

## **UBIQUITY INTERVIEW: KLEINROCK ON NOMADIC COMPUTING**

[Leonard Kleinrock developed the mathematical theory of packet-switching, the technology underpinning the Internet, while a graduate student at MIT -- a decade before the birth of the Internet which occurred when his host computer at UCLA became the first node of the Internet in September 1969. He is now at UCLA, where he is Professor of Computer Science. He has won numerous awards and prizes for his achievements.]

**UBIQUITY:** Let's see, where should we start? You've gotten so many awards it's surprising you haven't received one for poetry. You've written some very funny verses.

**KLEINROCK:** I certainly enjoyed writing them.

**UBIQUITY:** For example, "Ode to a Queue," which you wrote in 1989 for a symposium for the 20th anniversary of the ARPANET. Let's quote just a little from it:

"If you want to model networks  
Or a complex data flow  
A queue's the key to help you see  
All the things you need to know.

So the next time you feel lonely  
And wonder what to do,  
You'll soon feel fine if you join the line  
Of an analytic queue!"

**KLEINROCK:** You have really done your homework. Wonderful.

**UBIQUITY:** I guess it seemed right to find you had written poetry because it seems to me that the whole idea of packet switching flows from a poetic insight, a kind of metaphorical insight. Is that completely wrong?

**KLEINROCK:** No, it's not completely wrong -- there's an underlying philosophy of packet networks which permeates not only their technology but the way in which they are used. The philosophy I'm referring to is distributed control, i.e., no centralized authority. So, in a sense that's poetic, where you let things organize themselves in ways that work well according to natural laws that emerge. I think that contains a poetic quality, especially in the sense that that's the way nature organizes itself. Typically, nature evolves in an organic fashion according to the needs of the community and the species that live within that community in a non-centralized fashion. And so, part of the idea of sharing resources, of distributing control, of making sure there was resilience and recovery capability, of ensuring a repair capability in building these networks, was all part of the larger mosaic of packet switching.

**UBIQUITY:** Can you remember your absolutely original insight when you first thought, as God might think, "Let there be packet switching"?

**KLEINROCK:** Well, it was not a point in time. It was a process. And the process had to do with the MIT environment in which I found myself. I was working at the MIT

Lincoln Laboratory as well as doing graduate work at MIT. I was surrounded by computers and it was clear that sooner or later they would need to interact with each other and for that they would need to be connected to each other. At that time, there was no good technology to allow them to communicate with each other. So I decided to devote my PhD research to the task of developing a data networking technology based on sharing of communication resources. Now, at the same time I was thinking about that, the notion of time-shared computers was in the air. It was a similar idea, namely many people sharing one computer in a demand fashion, where you only get access to it when you need it and when you don't need it you release the computer to allow someone else to use it, a kind of shared medium. Well, the idea occurred to me that the shared data network communication medium should behave in the same way, where you only use the underlying network when you have something to send. When you have nothing to send you don't dominate those resources, but let other people share it. It wasn't an "aha" moment, it was an "of course" moment. I thought: Let's apply those demand-driven ideas to these very expensive networks that we need to build. And that idea occurred to me in the late '50s and very early '60s.

**UBIQUITY:** Various people contributed different parts to the overall vision. How do you see the parts interconnecting?

**KLEINROCK:** Well, yes, there were a number of people that have participated in the development of the underlying ideas. If we had to tease it apart, I think the way to describe it is that my early work from 1959-1962 at MIT was the creation of a mathematical theory of packet networks. Independent of that work, Paul Baran at Rand Corporation was developing the architecture of packet networks and developing similar concepts. Quite a bit later, Donald Davies considered the same idea in England at the National Physical Laboratory. His work was in the 1965 era, which was a few years later than the work that we had done. Now, there is a large idea here, the idea of packet switching. Some people think packet switching is simply the idea of taking messages, chopping them into smaller packets, sending them into the network independently, and then assembling them at the other end. But I think of that as the *narrow* definition of packet switching. The broader idea of packet switching is a packet-switched network that not only chops messages into packets but also uses distributed control for routing, uses a distributed architecture for fault-tolerance and makes sure there is demand access, where you only use the resource when you need it. Note that packet switching is just one way to accomplish that. You could also do it by polling, you could do it by message switching, you could do it by demand reservations. But packet switching happens to be a nice way to do it.

