



Trend Wars

Internet technology

Dejan Milojicic
Hewlett Packard Laboratories
1501 Page Mill Rd., MS 1U-14
Palo Alto, CA 94304
dejan@hpl.hp.com

As we've all seen, Internet technology has become a dominant factor in business, academia, and everyday life. E-commerce and e-services are fueling a market sea change, with practically any business anyplace nurturing its equivalent on the Internet. The change began with advertisements and news, followed by entertainment (with music and movies) and traditional businesses, such as retail sales, financing, brokerage, and auctions. Now, entirely new businesses and services are created on and for the Internet. As a consequence, Internet service providers and application service providers have created a revolutionary new model for driving new applications, hardware, and software development.

More globally, the recent AOL-Time-Warner merger revolutionizes our perception of future business, making us all wonder who will dominate tomorrow's businesses. If an Internet company that was lightly regarded—if not ridiculed and abused for its poor service just a few years ago—can buy one of the pre-Internet giants, where

will the changes stop?

On a different front, much of today's research revolves around the Internet. Networking teams explore how to bring big pipes that last mile to homes, while others explore operating systems within the context of Web servers or look to agent technology for help with Web interfaces and search techniques. Internet research has made itself felt most profoundly through startups. A number of university research projects ended up in products coming from startups, usually as products for the Web, such as Akamai or Inktomi. The open-source community is another relevant factor in this arena, while the Linux operating system and Apache Web server are the most widely represented systems in Web serving.

The Internet's impact on everyday life is also fascinating. The other day, my wife and I received a Web greeting card through Yahoo—animated, too! It is much easier to reach me by e-mail than by phone. Like most of my colleagues, I do not search for academic papers in libraries, but rather on the

Web. There, I also do spell checking, follow my favorite basketball teams on the NBA's Web site, and buy books and computers at Amazon.com and buy.com among many other sites. WebVan offers groceries online in the San Francisco Bay area and other companies offer clothing. "If it is not on out on the Internet, it ain't there" :-).

Unfortunately, the Internet has also opened some dark doors, including privacy invasion, pornography, and information abuse, to name a few—problems that existed in the presence of traditional technologies and that will have to be fought with traditional and new ways.

In this installment, we feature some of the people who have brought the Internet technical revolution to you: Erik Brewer of Inktomi and the University of California at Berkeley, Fred Douglass of AT&T Labs-Research, Peter Druschel of Rice University, Gary Herman of HP Labs, Franklin Reynolds of Nokia, and Munindar Singh of North Carolina State University.

—Dejan Milojicic



Eric Brewer

What were the decisive turning points for Internet and Web technology to become ubiquitous and pervasive?

For me, the turning point for the Internet was simply seeing lots of people have e-mail access. When I think about the turning point for the Web, I remember when I interviewed for a job in 1994 and I realized

that it was a big advantage to post my resume, even though not all my interviewers knew how to use the Web. In fact, I remember giving a talk at MIT where I tried to convince other students to use the Web before they went out and interviewed, because it would be an advantage to have some experience with it and because I felt it made it easier for the interviewers to figure out what it is the interviewee does.

Back then, the Web was a mixture of very formal things, like technical specifications, and informal things like people's home

pages. And of course, the need for a search engine was pretty obvious early on, given the Web's kind of wonderful anarchy. I also felt that although you could find things if people told you where they were, that was about the only way to find them. I felt automated search was fundamental to the Web's long-term success.

What are the next likely disruptive technologies in Internet space that might make new marks in the way we live and work?

The first is definitely wireless technology. I've spent a lot of the past five years working on wireless communications. Not so much the wireless part, but the services part. I'm a big fan of any device that has access, and usually that means wireless access, just for ease of use. I'm also a big fan of any kind of physical integration with the Internet. It takes a lot of forms, and FedEx uses one of the key ones: tracking. You can use the Internet to check where your FedEx package is anywhere along its route. That's a nice integration of the physical world with the virtual world. I think there's a lot more of that coming. People talk about more esoteric things like robots, but I actually think simple things such as sensors will become common. I like being able to check the traffic on the freeway using a camera so that I know first-hand how the traffic is, much more than trying to wait 10 minutes for the radio report to come on, which might or might not tell me what I want to know.

So you see more physical and real-world integrations?

Exactly. It would be nice if my stereo volume would go down when the phone rings. That's relatively easy to do technically, but that integration is nonexistent, or at least rare, in the present day. And even other kinds of physical integration such as Webvan, which is essentially Internet ordering and physical delivery. It's a simple but very powerful combination. I'd like it to be a lot less obvious where the virtual world ends and the physical world starts. I feel like that this will be a major transition for the next five years, and maybe one of the most rewarding.

What are the most important technologies that will determine the speed and the direction of the Internet's future?

Most people take fast networking for granted. I actually think it's more important to be continuously connected. So the reason I like DSL is not so much that it's faster, but because it's always on. And one of the reasons I like wireless technology is again, because it gives me access from any location at any time. The always-on principle is probably the most important thing that will push the Internet, because it will change the way people interact with it.

What about storage? Many people bring up storage as an important factor.

One of the most interesting changes on the Internet in the past few years has been the deployment of an incredible number of caches, which are essentially storage at the edge of the network to make the network faster and more reliable. It's been, I think, wildly successful as a technology. But it's inter-

Internet technology questions

1. In retrospect, what were the decisive turning points for Internet and WWW technology to become ubiquitous and pervasive?
2. What are the next likely disruptive technologies in Internet space that might make new marks in the way we live and work?
3. What are the most important technologies that will determine the future Internet's speed and direction?
4. Where do you see the roles of industry, startups, research labs, universities, "open source" companies, and standard organizations in shaping the future Internet?
5. What will be the major application areas dominating the Web?
6. What is the most controversial and unpredictable technology in the Internet space?

esting to me because that storage is transient. It doesn't really belong to anyone in particular. I like the notion that if I have storage in the network, I can access that storage from anywhere, using the always-on principle. That's why I expect most data over time to migrate into the infrastructure—because it's more reliable, more available, accessible from anywhere, and probably cheaper.

Where do you see the roles of industry, startups, research labs, universities, open-source companies, and standards organizations in shaping the Internet's future?

