

An Information-Centric Energy Infrastructure: The Berkeley View¹

Randy H. Katz², David A. Culler, Seth Sanders,
Sara Alspaugh, Yanpei Chen, Stephen Dawson-Haggerty, Prabal Dutta, Mike He,
Xiaofan Jiang, Laura Keys, Andrew Krioukov, Ken Lutz, Jorge Ortiz, Prashanth Mohan,
Evan Reutzler, Jay Taneja, Jeff Hsu, Sushant Shankar

Abstract: We describe an approach for how to design an essentially more scalable, flexible and resilient electric power infrastructure—one that encourages efficient use, integrates local generation, and manages demand through omnipresent awareness of energy availability and use over time. We are inspired by how the Internet has revolutionized communications infrastructure, by pushing intelligence to the edges while hiding the diversity of underlying technologies through well-defined interfaces. Any end device is a traffic source or sink and intelligent endpoints adapt their traffic to what the infrastructure can support. Our challenge is to understand how these *principles* can be suitably applied in formulating a new information-centric *energy network* for the 21st Century. We believe that an information-centric approach can achieve significant efficiencies in how electrical energy is distributed and used. The existing Grid assumes energy is cheap and information about its generation, distribution and use is expensive. Looking forward, energy will be dear, but pervasive information will allow us to use it more effectively, by agilely dispatching it to where it is needed, integrating intermittent renewable sources and intelligently adapting loads to match the available energy.

Key Words and Phrases: Smart Grid, Energy Networks, Supply-Following Loads, Slack and Slide

“The coming together of distributed communication technologies and distributed renewable energies via an open access, intelligent power grid, represents ‘power to the people.’”—Jeremy Rifkin [Rifkin]

1. Introduction and Motivation

1.1. The Energy Challenge

Today’s energy infrastructure is a marvel of the Industrial Age, yet it is showing its age. The Grid as it exists today is characterized by centralized generation via large plants, and a massive, centrally controlled transmission and distribution system. It delivers high quality power to all consumers simultaneously and is sized to service 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. Adding transmission capacity is expensive, and involves long construction delays. Local distribution, via step-down transformers, is also expensive in cost and efficiency, and is a single point of failure for an entire neighborhood. The system demands end-to-end synchronization, and lacks a pervasive mechanism for energy storage or buffering, thus complicating the integration of renewable generation sources, sharing among grids, or independent operation of subgrids during upstream outages.

Supplies are meticulously scheduled, or *dispatched*, to meet projected loads, which in turn are oblivious to the available supply. When demand exceeds supply, the only

¹ Research supported by NSF Grant #CPS-0932209, the FCRP MuSyC Center, Department of Energy/Lawrence Berkeley Laboratories, and industrial support from eBay, Fujitsu, Intel, Samsung, Siemens, and Vestas.

² Corresponding author: randy@cs.Berkeley.edu, 510-642-8778, 465 Soda Hall, CS Division, University of California, Berkeley, CA 94720-1776.

response is to shed load through brownouts and blackouts. Thus utilities are highly incentivized to mitigate energy consumption because of the expense of deploying new infrastructure to meet growing demands. The utility's strategy can best be described as one of *load-following supply*: the utility manages its portfolio of increasingly expensive generation sources—including if necessary entering a regional marketplace for energy—to scale its supply to meet its load (see Figure 1). The system is neither agile, making it difficult to exploit non-dispatchable renewable supplies, nor is it able to mitigate demand except through the blunt instrument of price signals, also known as *demand response*. Alternatively, some loads can be shifted to off-peak periods, through a process of *demand-side management*. Yet the average demand/consumer is a small fraction of the peak. For example, a 25 kWhr/day home draws on average less than 5% of its 100 amp service. Consumption correlations, like air conditioners on a hot day, drive demand beyond estimated aggregates. This can result in huge spikes in supply cost, and may trigger blackouts. Intelligence can be distributed to loads as well as the supply side. These are *supply-following loads*, that is, loads that adapt their behavior to the available supply.

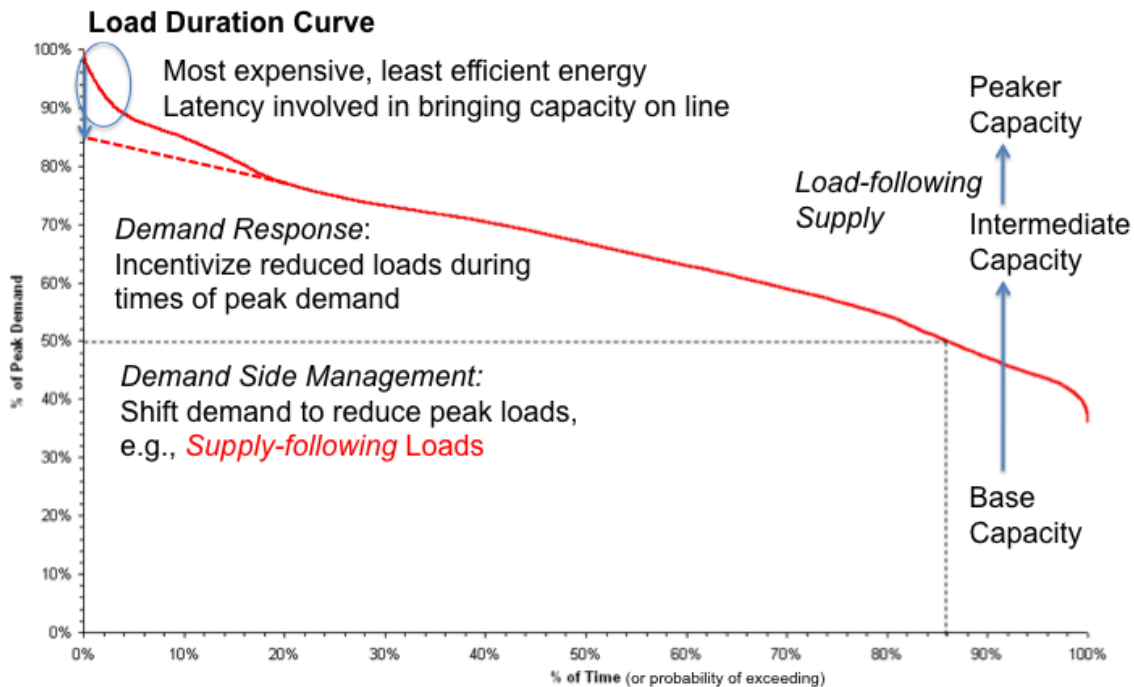


Figure 1. The Load Duration Curve

Information technology can improve the reliability, visibility, and controllability of the Grid. But making the Grid into a Smart Grid will require more than the intelligent metering currently being discussed in the context of demand response systems. The Internet suggests alternative organizing principles for a smarter Smart Grid. The Internet succeeded by pushing intelligence to the edges while hiding the diversity of underlying technologies through well-defined interfaces. Any device can be a source or sink of routable traffic and intelligent endpoints adapt their behavior to what the infrastructure can deliver in accordance with localized utility functions. While others have also made

these observations (e.g., [Gupta03, Keshav10]), we focus more deeply on the underlying information architecture of the Smart Grid, an intensive area of active research.

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 over time. The early network 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. This is similar to how the Internet, built in part on top of the telephony technology of ATM, has improved the phone network.

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 the current strategy of managing for the peak/worst case to focusing on the average case, with some headroom. This is analogous to statistical multiplexing in packet networks. The key is to use this headroom as an input for controlling generation, storage, and loads. Standardized intelligent “interfaces,” at the level of homes or even individual appliances, push intelligence to the edges to allow independent powered operation, distributed generation, and energy exchange. The architecture should allow aggregations of loads and supplies 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.