**UBIQUITY:** What were some other factors?

**KLEINROCK:** Yes. Large shared networks provide dramatic improvement in efficiency ((and I was able to demonstrate that mathematically)). So we see that distributed control; redundancy and reliability; chopping message into packets; and sharing large systems are all parts of the bigger fabric of packet switched networks. But returning to the narrow sense of packet switching (the idea of chopping message into packets), I first published a paper in April of '62 where I described the idea of chopping messages into a sequence of fixed size blocks (i.e., packetizing the message) and serving these blocks one at a time. And then when you finished your block, namely the packet transmission, someone else gets to send their packet -- it was the idea of a round-robin way of serving messages. So the idea of packetization

came early in my work, but I didn't elaborate on the idea of end-to-end reassembly of these individual packets. I was focusing on a single node. My concept was: Let's look at a single node, understand how that behaves, understand what kind of queuing and packetizing mechanism works there, and then expand that in my theory to the larger network.

**UBIQUITY:** What was Paul Baran doing?

**KLEINROCK:** Paul Baran was more interested in the engineering side of it. He was concerned with the end-to-end reassembly of these packets into messages. So, together, our ideas emerged independently, which is interesting. And these ideas laid dormant until Licklider, who was the first head of the computer research effort at ARPA, began to elucidate in the '62, '63 timeframe the idea of a galactic network. He had the idea of putting computers together in networks and letting man/computer symbiosis occur. He had no idea how to build such a network, and yet he was at MIT (on leave to the Defense Department at the time). He didn't realize that I had already published the technology to do it, and I was unaware of his grand vision as well. So, we three people participated in laying the groundwork in the very early '60s.

**UBIQUITY:** Then what?

**KLEINROCK:** The things that followed are that Davies did his work in England and coined the word "packet" for this technology. He actually implemented a one-node network, but the U.K. was reluctant to fund a multi-node network. Meanwhile, ARPA determined they wanted to put together a network. And the reason they wanted to put together a network, as you probably know, was not to protect the United States against a nuclear attack, but rather to provide for sharing of computers, applications, software and files among computer science researchers. They wanted to put these various specialized, high-performance computers in a network, so that if I wanted to use the capabilities of your computer I would access it through the network, run a job at your site, and bring the results back home. Once ARPA decided to create this network, that's when the whole implementation really gained some steam. And, as you know, in 1969 the first node was deployed at UCLA.

**UBIQUITY:** You mentioned J.C.R. Licklider, the psychologist.

**KLEINROCK:** Yes, J.C.R. Licklider. He was a most dynamic fellow. His MIT influence and connection is important. He had access to many people. He was a visionary as a psychologist and as someone interested in human-computer interaction. He was brilliant. He was one of the real heroes at the time, and he was the first director of the IPTO office within ARPA. The people who followed him as heads of IPTO were also magnificent. Ivan Sutherland came after him. Ivan was a classmate of mine at MIT, by the way. And after him came Bob Taylor, followed by Larry Roberts. Larry was also a classmate of mine. In fact, Ivan, Larry and I took our thesis defense at the same time because our doctoral committees had overlapping people. Claude Shannon was on each of our committees, for example. We had to demonstrate some of our work on a Lincoln Laboratory computer. All these distinguished faculty and the three of us students went out to Lincoln Lab and we showed them our wares. It was a wonderful and very exciting time. ARPA was a terrific engine of innovation, led by people who were themselves excellent researchers with ideas, with vision and with also the wisdom not to provide overbearing management and guidance, but rather to let the scientists they were supporting run with their own ideas without a lot of site

visits, without a lot of oversight management, without a lot of reports. "Here's a goal. Here's some money. Here's a vision. Let's make it happen."

**UBIQUITY:** Does that spirit still prevail in computer science, and still prevail in ARPA?

**KLEINROCK:** Sadly, no. You may have seen John Markoff's article recently, where he bemoaned the fact that ARPA is funding computer science with smaller dollar amounts and, secondly, funding projects that have short range goals and that are very applications-oriented and very tightly controlled, with limited vision. So, today ARPA is not the paragon of excellent funding of research that it once was. Now, within computer science itself there always have been people -- and there still are today -- who look to reach over the horizon and seek those things that are long-range, high-impact, high-payoff, and very risky. But the funding for those kinds of projects is very meager right now.