I think universities have to work on a much longer horizon, because a startup or an industry development group or industry research group will do anything that's on a short horizon better and faster. Startups, though, rely on timing. It might be that most startups fail because they're too early for the technology, not because they're too late. And though it's fine for the university to be too early, it's bad for a startup to be too early. For example, the Apple Newton was too early. It was a good idea in many ways—it even has some nice properties that are still not in the Palm Pilot—but their timing was off. So startups are about timing, and research labs are about having a vision that exists five years or farther out. And even if they don't achieve that vision, they should find lots of interesting things along the way.

Do you see long-term opportunities for open-source companies?

Absolutely. Most of the stuff we're doing at Berkeley has been and will continue to use some flavor of open source. Certainly large parts of Unix generally fit that model. I like it because it's about cooperation for the benefit of many people. And to the extent that I believe there's a lot of value in operating systems and in packaging systems that are robust, there will continue to be lots of room to add value for companies, even if their source code is public or if their source code belonged to somebody else originally.

What major applications will dominate the Web's future?

Finding information will continue to be the most common application. Variations of that will certainly be important;

e-commerce is essentially about finding information about products for businesses rather than for consumers. Communication will also dominate the Web's future, including instant messaging and chat. It's remarkable how many more people at Berkeley use instant messaging than people at companies do. So I think in some sense its impact is still on the way, even though something close to 40 million people use it already. Those classes of things—chat, instant messages, group white boards, group conferencing—are really going to be key application areas.

What is the most controversial and unpredictable technology in the Internet space?

I'm going to go out on a limb and say it's voice over IP—in part because I think it's not well defined. The notion that you can use packets to send voice I completely believe in, but the idea that you can do that and maintain the tight latency bounds that two-way voice requires means that you have to have a much more controlled environment than the public Internet. So voice over IP I believe literally; voice over the Internet I'm much more skeptical about. I think a lot of the trouble in that area will come from the fact that people aren't distinguishing IP versus public Internet. It's not the most controversial topic, but I certainly think it will be slower to come than people think.

What do you think about privacy?

I've been working in the privacy and security area as a researcher for several years now, and I believe that we should have a range of security and privacy options. This comes down to a need for an anonymous or pseudo-anonymous infrastructure, which hasn't really arrived yet. More importantly, it speaks for a need for very explicit rules or at least disclosure about what companies can do with the information they have about you. Most of that doesn't exist yet, and in the sense that it does exist, it's voluntary.

A privacy policy is a good start, but I don't think it's even clear what kinds of privacy policies there are, what classes they fit into, and what is and isn't acceptable behavior. There's also a notion that there is a fundamental right to privacy, and it might extend quite well into the digital world. For example, if I have data that I view as private, and I put it somewhere, and it's encrypted, is it subpoenaable? Does the government have a right to anything I store digitally? I'd like the answer to be "no" for the same reason that I think the things in my brain that I haven't written down yet are also inaccessible. This is a social discussion that'll take place over the course of the next decade, and it might have different answers in different cultures, which will be one of the interesting things to watch.

Eric A. Brewer is an assistant professor at the University of California, Berkeley in the Department of Computer Science. He cofounded Inktomi Corporation in 1996, and is currently its chief scientist. He received his BS in computer science from UC Berkeley, and his MS and PhD from MIT. He is a Sloan Fellow, an Okawa Fellow, and a member of Technology Review's TR100, Computer Reseller News' Top 25 Executives, Forbes "E-gang," and the Global Economic Forum's Global Leaders for Tomorrow. Contact him at 623 Soda Hall, UC Berkeley, Berkeley, CA 94720-1776; brewer@cs.berkeley.edu; www.inktom.com.



Fred Douglass

In retrospect, what were the decisive turning points for Internet and WWW technology to become ubiquitous and pervasive?

I think everybody agrees that Mosaic was the point at which the Internet went from being largely a research tool to being something that was used by the masses, because it provided a nice GUI and a way to get all the graphics that simple text and HTML didn't have.

What are the next likely disruptive technologies in Internet space that might make new marks in the way we live and work (agent technology, voice/text recognition, wireless, ...)?

I think that Mark Weiser had it right; it's a shame he passed away before he could see his vision realized. His vision of "ubiquitous computing" is finally at a point where we're seeing it actually happening. We're seeing big companies like IBM put a large effort behind it—what they call "pervasive computing." The requirements for ubiquitous computing, which are the cost model and the connectivity model, are such that everybody will have these computers in their homes. Voice and text recognition is another issue as far as ubiquitous access to computers. I'm hearing from people in AT&T Labs that speech recognition and speech synthesis are much more natural than before: the ability to interact with the computer in a kind of Star Trek mode is something else we'll see soon. And that, too, is going to make a huge difference, because if you can walk up to your computer and talk to it, then that brings in another level of ubiquity.

What are the most important technologies to determine the Internet's future?

Number one is broadband access to the home. Among my coworkers, virtually everybody has a cable modem if they're in an area that's served by cable modems. If they're not, then they have ISDN or DSL. One of the heads of the Yankee Group who spoke in front of AT&T about a year ago said, "Once you have a cable modem, you'd sooner give up your firstborn child than lose it." I use that quote because it's exactly what happened to me: I moved from an area where I had a cable modem to an area where one wasn't available, and I've been suffering for the last 10 months. I'm not giving up any of my children, thank you, but it really does revolutionize how you use the computer and the Internet from home. High-speed access also means both new applications and an ever-increasing load on networks and servers.

What about wireless? Are there any hard limits that could be crossed to provide a qualitatively better way of doing something?

It's all gradual. I don't think there's a threshold that we're coming up to, except for certainly the distinction between the POTS [Plain Old Telephone Service] line and anything else—when you jump from a POTS line to a cable modem or DSL or something. There's work on other things, like 3G, which are

attempts to make wireless connectivity available with the kind of access that would give you, let's say, multimedia when you're in the field—but that's down the road.

Where do you see the roles of industry, startups, research labs, universities, open source companies, and standards organizations in shaping the future Internet?

It's clear that startups are driving a lot of what's happening right now. We're seeing companies such as Akamai come in and have an enormous impact in a particular market in record time—and the time from startup to IPO is getting shorter and shorter. We see this in lots of different areas, where start-ups are taking the roles that industrial research labs and academia basically had to themselves until recently.

This also means that start-ups are providing a lot of competition to these environments in terms of recruiting the best talent. It doesn't mean that industrial research labs and academic departments don't still offer a lot and attract good people, because they do, but it's different. Start-ups tend to offer more immediate real-world impact, while labs and universities have more freedom to do long-term research.

What about open-source companies?