The “computing systems” analogy is valuable in conceptualizing the energy network. It is founded on the concepts of hierarchy and aggregation, layering, application programming interfaces (APIs) and protocols. Storage on the path breaks the synchronization between generation and loads, much as network buffers allow decoupling of senders and receivers. Critical service functions include resource allocation, load balancing, load shifting, and redirection. By analogy with the Internet’s smart edges and dumb pipes, an energy network can support smart generation and loads, but there are the challenges of reliability in this environment. Pervasive energy awareness is achieved through monitoring, modeling, and management.

In this paper, we review the current state-of-the-art in energy infrastructure and management, introduce some basic terminology, and identify opportunities for information technology oriented energy research. Our overarching goal is to provide an information system’s view of the next generation of energy infrastructure, a truly Smart Grid. The rest of this paper is organized as follows. Section 2 introduces the concept of *Energy Networks*. Section 3 describes recent trends in Energy Supply and Demand. Section 4 introduces Energy System Terminology. We present the concept of Deep Demand Response and Intelligent Loads in Section 5. We describe Energy Proportional Computing in Section 6. Our initial approach for an Architecture for Energy Networks is presented in Section 7. Our Summary and Conclusions are given in Section 8.

2. The Concept of Energy Networks

Pervasive information can fundamentally change how energy is produced, distributed and used. The crucial insight is to integrate information exchange everywhere that power is transferred. We call such an information-augmented system an *energy network*. The challenge is to exploit this information to construct a decentralized energy system that can match available supply to instantaneous demand on finer scales, be they geographical, logical aggregations, time grain, as well as all of these at once.

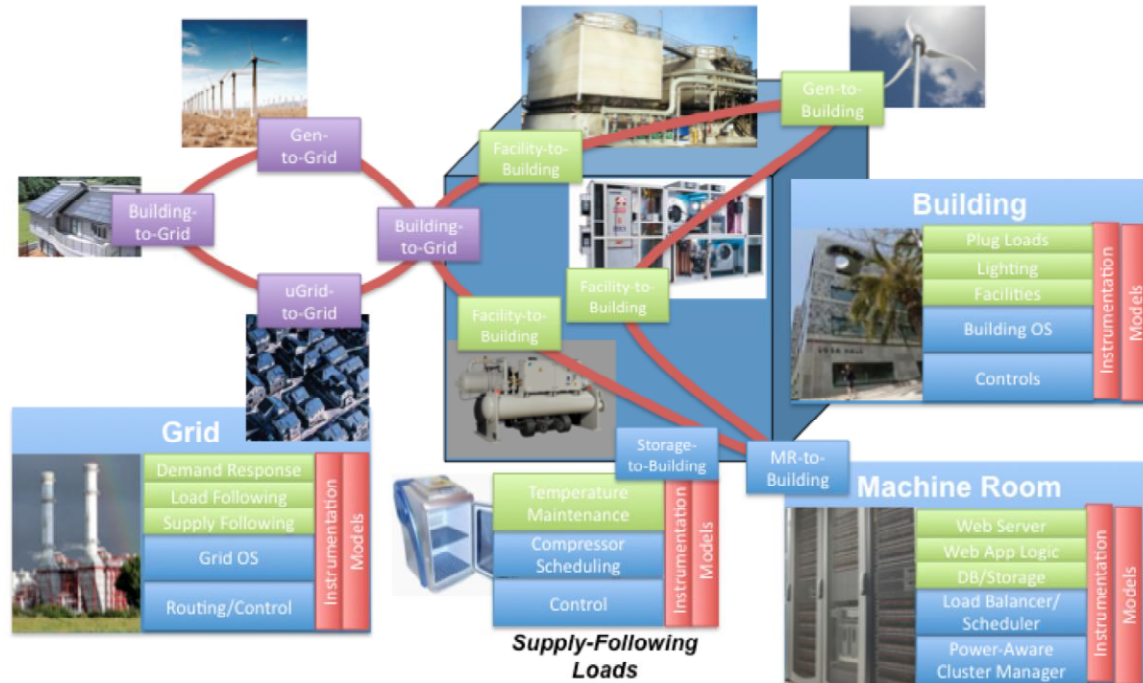


Figure 2. Energy Networks at Building and Grid Scales

An Information-centered Energy Network is an *information* overlay on the energy distribution system in its various *physical* manifestations, e.g., machine rooms, buildings, neighborhoods, isolated generation islands and regional grids. Figure 2 schematically presents the structure of such an energy network, with a particular focus on a building as an intelligent aggregation of loads, supplies, and even storage in its own right. It suggests the layers of software within loads like buildings and across the Grid necessary to implement its functionality. Such a network brings together pervasive information about energy availability and use, interactive load/supply negotiation protocols, and controllable loads and sources. It exists as a kind of control plane orthogonal to the physical network infrastructure over which power is actually moved. Such a network would enable an agile, distributed, and integrated energy management system that could buffer energy on the path to reduce peak-to-average energy consumption, moderate infrastructure provisioning, and encourage power-limited design and operation. Crucial to the network is the design and implementation of the components of the system that interface to the energy network.

3. Energy Supply and Demand Trends

There are several important observations about energy usage in the U.S. economy, drawn from publicly available data [DOE09, UBS07]:

- Only a tiny portion of the United States's current energy supply currently comes from renewable sources (1.1% of all energy supply, excluding hydro-electric and biomass). The largest single source of energy is petroleum (35% of all energy supply), which is primarily used to power vehicles. Other economies, albeit representing small populations, have seen a greater penetration of renewable energy sources, e.g., Denmark and Portugal.
- A large fraction of electricity generation is currently powered by coal (47% of total electricity generation), with smaller fractions supplied by nuclear (21%) and natural gas (19%). This is the largest use of coal in the U.S. economy (87% of total coal production), and the conversion of coal to electricity is a major contributor to green house gases, with associated implications for global climate impact.
- Losses are high in converting the energy intrinsic in coal and other sources to electricity delivered to loads. Up to 63% of the input energy is lost due to the physical process of "burning" sources to turn water into steam to turn turbine generators. A smaller contribution (2.6%) is due to transmission and distribution losses.
- Energy usage is more or less balanced among residential and commercial buildings (respectively 22% and 19% of all energy consumption), industrial plants and factories (30%), and vehicles (29%). Thus, buildings are about 40% of energy use, a kind of aggregated load that represents a near-term target for energy network deployment.
- Electricity generation currently consumes more energy than vehicles (respectively 41% and 29% of total energy consumption). It has been suggested that shifting to electric vehicles could significantly reduce the U.S. dependence on foreign oil. This is sometimes called "the million volt answer to oil." This would, of course, add significantly to the load and change the dynamics on the existing Grid, but also provides a potential source of distributed storage that could accelerate the deployment of a true Smart Grid.

4. Electrical Energy System Terminology

The traditional Electrical Energy System, also known as *the Grid*, consists of a three-tiered architecture of electrical energy generation, long-distance transmission, and regional distribution to local consumers of energy. We discuss each of these below:

Electrical Energy Generation: Electricity is generated from other forms of energy, typically by exploiting the principle of electromagnetic induction. This is accomplished by rotating a wire in a magnetic field to induce electricity through a device called a generator. The rotational forces come from many mechanical sources, such as water through a turbine, wind moving a blade, or the burning of a fuel source, such as nuclear, coal, or natural gas, to heat water into steam, which in turn is used to turn a turbine generator. The burning of fossil fuels, as well as the by-products of nuclear reactions, has implications for the environment, driving a strong desire to find alternative generation sources.




Figure 3. The Three-Tier Architecture of the Grid³

Transmission: This is the part of the system that delivers electricity at high voltages across long distances from remote generation sources for delivery to consumers. The transmission infrastructure can have bottlenecks: limited capacity for delivering energy to a particular region. This can result in large spikes in the cost of electricity when regional demand exceeds supply, whether due to generation or transmission capacity limits.