**UBIQUITY:** What about in industry?

**KLEINROCK:** The situation is very bad. The great research labs of the past are gone. IBM Yorktown is diminished. ATT/Bell Labs is gone. HP Labs is gone. PARC is not doing as well as it had been. It's a great shame. The U.S. is on a decline in many worrisome ways. The latest statistic shows that we are now sixteenth in the world in terms of broadband penetration and that is a sad turn of events. Now, it's true that some of the countries that have moved ahead are smaller countries with high-density populations, so it's a little easier to provide access to a larger fraction. Still, the trend is there. Our phone companies have not been at all innovative or forward-looking, but have instead been focused on rather short-term opportunities; they've been protecting their turf instead of rolling out broadband as they could.

**UBIQUITY:** Any hope?

**KLEINROCK:** There is a move now to get fiber and high-speed wireless out there, so I expect we're going to be turning around in terms of where we are on that list. But we stalled badly. And that stall may cost us in the long run. There may be long-term effects of the paucity of access. I couldn't get high-speed access in my home here until about three years ago. That's an embarrassment. I live in Brentwood, which is not a remote area. To be fair, it's not a high-density area, but nevertheless it's four miles from UCLA in a thriving metropolis. Yet I couldn't get DSL or cable until about three years ago. I finally gained access with a high-speed cable modem.

**UBIQUITY:** Why is that?

**KLEINROCK:** Because it's an economic question: if there's not enough people to buy the service they're not going to pull the lines and condition them, later to pull fiber to some point within reach of a number of phones. So A, it's all economics; B, the phone companies wanted to protect their T1 business, and DSL was a clear threat to displacing T1 with a much lower cost service with almost the same capability in terms of bandwidth.

**UBIQUITY:** What's in the future? You're a champion of nomadic computing. Tell us about that.

**KLEINROCK:** Nomadic computing is a technology allowing anyone to leave their office and still have seamless access to the same set of network services as they had at their office, wherever they go with whatever device they're carrying, regardless of the environment they enter. Now, that's a dream, of course, but the idea of providing this similar set of services to you wherever you go as easily as you access them from your home or corporate office is the goal of nomadic computing. Suppose I'm sitting at my office desktop computer accessing the Internet with a 100-megabit or gigabit per second wired Ethernet connection. Then, suppose I turn around to my conference table and start using a laptop with a wireless connection. I've just made a nomadic move. I have moved from one environment to another. My platform is different. My connectivity is different. The services I see may be different. The trick is to make that move, or a far more dramatic move over perhaps thousands of miles, as seamless as possible. When I just go from my office to my home that's another nomadic move. If I were to come to your office, for example, with my laptop or I go to ACM headquarters or some meeting or some hotel, when I come there with my laptop, my laptop is an alien in the networking environment it finds. My laptop expects to find its router. Maybe at my office it's behind a proxy. When it enters a new environment the router there says, "Who the heck are you?" The game is to make my system appear user-friendly in that alien environment so I can get the set of services I need. And that's a question of software, some hardware and some systems capabilities.

**UBIQUITY:** Is wireless the fundamental key?

**KLEINROCK:** Wireless plays a role here, because it's a communication technology that allows you to move in a wide variety of environments and still have connectivity, either sporadic or continuous -- but that's only one component of the nomadic equation, the connectivity.

**UBIQUITY:** What else is involved?

**KLEINROCK:** The rest is the ability to connect at all in terms of plug-and-play. If I can't find my router, my laptop may give up. Or if the router there says you're an alien, then you're not allowed here, and it'll throw me out. The point is to get around that barrier of inflexible protocols and software and to provide plug-and-play connectivity as well. The goal of nomadic computing is to be able to access high-speed bandwidth and services seamlessly, as effectively as you had back in your original corporate or home environment. And of course a lot of infrastructure is required to make that happen. It's beginning to be deployed: we now have a lot of plug-and-play capability. When you go to a hotel and you plug into the Ethernet connection there, it's a rather easy experience for you; there's a step of authentication, a step of providing a way in which you can pay, and there's a question of how much bandwidth you get. You may want a VPN back to your home or corporate environment. You may want just general Internet access. And you can begin to get that. But when you leave, say, a Wi-Fi hotspot now what do you do? Well, you'd like to be able to take advantage of the 3G system in a seamless fashion. That's beginning to happen as well. We're talking about handover within 3G and between 3G wireless, Wi-Fi and Bluetooth and WiMAX. Those technology interfaces are being developed.

