A Conversation with Peter Ford

Instant messaging (IM) may represent our brave new world of communications, just as e-mail did a few short years ago. Many IM players are vying to establish the dominant standard in this new world, as well as introducing new applications to take advantage of all IM has to offer. Among them, hardly surprising, is Microsoft, which is moving toward the Session Initiation Protocol (SIP) as its protocol choice for IM.

Providing us with the Microsoft perspective on IM is Peter S. Ford, chief architect for MSN Messenger. At Microsoft he has worked on Messenger, TCP/IP, IP security (IPsec), RSVP and QoS, voice over IP (VoIP), and Mobile Data. Previously he worked at MCI on Internet access and virtual private networks (VPNs), on the evolution of the National Science Foundation network to network access points (NAPs) and very-high-speed Backbone Network Service (vBNS), and at Los Alamos National Laboratory on high-performance computer networking and nonlinear systems. In an earlier life he was a systems hacker at the University of Utah and the University of Michigan. In the Internet Engineering Task Force (IETF) he cochaired the team proposing the use of connectionless network service (CLNS) as the candidate for IPv6. Ford has a bachelor of general studies degree from the University of Michigan.

Sparring with Ford in a discussion of IM is e-mail pioneer Eric Allman, chief technology officer and founder of Sendmail. Allman was an early contributor to the Unix effort at the University of California at Berkeley, where he received a master's degree in computer science in 1980. While there, he wrote the `syslog`, `tset`, the `_me` troff macros, and `trek`—along with Sendmail. Allman designed database user and application interfaces at Britton Lee (later Sharebase), where he worked with Ford on software for database machines. He also contributed to a neural network-based speech recognition project at the International Computer Science Institute and was CTO at Sift. He coauthored the “C Advisor” column for *Unix Review* magazine for several years and is a former member of the board of directors of the Usenix Association.

ERIC ALLMAN There are a lot of players in the instant messaging (IM) space today: AIM, MSN, ICQ, Yahoo, and

The IM world

ACCORDING TO

A MESSENGER

ARCHITECT

on down the list. It seems fairly obvious that having all these incompatible players is untenable. It

looks like things are going to converge on some combination of SIP (Session Initiation Protocol) and SIMPLE (SIP for Instant Messaging and Presence Leveraging Extensions) versus XMPP (Extensible Messaging and Presence Protocol)—whether that's all one, all the other, or some blend of the two is unclear. Do you believe those are the

only two players, and how do you think that convergence is going to happen?

PETER FORD I'm not sure the only two players are from the IETF (Internet Engineering Task Force). There's another forum among the mobile carriers called Wireless Village, and those guys are driving not only an architecture that's very similar to the other two architectures, but also a separate set of protocols. The wireless guys have a lot of clients, and as phones become more capable—with more memory, more CPU, better screens, and better input devices—they become a locus for yet another protocol and flavor, a mix of interoperability.

I do agree that you will see consolidation over time. You could probably argue that the e-mail experience that you and I and a lot of others lived through over the past 20 to 25 years is probably going to be repeated—perhaps more quickly than last time because the Internet makes that kind of evolution easier.

Does that mean 10 years, five years, two years? I couldn't predict. Quite frankly, the thing that fights against it being quickly this time around is that the communities operating with these mutually incompatible protocols are quite large. If you look at AOL's cloud or the MSN cloud or the Yahoo cloud, you're talking about fairly large, significant systems. To have them migrate and interoperate with standard protocols will happen, but it is going to take time.

I was involved with IPv6 (Internet Protocol version 6) for awhile back in the early '90s, and that's certainly taken longer than most would have anticipated. When you're talking about very large established systems, it takes time to move to new protocols.

EA Is the future of IM going to be driven by wireless carriers?

PF In addition to the PC IM systems and suppliers, you have the wireless carriers and their suppliers in the overall ecosystem for communications and IM. You have Nokia, Ericsson, Motorola, and others in that mix. Wireless platforms clearly are incredibly important as we move ahead, both in terms of what the carriers offer on top of 2G, 3G, and 4G networks, and in what's happening in the whole Wi-Fi space. That being said, I think the reason why the systems that we have for IM today are different from what you currently see in wireless systems and proposals is that the endpoint, or the power of the endpoint—the PC that people are sitting in front of—is so high and the functionality you get out of PC-based IM systems today includes new capabilities such as presence awareness

across the network with your peers.

