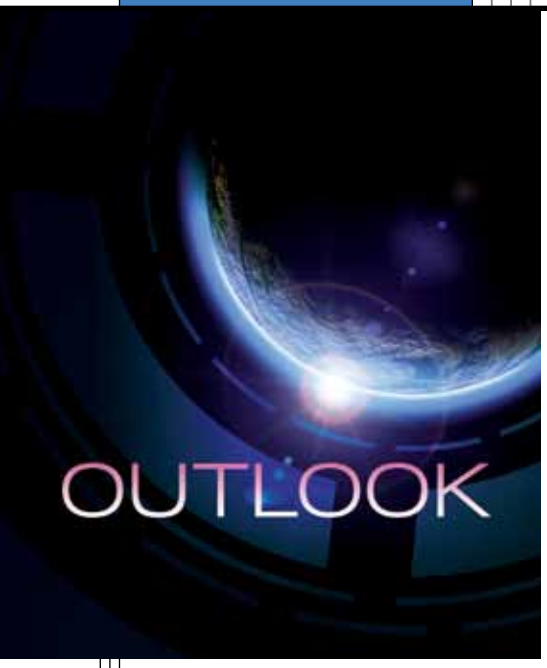


PROJECT GREENLIGHT: OPTIMIZING CYBER-INFRASTRUCTURE FOR A CARBON-CONSTRAINED WORLD

Larry Smarr, *University of California, San Diego*

Even with a variety of aggressive energy efficiency measures, the ICT sector's carbon emissions will nearly triple from 2002 to 2020. We must accelerate ICT energy efficiency so that we can increase the use of ICT in smart infrastructure capable of reducing global greenhouse gas emissions.

Everyone is now aware of the growing threat of global climatic disruption, but it's less well known that our information and communication technology (ICT) community can play a key role in this looming crisis. The Climate Group, on behalf of the Global eSustainability Initiative (GeSI)—a consortium of major IT and telecommunications companies—recently issued an informative new study, *Smart 2020: Enabling the Low Carbon Economy in the Information Age*, on this topic (www.theclimategroup.org). This report argues that in addition to making ICT systems more energy efficient, application of those systems to electricity grids, logistic chains, intelligent transportation, and building infrastructure could reduce global greenhouse gas (GHG) emissions by as much as 15 percent by 2020, compared with business as usual. This could be a critical element for enabling countries to meet their emission reduction goals.

First, I will review some key scientific results that illustrate just how far we have already come in changing the Earth's atmosphere and show that the climate is beginning to react. Then I will review ICT's role in GHG emissions and describe several advances that allow us to reduce future emissions.

CURRENT EARTH SCIENCE RESEARCH

Recent research has developed a probability distribution for the warming that we can expect from the carbon dioxide and other greenhouse gases already emitted since the beginning of the Industrial Age 250 years ago.[1] The most probable outcome, shown in Figure 1 from that study, shows that over time, about 2.5°C warming will occur as a direct result of our past emissions. However, to this point, we've seen only about a 0.8°C increase in warming, or only a third of what is going to happen. This delay has two major reasons: First, it takes about 50 years for the thermal equilibrium of the oceans to adjust; second, the aerosols that are being emitted, particularly in Asia, are cooling the Earth. Somewhat ironically, the rest of that warming will appear as we clean up current air pollution over the next few decades.

There is an emerging scientific consensus[2] on a variety of climate tipping points which begin to occur as the global temperature rises—for example, melting of the summer Arctic ice, the Himalayan glaciers, and the Greenland ice sheet. As Figure 1 shows, the current level of greenhouse gases has already committed us to serious environmental changes. Unfortunately, as we continue to add more GHGs

 Published by the IEEE Computer Society 0018-9162/10/$26.00 © 2010 IEEE

to the atmosphere, the peak of this curve moves further to the right, fostering greater disruption.

Arctic ice sheet

If this analysis is correct, we would conclude from Figure 1 that we should already be seeing indications that we are past the first climatic tipping point—melting of the Arctic Ocean summer ice. Figure 2 presents a summary of NASA satellite data on the Arctic ice sheet over the past three decades.[3] As the figure shows, from the 1980s to 2000, much of the ice was several years old, whereas in 2009 very little of this older, thicker, ice remains. The graph of the two-year and older ice shows why climate scientists predict that the Arctic may lose its several-year ice in the next five years, leaving only the annual ice, which is thin enough for ships to move through the Arctic Ocean. While this may be good for ocean transportation, this elimination of Arctic summer ice will cause dramatic climate changes all over the northern hemisphere.