**UBIQUITY:** How fast has nomadic computing been developing?

**KLEINROCK:** Nomadic computing has been developed considerably over the last 14 years since I started working on it in 1991, but it still does not provide a seamless interface. One of the problems is that the devices we carry around are varied and unique in many cases. And when you consider someone who's carrying around a cellphone, an e-mail pager, a laptop, a PDA, maybe a GPS unit, maybe a scientific watch as I do, they tend not to interoperate.

**UBIQUITY:** And people carry around all that stuff?

**KLEINROCK:** I'm carrying probably five speakers, two microphones, five different batteries, five clocks, three different databases, two or three different telephone databases. And, yes, they tend not to interoperate. So there's a move these days, as you've probably seen, toward convergence. They're trying to put the function of many of those devices into one device, a kind of handheld cell phone, PDA, maybe GPS, e-mail capability, paging system, et cetera. And a number of those devices are coming out. HP's come out with one. The Treo is a partial step there with multiple radio technologies converged as well. Now, this is a good thing. Of course, you can argue both ways, and I'll be happy to do that in a moment. But the ability to converge all of that technology into one device that performs all those functions for me is a big step forward. The thing that's amazing, which I don't think too many people realize, is the amount of content that's being delivered to these handheld devices. Do you know there's a multi-billion dollar industry in selling ring tones? You've probably never spent 99 cents to buy your ring tone, but the kids these days do that. It's a \$4 billion industry.

**UBIQUITY:** \$4 billion on ring tones! Amazing.

**KLEINROCK:** So somehow we're out of touch. During the SuperBowl if you paid \$1 you could put in a vote for the most valuable player right there online through your cell phone, because some of the sports networks were doing that, or get the immediate scores. And the content both in terms of interaction, in terms of the ability to download and upload images, the video that's coming across these devices, the content is just beginning to explode. Sports is a very large driver. Gaming is another large driver. As I've said, ring tones is one as well.

Moreover, it would be nice to have a seamless interface among the various radio technologies for various distances and bandwidths; you would then have a ubiquitous and continuous capability as you move around in your environment. That whole mosaic and framework is beginning to be developed right now.

**UBIQUITY:** Is there any difficulty converging voice and data on those gadgets?

**KLEINROCK:** Not from a technical point of view. Voice-over-IP is a big driver for these devices. And they're solving some of the latency problems and the quality of service problems. Fourth-generation (or 4G) phones are going to be up at high Ethernet speeds, for example, and WiMAX will be up to 75 megabits per second. So, you're going to have plenty of bandwidth. And the convergent device you carry around will render various kinds of content. The handheld device will be a *communicating, multi-function rendering device*; this device was earlier known as a cell phone.

**UBIQUITY:** Expand on that.

**KLEINROCK:** The device you carry around will render images and video to match the screen and other characteristics of that device. If I transmit an image from one handheld device to another it might be the wrong size, may not be the right set of colors, the orientation might be different, and the resolution might be different. So, you need some kind of rendering capability to match the device that's receiving it.

**UBIQUITY:** The Japanese seem to be using their cellphones quite a lot for buying things from vending machines and such.

**KLEINROCK:** Yes. Their vision is to adapt the cell phone to serve as your electronic wallet. And there's a lot of credibility for that. Today it is already used for getting Cokes out of a Coke machine. You may want some Wi-Fi capability right next to the Coke machines, so you can communicate with them to see the level of inventory and cash in those machines while at the same time use those locations as Wi-Fi hot spots as well as vending hot spots. So, what is this new handheld device? There have been different visions. The traditional view is that it was a phone; in Hollywood they think of it as a tiny television set; in Silicon Valley they see it as a PDA; and the game industry sees it as a GameBoy. It's the old blind men and the elephant problem. It can serve a lot of those functions. Actually, the cell phone is now being called the third screen from the entertainment point of view. The first screen was the movie screen, the second was the TV, and the third now is your cell phone screen.

**UBIQUITY:** Could you take a communication into a small gadget and then if you walk into your house shoot it up onto a wall screen?

**KLEINROCK:** Absolutely. That's where the rendering comes in. You have stored perhaps a high-resolution image in your phone, but on your phone you can only see it as a lower resolution. But the underlying high-resolution capability is still there. So when you transmit it to your television or your computer display it'll bloom out into the full resolution that's been stored.