I guess the biggest example of this is Netscape, which decided to make its browser open source. There was a long time when people had to basically write their own browser or use a very old version of Mosaic if they wanted a browser that they could modify in some fashion to, let's say, record traces or add some new functionality. So having Mozilla be open-sourced is wonderful for the Internet community.

It doesn't, however, seem like Netscape's being open-source has taken off in the same way as Linux. But even Linux is still a bit player in comparison to Windows, which is unfortunate. We're finally starting to see some cross-platform compatibility between the two with things such as Star Office and VMware, which will let you run Windows applications on Linux at a penalty. But that's a relatively recent development.

What major application areas will dominate the Web in the future?

I'm sure that there will be new applications down the road, but for the foreseeable future, today's applications are the same as future applications—e-commerce, pornography, interactive chat rooms, and games strike me as the big ones. As bandwidth to the end users—home users as opposed to corporate users—increases, then interactive video will certainly become more commonplace. And other sorts of things such as Internet telephony will become the norm as opposed to the exception.

How would Internet telephony impact big phone companies such as yours?

It's clear that the long-distance market, excluding the impact

of Internet telephony, is becoming more and more competitive, and the big phone companies are all making much less money at it. Of course, the costs are going down, too. Internet telephony was not even really on AT&T's radar for some time. Now AT&T, like everybody else, is moving to support it. In the end, telephony will be just one of a number of Internet services that we'll support, but certainly long distance as a money maker will have much less of an impact for AT&T and for the other big phone companies. Already, data dominates voice traffic on our networks.

What about charges on the Internet—for example, for anything including Internet telephony or any other applications?

I've questioned for a long time the model that everybody gets a sort of fixed pipe, where the end users don't pay based on what they do, but instead go on a flat rate. Ultimately there will be charges to give end users the best performance. You could sign up with some other network that charges you a little bit more but gives you better guarantees. We see this already, to some extent,

because at the ISP level there are settlement charges and things when there's an incompatibility between how much traffic one ISP is sending to another. But that doesn't tend to affect the end users.

How do you perceive security? What do you think will happen in the future with the right to privacy and governmental intrusion?

Advertising is driving much of what's happening on the net right now. For advertising to work, it needs some information about individual users. The idea is that when I go to a Web site, a company can make much more money if it can tailor its advertisements to me. At the same time, I really don't like the fact that people know all this stuff about me. There are special tools that intercept and anonymize cookies to try to keep people from finding out many of the details about a given person. There are also standards evolving regarding the guarantees a particular content provider can make as far as the privacy of the information provided to them.

The tension between all these companies throwing money at this, and a lot of these sites and the services that we have right now, wouldn't exist without advertising money. But as the tools evolve to separate that out—for example, eliminating advertisements completely—costs could rise. It's just like television, where it's possible to have a VCR that skips over advertisements. If commercials aren't being recorded, and if that becomes the norm rather than the exception, then the companies paying for commercials right now and providing broadcast or cable television would stop paying, moving us to a model where everything would be by subscription.

It's the same thing for Internet services. One model is that you go to a search engine and it shows you a bunch of ads while answering questions, and you don't pay anything. Another model

I question the model that everybody gets a fixed pipe, where the end users don't pay based on what they do, but instead go on a flat rate.

is that it doesn't show you a bunch of these ads, but you have to pay five dollars a month for the privilege of using the site.

In other words, advertisements are a good thing, as long as they are controlled?

Right now, because they're driving so much of this, they're a necessary evil to some. What I was trying to get at was that banner ads appeared sporadically, started becoming commonplace a few years ago, and are now ubiquitous. Because they've enabled so many things that people otherwise wouldn't have the time and money to support, they've actually done us all a great service. So we need a way to manage them.

I don't have a problem with banner ads when they're useful. To make them more useful, you need information about the people to whom you're providing them. The more you customize the ads, the more you base them on what it is that the Internet users are doing, and the more sensitized people are to the question of the lack of privacy. There are lots of people who are working specifically in this area, but I'm not one of them. I think you could have a whole separate round table on the merits of privacy and how it contrasts with what the providers are doing. But I think it's certainly of crucial importance to the evolution of the Internet.

Fred Douglass is the head of the Distributed Systems Research Department at AT&T Labs—Research. He has taught distributed computing at Princeton University and the Vrije Universiteit, Amsterdam. He has published several papers in the area of World Wide Web performance and is responsible for the AT&T Internet Difference Engine, a tool for tracking and viewing changes to resources on the Web. He chaired the 1998 Usenix Technical Conference and 1999 Usenix Symposium on Internet-working Technologies and Systems, and is program cochair of the 2001 Symposium on Applications and the Internet (SAINT). He has a PhD in computer science from the University of California, Berkeley. Contact him at douglass@research.att.com; www.research.att.com/~douglass.



Peter Druschel

In retrospect, what were the decisive turning points for the Internet and WWW to become ubiquitous and pervasive?

The key technologies that enable the Web developed over many years, and a lot of people made important contributions. The turning point in bringing Internet technology to the masses was the invention of browsers, HTTP and HTML. These were not fundamentally new technologies, but they made the Internet available to a broad audience by providing a convenient user interface. There's a lesson for us to learn here: the key to bringing these technologies to bear was ultimately not in the performance, as measured in transactions per second or bits per second, but usability and availability. That means simplifying all this pow-

erful technology, reducing it to intuitive user paradigms. That's essentially what enabled the Web. As a community, we don't have a good record in realizing that. For us, the command line interface was good enough. In retrospect, most people would agree it was inconvenient, but at the time there was not enough of a driving force for us to make it more usable. We just didn't see the enormous opportunity behind that.

What are the next likely disruptive technologies in the Internet space that might make new marks on the way we live and work?

These kinds of things are extremely difficult to predict, but it is probably going to be paradigms and technologies that make a dramatic difference in the usability and the ease of configuring and maintaining existing information technology. Agent technology and voice recognition and synthesis all have strong potential. I wouldn't be at all surprised if core technologies came out of this space that mark a turning point in the Internet.

What are the most important technologies that will determine the speed and the way of the Internet future, for example, fast networking, user interfaces, computer power (speed/memory/storage), and so on?

Fast networking, user interfaces, and computer power are clearly going to make a big difference and will play key roles in enabling future progress. However, I don't think that these things will, in and of themselves, bring about landmark changes. Such changes must come from new paradigms that can dramatically change or increase the availability and usability of existing technology. A few examples are automatic configuration of ad hoc networks, speech recognition and synthesis, and technologies that let a larger group of novice users take advantage of information technology through a device that behaves more like a household appliance than a computer. These must be devices that will not require people to perform tasks such as configuring and upgrading operating systems and applications, which are annoying and often too intimidating for a lot of people out there.