The water towers of Asia are melting

In May 2009, UC San Diego and the University of Cambridge held a three-day conference on the next tipping point in Figure 1: the melting of the glaciers in the Himalayan and Hindu Kush Uplift in Asia, which contains the largest amount of snow and ice outside the north and south polar regions.[4] This frozen water forms the "water towers of Asia,"[5] being the source of the great rivers from India to Southeast Asia to China—the Indus, Ganges, Brahmaputra, Mekong, Yellow, and Yangtze, among others—which carry the melting snow and ice to the ocean. Over the past decade, the glaciers have begun to melt very rapidly, impacting the water supplies of over a billion people. Also, temporary natural dams often form, backing up large lakes from the melting ice and snow and then giving way, causing flash floods that can destroy villages downstream.

So we see that there is significant evidence that current levels of GHGs are beginning to shift the climate as predicted in the tipping-point study. Unfortunately, the global GHG emissions are continuing to increase, implying even more profound climate shifts. Therefore it is worth taking a look at just how unusual the current level of carbon dioxide is compared to historical values.

Historical climate oscillations

It is true that the temperature and CO_2 levels of the atmosphere have oscillated over time. Some skeptics say that we are just experiencing another natural oscillation. To examine that hypothesis let's consider the last series of these oscillations. Figure 3 shows the past 800,000 years of oscillations in both CO_2 and temperature as derived from Antarctic ice cores.[6] The lowest point of these oscillations coincides with ice ages and the peaks with

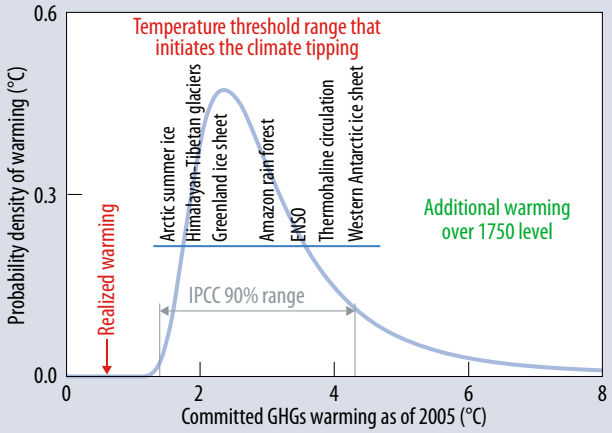


Figure 1. Temperature threshold range that initiates climate tipping. Earth has realized only one-third of the committed warming from previous greenhouse gas emissions; further emissions will move the curve to the right. Adapted with permission from V. Ramanathan and Y. Feng, "On Avoiding Dangerous Anthropogenic Interference with the Climate System: Formidable Challenges Ahead," *PNAS*, vol. 105, pp. 14245-14250, Copyright 2008 National Academy of Sciences, U.S.A.

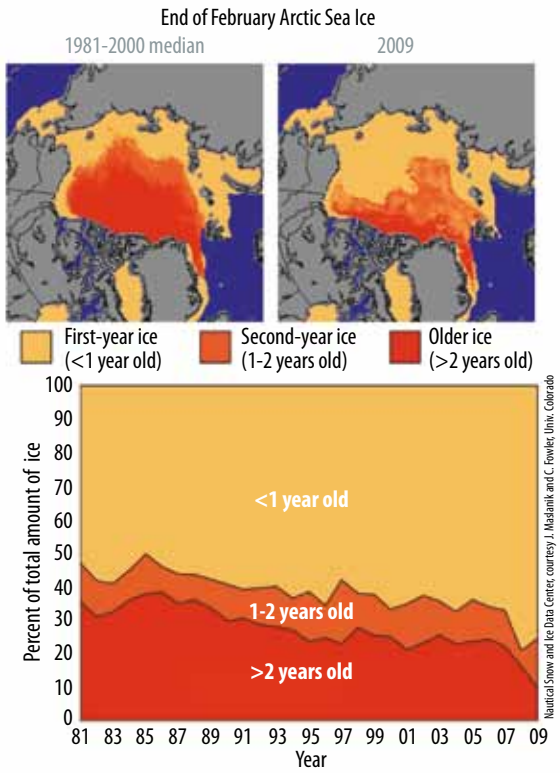


Figure 2. The multiyear Arctic ice sheet has diminished precipitously over the past decade, causing climate scientists to predict that the Arctic may become ice-free in the summer as soon as five years hence. Reproduced with permission, National Snow and Ice Data Center, courtesy of J. Maslanik and C. Fowler.