**UBIQUITY:** And that's no big deal?

**KLEINROCK:** No, because storage is getting very cheap these days. And if we do in fact provide the bandwidth within our homes and into our homes, then some of the storage can be out there in the Net. The question is where will all this storage be located? Will it be on the Net? Will your profile, your privileges, your preferences and your content be out on the Net or in your device? That's an open question. There won't be too much carried around in your handheld device because of the storage limitations, but to get a terabyte of storage in your laptop these days is trivial.

But there's a converse side to this. The discussion I just had was to say that considerable function and content is converging into one device. But do we really want one device?

**UBIQUITY:** Apparently not. Why not?

**KLEINROCK:** Look at teenagers who are running their iPods. One function they want on an iPod is a big button that allows them to pause the music. They don't want it buried deep down as a soft key in the menu somewhere. They want to be able to pause by pushing a button. Well, that's a special-purpose button that's not useful for a lot of other functionality and there is a question as to how much real estate should be devoted to that special function. They also like devices with more than two

batteries, so that when the iPod battery runs dry they can still call home on the phone. So I think we're moving very quickly towards convergence -- but then we're going to see a move back to divergence.

**UBIQUITY:** Give us an example.

**KLEINROCK:** Well, for example, right now many of us use Bluetooth earpieces with our cellphones. We put the phone in our pocket or in the automobile and we communicate via Bluetooth from an earpiece. That's a kind of divergence. A pacemaker that might be implanted in your body or things you can carry around in your belt or on a head-mounted display -- all those things are not encased in one device and as such represent a kind of divergence again. That means we need to communicate with our immediate local environment. And so now we're moving into another domain, and you've probably seen some of my writings called "Smart Spaces," where all the technology is not behind this screen I'm carrying around with me. Instead it lives in the physical environment in which I move around -- in my desk, in the walls around me, in my shoes, in my belt, in my eyeglasses. And these things collectively represent the local sphere in which I interact. My environment will be alive with embedded technology rendering it intelligent, or, if you will, a smart space. So whereas you may see convergence in a handheld device, that device will interact with other devices in its local environment. In other words, there's kind of a convergence of a lot of functionality into that device and a kind of divergence beyond that to devices it interacts with. We're probably going to see more divergence as we move forward after this strong convergence move.

**UBIQUITY:** Then nomadic computing is the next frontier?

**KLEINROCK:** Yes, and beyond nomadicity will be two other pieces. One is smart spaces, namely the environment of devices all around me that are intelligent and help me in my nomadic quest. The second is ubiquity, i.e., I want this capability everywhere. Nomadic computing means if I go from A to B, then I'm accepted as a friendly at B. And how many Bs are there? If there are only two or three Bs around that's not very useful. I want to be accepted everywhere. And that's where the notion of ubiquity comes in. Ubiquity also refers to the availability of access everywhere. As those three pieces get deployed, then, with a nomadic capability, I can be recognized when I get to other locations and this will occur in many, many places and those places will themselves be smart spaces, full of technology in the environment with which I can interact.

**UBIQUITY:** Give us another example.

**KLEINROCK:** I may not even want to carry my own device around with me. I may use any device that I see when I arrive there, which means my context and profile will be in the network and instantiated on whatever device I walk up to once I authenticate myself. That device I touch will become my personalized device for the time being while I'm using it. Now, that's a kind of grand vision, but I believe that's where we're heading. In fact, I'd like all of this technology to become invisible, just like electricity's invisible. When I walk into a room I expect electricity to be there. I don't care how it gets there and how it works. When I plug in I want electricity to appear. Similarly, when I walk into an environment I want all the Internet capability there and it should recognize me and my profile, etc. And it should be hidden from me in the sense of invisibility I just mentioned. I should be able to talk to the environment. It could respond with voice, or perhaps displays pop up or holograms.

Maybe there's a keyboard around. Maybe there's some kind of other input device, or output device. But I want it to be there when I get there and not have to import all of the technology with me. Indeed, we are talking about mass customization, a concept which is certainly enabled with today's technology.

**UBIQUITY:** That's quite a vision.

**KLEINROCK:** Yes. And it's going to take a while to occur. It's going to appear in certain places first, for example in some of the advanced laboratories. It will be used first by the early adopters. But eventually it'll be in our homes and our offices and the usual places we frequent, be it airports, hotels, convention centers and hopefully eventually wherever I choose to have it.

*[END]*

Source: Ubiquity 12-19 July 2005 <http://www.acm.org/ubiquity>