We have a long way to go in terms of truly allowing everyone, regardless of economics or educational background, to access the Internet. I think it's easy for us to overlook that. In spite of the fact that most of our friends and colleagues have Internet access, there is still a large fraction of the population who just cannot afford to buy a computer or pay monthly ISP charges. Making the technology truly ubiquitous in all parts of the world is going to be an important part in achieving broad coverage. That will require progress both in terms of the costs and usability to make it less expensive and intimidating to buy, use, and install an Internet access device.

Where do you see the roles of industry, startups, research labs,

universities, open-source companies, and standard organizations in shaping the future Internet?

All these organizations must play their respective roles; I cannot see how we can make significant progress without each of them. They are each in a unique position to realize certain opportunities to bring forward new ideas, invent and market the respective technologies, and develop the necessary standards that ensure widespread use.

How would you divide their space of influence? Have there been any significant overlaps?

The role of startups has amplified in recent years. This is the space in which a lot of existing base technologies and ideas have been crystallized and brought together with the respective energy and funding to make them marketable. At the same time, we also need to make progress in developing base technologies that are too long term or too risky for a startup company. That's where the research labs and universities come in. Universities also educate the next generation of technical leaders and technical personnel. The open-source movement ensures broad availability of base technologies, leverages the innovation and energy of the open-source contributors, and ensures diverse platforms.

Do you think there is any other area where the open-source approach can work?

I would like to see a similar movement in applications, in particular, core productivity applications—for example, alternatives to Microsoft Office. That will be of tremendous importance and could open the market and playing field, could open up doors for innovation. It remains to be seen if it is feasible for the open-source movement to take hold of the complexity and the level of sophistication in that application space.

What will be the major application areas dominating the future Web?

The major application areas are broadly electronic commerce, distance education, distance collaboration, and entertainment; those are the areas with the most potential. Identifying the specific applications within that space is going to be a big gamble and hard to predict. Those who have the right ideas in this space and know how to realize them can look forward to fame and fortune.

Games have been a very important factor around the evolution of technology. They draw a lot of new ideas and developments.

They do play an extremely important role, because they inspire the imagination and unleash a lot of creative energy that pushes forward innovation. They also have an important role in connection with the open-source movement, which leverages a lot of raw talent out there.

What is the most controversial and unpredictable technology in the Internet space?

Looking long term (several decades), it is the various types

of biocomputing that people are currently starting to get really interested in. Because of its potentially profound impact on computing and information technology, it will ultimately impact the Internet. In the short term, agent technology—automated intelligence that helps you configure information technology and networks among mobile users, automatically repair and heal disruptions and failures in networks, and automatically adapt to differences in the quality of service at any given time—has the potential to bring a new paradigm into play. Those technologies are most likely to profoundly impact the way we use information technology.

Peter Druschel is an assistant professor of computer science at Rice University in Houston, Texas. He received his MS and PhD in computer science from the University of Arizona. His research interests include operating systems, computer networking, distributed systems, and the Internet. He currently heads the ScalaServer project on scalable cluster-based network servers. Contact him at druschel@cs.rice.edu; www.cs.rice.edu/~druschel.



Gary Herman

In retrospect, what were the decisive turning points for Internet and WWW technology to become ubiquitous and pervasive?

For years, IP traffic was insignificant relative to voice traffic—it was minuscule, just a gnat. Then a number of things happened. The PC business paradigm pushed costs down to the point where the technology became affordable, and Web standards began to work across platforms to enable deployment of an information service with a global footprint. The standards also democratized the creation of information in a way that the videotext and even the online service providers were not able to do: people who had information could publish it, and people accessing information—even those with very narrow needs—could be served by a common infrastructure.

Prior to that, the cost of assembling information was high, and the value of any subset of information was relatively low, so information services struggled to be economically viable. Somehow, PCs, Internet connectivity, modem speed improvements, and a set of cross-platforms all happened at the same point in time. There were more entrants, lower barriers, more competition, and more alternatives—capitalism at its finest.

What are the next likely disruptive technologies in Internet space that might make new marks in the way we live and work (agent technology, voice/text recognition, wireless, ...)?

The most likely disruptive technology is cheap and ubiquitous wireless connectivity. Agent technologies have been pursued for years; voice-to-text has been pursued for years. Those things are hard problems that are addressed incrementally. Wireless seems to be the thing that creates the greatest number of options. Pervasive connectivity, pervasive computing, and the ability to have devices that relate with each other in ad

hoc ways—they create new options for how things work, how people work, and how systems operate.

So the next thing on the horizon could be that wireless becomes a commodity?

Wireless is a commodity. The ability to embed computing and wireless intelligence into everything is the next big change.

Speaking of commodity and client kinds of machines, does it somehow narrow and make obsolete the high-end machines? Once upon a time, people thought that big machines would never go away, but the push continues.

There are classes of information that involve large volumes of data, where big machines solve operational problems. But I think the percentage of big machines will continue to decrease. The one possible trend emerging is very large-scale data centers, which involve tens or hundreds of thousands of relatively simple machines.

The data-center-consolidation era dealt with the operational and administrative problem of having many small machines, which encouraged consolidation into large servers. If the multi-thousand CPU data-center problem is tamed from an operational and administrative point of view, then some of the operational benefits of having large, consolidated servers might go away, and they'd be left to storage and dynamic data problems, where centralization is necessary. You can't easily decentralize dynamic data and deal with the recovery and the administration of changing data. But there are still some big problems that are addressed well in single, logically centralized entities.

What are the drivers of the Internet?

User interfaces. I don't know that it's a technology as much as it's a design art.

The models that people use to deal with this world of computing—whether it's embedded computing or it's more explicit in the desktop sense—become the things that determine the rate at which the Internet permeates everyday life. I guess the counter argument to the importance of user interfaces as a gating factor in growth could be that the current Internet and client technologies are extremely complex, but the value has been so persuasive that people are motivated to overcome the complexity in ways that, five years ago, people would not have viewed as feasible.