This is different from what you get in the mobile-phone space today, where the working concept is “always on,” and you're always reachable. In the milieu of PC-based IM today you actually have a fairly rich presence system where you can indicate “I'm out to lunch,” or “I'm away for 15 minutes, I'll be right back.” So I think that one major difference is that you have much richer notification to your peers, your buddies, your contacts.

EA Apparently we've now got the three big players in the standards battles: Wireless Village, SIP, and XMPP.

PF Right. Just to be very clear, the three large networks don't use any of those for their current IM and presence—actually four networks, if you include ICQ.

■
EA AIM, MSN, Yahoo, and ICQ?

PF Yes, I'd say those are four fairly significant IM net-

works. They probably each have tens of millions if not over 100 million users. I know we (MSN) went over 100 million unique users a couple of months ago, in terms of people who are registered and login to use the service.

■
EA So MSN is not based on SIP at this point?

PF Not for presence and IM. At Microsoft, we designed and built the current MSN IM cloud that we ran as a

public service back in the late '90s. It uses a protocol called MSNP that we documented at the time as an Internet draft as part of the IMPP (Instant Messaging Presence Protocol) interoperability discussions in the IETF. SIP emerged from that round of discussions as the protocol choice, and we just released a product in July for Enterprise that uses SIP both for audio-video and for presence and IM. The Windows Messenger client that is released now actually is capable of talking to both the public cloud and to SIP servers that are deployed in Enterprise. If you were to look at a roadmap over several years, you would see a convergence onto SIP. We did add SIP for audio and video in the MSN IM network in 2001, including a PC-to-telephone integration feature. SIP has been the plan around here for several years.

EA Microsoft, Lotus, Sun, and Novell seem to have settled on SIP. Intel, H-P, Hitachi, Sony, and more or less the entire open source world is going toward XMPP, sometimes better known as Jabber. Are those two communities ever going to speak to one another?

PF I think so. Over time users will demand that, so you can imagine it working out in different ways. One way would be that a new standard emerges and everyone just moves to it. Or it could be, one of the existing standards—XMPP/Jabber versus SIP—basically becomes the winner. There may be interim solutions: I'd be shocked if vendors did not build, just as they did with e-mail gateways, ways to interoperate between those two communities.

There are two ways to interoperate. There is service-to-service interoperability, and then there's client unification, where people build clients that are able to talk to both, or several, protocols. I think to the one-protocol purists who live in the standards bodies, some of this is anathema to the way they think the world will work out, but the reality is we lived through this with e-mail, and we got to a fairly unified e-mail world without any disasters. In fact, the e-mail world is still evolving, and I think we'll see the same thing with IM.

EA Why did Microsoft choose SIP over XMPP?

PF I think it was from looking at the terrain at the time. If you go back three years, SIP had just recently been standardized for the purposes of signaling audio and video, so the notion of using SIP to signal other kinds of realtime communication like text-chat didn't seem implausible. In fact, it works.

And then there are a whole bunch of other externalities. For example, the 3GPP (Third Generation

Partnership Project), one of the communities that tries to standardize protocols for the international mobile telephony markets, elected to use SIP for mobile phones—for making calls, not for IM. Microsoft is still very bullish on voice over IP and the uses of common SIP signaling. We thought, "Gosh, wouldn't it be great if we had a common signaling protocol for all of these realtime sessions that we wanted to use and the protocol also allowed you to bundle those sessions together into a single session?"

You could have a session that was audio, video, and text. If you think about what a little screen on a PC would look like, you'd see your mom, you'd be talking to your mom, and you might be typing instant messages or dragging and dropping file transfers and pictures with your mom. That seemed like a very compelling scenario to us. We thought we could do it with SIP, so we built it.

And then we participated in the standards bodies, the IETF specifically, to work on the presence and IM text messaging specs. You see a lot of Microsoft people on those specs. At the time, David Gurle worked at Microsoft—he now works at Reuters on its IM network—as well as Christian Huitema, who used to be the chair of the Internet Architecture Board (IAB) and is an architect in the Windows Networking and Communications group. That's sort of a historical perspective of how we got there. It seems to work. It's one of these pragmatic approaches.

EA Those in the XMPP community seem to be claiming that XMPP is the mature protocol. It's got everything you need. SIP was designed, exactly as you said, as the signaling protocol that really doesn't have all of IM and doesn't have anything on presence, etc. There's a desperate urge to retrofit those in the form of SIP, which has not been standardized yet. Therefore, what's happening is all of you guys who are using SIP are doing proprietary extensions, and it's going to be hard to converge those. What's your take on that argument?