Distribution: This is the final stage of the delivery process, connecting the transmission system to the consumers and providing electricity at the appropriate voltages for the local loads.

Two emerging components of the energy system reflect the shift towards a more intelligent distributed architecture: *Distributed Generation* and *Energy Storage*.

³ Figure source at <http://peswiki.com/index.php/Image:Pathwaytopower.gif>.

Distributed Generation: Consisting of smaller generation sources than traditional generators, these are integrated closer to the loads. Examples include small-scale gas turbine generators, gas powered fuel cells, solar plants, and wind farms. They are intended to be used as an augmentation or enhancement to the traditional grid resources, providing more localized energy sources when the system as a whole becomes over subscribed or to bypass bottlenecks in the transmission system.

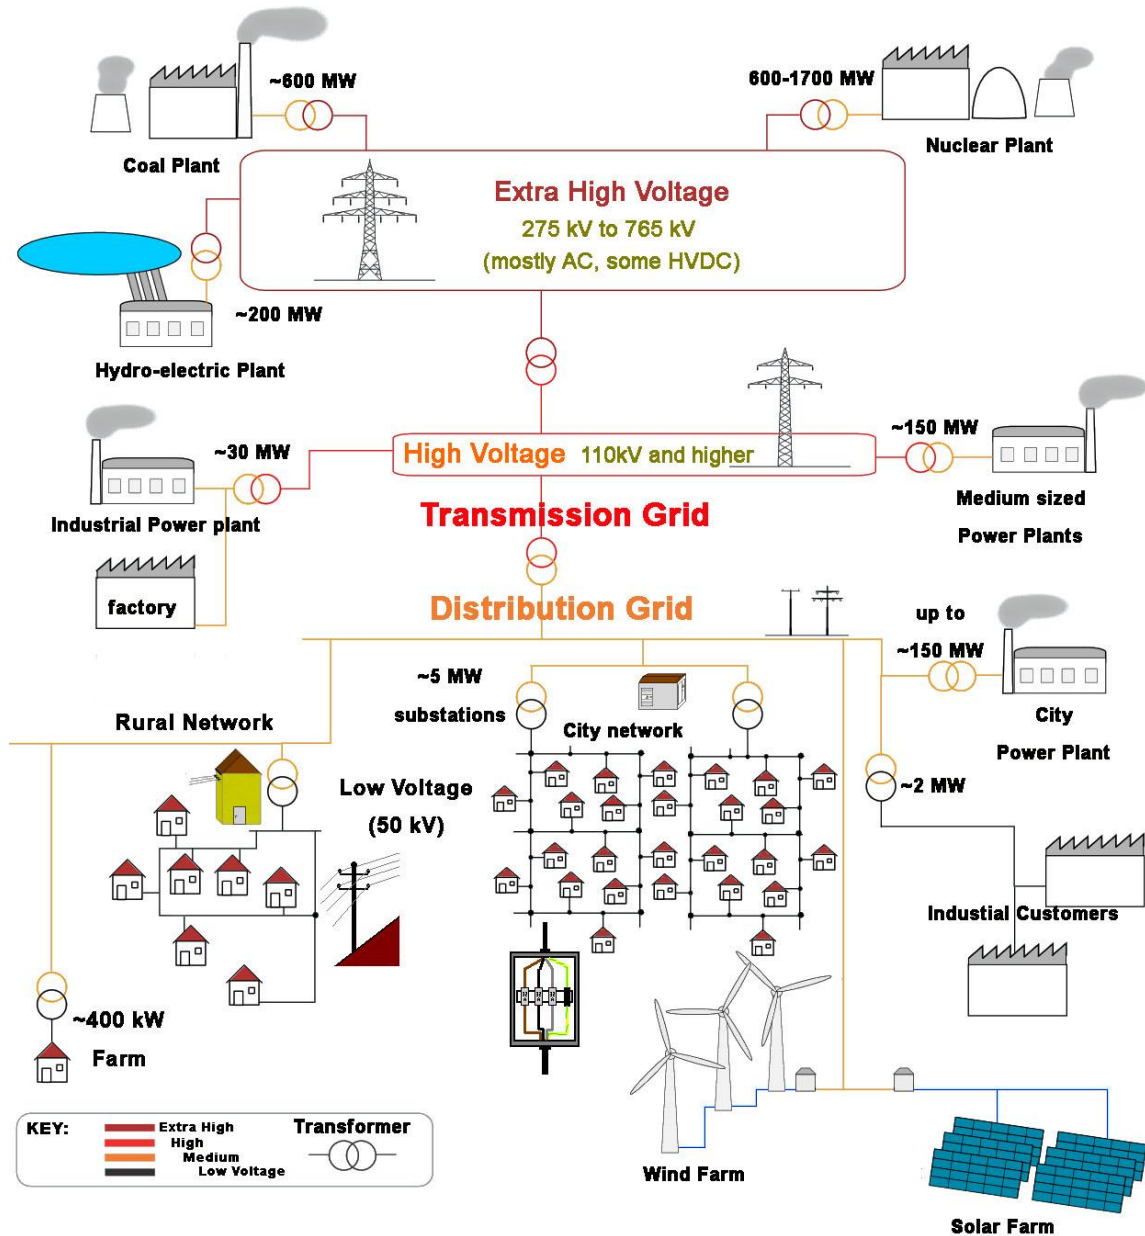


Figure 4. Detailed Grid Schematic⁴

⁴ Figure at http://upload.wikimedia.org/wikipedia/commons/5/56/Electricity_grid_schema-_lang-en.jpg

Energy Storage: This important component of the energy system allows energy production to be decoupled from its use in time. This is becoming increasingly important as difficult to schedule, renewable energy sources are becoming integrated into the Grid. These are known as *non-dispatchable sources*. Batteries exploit reversible chemical reactions to charge, hold, and supply energy as needed, yet tend to be expensive and have environmental impacts due to the kinds of materials used. Other approaches based on mechanical or thermal processes for energy storage are also possible. See Section 5.2 for how energy buffering can be implemented by slack and slide in intelligent loads.

Energy Markets: The energy system exists in a market place. If a utility has insufficient generation resources to meet its load demand, it can enter that marketplace to purchase additional supply at a market price. Similarly it can sell surpluses in its own capacity to other operators in its region. Some regions, like California, have deregulated their energy sector, limiting traditional utility companies to becoming energy distributors, and establishing an energy exchange that allows such distributors to purchase energy from a variety of energy suppliers [CalISO, Hedm09].




Figure 5. Smart Grid Schematic⁵

These trends in technology and markets, combined with a general societal desire to increase the efficiency of energy systems and to improve the environment by integrating more renewable sources in the generation mix, have yielded the Smart Grid concept.

Smart Grid: The Smart Grid exploits new technologies and architectures for managing the energy system, for improved efficiency and integration of alternative energy sources. A key enabler is pervasive monitoring and information communication between loads and generation sources, potentially to the level of individual consumer appliances. Monitoring and communication are essential for integrating non-dispatchable sources

⁵ Figure at <http://www.energy-daily.com/images/smart-grid-electricity-schematic-bg.jpg>.

with the Grid. The first step towards the Smart Grid involves the broad deployment of *Smart Meters*, with utilities communicating real-time prices to consumers. They, in turn, are expected to respond by reducing consumption when prices are high. The Smart Grid requires intelligent loads, beyond the meter on the building. *Demand Response* is an enhanced functionality for automated reduction of load demand in response to higher prices, by turning off non-critical loads or shifting their use in time. We describe such loads as *supply-following*, modifying their work function in response to energy supply signals. Enabling supply-following loads requires the Smart Grid to be as much focused on an information architecture based on communications and information exchanges between energy suppliers and consumers to better match supply and demand.