I know we've been looking at some of the work that people at places like MIT have done on these contextual, natural language interfaces using speech recognition. The problem there is the design complexity of the interface. The interface is very simple and intuitive, and there's a huge amount of effort in designing it to be that way. So if you were going to have a bus-stop interface or a McDonald's drive-up window interface, each of those contexts is narrow enough that perhaps you could engineer a very natural user interface. It's the engineering complexity that limits our ability to put the technology into those settings, because there are so many and they're so diverse. So

it's either the interface itself or the complexity of the interface design that limits the ability to push the Internet and computing technologies into those more mundane application areas.

Where do you see the roles of industry, startups, research labs, universities, open source companies, and standards organizations in shaping the future Internet?

This is a real challenge for many of the traditional models for innovation and the creation of commercial technology. In the Internet space, the required technical knowledge is taught in universities and used in everyday life. This lets startups aggressively attack every visible, commercially significant problem, and, obviously, they've become very significant suppliers of technology into the marketplace. Moreover, their success appears to be changing the motivations of faculty and students in universities. It certainly creates challenges for traditional corporate research labs.

Industrial research labs need to examine their traditional models for motivating and rewarding research staff. They need to examine their role as technology suppliers in contrast to the role played by startups. I think they need to reach farther into the future, speculate more, assume more risk, and try to create

what (HP CEO) Carly Fiorina has called "strategic dilemmas" for the corporation. That is, create significant future options unforeseen by the product development parts of the company.

The commodity aspect also doesn't help, because you no longer need substantially expensive equipment. You can do everything with off-the-shelf PCs and that kind of technology.

That also means that much of the innovation occurs at the business-model level. It isn't so much that there is some radical new technology, but rather that known technology has been applied to

solve a business problem in a very different way. It can be difficult for industrial labs to engage at that level, because if they remain in a corporate "ivory tower," they tend to be fairly remote from the practical applications technology in a business or industry context. To pursue innovation in such situations, the researchers will need to get into the front lines with the businesses that are looking to apply technology. I think Dorothy Leonard-Barton at Harvard has termed this process "empathetic design."

The challenge for universities is that they now have a conflict of interest, because if they're working in the Internet space, they typically don't get to work with a long-term horizon. If they do valid technology, they're able to personally bring it to commercial use by forming startups. So even at great universities, I fear that some of the faculty is learning to focus on very near-term work with a conscious eye toward commercial and personal financial gains. Universities, then, can be lured into abdicating their traditional role of working on big problems. It also changes the nature of the industrial relationship with the universities—large companies may view funding university research as potentially funding competition, as opposed to fund-

Even at great universities, the faculty is learning to focus on very near-term work with a conscious eye toward commercial and personal financial gains.

ing the creation of fundamental intellectual property of long-term significance. There's a whole different ecosystem that needs to evolve between startups, industrial labs, and universities in the Internet space because the rules have changed.

We'll see how open-source companies fare. They depend on an emotionally committed community of technical talent. Tim O'Reilly recently [see xml.com/pub/1999/10/tokyo.html; this is an xml.com article based on a Linux World Keynote] pointed out that almost everything that's great about the Internet was a result of an essentially open-source process, but it also pointed out why, once such technology becomes important commercially, forces work to destroy the open source's openness. So the democratization of technology—the multiple participants and the rabid competitiveness—has served to greatly accelerate the evolution and commercial application of Internet technologies. It would seem desirable to maintain that quality of rapid competition, but large, commercial interests naturally work to try to constrain it.

What major application areas will dominate the Web in the future?

Web-mediated personal communications will eventually displace the telephony-based model that has dominated personal communications to date. I'm not talking about voice over IP. I'm talking about interpersonal communications applications where the visual parts of the interface are tightly integrated with the aural or auditory parts of the interface, based on the Web, and without the relatively heavyweight use model and underlying infrastructure of telephony. We're starting to see a bit of this, but the infrastructure is not in place. You need persistent IP connectivity to have that work, and that's just starting to emerge.

Another big application area is to have the Web and the Internet start to permeate everyday things, so that things that are commonplace—automotive diagnostics, routine monitoring of appliances and mundane infrastructure services of all types—start to work in a transparent way. The important applications become embedded and transparent, so the user's not aware they're working.

One important use of Internet technologies has to do with the delivery of "passive" media (that is, traditional entertainment and information) to the mass market. The Internet is a potentially good way to push goods and services to consumers—so then there's an immense amount of money that now goes into broadcast, print, and other distribution channels that could flow into computing and communications infrastructure to enable the Internet to perform similar functions. It's hard to predict, because it deals with user behavior and people's ability to use a new distribution medium and new channels to accomplish those ends.

Do you think that it will also bring more charges for Internet use?

I think that the Internet will be more successful if the charging is carried in the cost of goods and services, rather than as an explicit fee you pay.

What about security and privacy? How is the government involved?

Because we live with pretty insecure systems today, we've just chosen to accept their vulnerabilities and the consequences. People's anxieties about privacy, however, are very powerful.

They're different in Europe than in the US. It's actually counterintuitive in some ways. I'd say the biggest issue for the US government deals with taxation. Too much of our governmental infrastructure depends on sales taxes associated with the physical presence of the purchaser and the distributor. Once that's all been virtualized, all the revenue dries up. That can't happen, so they're going to have to tax. I just don't know how they'll get around to doing it.

Gary Herman is director of the Internet and Mobile Systems Laboratory in Hewlett-Packard Laboratories, Palo Alto, CA, and Bristol, UK. His organization is responsible for HP's research on technologies required for deploying and operating the service delivery infrastructure for the future Internet, including the opportunities created by broadband and wireless connectivity. Prior to joining HP in 1994, he held positions at Bellcore, Bell Laboratories, and the Duke University Medical Center. He received his PhD in electrical engineering from Duke University. Contact him at Hewlett-Packard Laboratories 1U-19, 1501 Page Mill Road, Palo Alto, CA 94022; herman@hpl.hp.com.



Franklin Reynolds

In retrospect, what were the decisive turning points for the Internet and WWW to become ubiquitous and pervasive?

One of the first decisive events happened in the 1980s when Unix and TCP were made available on most workstations and minicomputers. The Berkeley Software Distribution was particularly influential. Instead of a proprietary solution like SNA or Decnet, different vendors began to deliver Unix- and TCP-based networks on many different platforms. TCP/IP was available with very little effort, which provided a business motive that was separate from any technical reasons for distributing and adopting the technology. Unix was particularly popular in universities and organizations affiliated with the US Department of Defense. The early Internet users were predominately in the DoD or universities. Before the widespread deployment of IP-based networks, Usenet was the most important large-scale Unix network. Electronic mail, news groups, and other Usenet applications were quickly developed for IP. The Internet grew rapidly and began to replace Usenet.