PF If I were to talk about who got where first, I think it's pretty clear that the SIP guys got started early in the presence game. SIP Registration is a form of presence. There was this earlier IMPP standards effort that yielded the requirement documents for IM and presence, and the SIMPLE (SIMPLE is the IETF group that extends SIP for IM and presence) guys kind of spun out of the IMPP efforts almost immediately. I think XMPP for IM and Jabber started a little bit later.

Standards take time. It's a layered architecture, and getting all the pieces of those architectures finalized and standardized, like any set of standards, is going to take

many, many years, especially in the more business-oriented IETF.

Twenty years ago these things would have just been built, and people would have thought of them as de-facto standards, whereas today the IETF requires going through Proposed, Draft, and then to Full Standard. That takes time. A lot of that time is getting working experience with building and deploying those protocols, and we are still in the early stages of deployment. For example, IPv6 took five-plus years in many cases to get things moved first into Proposed, then moved from Proposed to Draft, and you're only beginning to see the specifications for IPv6 becoming Full Standard.

I think we're going to have one of these interesting two-protocol situations that we've seen in the IETF before, much like OSPF (Open Shortest Path First) and IS-IS (Intermediate System to Intermediate System), both of them being discussed and worked on and subsequently deployed by ISPs. It is much like the early IPv6 work, where we worked on a couple of protocols, and then over time a decision was made within the IETF community that there's a de-facto winner and the bulk of the work will reside in one protocol family versus the other.

I doubt that kind of consensus on a single IM protocol will occur in the near term. There are a lot of standards dynamics, as well as application writers and ISVs (independent software vendors) voting with their feet.

I do think the telephony side of SIP does make it fairly compelling because there's a huge fleet of vendors out there that have solutions today that can interoperate with the telephone network. These solutions also interoperate with each other. I think you'll see SIP extended exactly the same way that XMPP is getting extended over time. So if there is a cool app that people want to run, you'll see one set of ISPs probably working with SIP and another set of people working in the XMPP world.

We at Microsoft feel pretty strongly about the SIP side of the game today. We think that the interoperability with the phone network matters because it's a very compelling user scenario, and so that's one area where we have our bet placed.

EA It's my understanding that SIP uses servers to find each other, but after that, it's all peer-to-peer. Correct?

PF It doesn't have to be. It's a policy decision. The first handshake—let's say the session invite that's going from A to B where A and B are each behind their respective SIP proxies—goes through A's outbound proxy to B's inbound proxy and then down to B. You can actually maintain that as the path that all communication goes through for signaling and for IM traffic as an example.

In that example, A may never actually discover B's IP address because B's SIP proxy actually provides a blinding function with respect to B's actual location inside that

administrative domain. In that case, you don't get IP address exposure.

For efficiency, you might want to do the policy checks in the first handshake of the invite and then any subsequent session traffic could go peer to peer if allowed. For the peer-to-peer cases you do get IP address exposure. We actually do that for things like audio-video traffic that

runs over RTP. We try to go peer to peer, especially for realtime traffic, because today SIP proxies by themselves are not designed to handle flowing RTP streams through them. The hard cases are when you have PCs connected to the Internet with NATs (network address translation), where you see if you have direct connectivity by attempting a connection peer to peer; and if that fails due to NATs/firewalls, you can revert to a proxy/relay in the cloud. It's kind of an ugly hack, but it's a way to check if you have efficient peer-to-peer connectivity. If you can talk to it directly with SIP after you've discovered the endpoint, the IP address, then you probably have a chance that the realtime media will flow. You can fall back to the proxy path only if direct connectivity fails, since proxies will introduce latency and cost to the realtime streams.

EA So what is the potential "peer-to-peerness" of it, say about scalability?

PF It's possible to build very fast SIP proxies. So, scalability is much less of an issue for signaling, basically because we have ever-improving silicon working on our side. I guess in the ideal, you would say a peer-to-peer protocol

would be more scalable because you'd have to deploy less intelligent infrastructure to handle every packet. That being said, we still believe people, especially enterprises, will deploy tons of proxies because the proxies are where the network managers and administrators will provide policy enforcement, logging, and monitoring.

