

# How Internet Concepts and Technologies Can Help Green and Smarten the Electrical Grid

Srinivasan Keshav  
David R. Cheriton School of Computer Science  
University of Waterloo  
Waterloo, Ontario N2L 3G1 Canada  
keshav@uwaterloo.ca

Catherine Rosenberg  
Dept. of Electrical and Computer Engineering  
University of Waterloo  
Waterloo, Ontario N2L 3G1 Canada  
cath@ecemail.uwaterloo.ca

## ABSTRACT

Several powerful forces are gathering to make fundamental and irrevocable changes to the century-old grid. The next-generation grid, often called the ‘smart grid,’ will feature distributed energy production, vastly more storage, tens of millions of stochastic renewable-energy sources, and the use of communication technologies both to allow precise matching of supply to demand and to incentivize appropriate consumer behaviour. These changes will have the effect of reducing energy waste and reducing the carbon footprint of the grid, making it ‘smarter’ and ‘greener.’ In this position paper, we discuss how the concepts and techniques pioneered by the Internet, the fruit of four decades of research in this area, are directly applicable to the design of a smart, green grid. This is because both the Internet and the electrical grid are designed to meet fundamental needs, for information and for energy, respectively, by connecting geographically dispersed suppliers with geographically dispersed consumers. Keeping this and other similarities (and fundamental differences, as well) in mind, we propose several specific areas where Internet concepts and technologies can contribute to the development of a smart, green grid. We also describe some areas where the Internet operations can be improved based on the experience gained in the electrical grid. We hope that our work will initiate a dialogue between the Internet and the smart grid communities.

## Categories and Subject Descriptors

J.7 [Computer Applications]: Computers in other systems, Industrial control

## General Terms

Design

## Keywords

Electrical grid, Green Networking

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright (c) 2011 ACM This is a minor revision of the work published in the ACM SIGCOMM Green Networking Workshop <http://doi.acm.org/10.1145/1851290.1851298>.

## 1. INTRODUCTION

Today’s electrical grid is a large-scale, heterogeneous, and widely-deployed system that is globally viewed as critical infrastructure. Unfortunately, this century-old system suffers from at least four major problems.

- First, a substantial fraction of the energy used to produce electricity is wasted due to generation, transmission, and distribution losses. For example, in the US, in 2008, this waste accounted for more than two-thirds of the energy produced [12].
- Second, a substantial fraction of electrical energy produced is from coal and natural gas, which contribute to global climate change. For example, in the US, in 2008, 65% of energy used to generate electricity was from these two sources [12].
- Third, generating plants are typically sized for the peak load because current end-devices are not elastic in their demands and their demands are highly correlated. This has a major impact on infrastructure cost. For example, in Massachusetts, 15% of the generation capacity is used less than 88 hours/year [13].
- Fourth, electrical storage capacity is very limited. Therefore, it is necessary to continuously match supply and demand, although both can change dramatically over time periods of a few hours.

The first two problems can be directly addressed by *distributed generation*—the production of energy from tens of millions of small scale sources, rather than a few hundred massive generators [8]. Distributed sources can be located closer to consumers, minimizing transmission losses. Moreover, these sources tend to be renewable (e.g., solar and wind), and have a negligible carbon footprint. Therefore, there is widespread consensus that the next-generation grid will feature distributed energy production from tens of millions of renewable-energy sources. This will substantially alter the nature of the grid, making it ‘greener,’ but also much more dynamic because many of these sources are time-varying. In parallel, the next-generation grid is likely to be ‘smarter.’ In a *smart grid*, grid operators might communicate with consumers to match supply to demand and to incentivize consumers to reduce their demand when the grid is under stress. This reduces the need to build out excess capacity that is mostly idle. We also expect more and

---

For conciseness, we refer to the electrical grid simply as the ‘grid’ in the rest of this paper.

more devices to handle elastic load and storage to be more prevalent. This will also significantly alter the nature of the grid. *The design and implementation of this next-generation smart, green grid will occur in the next decade and will create infrastructure that will last for the next century and beyond.*

In this position paper, we argue that the concepts and technologies pioneered by the Internet—the fruits of the past four decades of Internet research—can make fundamental contributions to the architecture and operations of the future grid. This is because the Internet and the grid were both designed to solve similar problems: that of matching geographically distributed suppliers (of information and energy) to consumers in near real-time. Symmetrically, we also identify a few areas where long-standing experience with the grid can be used to improve the operation of the Internet.

