

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

S. Keshav and C. Rosenberg

School of Computer Science and Department of Electrical and Computer Engineering  
University of Waterloo, Waterloo, ON, Canada  
DRAFT

## 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 demonstrate that 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 hope that our work will initiate a dialogue between these two communities.

## 1. INTRODUCTION

Today’s electrical grid<sup>1</sup> 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 three major problems. First, a substantial fraction of the energy used to produce electricity is wasted due to generation, transmission, and distribution losses (in the US, in 2008, this waste accounted for more than two-thirds of the energy produced [11]). Second, a substantial fraction of electrical energy produced is from coal and natural gas, which contribute to global climate change (in the US, in 2008, 65% of energy used to generate electricity was from these two sources [11]). Third, generating plants are typically sized for the peak load because consumers often do not have a strong incentive to reduce their

demand during times of peak demand, such as on a hot summer’s day.

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.’ In parallel, the next-generation grid is likely to be ‘smarter,’ using communication technologies to better match supply to demand and to incentivize consumers to reduce their demand when the grid is under stress. *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

<sup>1</sup>For conciseness, we refer to the electrical grid simply as the ‘grid’ in the rest of this paper.

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 [12, 13, 22, 9]. For example, there have been efforts to turn off network switches when they are not being used [12, 13]. 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 [9]. 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 [22]. Although these efforts are certainly commendable, servers and data centres accounted for only about 1.5% of the electricity used in the US in 2006 [20], which corresponds to less than 0.5% of all electricity produced [11]. In contrast, more than half of all electricity produced is lost in generation, distribution and transmission alone [11] (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 [3] 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 [23]. The current grid suffers from at least three 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 [11]. Similarly, in India, in 1998, transmission and distribution losses accounted for nearly 25% of all electricity generated [19].

- 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 [11]. 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<sup>2</sup> “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’, 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.

It is becoming increasingly apparent that the grid of the future will look rather different than the grid of today [17]. 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 integrate renewable energy sources [15]. Advanced storage technologies such as ultracapacitors and nanotechnology-based batteries can play an increasingly important role in smoothing the gap between supply and demand [2]. Energy storage technologies will allow the use of locally stored energy when the grid is stressed. In fact, 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

---

<sup>2</sup>“Pollution From Chinese Coal Casts a Global Shadow,” June 6, 2006.

- tens of millions of distributed and variable renewable energy sources;
- 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 base load 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.

We now discuss the fundamental similarities and differences between the Internet and the next generation of the grid, deferring to Section 5, our analysis 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.
- 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 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 (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 transmission that is stepped down in the distribution or access network.
- Both share a similar topology: hub and spoke in the distribution network and mesh in the core. This corresponds to the differing roles of these components.
- 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 [2]. 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.
- 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. Arrangements between the central components and the access networks are complex and usually embody some notion of a Service Level Agreement (SLA).
- 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 infrastructures. 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.

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 head-

ers. 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.

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. GREENING AND 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 [10] appear to be identical to those of the grid. Moreover, the problems discussed in Section 3, that of aggregation of stochastic sources, distributed resource management, multiple time scales of con-

trol, and user incentivization, are similar to that faced in the Internet.

Internet researchers have studied, and, to first order, solved, these problems in the last four decades. For example, fundamental advances in the modeling 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<sup>3</sup>. 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** Generation, transmission and distribution losses account for more than two-thirds of the energy used to produce electricity. Therefore, it is critical to reduce the distance traveled 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. 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.

**Tomography** A second reason for transmission losses, especially in developing countries, is the illegal tapping of transmission lines. In India, for example, illegal tapping was the primary reason for transmission and distribution losses. Reference [19] 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

---

<sup>3</sup>This listing is by no means exhaustive, but suggests the potential of this approach.

similar techniques, making it easy to detect illegal tapping.

**Modelling of stochastic 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 [18]. 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 [14] 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 they are 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 smoothen their aggregate load profile.

**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).

**Simulation and Emulation** Internet designs are typically evaluated using detailed simulations and emulations. Over the past two decades, large-scale simulation and emulation techniques have been developed to study the Internet [1, 21, 16]. We believe that similar techniques will allow us to understand the qualitative behaviour of the grid, and to explore the effect of potential changes to the system.

## 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] Ns2. <http://www.isi.edu/nsnam>, 2005.
- [2] Electricity storage technologies. <http://www.electricitystorage.org/site/technologies/>, 2010.
- [3] Ideas project. <http://www.ecs.soton.ac.uk/research/projects/618>, 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] 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.
- [10] D. Clark. The design philosophy of the DARPA Internet protocols. In *Proc. ACM SIGCOMM*, 1988.
- [11] Energy Information Administration of the US Department of Energy. Annual Energy Review 2008. *Report DOE/EIA-0384(2008)*, 2008.
- [12] M. Gupta and S. Singh. Greening of the Internet. In *Proc. ACM SIGCOMM*, 2003.
- [13] B. Heller, D. Underhill, S. Seetharaman, N. McKeown, P. By, and A. Mishra. ElasticTree: Saving energy in data center networks. *Proc. ACM NSDI*, 2010.
- [14] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proc. ACM SIGCOMM*, 2004.
- [15] B. Lee and D. Gushee. Electricity storage: The Achilles’ heel of renewable energy. *Chemical Engineering Progress*, 104(3):29, 2008.
- [16] J. Lepreau. The Utah Emulab Network Testbed.
- [17] 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.
- [18] R. Mazumdar. *Performance Modeling, Loss Networks, and Statistical Multiplexing*. Morgan and Claypool, 2009.
- [19] P. Purkayastha. Power Sector Policies and New Electricity Bill: from Crisis to Disaster. *Economic and Political Weekly*, 36(25):2257–2262, 2001.
- [20] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the electric bill for internet-scale systems. *Proc. ACM SIGCOMM*, 2009.
- [21] L. Rizzo. Dummynet: a simple approach to the evaluation of network protocols. *ACM SIGCOMM Computer Communication Review*, 27(1):31–41, 1997.
- [22] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez. Greening the internet with nano data centers. In *Proc. ACM CoNEXT*. ACM, 2009.
- [23] A. Von Meier. *Electric power systems: a conceptual introduction*. Wiley-IEEE Press, 2006.