There are things such as HIPPA (the Health Information Privacy Protection Act) and SEC (Securities and Exchange Commission) requirements on certain industries, where having that kind of proxy-based routing is actually important from a management and compliance perspective.

EA The ability to queue messages, of course, is one of the great things about e-mail. You don't have to be there right now to read it. Do you see any kind of queuing happening in IM along the lines of "Gee, when I log in next, I'll see any messages that came in for me in the meantime"?

PF A lot of us call that the offline messaging scenario. Offline basically says you're not available. Where does the instant message go if you are offline? You could either queue in an intermediary node or you could actually queue at the source. Typically, SIP, as it's designed today, is pretty much an end-to-end protocol.

So the responsibility for queuing and reissuing a retransmit probably would reside on the endpoint if you were to build offline IM with SIP as currently specified. Alternatively, when users are offline, they can have a SIP-based proxy running on their behalf that picks up the inbound IM and queues it for later delivery when the user comes online. I do think that you're going to see interesting application intermediaries built that will provide the kind of functionality you're talking about. End users are beginning to ask for it. They get jazzed about the idea of being able to start an IM session with somebody, then if that person goes offline at some point, the message being sent would be saved and retrieved at a later time. Other exciting applications such as in-network language translation could also use that architecture.

The current IM clouds really aren't built for that. By building smarter endpoints, it's pretty easy to get that

kind of functionality. Then, also, in the MSN messenger network, we don't actually go peer to peer with text messaging. We always go through an intermediate system we call a switchboard. You could put something like a queuing mechanism there. We don't do queuing today. The interesting thing is that since you have some semblance of presence flowing between the two ends, it's more than likely that your endpoint application could determine whether to use more of a queued protocol or application stack rather than an IM stack. There's nothing to say that just doesn't become mail, right?

Do I think that IM and e-mail will become the same protocol in the end? I think probably not. Do I think that you'll have end-user or user-agent experiences that basically blend the two together? I think it's inevitable.

I could be wrong here, but that is my guess.

EA Is IM to e-mail the same as phones are to voice-mail? Is there something distinctly different there?

PF Interesting question. I guess I would say, from a user's perspective, the answer is probably yes. From a technical protocol perspective, however, the answer is no, because most voice-mail systems are

really just different user-agents, hanging at the recipient end of the connection. The protocols that are running are still the voice protocols, and you get the worst possible experience.

If you call someone's VM system, you have no idea what the UI is. If it's in a foreign country, it's talking to you in a language that you probably don't know. It's just horrible. If those of us who have lived in the e-mail universe for several years would have built voice-mail, we would have made it the caller's user-agent's responsibility to manage, building up the piece of voice-mail and then sending it across with some sensible store-and-forward protocol, for a couple of reasons.

One reason is that coding efficiency would be much higher. You wouldn't have to have this realtime requirement for carrying the bits across that you still have with voice-mail systems today. The second thing is, users could work with a client or a client user-agent that was familiar to them, so they could edit it and they'd have commands

that they understood. They might have precanned messages, things of that sort, and you'd just have a much better user experience for the caller and the voice-mail recipient. So in that sense, I'd argue that the architectures for VM and IM are quite different.

EA Another issue that enterprises have had is essentially letting their confidential data go over an insecure mechanism—in particular, financial people using Yahoo IM to talk to their clients means there's potentially very sensitive information that's going through a third-party's host, all cleartext, etc. How's the IM world coping with this for the future?

PF It has the same ideas about obtaining privacy and security, but it's struggling much the same way that the e-mail community has. It's hard to build secure e-mail that's as easy to use as existing products, so there's a lot of work involved just from the user experience and the systems to make it simpler.

There's been some fair progress on that. The protocols are kind of just now coming together. The e-mail space has had several false starts on protocols for secure mail. I think the IM world is a little bit better off because it can look back at that experience. But at the same time, you have challenges. The IM world is probably in a little bit better shape than the e-mail world. The reason I say that is, in most of the IM systems you actually have some notion of who the user is, who's logged in, unlike e-mail systems, which were really designed for broad open communication and didn't worry too much about user authentication up front.

The e-mail world is busily grappling with that as a result of spam. But once you have authentication in the system, it usually is easier to get to some kind of private communication channel or end-to-end package, based on that authentication.

Our corporate product actually uses SSL for enhancing the privacy of the channel, and it won't surprise me that in areas where you actually want full end-to-end assurance of the security of the package, you'll use end-to-end encrypting techniques as well.

