systems dissipate it by interrupting us and giving us opportunities to get sidetracked.

- Concerns about privacy are a barrier. An S&R application will make individuals and organizations more effective; however, the company that hosts the application will know the most important aspect of its users — their goals.
- Security is a major hurdle. Widespread use of sensors and responders gives hackers multiple points of entry into S&R systems. The systems that form the backbone of critical services such as food, water, energy, and finance are likely to have common components; successful attacks or errors in these components will have devastating consequences.
- S&R systems enable efficient use of limited infrastructure, such as electric grids and roads, by distributing demand over time and reducing peak congestion. As a consequence, the infrastructure operates close to capacity much of the time, and an increase in demand can take it over the edge and bring the system down. Resilience requires some spare capacity as well as S&R technology.

Society will feel the impact of S&R technologies in many ways. S&R systems will let people conduct a variety of new services from anywhere. They'll let nurses in Manila monitor senior citizens in Manhattan, and engineers in Bangalore monitor intrusion into buildings and networks in London. S&R technologies will accentuate the digital divide; those who master the technology will function better at school and work than those who don't.

The next 10 years will see rapid development of S&R technologies in applications that touch the daily lives of people across the globe.

K. Mani Chandy is the Simon Ramo Professor at the California Institute of Technology. His research interests include sense and respond systems, event-driven architectures, and distributed systems. He blogs at senseandrespond.blogspot.com, and he has coauthored a recent book called *Event Processing: Designing IT Systems for Agile Companies* (McGraw Hill, 2009). Chandy has a PhD in operations research and electrical engineering from the Massachusetts Institute of Technology. Contact him at mani@cs.caltech.edu.

The Play's the Thing

R. Michael Young
North Carolina State University

For most of the 20th century, our entertainment media — film, music, novels, and TV — were happily non-interactive. But a significant shift in the past 30 years toward interactive entertainment has built the computer game industry into a powerhouse that generates more than US$19 billion in annual revenue worldwide, rivaling both music sales and box office receipts. For most of this industry's history, games were primarily designed to be played alone, but even this has changed, with the single-player focus shifting in the past five years to exploit the increase in broadband availability and include additional players.

As computer and console games continue to exploit new services available via the Internet, the design of gameplay itself will correspondingly change. These changes will expand the already powerful social and cultural roles that games play as well as enable the development of new core game technologies involving 3D graphics, real-world/augmented reality interfaces, and artificial intelligence.

Playing in the Cloud(s)

From a market perspective, it's the players' desire for social connectivity that will drive the coming shift to networked gameplay. Already, developers of major game titles are marginalizing their single-player modes and focusing their development efforts on enhancing their networked multiplayer offerings. In fact, some high-profile games are now designed exclusively for online play. Although the shift toward network-enabled games is currently motivated by the desire to enhance gameplay with a social element, the added computational power the shift brings has much broader significance.

We can categorize the kinds of innovations we'll see in game development as a result of the increased access to network services as belonging to one of two types: those that make current high-end game capabilities available across a range of hardware and those that bring new game capabilities into existence. In the former category, we're already seeing early steps to provide compute-intensive game services via the cloud — for instance, by shifting the graphics-rendering process from a player's PC to cloud-based render farms. In these approaches, a game's high-end 3D graphics are produced on remote servers and streamed as video to lightweight clients. In general, approaches like this will add new value to games by migrating conventional game computation from the player's machine to high-end servers, effectively raising the bar for compute power across all users. It will also allow, for instance, high-end virtual worlds, educational simulations, and serious games to run on low-end hardware in schools that lack modern computer laboratories and in the homes of families who can't afford today's high-end hardware.

Even more significantly, this shift to the cloud will provide access to compute services that will enable new types of intelligent tools to add value to games in ways we've only begun to explore. Exciting new techniques are currently being developed that let game engines create game content on-the-fly rather than requiring it to be crafted by hand and compiled into a game at design time. These methods, collectively referred to as *procedural content generation* (PCG), leverage computational models of in-game phenomena to generate content dynamically. Ongoing PCG research projects seek to build systems that can automatically create entire cities, forests full of diverse and unique trees and plants, and novel game character bodies that move smoothly according to anatomical and physiological constraints.

General methods for PCG are computationally costly and so have seen commercial use only in very limited contexts. By moving these functions to the cloud, PCG techniques bring this new functionality to the game client software at almost no cost. Furthermore, the use of cloud-based servers for PCG will promote the development of even more transformative uses of content generation, including complex character dialogue, dynamic 3D camera control, and complex and adaptive story generation. In the future, games that use PCG on remote servers will tailor each player's session to his or her preferences, goals, and context. Each city street you race down, each thug you interrogate, each quest your raiding party embarks on will be created on the spot to provide you with a carefully crafted entertainment experience.

Taking It to the Street

One of the most significant changes in interactive entertainment will arise from the combination of network-centric game services with powerful, pervasive, and location-aware handheld computing platforms and smart phones. This powerful combination will break down the boundary between play and many other aspects of our lives, making entertainment not just accessible during our leisure time but an integral part of our work, social life, shopping, and travel. Thanks to GPS, games running on mobile platforms will not only know who you are, but where you are, letting designers adjust a game's content and challenges to the physical/geographical space in which you're playing.

By relying on network services to manage a game's state, games will be designed to seamlessly slide from cell phone to game console to work PC to home media center as players move from context to context during the day. Social gameplay will be further enhanced by designing games that take into account other players located in the same physical space — for example, when riding on a city bus or touring a foreign city. Services that facilitate the easy creation of and access to location-specific data will make game content creators out of local governments, merchants, civic groups, and individuals. In the near future, your game will adapt to the political history of the village you're driving through, the goals of the anonymous player who's sharing your subway car, and the sale on khaki pants at the Gap that you just walked past.

The two network-centric aspects of games described here — the power of cloud computing and pervasive, location-aware connectivity — will change not just gameplay but will also alter the boundaries between entertainment and what we've traditionally thought of as more serious computing contexts. I expect to see a stronger integration of virtual spaces, information spaces, and real-world spaces. The pervasive nature of online interactive entertainment will push the interface metaphors and user experiences found in games into the broader context of computing. It's clear that those broader contexts will change as a result. The challenge for game designers is to figure out how the broader contexts will, in turn, change games.

R. Michael Young is an associate professor of computer science and the co-director of the Digital Games Research Center at North Carolina State University. He has a PhD in intelligent systems from the University of Pittsburgh and is a member of the IEEE, the ACM, and the Association for the Advancement of Artificial Intelligence. Contact him at young@csc.ncsu.edu.

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.