Our main contributions are:

- Identifying the fundamental similarities and differences between the Internet and the electrical grid;
- Identifying areas where the grid can be made smarter and greener using Internet concepts and technologies;
- Identifying areas where grid design principles can be used to improve the operation of the current Internet.

The rest of the paper is laid out as follows. We start with a survey of related work in Section 2. In Section 3 we describe the current status of the grid and some trends that will affect its future evolution. Section 4 compares the Internet and the grid. This motivates Section 5 where we identify areas where Internet technologies can be used to green and smarten the grid. Symmetrically, Section 6 identifies areas where grid design principles can be used to improve the Internet. We conclude in Section 7.

## 2. RELATED WORK

The terms ‘Green IT’ and ‘Green Networking’ have been used in literature primarily to refer to the reduction of electrical energy used in the IT infrastructure [14, 15, 24, 10]. For example, there have been efforts to turn off network switches when they are not being used [14, 15]. Similarly, Bolla et al study how to reduce the energy cost of a software router by dynamically adapting the number of active processor cores to the load [10]. There has also been a proposal to offload services from data centers to set-top boxes that are already on, so that servers in the data center can be turned off [24]. There has also been a recent proposal to source electricity from the cheapest source [22]: this does not reduce the amount of electricity used, but only its cost. Although these efforts are commendable, servers and data centres accounted for only about 1.5% of the electricity used in the US in 2006 [22], which corresponds to less than 0.5% of all electricity produced [12]. In contrast, more than half of all electricity produced is lost in generation, distribution and transmission alone [12] (two thirds was lost in the US in 2008). Thus, we believe that it is more important to focus on reducing transmission and distribution losses than to reduce the energy consumption of the Internet itself.

Recently, some research groups have focused on the use of Internet technologies to improve the grid. The IDEAS project [2] at the University of Southampton studies the decentralized control, operation and management of future generation electricity networks. The Spotlight [7] and PEIR

[6] research projects at UCLA focus on the use of sensors and smart phones to reduce the carbon footprint of daily activities. These projects primarily focus on the use of the Internet as an information overlay, rather than the use of Internet concepts and technologies themselves as a way to evolve the grid.

Our vision is most closely aligned with that of the Local project at UC Berkeley [5] that seeks to apply Internet design principles to the design of the next-generation grid. However, we believe that this paper is the first to lay out the similarities and differences between the Internet and the grid, and the first to demonstrate the use of specific Internet concepts and technologies for the next-generation grid.

## 3. THE GRID: STATUS AND EVOLUTION

The current electrical grid has a mesh-like core interconnecting a few hundred electrical energy producers supplying energy to hundreds or thousands of distribution networks that provide electrical energy to industrial and domestic consumers [25]. The current grid suffers from at least four major problems:

- Generation, transmission and distribution losses are pervasive. For example, in the US, in 2008, more than two-thirds of the energy used to produce electricity was wasted [12]. Electrical energy is wasted even after production. For example, in the US, in 1995, about 7% of electricity was wasted due to transmission and distribution losses [23]. Similarly, in India, in 1998, transmission and distribution losses accounted for nearly 25% of all electricity generated [21].
- Traditional energy generation tends to be carbon intensive: in 2008, non-carbon sources such as nuclear and hydroelectric power produced less than 30% of the energy used for producing electricity in the US [12]. Electricity generation is the primary consumer of coal in the US (using 91%) and therefore a major cause of the US carbon footprint [4]. Although countries such as France and Japan produce nearly 70% of their energy from non-carbon sources, these tend to be exception rather than the rule. According to the New York Times “Every week to 10 days, another coal-fired power plant opens somewhere in China that is big enough to serve all the households in Dallas or San Diego.”
- Grid endpoints are mostly ‘dumb’ and inelastic, unable to learn of the status of the grid, and therefore unable to reduce consumption when the grid is stressed. Therefore, electricity generation has to be built for peak load, which can be substantially higher than the average load.
- Electrical energy storage is scarce. Therefore, supply and demand must be continuously matched, and operators must maintain ‘spinning’ reserves that can be brought online at a few minutes’ notice.

