The Growing Interdependence of the Internet and Climate Change

Larry Smarr
University of California, San Diego

As proven by the global attendance at December's UN Climate Change Conference 2009 (http://en.cop15.dk/), more attention is being paid to the components of our society responsible for the emission of greenhouse gases (GHGs) and how to reduce those emissions. The global information and communication technology (ICT) industry, which includes the Internet, produces roughly 2 to 3 percent of global GHG emissions, according to the Climate Group's Smart2020 report (www.smart2020.org). Furthermore, if it continues to follow a business-as-usual scenario, the ICT sector's emissions will nearly triple by 2020.

However, the Climate Group estimates that the transformative application of ICT to electricity grids, logistic chains, intelligent transportation, building infrastructure, and dematerialization (telepresence) could reduce global GHG emissions by roughly 15 percent, five times ICT's own footprint! So, the key technical question before our community is, can we reduce the carbon intensity of Internet computing rapidly enough that even with its continued spread throughout the physical world, the ICT industry's overall emissions don't increase?

This is a system issue of great complexity, and to make progress we need numerous at-scale testbeds in which to quantify the many trade-offs in an integrated system. I believe our research university campuses themselves are the best testbeds, given that each is in essence a small city, with its own buildings, hospitals, transportation systems, electrical power generation and transmission facilities, and populations in the tens of thousands. Indeed, once countries pass legislation for carbon taxes or "cap and trade" markets, universities will have to measure and reduce their own carbon footprints anyway,[1] so why not instrument them now and use the results as an early indicator of the optimal choices for society at large?

As discipline after discipline transitions from analog to digital, we'll soon find that when the carbon accounting is done, a substantial fraction of a campus's carbon footprint is in its Internet computing infrastructure. For instance, a major carbon source is data center electrification and cooling. Many industries, government labs, and academics are working to make data centers more efficient (see http://svlg.net/campaigns/datacenter/docs/DCEFR_ report.pdf). At the University of California, San Diego (UCSD), our US National Science Foundation-funded GreenLight project (http://greenlight.calit2.net) carries this work one step further by providing the end user with his or her application's energy usage. We do this by creating an instrumented data center that allows for detailed real-time data measurements of critical subcomponents and then making that data publically available on the Web, so that the results can guide users who wish to lower their energy costs.

This is more complex than you might think at first. Any given application, such as bioinformatics, computational fluid dynamics, or molecular dynamics, can be represented by several algorithms, each of which could be implemented in turn on a variety of computer architectures (multicore, field-programmable gate array, GPUs, and so on). Each of these choices in the decision tree requires a different amount of energy to compute. In addition, as UCSD's Tajana Rosing has shown, we can use machine learning to implement various power[2] or thermal[3] management approaches, each of which can save up to 70 percent of the energy used otherwise in the computations.

Another strategy to reduce overall campus carbon emissions is to consolidate the clusters and storage systems scattered around campus in different departments into a single energy-efficient facility and then use virtualization to increase the centralized cluster's utilization. We could also use zero-carbon energy sources (solar or fuel cells), which produce DC electricity, to drive the cluster complex, bypassing the DC to AC to DC conversion process and reducing the operational carbon footprint of campus computing and storage to zero.

As we reduce the carbon emissions required to run Internet computing, we can extend the Internet into new functions, such as instrumenting buildings for their energy use and eventually autonomously controlling building systems in real time to reduce overall energy use. An example is the research performed in UCSD's Computer Science and Engineering building by Rajesh Gupta and his colleagues, who found that roughly 35 percent of the building's peak electrical load is caused by PCs and servers. His team's research also showed that intelligent sleep-state management could help avoid a large fraction of this Internet computing electrical load (www.usenix.org/events/nsdi09/tech/full_papers/agarwal/agarwal_html/).

Another application of Internet computing to avoid carbon emissions is *dematerialization*, such as using Internet video streaming to reduce air or car travel to meetings. At Calit2, we use a variety of compressed high-definition (HD) commercial systems such as LifeSize H.323 videoconferencing (approximately 1 to 2 Mbps) or high-end systems such as Cisco's Telepresence system (approximately 15 Mbps). However, we're also experimenting with uncompressed (1,500 Mbps) HD (developed by the University of Washington's Research Channel) or with digital cinema (four times the resolution of HD), which requires 7,600 Mbps uncompressed! These higher-bandwidth video streams are used over dedicated optical networks (such as CENIC, Pacific Wave, the National LambdaRail, Internet2's Dynamic Circuits, or the Global Lambda Integrated Facility, all operating at 10,000 Mbps).

We can extend the notion of virtual/physical spaces from simple face-to-face meetings to creating collaborative data-intensive analysis environments in which whole rooms are "sewn together" using the Internet video streaming technologies mentioned earlier. Calit2 is an institute that spans two University of California campuses, San Diego and Irvine, separated by a 90-minute drive. We recently started using HD streaming video to link our two auditoriums together for joint meetings, such as our all-hands meetings. Previously, we needed dozens of people from one campus to drive to the other campus for such a meeting.

Another example that focuses more on research is how Calit2 in San Diego and the NASA Ames Lunar Science Institute in Mountain View, California, have both set up large tiled walls (displaying tens to hundreds of megapixels) called OptIPortals and then used the CENIC dedicated 10-Gbps optical networks to couple their two rooms with streaming video and spatialized audio. This lets researchers at both ends explore complex lunar and Martian images taken by orbiting or surface robotic craft. Each side can control image placement and scaling on the other's wall, so team brainstorming is as easy as if both sides were in the same physical room. We use this on a weekly basis, avoiding a significant amount of plane travel and the carbon emissions that it would otherwise produce.

These ideas are just the tip of the iceberg of how we can turn our research universities into living laboratories of the greener future. As more universities worldwide begin to publish their results on the Web, best practices will quickly develop and lessons learned can be applied to society at large. This is essential because the world must act this coming decade to make drastic changes in the old "high carbon" way of doing things and transition to a new "low carbon" society if we're to avoid ever worsening global climatic disruption.

References

1. B. St. Arnaud et al., "Campuses as Living Laboratories for the Greener Future," *EDUCAUSE Rev.*, vol. 44, 2009, pp. 14–32.
2. G. Dhiman and T. Simunic Rosing, "System-Level Power Management Using Online Learning," *IEEE Trans. Computer-Aided Design*, vol. 28, no. 5, 2009, pp. 676–689.
3. A. Coskun and T. Simunic Rosing, "Proactive Thermal Management," to be published in *IEEE Trans. Computer-Aided Design*, 2009.

Larry Smarr is the Harry E. Gruber Professor in the Department of Computer Science and Engineering in the Jacobs School of Engineering at the University of California, San Diego, and director of the California Institute for Telecommunications and Information Technology, a UC San Diego/UC Irvine partnership. His research interests include green IT, global telepresence, and microbial metagenomics. Smarr has a PhD in physics from the University of Texas at Austin. Contact him at lsmarr@ucsd.edu or follow him as lsmarr on Twitter.