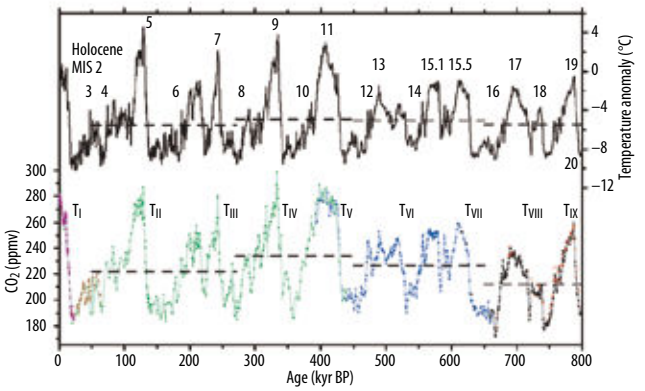


Figure 3. The measured changes in atmospheric CO_2 and temperature from ice cores over the past 800,000 years. Reprinted with permission from *Nature*, Nature Publishing Group, 15 May 2008, vol. 453, pp. 379-382.

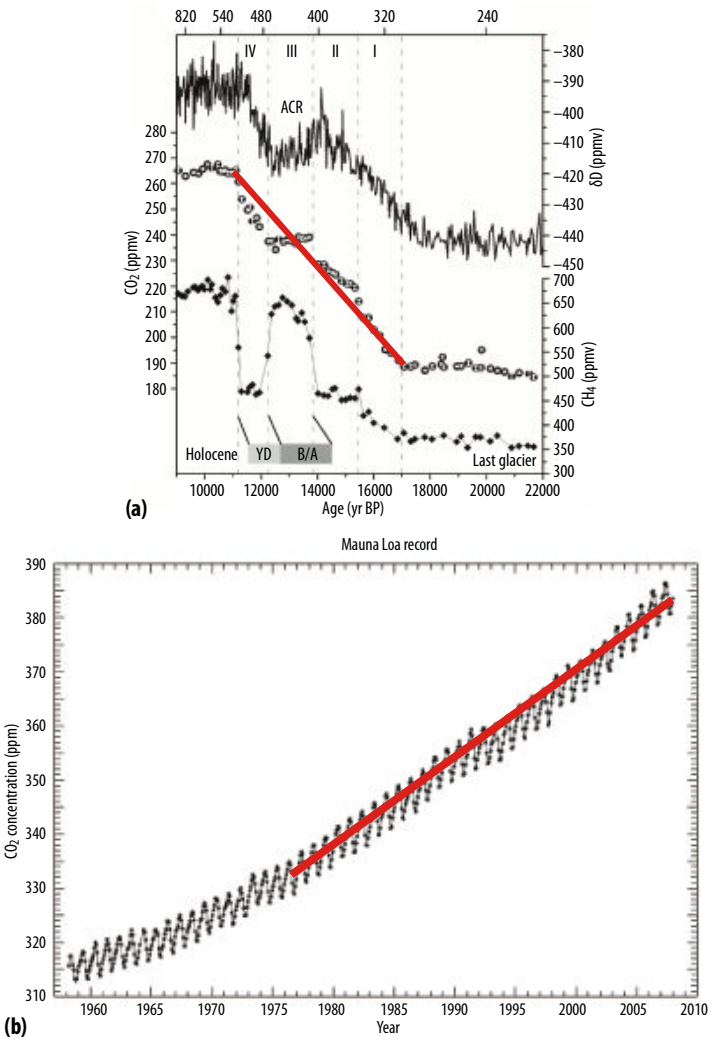


Figure 4. Time rate of change in the atmosphere's CO_2. (a) The middle curve shows the increase in carbon dioxide emissions since the last ice age. Reprinted with permission from *Science*, vol. 291, 2001, pp. 112-114; (b) increase in emissions over the past 50 years; http://scrippsco2.ucsd.edu/program_history/keeling_curve_lessons.html.

interglacial periods, such as we have been in for the past 10,000 years. As Figure 3 shows, the amount of carbon dioxide in the atmosphere has oscillated between about 170 and 300 parts per million (ppm), while the temperature has oscillated as much as 12 degrees centigrade for many ice age/interglacial cycles.