5. Deep Demand Response

5.1. Intelligent Loads Meet Intermittent Supply

On the demand side, intelligent loads, and on the supply side, the larger penetration of renewables, will be crucial to the Smart Grid's success. There is a special opportunity to combine both. The tipping point may well be the introduction of electric vehicles (EVs) on a large scale, which will place considerable demands on the Grid. EVs also represent a new revenue stream for the electricity industry as vehicles migrate away from petroleum, presumably sufficient to underwrite new infrastructure.

The average quick charging load of an EV is comparable to that of a typical residence. More intelligent charging, spreading demand over time or sequenced in cooperation with other loads in the neighborhood, allows better integration of this new load with the Grid by reducing its effect on peak demand. EVs also offer the possibility of stabilization, by acting as Grid storage to absorb intermittent energy generated by renewables while locally sourcing energy during times of high demand. However, in current technologies, battery lifetimes are affected by charge-discharge cycles, which mitigates their usability in this role.

A regulatory push to increase renewables in the source mix of is well underway in the U.S., with 24 states and the District of Columbia having policy in place to achieve a Renewables Portfolio Standard (RPS) of 10-45% of total energy consumed by 2015-2030. These states account for well over half the total U.S. electricity consumption, with more states setting voluntary goals. California's RPS calls for renewables to comprise 33% of its energy mix by 2020 [ARB, DoE10].

5.2. Slack

These trends suggest the need for a new kind of *deep demand response*: combining energy information with new kinds of control systems and sensor/actuator networks to enable more intelligent loads. This allows energy systems to match supply and demand dynamically, down to the appliance level, in real-time. The keys are predicting the output of renewable sources and controlling the consumption of loads in response to these. The challenge lies in meeting the work function “quality-of-service requirements” of the load,

while adapting to any variations in the source. To reason about load adaptability, we introduce *energy slack*: the amount of time an energy-consuming operation can be advanced or delayed while still performing its intended function. To model source variability, we consider the cumulative distribution function (CDF) of the change in output.

To illustrate deep demand response and slack, consider wind as the renewable source and a refrigerator as supply-following load. A wind source exhibits extremely high variability over time. As an example, consider a power profile that changes by less than 5 kW over a 5 minute window 95% of the time, yet can source up to 25 kW over a 60 minute window 95% of the time. If this kind of information were available to loads, an intelligent refrigerator could advance the start time of its cooling cycle to consume excess energy (pre-cooling), or it can increase a temperature set point slightly to reduce its energy consumption.

To make decisions in sculpting energy loads in this way, it is important to have metrics to express the ability to change the load. Fine-grained alignment of supply and consumption requires controllers that make continuous decisions based on information about load schedule as well as energy availability. For a load to be sculptable,

- Its consumption schedule involves choices on when to consume energy;
- There is some capacity to store energy in the system, which we call *slack*. Precooling is an example of a mechanism for increasing slack in a physical system;
- There is an ability to schedule its energy consumption, which we call *slide*. Running a washing machine at an off-peak time is a conventional example of slide. Less conventionally, slack can enable slide, for example, by deferring a cooling cycle in the refrigerator. Harnessing slide in physical systems like heaters, air conditions, and refrigerators is a rich arena for making loads more intelligent and adaptable.

Slack is the potential of an energy load to be advanced or deferred without affecting earlier or later operations or outcomes. In critical path analysis, the term refers to the scheduling flexibility in a task that keeps it off the critical path. In our context, we measure slack in units of energy, and apply it to the operation of physical systems, where the goal is not completion time, but rather reshaping the profile of how energy is consumed. More simplistically, energy slack closely follows a credit analogy. As energy is used by the system to do work—in the case of a refrigerator, as the compressor operates—slack is credited. Then, during the natural warming cycle, the credits are drained as slack is slowly reduced. Slack can remain in a system across cycles in the form of stored energy. In the case of a battery, this is the capability of the battery to introduce energy into the system later. For a refrigerator, this is the thermal energy stored in the contents of the refrigerator, whether it is food, liquids, or even air. Its consumption manifests itself in affecting the cooling schedule of the refrigerator's compressor.

6. Energy Proportional Computing

6.1. Introduction

The energy consumption of datacenters has received a great deal of attention, from massive web services like Google and Amazon located in proximity with massive electrical generation, to enterprise datacenters on urban grids, to building machine rooms. The EPA estimates that U.S. datacenters consumed 61 billion kilowatt-hours in 2006, at a cost of about \$450 million. This constitutes about 2% of US electricity consumption and has been the fastest growing consumption sector, representing an alarming growth trend [EPA, Mitchell-Jackson03].

An influential 2007 study examined the carbon footprint associated with information technology (IT) equipment extrapolated to the year 2020 [ClimateGroup08]. In 2007, IT equipment generated 1.43bn tons CO₂ equivalent, which represented 2% of the world's greenhouse gas footprint. This is comparable to the footprint of commercial aviation. This is expected to grow to 4% by 2020. The breakdown by equipment type extrapolates the category of PCs, peripherals and printers to yield 820m tons CO₂, or 57% of the total. Telecoms infrastructure and device yields a further 360m tons CO₂, or an additional 25%, while datacenters are expected to generate a further 260m tons CO₂, or 18% of the total. The compounded aggregate growth rate per sector from 2002 to 2020 is 5% in telecommunications, 7% for datacenters, and 5% for PCs and related equipment. Thus, datacenters represent the fastest growing contributor. In developing this estimate, numerous assumptions about the effectiveness of server consolidation and energy efficiency at the node and datacenter facility levels were assumed, so this should be considered as a lower bound to the actual contributions.

Much of the focus has been on the ratio of total datacenter consumption to that consumed by the computing equipment, called *Processor Utilization Efficiency* (PUE). Typical, large data centers operate at a PUE of 2-3 so the energy required to deliver power to and remove the heat from the servers is equal to that used by the servers. Considerable design efforts focused on improved heat exchange, airflow, and power distribution has reduced this overhead and demonstrated PUE of 1.2-1.4. PUE is simple to measure and track by metering the whole center and the PDU output, but it fails to recognize that real measure of effectiveness is not the power consumed by the servers, but the work accomplished.

Another key observation, made by engineers at Google, is that current computer systems are not power proportional [Barroso07]. At the system level, power consumed does not scale down with system utilization. In practical terms, it is surprisingly hard to achieve high levels of utilization of typical servers. *Energy Efficiency*, defined as:

$$\text{Energy Efficiency} = \text{System Utilization} / \text{Power}$$

should approach 0 as utilization approaches 0, but does not. Even a relatively energy-efficient server still consumes about half of its peak power when doing virtually no work.

Interestingly enough, many physical systems exhibit this poor behavior, including air conditioning and power distribution systems. These have been optimized for peak efficiency in case of rarely encountered peak use, yet operate proportionally inefficiently in the more common case of lower usage. Such a system does not exhibit the critical principle of *Doing Nothing Well*.

In fact, the utilization of the servers in datacenters is typically only about 25%. For example, in a study of over 5,000 of its servers, Google found that the average CPU utilization for most servers is between 10% and 50% of maximum utilization. This should not be surprising, since reasonable capacity planning must provide headroom to accommodate transient bursts in demand and potential future growth.

6.2. Architecture of Internet Datacenters

With enormous amounts of computing located under a single roof, the design of an Internet Datacenter must critically balancing the dense computing infrastructure with the needed facilities for power distribution and cooling. Within the datacenter, power passes through a distribution tree of Mains-UPS-PDU-Computer Racks-Computer Equipment. The HVAC System provides cool air to maintain the equipment at a proper operating temperature. A modest amount of the power budget goes to lights, conditioning office spaces, and so on, but the overarching loads are the computer equipment, the air conditioning system, and the power distribution infrastructure itself.

6.2.1. IT Equipment in the Datacenter

Energy consumption by information technology equipment is especially important because other elements of datacenter energy consumption scales with it, such as air-conditioning and power distribution facilities. While a PUE of 2 to 3 is considered good, Google has achieved PUEs as low as 1.2 in some of its datacenters [Google10].