It is becoming increasingly apparent that the grid of the future will look rather different than the grid of today [19].

“Pollution From Chinese Coal Casts a Global Shadow,” June 6, 2006.

The need to reduce carbon footprint has greatly increased investment in heterogeneous renewable sources of energy such as water, waves, wind, and sun. These tend to be inherently small-scale and distributed [8]. Synergistically, in an effort to reduce transmission losses, energy production is increasingly becoming more distributed, with numerous small-scale sources of production located close to sources of consumption. Thus, instead of a few large carbon-based producers, the future grid will have tens of millions of renewable-energy based distributed energy producers.

Renewable energy sources tend to be highly variable: a wind generator stops producing power when the wind stops blowing and solar cells do not work at night. However, in the current grid, there are inadequate provisions for energy storage, making it difficult to provide reliable power from renewable energy sources [18]. Advanced storage technologies such as ultracapacitors and nanotechnology-based batteries can play an increasingly important role in smoothing the gap between supply and demand [1]. Energy storage technologies will allow the use of locally stored energy when the grid is stressed. Indeed, stored energy can even be fed back into the grid.

Finally, given the availability of cheap communication technologies, it becomes easier for consumers to learn of grid status and use this to guide their consumption.

Thus, the current grid will evolve to a future grid that features

- tens of millions of distributed and variable renewable energy sources
- more flexible (elastic) end-devices
- vastly more storage
- the use of communication technologies both to allow precise matching of supply to demand and to incentivize appropriate consumer behaviour

This evolution will not be easily achieved. Aggregating inherently stochastic sources to achieve reliable electricity supply is a non-trivial problem. The unprecedented increase in the number of energy generation sources will require large-scale decentralized control operating simultaneously at time scales ranging from seconds to months. Assuring stability of this control is a challenging open problem. Similarly, introducing storage into the grid dramatically increases its complexity: efficiently managing geographically-distributed online and stored energy sources will require both sophisticated algorithms and a balancing of conflicting incentives. Finally, consumers need to be informed about the current status of the grid in a comprehensible format and given incentives to reduce their consumption.

Our key insight is that these problems mirror those in the Internet. Therefore, solutions to these problems pioneered in the Internet may be applicable, with some modification, to the grid. We now discuss the fundamental similarities and differences between the Internet and the next generation of the grid, deferring to Section 5 our discussion of how Internet concepts and technologies can be used to address these problems.

## 4. SIMILARITIES AND DIFFERENCES

### 4.1 Similarities

Both the Internet and the grid share some fundamental similarities, as discussed below.

- Both the Internet and the grid are critical to society, and represent large-scale, expensive, and highly-evolved systems. The grid is designed to operate at 99.999% reliability, which is a somewhat higher standard than that used for the Internet.
- Historically, both arose from a complex process that included a bottom-up federation of decentralized local providers and the top-down creation of a national infrastructure. Like LANs, local electricity providers originally provided services to a few tens of customers. Grids were formed in an effort to increase diversity and reduce downtime. At the same time, recognizing the value of the infrastructure, governments created both national grids and telcos, legislating a monopoly by fiat.
- Both the Internet and the grid essentially meet distributed demands, for energy and information respectively, from distributed suppliers. The common goal is to match supply to demand.
- Consumers of both systems are generally unaware of the underlying complexity. They pay either a flat fee or a usage-based fee and expect uninterrupted service.
- Both have a three-part structure comprising generation, transmission, and distribution. The generation part corresponds to electricity generation in the grid and content servers in the Internet cloud. Generators are interconnected by a mesh-like transport tier, corresponding to the long-haul transmission network in the grid and long-haul links in the Internet core. Finally, consumer demand is aggregated by a hub-like distribution tier, corresponding to local distribution companies in the grid and tier-two ISPs in the Internet.
- Both provide consumers with services in return for payment. In the case of the electrical grid, services correspond to consumption of electricity by end-devices such as toasters, air conditioners, and ovens. In the case of the Internet, services are (a wide range of) applications that ride over TCP/IP.
- Both differentiate between the core network and the edge (also called distribution or access) network. The core Internet is optimized for high capacity and often uses a packet format (such as MPLS) that is different from the packet format at the edge (IP). The core electrical grid is optimized for minimizing transmission losses and uses high-voltage AC or (rarely) DC transmission that is stepped down in the distribution network.
- Both embody several degrees of control: from very strict control in the core, to a degree of heterogeneity in the distribution network, to a lightly-controlled environment within a home or building.