I remember that in the mid 80s a lot of people thought TCP was "not quite real" and the OSI protocols were "industrial strength." Many governments, including the US government, announced plans to make OSI support mandatory for future procurements. But it turned out that OSI's technical advantages were not compelling (some would argue that they did not exist). The Internet's rapid growth and the proof of the underlying technology's interoperability and scalability created an insurmountable barrier to the widespread deployment of OSI protocols.

Linux also bears some contributions to the Internet because all the Web servers are hosted by it.

I think that very few people would have predicted the current commercial success of Linux. Certainly, I never would have guessed that Linux would enjoy the kind of success it has.

The ability to use Linux in a commercial environment would have been scoffed at by most people five years ago.

I don't think we should understate the importance of Tim Berners-Lee and the World Wide Web. The Web has popularized the Internet. However important computer-to-computer communications become, the Web has made interactive use of computer networks interesting to millions of people. If the Web had not been invented, however, something else would have come along. It is easy to forget how quickly things such as Usenet, Gopher, Archie, and online bulletin boards gained popularity before they were all swept away by the Web.

Perhaps, the most important factor in the Internet and the WWW becoming ubiquitous has been the continuing improvements in performance and reduction in price of personal computers and network technology. As the price comes down, it becomes reasonable for more consumers to test drive the Internet and the Web. User population growth provides opportunities and incentives for the development of new uses of the Internet. These new applications and services attract more users, which in turn.... Well, you get the idea.

What are the next likely disruptive technologies in the Internet space that might make new marks in the way we live and work (agent technology, voice/text recognition, wireless, and so on)?

My favorite candidates include large-scale wireless networks, such as third-generation cellular, and small-scale wireless networks, such as Bluetooth. I also think really cheap mass storage and tiny, ubiquitously deployed, embedded computers coupled with small-scale wireless networks will be important. These technologies will make possible new applications that change the way we use computers and networks.

Another interesting technology with tremendous potential is mobile code. Possible uses range from low-level network infrastructure, such as active networks, to high-level application platforms, including mobile-agent systems. If we include Java applets and JavaScript, then applications based on mobile code are already an important part of the Web. The most important characteristic of the various mobile code platforms is the ability to deploy new behavior or functionality on demand. There are still some hard problems associated with security and performance, but as mobile code technologies mature, their influence will grow.

What will be the major application areas dominating the future Web?

In the near term, the next five years, there are some safe bets: Digital network convergence (telephony and data and broadcast networks) and electronic commerce will undoubtedly continue to grow. Going a little further out on a limb, the Web will revolutionize education and training. Not that the techniques and technology used will be dramatically different than what we have today. What will be important is that it will be available to everyone—especially for the economically disadvantaged.

If we have really cheap mass storage, we will start to record everything. Digital cameras and recorders will be everywhere and will not be limited by storage capacity, because they will use the wireless network to ship the data to storage farms. The motivations for recording everything will include personal security, liability management, curiosity, and so on. We will

record the office, classroom, shopping mall, lawyer's office, police on the beat, construction sites, operating rooms, nursing homes, playground, and almost anything else.

Pervasive wireless networks will make it possible for mobile users to be always on and always connected. Ubiquitous embedded-computing devices will make it practical to deploy smart environments, which will lead to personal digital assistants becoming the most important Internet access device for most people. PDAs will provide personal information management applications, games, Web access, storage farm access, video and audio recording, communications (telephony and broadcast), and universal remote control. Think Star Trek communicator and tricoder.

Where do you see the roles of industry, startups, research labs, universities, open-source companies, and the standard organizations in shaping the future Internet? What do you think about this gold rush in the Silicon Valley? Will it improve the future Internet, or is it just muddying the water?

It seems to me that there has been a gradual shift away from corporate research. Fewer companies fund significant internal research groups, and the work of most industrial research groups is more akin to advanced product development than research. Most research is conducted in academic settings. Startup companies seem to be one of the preferred ways to bring research results to market. Even large companies that own a crazy idea sometimes spin it off into a startup, limiting their risk. Successful startups are then candidates for acquisition.

A difference between the startup craze of the 1980s and today is the emphasis on acquisition. Companies increasingly grow by acquiring technology and products rather than developing them. I think this trend will continue.

What is the most controversial and unpredictable technology in the Internet space?

Security and privacy-related technologies are already the most controversial, and I expect the controversy to intensify in the future. There are many issues, for example, should a government be able to eavesdrop on your Internet use? Should your employer? During the November 1999 Internet Engineering Task Force meeting, there was a vote on the issue of wiretapping the Internet. The idea of adding support for wiretapping to the core Internet protocols turned out to be very unpopular, but it is unlikely that this has settled the issue. Another example is the amount of information about you available on the Internet and who can access it. What are reasonable privacy policies? How should privacy contracts be negotiated and what are the enforcement mechanisms? I recommend an article in the January 2000 *Harper's* (Mark Costello et al., "The Searchable Soul: Privacy in the Age of Information Technology," pp. 57–68.) for a thoughtful discussion of privacy and the Internet.

Franklin Reynolds is a senior research engineer at Nokia Networks and works at the Nokia Research Center in Boston. His interests include ad hoc self-organizing distributed systems, operating systems, and communications protocols. Over the years he has been involved in the development of various operating systems ranging from small, real-time kernels to fault-tolerant, distributed systems. Contact him at franklin.reynolds@nokia.com.



Munindar Singh

What were the decisive turning points for Internet and Web technology to become ubiquitous and pervasive?

The key turning point has got to be the appearance of useful interfaces such as Mosaic and of enough nodes being around on the Internet to make it worthwhile for you to use the Internet to find something. Tools like Mosaic made it simple for you to find information. Many people could see the commercial possibilities and that's where it took off.

What bottlenecks do you foresee?

I don't think that the future of the Internet depends on fast networking. There's enough bandwidth to go around, at least for now. In computing power, yes, more development is needed. We need better Web servers and so forth. But I don't see that as a bottleneck.

The bottleneck is going to be the user interface. Right now, the Internet does well up to a point, but it is certainly not very good in terms of how people interact with each other, how they interact with services, how they find services. The Internet doesn't give you a semantics, and you're stuck with using keywords. Keywords might be acceptable in a circumscribed domain. They are much less effective as the domain grows, when precision declines.

What are the next likely disruptive technologies in Internet space that might make new marks in the way we live and work—agent technologies, voice/text recognition, and wireless?