Consider a datacenter with a PUE of 3.3, one in which for every dollar spent on power for its IT equipment, another \$2.3 is spent on the power and cooling infrastructure. In such a datacenter, typical numbers for energy consumption are as follows. For the physical infrastructure of cooling and power equipment, the 70% of the datacenter energy footprint breaks down as: the Computer Room Air Conditioner (CRAC) consumes 9%, the humidifier 3%, the chiller 33%, the uninterruptable power supply (UPS) 18%, the Power Distribution Unit 5%, the lighting 1%, and the transformers and related power switchgear 1%. For the 30% that is the IT load, 33% can be attributed to the CPU, 30% to DRAM, 10% to disks, 5% to networking equipment, and 22% to other backend servers and appliances [Hamilton10].

Some simple rules can be followed to minimize cooling costs. First, raise the datacenter temperatures. Highly redundant architectures combined with practical experience suggest that traditional machine rooms have been overcooled. Second, maintain tight control of the flow of cool and hot air via short path while avoiding mixing with short paths. Third is to exploit *airside economization*, using cool outside air to augment and thus reduce the

need for conventional air conditioning. Waterside economization turns off the air-conditioning when it is not needed. The best current designs have water cooling placed close to the load, but avoid direct water cooling.

6.1.2. Datacenter Power Distribution

The prime figure of datacenter sizing is its power rating: the maximum load it presents to the utility. From this we can derive the number of computers and other information technology equipment it can support.

First, let us consider the conversion losses in a typical datacenter. Moving power from high voltage transmission lines (115kV) to the local distribution infrastructure (13.2kV) typically encounters a 0.3% loss. Power then passes through the datacenter's *Uninterruptable Power Supply* (UPS), still at 13.2kV, where it sustains a further 6% loss. The purpose of the UPS is to protect the datacenter from disruptions in the power supplied from the Utility. It is usually implemented as large battery banks. When combined with a local generation source like a diesel generator, this makes it possible for the datacenter operate independently of the Grid for a period of time. *Power Distribution Units* (PDU), operating at 480V, feed power throughout the datacenter, incurring a further 0.3% loss. At the next stage, power passes through the *Power Panel* at 208V, with a further 0.3% loss. The final step takes the power through the *Power Switch Gear and Conductors* to deliver it finally to the computers in their racks, losing an additional 1%. Thus the total power delivered after these losses is approximately 92%: a 1 MW datacenters can provide only 920kW of power to its computers and other facilities.

Some general rules are useful for minimize power distribution losses. First, avoid conversions, such as making use of fewer voltage transformation steps as well as a more efficient UPS. Google gets rid of the need for a UPS by distributing batteries to individual server nodes, dramatically reducing the inefficiencies of conventional large-scale battery-based UPS systems. Another possibility is to increase the efficiency of conversions, particularly at the level of the server node. This can be achieved By better matching voltage regulators to their loads and using more efficient parts, more efficient power conversion can be achieved. Commodity server designs rarely make these choices. Other efficiency improvements can be achieved by distributing high voltages as close to the load as possible, as well as DC distribution.

A critical observation is the principle of distributing power and cooling infrastructure via smaller, better utilized components located close to where they are used.

6.1.3. Datacenter IT Equipment

The typical building blocks of a large-scale datacenter are server racks with 10-80 nodes, with 20-60 such racks per PDU. For a given power sizing, the goal is to deploy as many machines as possible. In this endeavor, it is important to distinguish between nameplate and actual peak power. It is rare to be able to fully utilize all components of the compute node simultaneously. Google engineers have reported that a typical compute node with a

nameplate power of 213 W was measured under a heavy workload as consuming only 145 W, about two-thirds of its rated peak. Furthermore, in practical terms, it is difficult to drive such systems to anything approaching even these levels of utilization. These observations are essential for proper sizing of the number of nodes to the deployed power infrastructure.

In general, it is worthwhile consolidating computation on fewer machines if possible, while placing the remainder in low power sleep states. It is better to have one computer at 50% utilization than five computers at 10% utilization. In the design of the compute node, issues of system balance must be carefully considered. It does not make sense to deploy too high performance a CPU if the other components of the system, like memory or disk, represent the performance bottleneck.

6.2. Datacenter in a Box

A new form of system packaging has emerged, consisting of containerized datacenters. These systems represent a new kind of mechanical-electrical co-design. The idea is to just add power, chilled water, and network connectivity to create a datacenter building block. Manufacturers include Rackable Systems [Rackable] and Sun Microsystems [Sun]. Systems like these have been used by both Google and Microsoft in their datacenter deployments.

Many considerations have led to modular datacenters. The first is the faster pace of infrastructure innovation, now typically on a three-year cycle. Modular design allows a container to be swapped out and replaced with an updated model. The second is efficient scale down, as compute density and network connectivity increase. The third is service-free fail in place. 20-50% of system outages are caused by administrator errors, and it is better to allow components to fail and remain in place. Once a critical number have failed, the whole container can be replaced. The final consideration is the way in which containers allow the datacenter to grow incrementally.

Some parameters of Sun's pioneering Black Box container are as follows. It provides power and cooling for 200 KW of racked hardware. It provides external taps for electricity, network, and water. A 20 foot x 20 foot x 20 foot container contains 7.5 racks of approximately 250 Servers, 7 TB DRAM, and 1.5 PB disk [Sun].

6.3. Case Study: Energy-Proportional Cluster Management

For web service workloads, the amount of work is primarily determined by the rate of user requests, which can vary drastically over time. Web service operators must at the very least provision for the observed peak, taking "... the busiest minute of the busiest hour of the busiest day and build[ing] capacity on that." However, most clusters provision for far more than the observed peak in order to provide a safety margin against flash crowds. A typical rule of thumb is to provision for twice the maximum average load over a moderate window experienced in the last planning cycle. In a one-week request trace obtained from Wikipedia, we observed the peak demand over a minute that was 1.6

times the average rate, though peaks as large as 19 times the average have been observed. A natural consequence of the gap between peak and average requests rates in workloads, amplified by overprovisioning, is that much of the time many servers sit at low levels of utilization. In an energy efficient data center, not only should the PUE be small, the power consumption should be proportional to the rate at which work is performed.

Unfortunately, modern server platforms are very far from power proportional despite substantial improvements in power efficiency of the microprocessor, including Dynamic Voltage/Frequency Scaling (DVFS) and the introduction of a family of sophisticated power states. Even for specially engineered platforms, the power consumed when completely idle is over 50% of that when fully active, and idle consumption often over 80% of peak for commodity products. Thus, the service power consumption is essentially proportional to provisioned capacity, not the request rate. To make the service power proportional, we must harness the idleness for energy savings.

In this subsection, we consider the design and implementation of power proportional clustered services constructed from non-power proportional systems [Alspaugh10]. The basic approach is to put idle servers to sleep and wake them up when they are needed. Thus, capital equipment spending tracks peak demand, but energy consumption tracks delivered service. However, realizing this simple goal presents several challenges. The active work must be coalesced into a subset of nodes, so that a significant amount of coarse-grained, i.e., entire node idleness, is obtained. But enough capacity must be available to absorb bursts in demand. Servers exhibit a dramatic increase in delay or loss rate as the requested load approaches the maximum delivered load, which traditional services avoid by having all the resources running all the time. We dynamically provision active nodes to track demand, while providing headroom for bursts. Interestingly, the solution is very similar to approach used by electric utilities to match power generation to time varying demand without any explicit protocol between consumers and providers. A set of baseline resources is supplemented by an array of intermittent resources. Intermittent resources are brought on-line or off-line in accordance with estimated demand in the near future. Each of those resources maintain a certain amount of “spinning reserve” to absorb transient bursts and to cover the ramp-up delay in spinning up additional resources. Each seeks to operate at sufficiently high utilization to be efficient, while leaving enough reserve capacity to be reliable.