- Both support some degree of storage. In the grid, this corresponds to technologies such as batteries, gyroscopes, and pumped storage in hydroelectric dams [1]. In the Internet, this corresponds to buffers in routers and content caches.
- Both provide a simple API between the infrastructure and users, corresponding to IP in the Internet and the three pin plug in the grid. This allows independent innovation.
- Both require distributed resource management: geographically decentralized demands must be met by geographically diverse sources of supply.
- Both have a tremendous investment in legacy protocols and components. Given their current scale, the protocols and APIs have ossified. Therefore, any changes to be made in them necessarily have to be legacy compatible.
- Both codify experience in the form of open standards that permit interconnection and interoperation. In the Internet, this corresponds to RFCs published by the IETF. In the grid, standards are codified by the International Electrotechnical Commission [3].
- Both have established a tenuous balance between centralization and decentralization. Local utilities, just like ISPs, provide customer care and access. Like ISPs, they have the ability to make decentralized decisions. On the other hand, Tier-1 ISPs, like national grids, are centrally controlled and are typically oligopolies or monopolies.
- Agreements between the central components and the access networks are complex and usually embody some notion of a Service Level Agreement (SLA). For example, in the grid, interconnection standards for distributed energy sources are codified in ANSI/IEEE standard 1547 and IEC standard TC8, and in the Internet, interconnection standards between ‘user networks’ and ‘service networks’ are governed by RFCs such as RFC 5865.

We found it helpful, given these similarities, to *think of the grid as a content distribution network designed to distribute a single content stream.*

## 4.2 Differences

While sharing many similarities, the electrical grid and the Internet differ in some fundamental ways. One major difference is that packets allow metadata in the form of headers. This allows the packet to have a type and a destination. This allows the Internet to support multi-commodity flows, whereas the grid supports a single commodity. More importantly, the Internet allows point-to-point flows, which are difficult to accomplish (if even feasible) in the electrical grid. For instance, in the grid, it is difficult to meet demands with fine granularity (think of ‘addressing’ a single light bulb when sending electricity down a wire). This is trivial in the Internet.

A second set of differences is in the directionality of flows. In the Internet, flows are bidirectional, though downstream flows dominate. In contrast, in the grid, almost all flows are downstream, with negligible upstream flow (though this will

change with the introduction of distributed generation). On a related note, Internet flows are more intelligent, in that a consumer of a flow can dictate precisely how much it wants to consume, and from where. In contrast, electrical grid endpoints are ‘dumb.’

The two infrastructures differ in the time-scales involved in control. Flash-crowds for content in the Internet can be responded to in a few seconds. In contrast, rapid increases in demand in the grid are not easy to meet: bringing a nuclear reactor online can take a few days, and even hydroelectric turbines need nearly an hour to be brought online.

Long-haul transmission in both networks scale differently. In the Internet, once a fiber-optic link has been installed, adding more capacity corresponds to simply adding more wavelengths to the optical transducers at the ends of the link. This allows rapid scaling up as demand ramps up. In contrast, electrical transmission lines are limited by the conductors in place, and adding more capacity is tremendously expensive.

Services in both networks differ in their range. The Internet supports thousands of applications while the grid will probably remain a single service network.