User interfaces will need to improve a lot, especially long-lived interactions between the user and something on the other end—the Internet. In terms of specific technologies, all of the technologies that you mentioned will have an increasing role to play.

However, I would say wireless technologies are going to have the biggest impact on how we think of the Internet. Right now, we conceptualize the Internet as wired. We can remain connected for long periods of time, and those of us with dedicated modems or LANs can stay connected almost perpetually. But wireless connectivity is generally unreliable and expensive, and also of a low bandwidth. That changes the way we look at things, even from the technology standpoint. Several problems that appear solved right now with the wired technologies will reemerge in a wireless setting.

If you can't even take for granted that you will stay connected, you have to think of different computational metaphors altogether. For example, the Jini specifications (from Sun) that came out publicly about a year or so ago have this notion of leasing. As its defined, leasing is an attractive notion for open systems in which you get a resource, but you get a resource only for as long as a lease lasts. When the leases expire, unless you renew them, they soon go away. The way they set it up, leasing doesn't consume resources forever, but the leases must be renewed in the order of several milliseconds. This is demanding in a wireless setting, and you might want to have

Security next

The next Trend Wars installment will focus on security.

We have invited Dan Geer, Li Gong, Marcus Ranum, Clifford Neuman, and Mary Ellen Zurko to participate.

If you have any specific questions that you want our guests to answer, please send them to cbaltes@computer.org.

some other way of handling disconnected operations.

Also, wireless technology makes more demands on the user interface itself. Wireless can make you mobile, but only if don't carry around a big monitor! And if you have a PDA, there's a small screen and low bandwidth that is susceptible to broken connections. That gives you an entirely different Internet experience. Increasingly, people are going to want that technology. In Italy, there are more wireless telephones than there are wired telephones, and in Japan the number of wireless telephones exceeds the number of desktop computers. Some studies suggest people like telephones a lot more than computers. So, wireless is the way of the future, and it is going to shake things up quite a bit. I think agent technology will provide solutions to many of the problems caused by going to wireless.

Do you think also on the programming paradigms side there will need to be a parallel thing? So far, we have kind of a nomadic computing where you would have laptops that move and periodically connect and disconnect. But would wireless require much more dynamic disconnected behavior?

Yes, I think so. And I think it would change the way we set up our desk systems. They would have to be a lot more open, a lot more context sensitive. With a laptop, you pretty much have a fixed context. Maybe the numbers you dial out change a little, but for practically everything else, a laptop is just like a desktop.

We really must think of more dynamic changes in context, say, where you're walking down the street and you want to know what's around you at that time. User expectations will be different, and at the same time the demands on the systems that we build will be different, as will the techniques that we use to build the systems.

Do you think that wireless will justify the need for more mobile agents? I recall somewhere your statements that there is no need for that.

There's a distinction between mobile objects and mobile agents. In some cases, mobile objects might be a good technology to have. For example, if you have a long-lived computation, survivability of the computation beyond the lifetime of the hardware is important. Clearly, it's good to have mobile objects in those settings. In other settings, and in particular where mobile agents are involved, I don't think they really buy you anything else. You just need more sophisticated communication standards.

Another significant thing that's going to happen is in adding semantics to the Internet. That's a difficult problem that's not going to go away. The Internet's success in connectivity only

Technical Committee on Operating Systems, Applications and Environments (TCOS)


<http://www.tcos.org>, [tcos-\[announce,discuss\]@tcos.org](mailto:tcos-[announce,discuss]@tcos.org)

IEEE-CS TCOS is a large membership organization (over 3,500 members), an affinity group of developers and users interested in Operating Systems (OS), applications and environments. TCOS is to IEEE-CS what SIGOPS is to ACM (though TCOS membership is free).

TCOS is involved with theoretical and practical aspects of OS design and implementation, including system organization, resource allocation policies, measurement, performance evaluation, and system reliability. It is also involved with OS aspects affecting system interface, the completeness of services, and portability of applications and environments.

There are two main types of activities of the TCOS of benefit to members and to community in general: the organization of OS-related events (see the list below), and gathering and distributing OS-related information, mainly on the Web site. We publish a Newsletter a few times a year.

TCOS collaborates with IEEE-CS publications, such as IEEE Concurrency. Examples include making some of the magazine



materials available on the TCOS Web site and pursuing joint publications. TCOS seeks to actively cooperate with other organizations such as ACM SIGOPS and USENIX and has been cosponsoring several events with them

TCOS sponsors The Hot Topics in Operating Systems (HOTOS) Workshop which is its main event. TCOS cosponsors the USENIX OSDI Symposium, Virtual Enterprises and Mobile Technologies, Workshop on Mobile Computing Systems and Applications (WMCSA), and WEISS (see last page).

TCOS membership is free. You can register on the Web: through IEEE-CS (on www.computer.org/members/, follow *Other Services*) or on <http://www.tcos.org>.

We are actively looking for new ideas, but even more for new volunteers to make the ideas happen. Please take a look at our member survey and tell us what else you would like TCOS to offer. The form is available at:

<http://www.tcos.org/forms/feedback.html>

makes the semantic problem harder, because, again, you want to find things that are relevant and meaningful, and there's no current way of handling that. That would be a good place for agent technology. If agents can understand what the user needs and help find what the user wants to read or buy or whatever more effectively than if the user had to go unprotected into the Internet, then technology of that nature will develop.

Where do you see the roles of industry, startups, research labs, universities, open source companies, and standard organizations in shaping the future Internet?

The Internet has to become professional. The way things have been done so far; it's very amateurish. Any kid in any garage can start a new company and do stuff with the Internet. Someone comes up with a design and they make it work. Because there's nothing else to compare it with, we are happy that it works. But it's as if people are doing things off the top of their heads: for all its success, the work sounds ad hoc.

In the next few years, the Internet is going to become more professional. It's going to be done by real computer scientists, as opposed to the physicists in their spare time! It will have a different kind of a flavor to it. That might be controversial, but the role of computer scientists in the Internet is going to grow significantly.

For example, look at the recent success of companies like Akamai. That's the kind of thing we'll see more of—hard core computer scientists applying graph theory, for example. The Internet is reaching the fundamental limits of how far the ad hoc approaches will go. Soon, we'll have to do things more carefully, more as engineering, more studied.