Our basic result is summarized in Figure 6, which shows the second-by-second request rate of a Wikipedia trace over a week relative to the energy-consuming resources provisioned to service that load. The diurnal pattern is apparent, as is the stochastic fluctuation in demand. Traditional capacity planning would provision for approximately 10,000 requests per second. The power-aware cluster manager presented here, which we call *NapSAC*, performs Server Actuation and Control on a heterogeneous set of machines. It dynamically provisions enough machines to serve the requested load with just enough headroom to accommodate the bursts and allow for re-provisioning to increased demand. The jagged pattern shown in the graph reflects the capacity of the individual machines engaged to serve the load. *NapSAC* energy consumption is proportional to the work performed at essentially a joule per request, despite large variation in the request rate,

while static provisioning is only efficient near the peak load. Furthermore, the fine-grain overprovisioning required to track the load without incurring a significant congestion penalty is small compared to the worst-case provisioning used in capacity planning.




Figure 6. The Efficiency of Dynamic Provisioning

A key observation is that most clustered systems employ a master-worker pattern where the master is responsible for load management. It is natural to extend the master to manage the power state of the workers and to assign work in a manner that is consistent with its power provisioning.

Leading server products do not actually support going into standby and resuming, other than by actuating the AC power line, so we examine the use of desktop, mobile, and even embedded platforms for use in service applications. The Wake-on-LAN and low-power sleep states available on these platforms are sufficient mechanisms for this design. Under the 2x provision rule, we are able to achieve 90% of the savings that an optimal scheme

would achieve. With our design we are able to reduce energy consumption while maintaining acceptable response times for a web service workload based on Wikipedia.

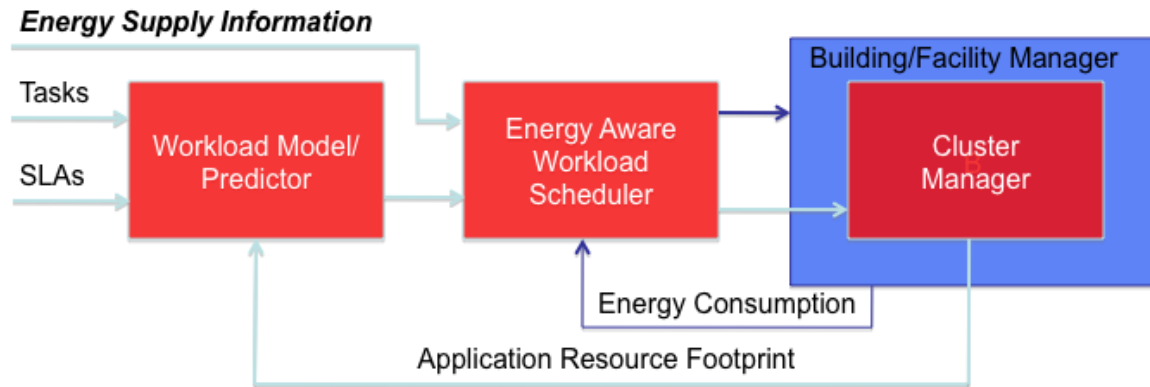


Figure 7. Energy Aware Cluster Manager

Figure 7 shows a block diagram for an energy-aware task scheduler. The cluster manager moves processors from low power sleep to active states and back again, while monitoring processor resource utilization and energy consumption. A predictor tracks the arrival of work over various time scales to predict forward in time the workload. Energy supply information, perhaps represented as prices, influences how and when work is scheduled on the cluster. For example, when energy is cheap, accelerate the execution of batch tasks; when it is expensive, defer such tasks to a later time when energy is less expensive.

Figure 8 shows how such an energy aware scheduler responds to a step impulse of a workload. A predictor tracks the rise and fall of activity to predict what the demand will be forward in time. The controller then brings processors out of their low-power sleep states to service this workload. When the workload exceeds the capacity of the active processors, the shortfall will affect performance, such as longer service time for requests. When capacity exceeds the workload, some additional energy will be consumed, but if the provisioning closely follows the workload, the surplus can be minimized.

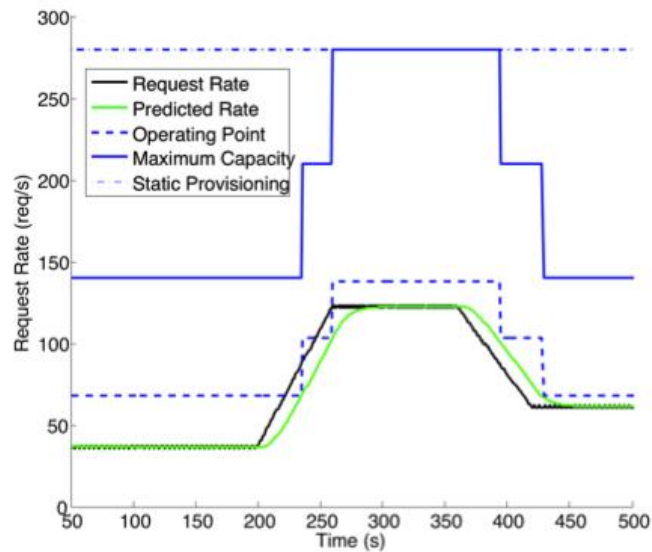


Figure 8. Close Provisioning Matching of the Workload to Resources

7. Towards A Layered Architecture Model for Energy Networks

7.1. Introduction

A key element of an information-centered architecture for energy systems must incorporate the logical separation of energy and information flows. This is analogous to the separation of the data and control planes in conventional networks.

We are at too early a stage to be able to fully specify how such an architecture should be constructed. Nevertheless, we offer the sketch of a solution. Energy storage and buffering must be first class components of the energy architecture, to make it possible as a first class component, to break the synchronization between sources and loads. An important capability is the ability to forecast energy availability and usage, and to use this as the basis for negotiation between energy suppliers and consumers. Exchanging such information, in appropriate levels of aggregation, will allow adaptation and cooperation on both the supply and load sides. With the appropriately specified interfaces, incremental deployment should be possible. And an information-centered approach should enhance reliability, by enabling better adaptation and agility among sources and consumers.

To be able to manage the energy systems, pervasive monitoring and modeling of loads and supplies are necessary. Smart aggregation and disaggregation of information protects privacy of individuals while also reducing the sheer volume of information that needs to be exchanged. Sources and loads can be aggregated into virtual associations. In effect, (bidirectional) bits follow power flows, thus making it possible to exploit information to match sources to loads, manage buffers, integrate renewables, signal demand response and take advantage of locality.

Traditional energy systems collect and act upon information centrally, usually under the control of a Utility company. An alternative view is possible, based on distributed information and control. Supplies are intelligent, and can communicate their forward-looking profile of energy availability. Loads are also intelligent, shaping their work function to the energy availability signals they receive from suppliers.

Within such an architecture, energy storage plays a role as a load as well as a supply. Control algorithms must determine, given availability or pricing of energy from the Grid as well as from distributed generation sources, whether to charge storage for future use or discharge it for current use. Beyond conventional storage, such as batteries, energy can be buffered by essentially scheduling its work. By consuming energy now, we advance the work function of the load, introducing slack that allows energy consumption to be deferred later. In effect, we are buffering energy.

Controlling this system admits of many optimization approaches. In addition to traditional energy system metrics of peak reduction, minimum energy consumption, and responsiveness to changing energy conditions, we are interested in stability and convergence properties of the system. The best distributed control system remains an area of active research investigation.

7.2. The Building Operating System

Multiple studies have indicated that commercial building energy consumption can be reduced 10-50% should better building energy management systems be available. Although some new commercial systems have started to break out of the "silo" approach to building management software architecture, the predominant paradigm is still monolithic software managing all facets of a building's energy spend. A better approach is a flexible, open, service-based architecture to solve the problems inherent in the older system and prototype all of the core functionality. Our evaluation proceeds on real metrics like problems identified, portability between buildings, and ultimately, energy savings.