Although both networks provide some degree of storage, the ease of adding and managing storage is very different for the two networks. Memory and storage costs in the Internet are low and continue to decline every year. In contrast, electrical storage is scarce and expensive. Electrical storage costs are not declining at the same rate as that of digital storage.

Finally, the Internet and the grid differ in their predictability. Grid demands are highly predictable based on time-of-day, day-of-week and so on. This allows grid operators to plan for operations days or even weeks in advance. This is usually not possible for Internet providers, where even core network traffic can fluctuate dramatically and unpredictably over periods of hours.

These facts fundamentally differentiate the two networks. It is instructive to think of the electrical grid as a content distribution network for a single file. Given our goal of bringing together the Internet and the electrical grid, we will focus primarily on their similarities, while keeping these differences in mind, so as not to be led too far astray.

## 5. SMARTENING THE GRID

Given the many similarities between the grid and the Internet, we believe that technologies and principles pioneered and proven effective in the Internet can be applied to green and smarten the grid. Interestingly, the fundamental and secondary design goals of the Internet [11] appear to be nearly identical to those of the smart grid. Loosely paraphrased, for both networks the fundamental goal is to effectively integrate a number of separately administered existing networks into a common utility using ‘gateways.’ The common secondary design goals are: a) to tolerate loss of individual components, b) to support different underlying network types, c) to allow distributed resource management, d) to be cost effective, e) to permit easy endpoint attachment and f) to be accountable for resource usage. Moreover, the problems discussed in Section 3, that of aggregation of stochastic sources, distributed resource management, multiple time scales of control, and user incentivization, are similar to that faced in the Internet.

Internet researchers have studied, and, to first order, have

solved, these problems in the last four decades. For example, fundamental advances in the modelling of aggregate stochastic sources allow operators to size long-distance links with high confidence. The Internet has always implemented highly-scalable decentralized control: Internet control algorithms such as scheduling, flow control, traffic management, and topology design span many orders of magnitude in time and have been found to be both robust and efficient in practice. Large-scale distributed algorithms for cooperative caching and resource placement in content distribution networks have addressed the distributed resource management problem for more than a decade. Moreover, there is a rich body of work studying incentives for good behaviour in the context of the Internet. Therefore, we believe that Internet concepts and technologies can be used to study and alleviate these problems.

Given this overall context, we now present details of some specific solutions pioneered in the Internet that can be carried over to the grid. This listing is by no means exhaustive, but suggests the potential of this approach. We recognize that the grid and the Internet differ in many ways so the translation, though promising is not straightforward.

#### **Local matching of supply to demand in a microgrid**

Transmission and distribution losses account for 7-25% of all electricity produced. Therefore, it is important to reduce the distance travelled by electrical energy. This is isomorphic to the need to reduce delay to serve content on the Internet. Therefore, ideas from content distribution network can be applied to the grid. Specifically, we propose the use of caching (battery piles at sub-stations) and peer-to-peer matching of supply and demand (using a *microgrid* [17]). Caches can hold excess electricity generated by a home or business using solar or wind energy in a neighbourhood during periods of slack demand. These could be released when demand increased at a later time. Demands could also be met using peer-to-peer matching, i.e., by plug-in-electric cars parked in neighbours' garages. An example of an electrical interconnection network necessary for two-way electricity flow is described in Reference [9].

**Tomography** A significant cause of transmission losses, especially in developing countries, is the illegal tapping of transmission lines. In India, for example, illegal tapping is the primary reason for transmission and distribution losses. Reference [21] reports a loss of 22-24% in the period 1995-1998. Tapping is usually discovered using time-domain reflectometry, an effective but painstaking procedure. Internet tomography reveals the traffic matrix of a network by monitoring a few aggregation points. We believe that it may be possible to reveal the fine structure of the grid using similar techniques, making it easy to detect illegal tapping.