In contrast, in 2010, the CO_2 will reach 390 ppm, and an MIT study indicates that global economic growth in a "business as usual" scenario would raise this to approximately 900 ppm by the end of this century.[7] Clearly, we are entering levels of CO_2 that the Earth's atmosphere has not seen for a very long time.

Next, let's consider the time rate of change of the atmosphere's CO_2. Figure 4a shows the warming transition at the end of the last ice age, starting about 17,000 years ago, during which the CO_2 level rose from about 190 ppm to 270 ppm, or about 80 ppm, in 6,000 years.[8] This yields an average rate of change in the "natural" rise of carbon dioxide of 1.33 ppm per century. Turning to modern times, the Keeling curve (http://scrippsco2.ucsd.edu/program_history/keeling_curve_lessons.html) measured at the Mauna Loa Observatory, where researchers have been measuring carbon dioxide in the atmosphere since 1958, demonstrates that the atmospheric CO_2 level has increased 50 ppm during the past three decades. The slope of the curve indicates the current rate of change is approximately 1.6 ppm per *year*, over 100 times faster than the "natural rate of warming" experienced during the rise from the last ice age to our current interglacial period.

So we see that both the absolute value of CO_2 and its time rate of change are radically different today from historical oscillations, meaning that the Earth's climate is dramatically out of equilibrium in an unnatural way. More unsettling still, a recent study by Shell Oil[9] suggests that with very aggressive global efforts, the level of CO_2 might be held to "only" 550 ppm by 2100, 83 percent higher than the Earth's atmosphere has seen in 800,000 years. Susan Solomon, one of the world's leading atmospheric scientists, has carried out calculations showing that global warming will only slowly decline over the following 1,000 years.[10,11] Clearly, we are facing an unprecedented challenge in this new century.

HOW CAN THE ICT COMMUNITY HELP?

The Smart 2020 report estimates that our ICT industry contributed about 2 to 3 percent of the

total global GHG emissions in 2007, growing at a compounded rate of approximately 6 percent, even assuming efforts to lower the industry's carbon intensity over the next decade. This means that the total emission will roughly triple between 2002 and 2020. The graphs in Figure 5 include methane, nitrous oxide, and other greenhouse gases, combined into a "CO_2 equivalent" figure. The dark green in the figure shows the life cycle emissions that are associated with making our equipment and then disposing of it. The light green represents the carbon dioxide equivalent associated with generating the electricity needed to operate and cool all our ICT equipment. In forming its 2020 projections, the Smart 2020 study takes into account both the likely technological improvements in energy efficiency over the next decade, as well as detailed projections of adoption rates over various forms of ICT around the world.

The Smart 2020 study shows that all but 14 percent of the ICT emissions in 2020 will occur outside the US and Canada, with China alone emitting twice this level. Clearly, the efforts to reduce the ICT emission intensity will require a global effort.

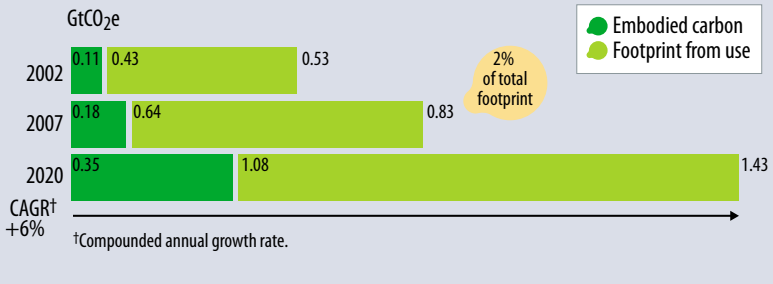


Figure 5. ICT carbon footprint worldwide. ICT emissions are increasing at 6 percent annually, with most of the increase in developing countries. Adapted with permission from the Smart 2020 report.