A building operating system provides a context and runtime for other software that runs "on the building." We envision numerous vendors competing on delivering the best solution for a particular building management task, but running alongside software from other providers. Pervading all of this are core pieces of functionality, providing five key backplane service abstractions: *sensor and actuator access*, *access management*, *metadata*, *archiving*, and *discovery*. Executing on this platform are building-scale applications: *batch and real-time analytics*, *supervisory control loops*, and *individualized energy feedback*. These building services are described below.

Sensor and Actuator Access: Imposing a uniform data model and access method on the diverse instrumentation found in a typical commercial building is a critical enabling piece of technology. We define a uniform data model for sensors and actuators, in which devices expose multiple sense points and channels. Using this decomposition, we define a simple set of objects and properties required for interpreting nearly any sensor; the result

is the *Simple Measurement and Actuation profile*, or sMAP [DawsonHagerty10]. It is possible to expose a huge diversity of underlying sensors as time-series.

Archiving: Data is not commonly used in real-time but rather it is stored for later analysis. We have experimented with several different archival backends for storage: a custom file-based engine with a simple query language on top, SQL-based stores, and NoSQL document stores. In general, we have found these solutions to be acceptable for the current amount of data generated (3000 points at 20 second resolution), although doubts exist as to their ability to scale to larger numbers. Additionally, it seems that it will be necessary to degrade historical data at some point instead of storing it “forever.” Alternatively, horizontal scaling can be employed to scale as the amount of data grows. The take-away, however, has been is that existing data management solutions are acceptable for this form of time-series data when properly employed and with the addition of a simple RESTful interface can be easily integrated into our architecture.

Metadata: A key piece of infrastructure is the metadata server. Potential examples of metadata include building models, locations and types of sensors and actuators, and the logical entity relationships between these devices. This type of information is key to automatically drive analyses of the building data. Two prototypes of this system exist. The first is a simple application built on a heavily-de-normalized relational schema. It was designed to capture the electrical distribution design of Cory Hall, the home of Electrical Engineering at Berkeley, which necessarily forms a tree. Significant metadata concerning the sensors in the system has been entered into this system, to the point at which we can consider re-normalizing the schema and thus deriving a data model rich enough to capture the diversity of the real building systems. A second prototype system, StreamFS (formerly IS4 [Ortiz10]), is much further along at developing RESTful machine-to-machine interfaces for exposing the building metadata to other consumers. It treats metadata as a set of file-like structures, where the file paths refer to a physical resources and the associated metadata is kept as a JSON document associated with that resource. StreamFS also provides versioning of the file system structure in order to keep track of changes in the building or deployment over time; an important innovation so that the appropriate metadata for any sensor reading can be reconstructed. The goal is to encode enough information into a uniform data model so that metrics such as kW/ft³ and power breakdowns between different device classes can be easily computed.

Access Management: A crosscutting design requirement is the need for security. Depending on the part of the system, there may be a need for authentication, integrity checking, or confidentiality. For instance, certain data is public but consumers need to check that it has not been altered since production in transit; they must be able to verify its integrity. Other sensors may allow a malicious user to discover when a person is in their office and thus invade their privacy. Thus access to data should be contingent on authentication and confidentiality guarantees. Multiple systems exist which have solved these problems in other areas. Kerberos and PKI are the most relevant in this context. One key concept is that of *principals*: identities that receive capabilities. In a building operating system, both computer users and pieces of software (agents) are principals. Another is *role*, as defined by role-based access control. In this paradigm, the capabilities

granted to a principal are determined by the role they play. A principal can play multiple roles. We propose to define a set of roles for a building operating system, using well-known cryptographic primitives for enforcement. By defining a new HTTP authentication mechanism that understands and enforces these access restrictions, we construct a modern version of "kerberized" web service protocols that carries the traffic for the building operating system.

Discovery: The final key component tying together all others is discovery. At the minimum, a client must be able to enumerate the access management server, although the ability to enumerate other services may itself require a particular role. There are many ways to implement this. Predominant approaches are DNS based, either using Zero configuration-like multicast DNS, or a centralized dynamic DNS system. Services receive names that resolve to a number of records containing details on how to contact them. Although DNS names are themselves hierarchical, we do not use the DNS hierarchy to capture a significant amount of metadata concerning the services. Although it may be convenient to do so in some cases, the full metadata is too complex to be captured by a simple hierarchy. We use either dynamic DNS, where services coming online identify themselves to a DNS server and update their records with appropriate information, or a simplified version of a custom web-services protocol to accomplish essentially the same thing; the decision be made based on the difficulty integrating DNS access control mechanism (TKIP) with our new authentication architecture.

Significant challenge remains in extending such a system to accommodate multiple building or entities within the same system. It is entirely reasonable to presume that a single campus or office park would run only a single "instance" of the same operating system. The boundaries between systems are organizational rather than physical. Thus, key concepts from the DNS architecture can play a role such as delegation and replication.

7.3. Underlying Layers for Monitoring, Modeling, and Management

Building-Scale Web Services Architecture

As traditional building instrumentation becomes more networked, more diverse, and more pervasive, and as physical information is widely used in a broad range of applications, networked sensors are evolving into tiny embedded information servers. Web Services can enable the integration of diverse sources of physical information, as do diverse conventional sources. Uniform, machine-independent access to self-describing physical information obviates the innumerable drivers and adapters for specific sensor and actuator devices found in industrial building automation, process control, and environmental monitoring solutions. Physical information is in more challenging than conventional data because its interpretation is dependent on the behavior and context of the particular sensor or actuator. Also, the diversity of sensors is large. In many cases, the representation, transportation, and storage of this data must be efficient. While it is easy to wrap readings in XML and transport them over HTTP, it is challenging to get widespread agreement on a simple, easily understood solution.

An approach based on a simple representation of measurement information and actuation events, founded on modern RESTful web service techniques, allows for arbitrary architectural composition of data sources, freeing application designers from tight frameworks and enabling widespread exploration of the sensor-application arena. We focus on defining a particular web service for exchanging physical data. Success means becoming a widely accepted interchange format for sensor data on the Internet.

The design space for a web service for physical information consists of three areas:

- Metrology the study of measurement; what is necessary to represent a datum.
- Syndication concerns how a data is propagated out from the sensor into a larger system.
- Scalability relates to the range of devices and uses the service can support, from small, embedded systems to huge Internet data centers.

Each of these concerns presents a set of design decisions. We focus on representing and transmitting physical information.

A prototypical interaction between a client and an instrument exposing a sMAP interface is as follows. The client contacts the device over HTTP and discovers what resources, and thus what data is provided by the instrument. The client first determines what the units of a particular channel and measurement point are, before taking a sample by fetching the reading resource. Finally, the client arranges for period reports to be delivered to him. Since all communication references a schema, the client can also fetch the schema to validate the data.

sMAP forms the basis for a build to electric grid testbed which provides fine-grained, multi-resolution sensor data about the building under inspection. The project currently provides around 2000 distinct measurement channels. These channels monitor electricity consumption, environmental quality data, HVAC parameters, weather data, and more. An early problem we confronted was how to integrate all of these sensors to allow “building applications” to be possible; the answer was sMAP. sMAP is designed to unify access to legacy instruments with a consistent representation and up-to-date technological underpinnings. While many recognize the need for such integration, much of the work is present in proprietary products and seeks to interoperate only within the particular industrial sector. By presenting a simple but carefully considered specification along with a substantial amount of data, we can bootstrap the process of making all physical information universally accessible.