**Modelling of stochastic generating sources** There is a rich body of work on Internet traffic modelling and analysis. A deep theoretical basis has been developed to study stochastic sources and aggregates of stochastic sources [20]. These stochastic models and analysis are critical in understanding the behaviour of renewable energy systems such as windmills and solar panels, whose output varies over time. Specifically, these

models allow us to intelligently aggregate such sources with sufficient diversity so that base load can be met with high confidence.

**Delay-tolerant networking** A recent development in Internet evolution is the study of networks where data is carried physically on data mules to provide delay-tolerant networking. This is similar, in principal, to the carrying of battery packs in a car. Considering that a typical home consumes about 20 kilowatt-hours (KWh) of energy, and that a plugin electric car, such as the Chevy Volt already has a battery pack of similar capacity, using cars to physically move energy using DTN routing [16] may be feasible, and even attractive.

**Latent indexing** A recent development in the Internet is to collect monumental amounts of usage data and use data mining techniques to extract useful information from them. For example, every hyperlink on a website and every click on a hyperlink tells a search engine about user preferences. These can be used to rank web pages and to fine-tune search results. In the same way, every click of a light switch also reveals user preferences: for example, if a user turns on an air conditioner during peak-pricing, this reveals that she is demand-inelastic. We envision the monitoring and harvesting of this switch-stream, its processing using algorithms developed for signal processing, and its analysis through techniques pioneered in data mining.

**Incentive-compatibility** Internet researchers have long studied the problem of incentive compatibility: how to make end users behave in a way that is beneficial to the network. Specifically, game-theoretic formulations of user behaviour, along with the principles of mechanism design, have given insights into how user behaviour can be made to satisfy a preferred social welfare function. These insights and analyses are applicable to the grid, where consumers will be expected to reduce usage during times of grid stress, and cooperate with their neighbours to smooth their aggregate load profile.

**Congestion control** The electrical grid congests when the load exceeds supply. This can happen due to a short circuit, which creates an infinite demand, or due to highly correlated demands, for example due to air conditioner use during hot weather. We believe that proactive and reactive flow control algorithms developed for the Internet may be applicable to the grid. For example, time-division scheduling of air-conditioners in a neighbourhood would decorrelate their demands, reducing the peak load on a substation. The rapid deployment of smart meters makes such congestion control schemes practicable, even in the near term.

**Distributed algorithms for control** The Internet embodies control algorithms operating at the time scale of nanoseconds (scheduling), milliseconds (flow control), minutes (routing), hours (traffic management), and months (topology planning). These algorithms are distributed and found to be robust. We envision the application of these types of algorithms to the grid to deal with control problems of source impedance matching

(milliseconds), fault containment (seconds), dispatch (minutes), and topology planning (years).

## 6. WHAT THE INTERNET CAN LEARN FROM THE ELECTRICAL GRID

Our study of the grid has led to the realization that it is not just the grid that can benefit from Internet concepts and technologies. In fact, the Internet has much to learn from the grid. This is not the primary focus of our work, so we will merely touch upon a few general ideas:

**Peak-load pricing** The grid has long used peak-load pricing, where consumers pay more for electricity during peak periods. Given the effectiveness of this approach, and the fact that consumers are already trained to deal with two-part pricing schedules, it seems that ISPs could reduce the capacity of their network links, if necessary, by using a similar approach.

**Fault isolation** The grid has invested heavily in the ability to contain faults using electrical breakers. This contains faults within small regions, leaving the overall grid unharmed. A similar approach could be used in the Internet, where traffic triggers or other warnings of danger could isolate flows or even entire ports from the Internet. This may make the Internet more tolerant of faults.

**Load forecasting** The grid depends on accurate load forecasting. Utilities use weather and usage models to predict demand a day, week, and months ahead. These techniques could be used to load forecasting for the Internet.

Again, we note that this list is not meant to be exhaustive, but merely illustrative of the many Internet research areas that open up by a careful study of the grid.

## 7. CONCLUSIONS

The century-old grid will see a major face lift in the next decade. Based on a careful examination of the Internet and the grid, we have shown that the concepts and techniques pioneered by the Internet are directly applicable to the design of the next-generation grid. Moreover, we have identified specific areas where the Internet community can work with the electrical grid research community. We hope that our work will initiate a dialogue between these two communities.