You'll see products like that emerge. In fact, there are people who built plug-ins into our clients to provide some of that. People have figured out how to piggyback on top of existing clients and technology, to provide some of those kinds of encrypted conversations.

EA One of the approaches that people are talking about

for e-mail spam is the extensive use of whitelisting. This has been compared to buddy lists in IM, where I gather the very common style of usage is you're not available to people you don't already know. In other words, to meet somebody you've never met before, you need an introduction, which the high priests of etiquette would probably like, but it doesn't really fit into what I think of as the Internet world where you can talk to anybody at anytime. Do you see that changing the way we communicate and the way we think about messaging and mail?

PF I do. I think one of the reasons people like IM is that several IM systems, ours included, do use whitelisting as a basic part of the system. When I decide to add you to my contact list, you have to basically say yes before I can start IM'ing with you. The natural question, of course, is how do you bootstrap that entire process? The current common answer is e-mail, because e-mail doesn't have that kind of enforcement.

At some level you could argue that e-mail, and maybe the future of e-mail, is the open-for-everybody protocol and system. You essentially bootstrap either to another flavor of e-mail, where you use extensive whitelisting, or to instant messaging as a different application, but it's still a messaging application where the whitelist is enforced.

You still have to have some low-level bootstrapping protocol that's pretty open. Another question is, of course, how do you make that protocol less spam-capable, because people use that open conduit to you. People will use those low-level protocols to spam you. There may be another set of mechanisms that you have to put in place. Basic authentication of the source may be enough, much like, in some cases on the Web, asking for people's credit cards just so you can prove that they're human beings whom somebody else is going to vouch for. You don't charge them anything; you just want to know for sure that the person signing up on your site is a bona-fide human and not a robot or spam-o-matic. This is just one possibility in an entire spectrum of techniques to mitigate spam.

EA So you think e-mail is destined to become the singles bar of the Internet?

PF (Laughs) Well, that and chat.

EA Chat is probably closer to it.

PF Chat is certainly closer in terms of the interaction style. And it's interesting, that's actually one of the differences between chat and IM. Typically, IM is much more focused on interacting with people who you know in

some way or trust in some way, whereas a lot of chat is really more about the singles bar or the lobby of a hotel where you just serendipitously meet someone and start up a conversation. Many people meet in chat and then as they learn more about each other they take their one-to-one personal conversations into IM.

There's this whole social dynamic of what protocols you use and whether they are the same protocols or different. With IM, typically, as I said, you have some notion of an authenticated identity. That usually means you know who they are, whereas with chat and e-mail—in the sense of people just getting some random name on Hotmail or Yahoo—you may not really know who they are, and interestingly enough, you might not care to know exactly who they are.

There's room for both. In fact, I'd argue you need both because part of human interaction is meeting new people whom at the very beginning you know very little about. So how do you make that work? You're not going to demand a person's Visa card just to have a conversation.

EA Would you care to give a prediction as to what the messaging world in general is going to look like in two years and what it's going to look like in five years?

PF I'm pretty optimistic that a lot of the spam issues will be resolved, both in the IM and the e-mail worlds in the next two years because I think we've reached a crisis with respect to the spam problem. I actually think things like whitelisting, which is ubiquitous in IM, will actually become far more common in the e-mail world.

Coming up with a way that will allow for the—I don't want to say the anonymous but basically the un-introduced—person being able to communicate with you, that's probably a five-year, maybe even a five-year-plus, problem.

The sort of multimedia experience around messaging will become far more pronounced, even in the next couple of years. A lot of people around Microsoft worked on NetMeeting and other efforts to build a PC- and Internet-based follow-on to Picturephone, and I think we're probably at the point now where for very low cost and with very high fidelity and great user experience, you can have conversations with people with video and audio, simultaneously sharing pictures and applications—either sharing a photo album in the consumer space or working with a PowerPoint presentation together in a collaboration enterprise. I think that's going to be incredibly common in the next couple of years as broadband Internet gets deployed.

As we come up with protocols that do a good job of signaling and establishing that kind of connectivity, I think that will translate probably initially in the IM world because those are the scenarios that a lot of us are focusing on. But it'll translate over time into the e-mail world as well, where you'll see a lot more audio e-mail, a lot more video e-mail, and a lot of other graphical content. A lot of that is going to come from things like tablet-like PCs, where people actually write by hand. We just built something like that into our IM product, and I'm just amazed by watching kids draw pictures and send them. So, the Picturephone, presented at the 1964 World's Fair, is probably right around the corner—a couple of years out. It's going to drive tremendous bandwidth requirements in the networks.