So, it's actually a good time for computer scientists—people with interest in concurrency and knowledge of concurrency, among other areas. They will stand apart from those who are unqualified in the details of computer science, those whose only computer science qualification might be the abil-

ity to write a script. In that sense, universities will have an increasing role to play. The entry requirements will be higher, so people will need more university training. I suppose startups will continue to be the place where good ideas come out, but better-qualified people will more often found the startups.

Do you think that the current negative trend of people not going for PhDs will reverse again?

Not right away. The economy is booming and people can achieve a lot of things in industry, so they will continue to do that. Universities are going to have a tough time retaining faculty and PhD students for another few years. However, as the Internet becomes more professional, people will realize that they need more formal education as opposed to taking off with a minimal courses in programming, which frequently seems to work these days.

But the PhD question, that's difficult. Some students will get enough experience doing a good masters, as opposed to necessarily sticking around for a PhD. Universities that are entrepreneurial and find interesting challenges for students will do well. Universities that retain a traditional orientation, resisting dealing with companies, will generally not do as well.

Again, as the Internet becomes more professional, people will want more standards. But like in the rest of computing, standards have often brought up the rear.

There is this old saying that standards should come in the middle of the technology development. If they come too early or too late, it's not as good.

That's right, yes. We certainly don't have the too-early problem right now.

What major application areas will dominate the Web in the future?

Two kinds of applications will dominate: e-commerce, of course, and then personal technologies. By personal tech-

"From the Trenches", The First Workshop on Industrial Experiences with System Software (WIESS'2000)

Co-located with OSDI 2000 October 23-25, 2000, Paradise Point Resort, San Diego, CA

<http://www.usenix.org/events/osdi2000/wieiss2000/>

Co-sponsored by USENIX, IEEE-CS TCOS, and ACM SIGOPS (pending approval)

Important Dates

Paper Submissions due: Mon, May 15, 2000
Notification to authors: Thu, June 29, 2000
Camera-ready papers due: Tue, August 31, 2000
WIESS'2000: Sun, October 22, 2000

Overview. WIESS'2000 will feature short papers (5-10 pages), abstracts, and posters on designing, implementing, and using industrial system software and applications. It is an attempt to allow people from "the trenches" to present their work. SOSP and OSDI conferences feature very high quality technical programs that have also "raised the bar" for paper acceptance. It has therefore become more difficult for authors of papers presenting industrial experience to be included in these programs. WIESS is an attempt to complement SOSP and OSDI, and focus primarily on industry results of immediate benefit and use rather than long-term research.

Submission Guidelines. WIESS will focus on papers that draw important conclusions from practical experience in developing and using system software solutions. War stories, outrageous conclusions, and negative results are especially welcome. Topics of interest should include, but not be limited to: Operating Systems, Distributed Systems, Real Time & Quality of Service, Embedded Systems, Security and Privacy, Networking, Internet, Web-Based Technologies, Programming Environments and Tools, Fault Tolerance and High Availability, Middleware, Appliances and Personal Digital Assistants, and System Administration.

Papers will be valued for the relevance and usability of the work more than for the presentation. The best 20-30 papers will appear in the conference proceedings. The other accepted papers will appear as short abstracts. Most papers (except those that clearly did not meet the above criteria) will be provided an opportunity to present their work as a poster at WIESS 2000. Papers should be submitted using Web:

http://tesla.hpl.hp.com/wieiss/forms/authpaper_reg.html

Submitted papers should be 5-10 pages, single-spaced 8.5x11 inches, including figures, tables and references. Only printable PS and PDF (preferably) will be accepted.

Program Committee:

Gaurav Banga, NetAppliance
Eduard Bugnion, VMWare
Rob Gingel, Sun
Fred Glover, Compaq
Ira Greenberg, Oracle
Larry Huston, Intel
Rodger Lea, Sony
Udi Manber, Yahoo
Franklin Reynolds, Nokia
Toshi Sakuraba, Hitachi
Indira Subramanian, HP
Franco Travostino, NortelNet
Richard Wheeler, EMC
Mark Brown, IBM
Dejan S. Milojicic, HP Labs, Chair

Steering Committee

Jean Bacon, Cambridge University, SIGOPS Member
Andrew Hume, AT&T Research, USENIX President
Valerie Issarny, INRIA, ACM SIGOPS Vice Chair
Mike Jones, Microsoft Research, TCOS Vice Chair
Marshall Kirk McKusick, Self Employed, USENIX Member
Dejan S. Milojicic, IEEE-CS TCOS Chair

nologies, I mean things that help people go about their daily lives. For example, technology can help people find information, help people find other people, help people do community work. Many of the noncommercial applications will be of the personal technologies variety.

What is the most controversial or unpredictable technology in the Internet space? Is there something that some people believe in, but that others violently oppose, for example? One example might be privacy. It's not really technology, but it's a topic that's being discussed a lot lately.

That's a good example. With security, people are either extremely careless about it or they're really paranoid about it—with nothing in the middle. Maybe standards will get established and people will be able to check their privacy requirements automatically, making sure that they preserve those.

How do you see development of the Internet infrastructure over the years, with respect to continents, governments, and so forth?

It will spread. Wherever there's an economic motive involved, it will generally spread fairly well. Even in a country that isn't as fully wired as the US, if the government sees advantages, it will expand the reach of the Internet. Countries like North Korea might not, because of extreme politics and because the people running the country might not see the economic advantages of having net access for everybody. But vir-

tually all other countries will see a growth—even regimes that by our standards are not very nice will support the growth of the Internet, perhaps with some kinds of control over it.

That brings up another controversial point—content-rating services. A recent *CACM* had an article about how simple keyword matches seem to eliminate good content more than bad content. For example, they might eliminate Superbowl XXX because of the letters "X-X-X." Although content-rating services are controversial, content rating by consumers themselves by having ways to manage reputation might work. I have a personal interest in technologies for reputation management.

Again, the success of content-rating by individuals might depend on the kind of society you live in. In the US, we can get away with expressing our opinions, although at the risk of a lawsuit. Other countries might not appreciate the dissemination of certain kinds of information, for political or religious reasons. Such societies might not appreciate the right of ordinary citizens to publish ratings.

Munindar Singh is an assistant professor in computer science at North Carolina State University. His research interests are in multiagent systems and their applications in e-commerce and personal technologies. Singh received a BTech from the Indian Institute of Technology, Delhi, and a PhD from the University of Texas, Austin. His book, *Multiagent Systems*, was published by Springer-Verlag in 1994. Singh is the editor-in-chief of *IEEE Internet Computing*. Contact him at singh@ncsu.edu.