## 8. REFERENCES

- [1] Electricity storage technologies. <http://www.electricitystorage.org/site/technologies/>, 2010.
- [2] IDEAS Project. <http://www.ecs.soton.ac.uk/research/projects/618>, 2010.
- [3] International electrotechnical commission. <http://www.iec.ch>, 2010.
- [4] Lawrence Livermore National Laboratory Energy Flow Diagram. <https://publicaffairs.llnl.gov/news/energy/energy.html#2008>, 2010.
- [5] LoCal Project. [http://local.cs.berkeley.edu/wiki/index.php/Main\\_Page](http://local.cs.berkeley.edu/wiki/index.php/Main_Page), 2010.
- [6] PEIR Project. <http://peir.cens.ucla.edu/>, 2010.
- [7] Spotlight project. <http://nesl.ee.ucla.edu/project/show/59>, 2010.
- [8] K. Alanne and A. Saari. Distributed energy generation and sustainable development. *Renewable and Sustainable Energy Reviews*, 10(6):539–558, 2006.
- [9] A. Arulampalam, M. Barnes, A. Engler, A. Goodwin, and N. Jenkins. Control of power electronic interfaces in distributed generation microgrids. *International Journal of Electronics*, 91(9):503–523, 2004.
- [10] R. Bolla, R. Bruschi, F. Davoli, and A. Ranieri. Energy-aware performance optimization for next-generation green network equipment. In *Proc. ACM SIGCOMM PRESTO workshop*, 2009.
- [11] D. Clark. The design philosophy of the DARPA Internet protocols. In *Proc. ACM SIGCOMM*, 1988.
- [12] Energy Information Administration of the US Department of Energy. Annual Energy Review 2008. *Report DOE/EIA-0384(2008)*, 2008.
- [13] P. Giudice. Our energy future and smart grid communication, Testimony before the FCC Field Hearing on Energy and Environment. [www.broadband.gov/fieldevents/fh\\_energy\\_environment/giudice.pdf](http://www.broadband.gov/fieldevents/fh_energy_environment/giudice.pdf), 2009.
- [14] M. Gupta and S. Singh. Greening of the Internet. In *Proc. ACM SIGCOMM*, 2003.
- [15] B. Heller, D. Underhill, S. Seetharaman, N. McKeown, P. By, and A. Mishra. ElasticTree: Saving energy in data center networks. *Proc. ACM NSDI*, 2010.
- [16] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proc. ACM SIGCOMM*, 2004.
- [17] R. Lasseter and P. Paigi. Microgrid: A conceptual solution. In *2004 IEEE 35th Annual Power Electronics Specialists Conference, 2004. PESC 04*, volume 6, 2004.
- [18] B. Lee and D. Gushee. Electricity storage: The Achilles' heel of renewable energy. *Chemical Engineering Progress*, 104(3):29, 2008.
- [19] E. Lightner and S. Director. Evolution and Progress of Smart Grid Development at the Department of Energy. In *DOE presentation at FERC-NARUC Smart Grid Collaborative Workshop, Washington, DC*, 2008.
- [20] R. Mazumdar. *Performance Modeling, Loss Networks, and Statistical Multiplexing*. Morgan and Claypool, 2009.
- [21] P. Purkayastha. Power Sector Policies and New Electricity Bill: from Crisis to Disaster. *Economic and Political Weekly*, 36(25):2257–2262, 2001.
- [22] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the electric bill for internet-scale systems. *Proc. ACM SIGCOMM*, 2009.
- [23] U.S. Climate Change Technology Program. Technology Options 2005. [www.climatechange.gov/library/2005/tech-options/tor2005-132.pdf](http://www.climatechange.gov/library/2005/tech-options/tor2005-132.pdf), 2005.
- [24] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez. Greening the internet with nano data centers. In *Proc. ACM CoNEXT*. ACM, 2009.
- [25] A. Von Meier. *Electric power systems: a conceptual introduction*. Wiley-IEEE Press, 2006.