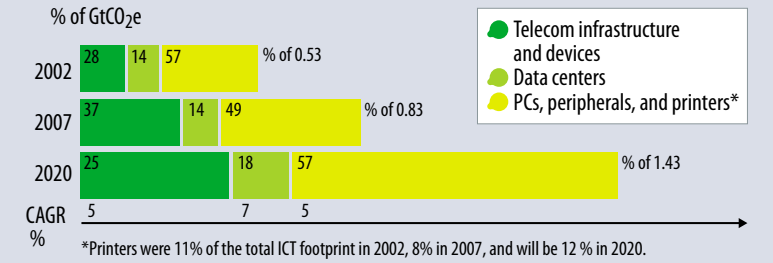


Figure 6. ICT carbon footprint by industry sector. The number of PCs (desktops and laptops) globally is expected to increase from 592 million in 2002 to more than 4 billion in 2020. Adapted with permission from the Smart 2020 report.

Which ICT sectors?

The report also divides the emissions into the component parts: emissions resulting from the fixed and mobile telecommunications/Internet infrastructure, data centers, and the edge of the network. Much of the attention in green IT discussions focuses on data centers, whether located in academia and industry or forming the "back end" of the Internet such as those deployed by Google, Amazon, Yahoo, and Microsoft. This makes sense, since these "superclusters" are measured in hundreds of thousands of PCs. Yet the Smart 2020 report shows that this only adds up to less than 20 percent of the total emissions in 2020. The majority (57 percent) will come from the Internet's edge: PCs, peripherals, and printers. This is because of the enormous scale as China and India rapidly adopt PCs. By 2020, the report estimates there will be 4 billion PCs in the world. So the vast number of PCs is going to dominate this problem.

Cleaning up the edge

Addressing the problem of power management in edge devices needs to start at the system level, focusing on the integration of the hardware and software architectures (http://scipm.cs.vt.edu). It requires coordination across processing, communications, and networking in a modern environment that includes PCs, servers, laptops, and smart phones. All of these devices have complex architectures, including multiple radios, ASICs, microprocessors, DSPs, memories, batteries, AC/DC converters, disks, and displays. This makes monitoring and management of energy a hard problem.

There is wide variance in a component's energy consumption depending on whether it is asleep or active, as much as 6 to 10 times, and radios have an even wider variance. So how can we exploit this for improved energy efficiency? One strategy is to deploy the devices that use the least energy to shut down the bigger ones when they're not needed. For instance, it's possible to coordinate between radios—Wi-Fi, cellular Internet, Bluetooth, Zigbee, and so forth—and use them to page each other to keep the system energy efficient.

The real help here is continued miniaturization, so that there is room to add sensors, control systems, and actuators throughout the system layout to provide data that can feed energy algorithms which attempt to attain the thermal limit of what the devices can achieve.

As UCSD's Rajesh Gupta has shown, from an algorithmic point of view, a power-aware architecture either shuts down a component using dynamic power management or slows down using dynamic frequency scaling, or both. As a demonstration of this, working with Microsoft Research,

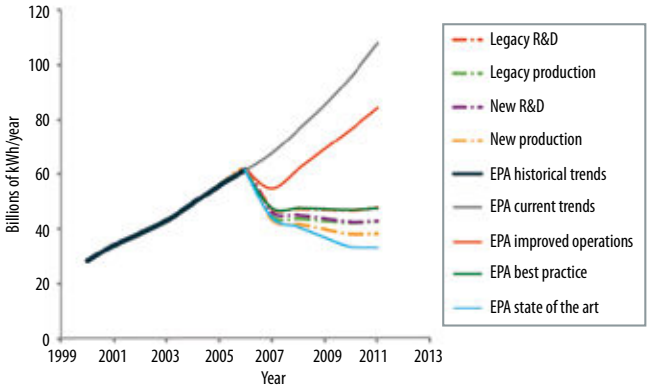


Figure 7. Annual electricity use in US servers and data centers. Source: *EPA Report to Congress on Server and Data Center Energy Efficiency*, 2007.

Gupta and his team developed the Somniloquy architecture (http://mesl.ucsd.edu/yuvraj/research/documents/somniloquy-NSDI09-yuvraj-agarwal.pdf). An implementation of this approach, housed in a standard USB device and inserted into a ThinkPad laptop, manages radios and hardware components so that it's possible to power down to only 1W doing normal work and achieve 63 hours of battery life—as opposed to the four to six hours available with normal (16W) or low (11W) power strategies. Either widespread use of such devices or engineering this capability into edge devices themselves could have a major impact on reducing the estimated 2020 emissions for ICT.