I think the demand for bandwidth, five years out, is going to be phenomenal. We're talking about cameras and phones and stuff like that. That's just going to become much more commonplace—both in the consumer market and in business. Take, for example, the ability to take a snapshot of yourself and say, "Hey, remember me? We met at the IETF and we were talking about the protocol for XYZ." That kind of interaction is going to be greatly facilitated in these messaging products.

I think the flip side is that this whole notion of controlling your communications—and people controlling inbound connectivity—is probably going to be the big issue. And I think the spam issue is only one element of that kind of control.

You're also going to want to have that kind of control in terms of the people you work with and the people you play with. Building those social protocols for when you are going to allow your spouse to interrupt you when you're at work and things of that sort, and getting all that stuff right, is going to be quite important. That's probably a five-year problem.

Artificial intelligence techniques—things tied to your schedule, things tied to your friends-of-friends-of-friends network—will become important in terms of the way you decide how to communicate with people and what communication you accept and when you accept it.

I see tremendous amounts of evolution in what you and I would call user-agents. The current IM clients and the current e-mail clients are just going to evolve like crazy in the next five years. We probably won't recognize them five years from now.

The explosion of people trying to communicate with you, using instant messaging and e-mail, is going to grow tremendously. Having systems that can manage that in a human-friendly collaborative manner is going to be

critically important as we move ahead.

I'm very optimistic about that. At some level, it sounds like, "Oh my god, we're all going to drown in an e-mail sea," and I think that filtering technologies have come around very quickly in the last two years. I've been very impressed by how quickly people have addressed spam, and that's because it's so important. I think e-mail was the killer app of the Internet, and messaging still is the killer app. E-mail plus instant messaging are part of that whole messaging milieu. They probably will be for a long time.

I'm one of the people who still believe that e-mail is as important if not more important than the Web in the Internet. The Web-heads of course would say, "No, no, the Web is the most important thing," but I'm a big believer that person-to-person messaging, whether it be e-mail or IM, is probably still the driver. Clearly, both are important. I give the nod to messaging because people can be closer to the people they care about, and it makes it easier for them to work with the people they need to work with. I don't think that's going to go away. I think that's just going to get amplified by ways that we can't even imagine today.

Someone told me that they would like the client on their PC to be much like Della Street, Perry Mason's secretary, so that you'd have someone that knowledgeable about you, that helpful to you, facilitating your communication with the outside world.

EA Somebody who, when you say, "Could you please do this?" will say, "Yes, I already did."

PF Yes, and the mail that you have on your desk to read or on your desktop to read would be in priority order. We're getting there. Stuff like that I think is going to take longer because it's going to require some of the user identity and persona issues to get sorted out, deep language

understanding and knowledge bases, and the protocols are going to need to mature in the sense that they need to support that kind of authentication and authorization functionality.

There's also going to be more negotiation in the protocols, in the sense that they may have this conversation:

"Hey, Eric, can I talk to you now?"

"Well, I'm busy right now. Could you talk to me in 15 minutes?"

"I really need to talk to you immediately because it's about this important business decision."

"OK, can I call you back in two minutes?"

Imagine that communication going on in PC-to-PC messaging space. I think that kind of functionality is going to get folded in as well, and that to me is very exciting because the thing about IM is the immediacy. You really feel like you're communicating with a person on the other side.

I don't know if you use talk (an early Unix-based IM program) at all anymore. You said you're not a big IM user. There was a certain pleasure—at one o'clock in the morning in a machine room at the computing center—in launching up a talk session with someone to help you understand how to rebuild the G partition on your system. That was different from using e-mail. That captures some of the difference between IM and e-mail today. I think that's important.

EA Good point. I remember those late-night sessions, too. And you're right, they were really cool. **Q**

Visit our Web site for the full Peter Ford interview:
<http://www.acmqueue.org>

LOVE IT, HATE IT? LET US KNOW:

feedback@acmqueue.com or www.acmqueue.com/forums

© 2003 ACM 1542-7730/03/1100 \$5.00.

The explosion of people trying to communicate with you, using instant messaging and e-mail, is going to grow tremendously.