Figure 9: Monitoring Layer

Figure 9 shows the lowest, physical level of the building operating system, consisting of the pervasive sensors [Jiang09] that monitor the consumption of energy and the environmental conditions of the physical space being managed. These provide critical inputs about the current state of the physical world, to drive the development of models that in turn form components of the necessary building-scale control systems.

Figure 10 takes the perspective up to the networking level, in which the sensors form a self-organizing network from which their measurements can be collected and aggregated for use by higher layers of processing [Dutta05, Dutta05]. Services are needed here for hiding the heterogeneity of the underlying sensors and the data they collect. Collection and processing is not limited to the physical space of the building. The collections can be shared across the Internet, whether it is to aggregate the collection for campus-scale purposes, or to create shared repositories from which more flexible models can be derived.




Figure 10: Networking Layer

Figure 11 offers the highest layer, which we have described as the building operating system. This level performs the critical resource allocation tasks that balance consumption, efficiency, and cost of energy to achieve the desired qualities of the

environment within the building. This where monitoring data and extracted models come together to provide the foundation for control systems that can actually shape the behavior of the major facilities within the building and their energy consumption. With the necessary open interfaces defined, such an operating system view would enable whole new kinds of applications, user interfaces, and services to be developed for energy-efficient buildings of the future.




Figure 11: Modeling Layer

8. Summary and Conclusions

In this paper, we have discussed energy consumption in information technology equipment and how it serves as a first model system of a supply-following sculptable load. The key idea here is that of *energy proportional computing*: consume energy in direct proportion to the amount of work being performed. A related observation is that most systems are designed to be energy efficient at high levels of utilization, yet few systems actually operate at these levels. This suggests a new challenge: design systems to *doing nothing well*. Internet Datacenters are a particularly interesting design point, with tens of thousands of computers under one roof, a computing facility co-designed with a physical building to achieve low power utilization efficiency metrics to minimize the overhead of delivering power and cooling to the processor complex. Containerized datacenters further suggest new packaging opportunities for the better optimization of computing, communications, power, and mechanical subsystems.

Our group at Berkeley is designing *LoCal*, a prototype scalable energy network [He08, LoCal]. Considering the inherent inefficiencies at all levels of electrical energy distribution and usage, *LoCal* investigates how information exchange can better support an energy system with close integration of energy generation, storage, and consumption. We are developing energy information interfaces and exchange protocols at a variety of scales including the machine room, building facility, energy storage, building, and grid generation levels. We feel it is particularly important to explore intelligent loads beyond datacenters to include the whole building and aggregations of buildings. The concept of *deep demand response* and *supply-following loads* focus on how energy-aware loads can modify their work functions in response to information about energy availability and costs.

LoCal is based on a layered architecture founded *sensing* to obtain information about energy consumption, *actuation* to affect energy consumption in component facilities, *decision making* based on inputs from sensors and generating outputs to actuators, distributed among multiple sensing/actuating/decision making entities, all of whom communicating through well defined APIs. LoCal makes it possible to turn non-dispatchable sources into sources whose energy can be matched to agile, supply-following loads. This is an alternative for integrating distributed energy resources, by developing loads that can closely follow varying supplies.

These systems demonstrate the ways in which information can be exploited to yield a more efficient energy system. While new energy sources and storage technologies are of crucial importance, this paper has shown the essential value of an underlying information architecture for the Smart Grid. Information exchange and better awareness of energy availability and load profiles can be used to more accurately match supplies to loads, integrating renewable sources, achieving higher levels of overall energy efficiency, and avoiding the need to grossly overprovision the energy system.

References

[Alspaugh10] S. Alspaugh, L. Keys, A. Krioukov, P. Mohan, D. Culler, R. H. Katz, "NapSAC: Design and Implementation of a Power-Proportional Web Server," First ACM SIGCOMM Workshop on Green Networking, New Delhi, India, (August 2010).

[ARB] California Air Resources Board, <http://www.arb.ca.gov/cc/cc.htm>.

[Barroso07] L. A. Barroso, U. Holze, "The Case for Energy-Proportional Computing," *IEEE Computer*, V 40, N 12, (December 2007), pp. 33-37.

[CalISO] California Independent System Operator, <http://www.caiso.com/>.

[ClimateGroup09] The Climate Group. SMART 2020: Enable the Low Carbon Economy in the Information Age. See www.smart2020.org/.

[DawsonHaggerty10] S. Dawson-Haggerty, X. Jiang, G. Tolle, J. Ortiz, D. Culler, "sMAP—a Simple Measurement and Actuation Profile for Physical Information," ACM Sensys 10, Zurich, Switzerland, (November 2010).

[DOE09] U.S. Energy Information Administration, "Annual Energy Review," U.S. Department of Energy, 2009. See <http://www.eia.doe.gov>.

[DOE10] U.S. Energy Efficiency and Renewable Energy, "State Activities and Partnerships," U.S. Department of Energy, 2010.

[Dutta05] Prabal Dutta, David E. Culler: *System software techniques for low-power operation in wireless sensor networks*. [ICCAD 2005](#): 925-932.

[Dutta08] Prabal Dutta, David Culler, "Practical Asynchronous Neighbor Discovery and Rendezvous for Mobile Sensing Applications," ACM SenSys 08, Raleigh, NC, (November 2008).

[EPA07] Environmental Protection Agency, 2007.

[Fan07] X. Fan, W-D. Weber, L. A. Barroso, "Power Provisioning for a Warehouse-sized Computer," ISCA'07, San Diego, CA, (August 2007).

[Google] <http://www.google.org/powermeter/>.

[Google10] "Data Center Efficiency Measurements," Google Inc. 2010.
<http://www.google.com/corporate/green/datacenters/measuring.html>.

[Gupta03] M. Gupta, S. Singh, "The Greening of the Internet," ACM SIGCOMM 2003, Karlsruhe, Germany, (August 2003).

[Hamilton10] James Hamilton, <http://perspectives.mvdirona.com/>.

[He08] M. He, E. Reutzel, D. Culler, R. H. Katz, S. Sanders, "An Architecture for Local Energy Generation, Distribution, and Sharing," Proc. IEEE Energy 2030 Conference, Atlanta, GA, (November 2008).

[Hedm09] K. Hedman, R. O'Neill, E. Bartholomew Fisher, S. S. Oren, "Optimal Transmission Switching with Contingency Analysis", *IEEE Transactions on Power Systems*, V 24, N 3. (May 2009), pp. 1577-1586.

[Mitchell-Jackson03] J. Mitchell-Jackson, J. G. Koomey, B. Nordman, M. Blazek, "Data Center Power Requirements: Measurements from Silicon Valley," *Energy (Energy)* ISSN 0360-5442, 28(4):837 - 850, 2003.

[Jiang09] Design and Implementation of a High-Fidelity AC Metering Network, *Xiaofan Jiang, Stephen Dawson-Haggerty, Prabal Dutta, and David Culler*, 8th international conference on Information processing in sensor networks, (IPSN SPOTS 2009) April 2009 (to appear).

[Keshav10] S. Keshav, C. Rosenberg, "How Internet Concepts and Technologies Can Help Green and Smarten the Electrical Grid," ACM Sigcomm Workshop on Green Networking, New Delhi, India, (August 2010).

[LoCal] <http://local.cs.Berkeley.edu>.

[Ortiz10] Jorge Ortiz and David Culler. "A System for Managing Physical Data in Buildings." Technical Report No. EECS-2010-128, EECS Department, University of California, Berkeley 2010.

[Rackable] See http://www.sgi.com/products/data_center/.

[Rifkin] Jeremy Rifkin, "The Third Industrial Revolution: Leading the Way to a Green Energy Era and a Hydrogen Economy," <http://www.foet.org/lectures/lecture-hydrogen-economy.html>.

[Sun] See <http://www.sun.com/service/sunmd/>.

[UBS07] UBS Research Focus, "Climate Change: Beyond Whether," (January 2007).