Data centers

The professionals who run data centers have made major improvements in energy intensity in the past four years. As an example, consider the findings of the Data Center Demonstration Project, launched by the Silicon Valley Leadership Group and Lawrence Berkeley National Laboratory, with 17 case studies.[12] As Figure 7 shows, they find that with best-practice methods, the inexorable rise in the kWhr/year that had been taken for granted can be reversed and indeed rapidly decreased. Fundamentally, this involves a series of techniques that move traditional data centers from a strategy of cooling the entire room to one of cooling just the heat-generating processors.

One innovation for small data centers is to enclose the racks in a box not much bigger than the equipment. For instance, Sun Microsystems has created a Modular Data Center (www.sun.com/products/sunmd/s20) out of an international cargo container that can hold seven racks with both air and chilled-water cooling, adding sensors that monitor temperature and energy and allow for active management of disks, CPUs, routers, and so on. The NSF-funded GreenLight project at UCSD (http://greenlight.calit2.net) has purchased two of these data centers to explore methods that go beyond simple physical cooling to the role of software and the applications themselves for increasing energy efficiency.

To provide a realistic load on the system, Calit2 has pulled together a wide range of computational science applications. For example, an end user doing metagenomics can use this service-oriented architecture to run an application remotely, choosing a variety of algorithms and running them on different computer architectures—multicore, GPU, FPGA, and so on—each combination of which has a different energy profile and turnaround time that can be measured and published in an open fashion on the Web. Project researchers are also developing middleware that automates the optimal configuration of hardware and software.

As Tajana Simunić Rosing and her colleagues at UCSD have shown,[13-15] dynamic power management coupled with machine learning based on the outcomes of the sensors and performance counters can control voltage and frequency, achieving up to a 70 percent energy savings for a certain class of workloads. For dynamic thermal management, machine learning can predict that a certain algorithm, say a graphics algorithm, generates significant heat when it reaches the GPU, so the system can precool the GPU to immediately transfer the heat—with up to a 60 percent reduction with no performance hit. These are pioneering examples of how thinking about the interaction of software and hardware can achieve even higher levels of energy efficiency than just thinking about the hardware.

Another example is virtualization. On a typical campus, departments have compute clusters located in poorly air-conditioned rooms and often are only computing a small fraction of the time, even though the electricity to run and cool them is running 24/7. Virtualizing the workload to run on a larger system that is enclosed in an energy-efficient environment and that runs 80 percent of the time allows many more calculations for the same amount of energy. If the end user's laboratory is connected to the centralized cluster with a 10-Gbps clear channel optical fiber as in the GreenLight project, there will be no more latency than if the cluster was in the user's lab, yet the campus will be spending less on energy and lowering its carbon footprint.

Finally, if the energy can be generated in a manner that does not produce carbon emissions, then we end up in the best of all possible worlds—computing with zero carbon emission. Since campuses are beginning to install zero-carbon energy sources such as solar panels or fuel cells, why not use them to power the data centers? Even better, since these sources produce DC, we can use the power source to directly power the computer (which natively runs on DC) and save the wasted energy that goes into AC/DC conversion.

At UCSD, we are installing 2 megawatts of solar power cells, and next year we're going to launch a 2.8-megawatt fuel cell that liquefies methane produced at the Point Loma waste treatment plant and uses it as fuel. This process produces no carbon dioxide—in fact, it recycles the methane that would normally be released into the air—and we could run 10 or 20 Sun Modular Data Centers from this one fuel cell. As part of our GreenLight project we are exploring this option with Lawrence Berkeley Laboratory.

Whether in a laptop or a data center, we're wasting a large fraction of the energy we are using to power and cool ICT devices because we haven't focused on how we can more efficiently perform our calculations. However, the many experiments around the US and the world are a positive sign that this will change soon.

All of this effort is intended to develop ICT components that use less energy, so that we can use more ICT to build out smart infrastructure in electric grids, transportation systems, logistic systems, and buildings. The Smart 2020 report shows that such applications of ICT could reduce global carbon emissions by five times the amount of the entire ICT sector. We have a great opportunity in academia to explore these possibilities and transfer our innovations to society at large, because our campuses are essentially small cities, which can be thought of as testbeds[16] for exploring changes that lead to a greener future. This process is already under way at many campuses (www.presidentsclimatecommitment.org).

With academia becoming first movers, they can drive innovations that will be transferred to the market and applied at scale, helping to speed society's transition from a high-carbon to a low-carbon economy.

References

1. V. Ramanathan and Y. Feng, "On Avoiding Dangerous Anthropogenic Interference with the Climate System: Formidable Challenges Ahead," *Proc. Nat'l Academy of Sciences* (PNAS 08), vol. 105, no. 38, pp. 14245-14250, 2008; www.pnas.org/content/105/38/14245.full.pdf+html.
2. T.M. Lenton et al., "Tipping Elements in the Earth's Climate System," *Proc. Nat'l Academy of Sciences* (PNAS 08), vol. 105, 2008, pp. 1786-1793.
3. C. Lombardi, "NASA Images Show Thinning Arctic Sea Ice," CNET News, 7 Apr. 2009; http://news.cnet.com/8301-11128_3-10213891-54.html.
4. "Ice, Snow, and Water: Impacts of Climate Change on California and Himalayan Asia;" http://esi.ucsd.edu/esiportal/images/cambridge/ucsd_cambridge_gwi-conference_05-2009_press_declaration_final.pdf.
5. http://maps.grida.no/go/graphic/water-towers-of-asia-glaciers-water-and-population-in-the-greater-himalayas-hindu-kush-tien-shan-tib.
6. D. Luthi et al., "High-Resolution Carbon Dioxide Concentration Record 650,000-800,000 Years before Present," *Nature*, vol. 453, 2008, pp. 379-382.
7. A.P. Sokolov et al., "Probabalistic Forecast for 21st Century Climate Based on Uncertainties in Emissions (without Policy) and Climate Parameters," *J. Climate*, vol. 22, no. 19, pp. 5175-5204, 2009; http://globalchange.mit.edu/pubs/abstract.php?publication_id=99.
8. E. Monnin et al., "Atmospheric CO_2 Concentrations over the Last Glacial Termination," *Science*, vol. 291, 2001, pp. 112-114.
9. Shell Oil, "Shell Energy Scenarios to 2005"; www-static.shell.com/static/public/downloads/brochures/corporate_pkg/scenarios/shell_energy_scenarios_2050.pdf
10. S. Solomon et al., "Irreversible Climate Change Due to Carbon Dioxide Emissions," *Proc. Nat'l Academy of Sciences* (PNAS 09), vol. 106, no. 6, pp. 1704-1709; www.pnas.org/content/106/6/1704.full.
11. D. Archer, *The Long Thaw: How Humans Are Changing the Next 100,000 Years of Earth's Climate*, Princeton Univ. Press, 2009.
12. Silicon Valley Leadership Group, "Data Center Energy Forecast," 29 July 2008; http://svlg.net/campaigns/datacenter/docs/DCEFR_report.pdf.
13. G. Dhiman and T. Šimunić Rosing, "System-Level Power Management Using Online Learning," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems,* May 2009, pp. 676-689.
14. A. Coskun et al., "Static and Dynamic Temperature-Aware Scheduling for Multiprocessor SoCs," *IEEE Trans. Very-Large-Scale Integration Systems*, vol. 16, no. 9, pp. 1127-1140.
15. A. Coskun et al., "Utilizing Predictors for Efficient Thermal Management in Multiprocessor SoCs," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 28, no. 10, pp. 1503-1516.
16. B. St. Arnaud et al., "Campuses as Living Laboratories for the Greener Future," *Educause Rev.*, vol. 44, no. 6, 2009, pp. 14-33.

Larry Smarr *is the Harry E. Gruber Professor in the Department Computer Science and Engineering of the Jacobs School of Engineering at the University of California, San Diego, and is the founding director of the California Institute for Telecommunications and Information Technology, a UC San Diego/UC Irvine partnership. His research interests are in high-performance cyberinfrastructure, climate change, and green IT. He received a PhD in physics from the University of Texas at Austin. He is a member of the National Academy of Engineering and a Fellow of the American Physical Society and of the American Academy of Arts and Science. Contact him at lsmarr@ucsd.edu and follow him on twitter as lsmarr.*

Vint Cerf and Munindar Singh, guest editors for *IEEE Internet Computing*'s January/February special issue, have invited a dozen computing luminaries to discuss their predictions for the Internet's future over the course of this new decade. For additional information, see computingnow.computer.org.