5

Computing and Communications Unchained: The Virtual World

Leonard Kleinrock
John Major

Leonard Kleinrock

I am tempted to suggest to the information technology users of the world, "Unite and throw off your chains." So let's do that.

I am going to talk about "Computing and Communications Unchained: The Virtual World." We could have called it the "virtual universe," but we thought we would restrict our focus. After all, academics are narrow (wasn't that the word used this morning?). I will talk about the virtual world of the nomad, discussing some of the issues and the technology. John Major will then talk about the applications in which you begin to see nomadicity happen.

So let us start at the right place, a dungeon. Most of us associate our computers, such as they are, with some kind of desktop device, possibly connected to a server, located down in somebody's dungeon. You never see it. You are rigidly attached to that architecture.

In fact, most of us are nomads. I do not know how many laptops there are in this room, but there are more in your hotels or automobiles. We travel everywhere. We travel to our office, home, airplane, hotel, automobile, branch office, and bedroom. My wife will not let me use a laptop in bed any more. What bothers her is the noise of the keyboard: it is forbidden.

Of course, we also travel to places like this symposium. When I go on the road, I usually travel with a laptop computer, pager, cellular telephone, and personal digital assistant (PDA). It used to be that shelf space in a bookstore or counter space in the supermarket was a precious commodity that all merchandisers would fight for. Today, they are fighting for waistline space to hang all this stuff on. You can rent out square inches of your belt these days. When I load up and go out into the world, I feel like Pancho Villa going into battle with the U.S. cavalry.

So what is nomadicity? In my mind, it is basically the system support needed to give all kinds of capability to nomads—no matter where they go—in an integrated, transparent, and convenient fashion. You should not have to suffer because you are moving around. Today, you suffer a great deal, so let us talk about what is needed and what some of the issues are.

Why should we care about nomadicity? There are many reasons. It is here right now; the users see it. We all move around and experience problems with synchronization, updating, access, weight, and more. I think there is a major paradigm shift in computing. It may not be as important as information warfare, but I think this is a dominant trend, not just a tangential issue. This is how people are now and how they will be using information

What are the components of the system design you have to worry about? Well, here is a list of the usual components: bandwidth, latency, reliability, error rate, delay, storage, interoperability, user interface, and cost. No surprises here; these are the usual suspects that we always worry about. However, there are some other things you have to worry about when you move to a nomadic environment: physical size, weight and processing power (how many Pentiums can I put on my lap?), battery life, communications, interference in the radio world, and damage. Moreover, these technologies are portable, which makes them easy to lose and subject to theft. I am sure you have read about the latest European scam that has now hit the United States. You are at an airport. You pull your luggage with the laptop attached. You reach the x-ray unit at the security gate. One person goes through the x-ray machine and his cohort is directly in front of you. When the cohort gets there, he fumbles and takes a few minutes to get through. Meanwhile, your laptop is on the way through. The first person gets it and disappears, and you have lost your laptop. This is a very easy way to lose a machine. Do not take your laptop to airports unless you watch where it goes.

Nomadicity exacerbates several concerns: disconnectedness; variable connectivity, either because the world does it for you or because you choose to move and use some other communications medium; latency, which varies significantly; variable routes in virtual circuits as you move around; variable requirements that you put on the system—what you expect and need; replication of resources such as files, machines, databases, and applications because you are moving; and foreign languages.

When you go to a new environment, you have to locate the power supply. Where is the modem? Where is the high-resolution screen or printer? You must become familiar with the environment. Conversely, the environment has to become aware of you. It must know that you are there. It should know your profile and what you want and should send the things to you that you expect. As the bandwidth and the platform capabilities change, the system should adapt what it sends to you. For example, it should compress some things and not send you high-resolution video. Perhaps it will send only the name of the movie, if all you have is low bandwidth. This is a major issue. Most of all, nomadicity is one of the ultimate problems in distributed systems.

What should we do? First, develop the systems architecture and network protocols. We need to interoperate between the wireless and the wired worlds to handle the concerns of unpredictable user, network, and computing behavior and to provide graceful degradation in all of this (simple problems!).

Second, we should develop a nomadicity reference model. We have heard today about one of the reference models that CSTB produced (Figure 5.1).[1] It has various names, the hourglass-shaped open data network (ODN), with a control level referred to earlier today as the bearer service. The ODN could provide a reference model. Suppose I wanted to send some e-mail. In one environment, I might connect to a cellular bandwidth source. A few minutes later, I might find that I can use a modem. Maybe I will walk into an office and get a 10-megabyte local area network, and maybe I will get a 150-megabyte ATM (if that ever appears in the office). As I move among these choices, what do I have to do? Well, right now I have to do a lot of things. I have to plug in a different PCMCIA card, reconfigure, reboot, and put down some IP addresses. I do not ever want to have to do this. I want something to do it automatically for me. This is one of the things that a nomadic system design should be able to take care of.

Another architectural model might be a more standard one—network infrastructure, support, and middleware. Hopefully, we will achieve integrated nomadic support as opposed to velcro integration, which we have today in many of our belt-laden devices. We happen to have one of these models at the University of California at Los Angeles. We filled in some of these architectural pieces (Figure 5.2). It has happened in a number of locations where you provide connectivity management, some file synchronization, and update schemes.

What else needs to be done? You have to understand how the system works and develop performance models. There are a lot of choices here. You can make a mathematical model, but the mathematics is not that strong. You can do numerical evaluation, but then you run into an exponential explosion of computation. There are iterative solutions, but then does the darn thing converge? Simulation—it is hard to search a large space. Emulation—it is

[1] Computer Science and Telecommunications Board, National Research Council. 1994. *Realizing the Information Future: The Internet and Beyond.* National Academy Press, Washington, D.C., Chapter 2.

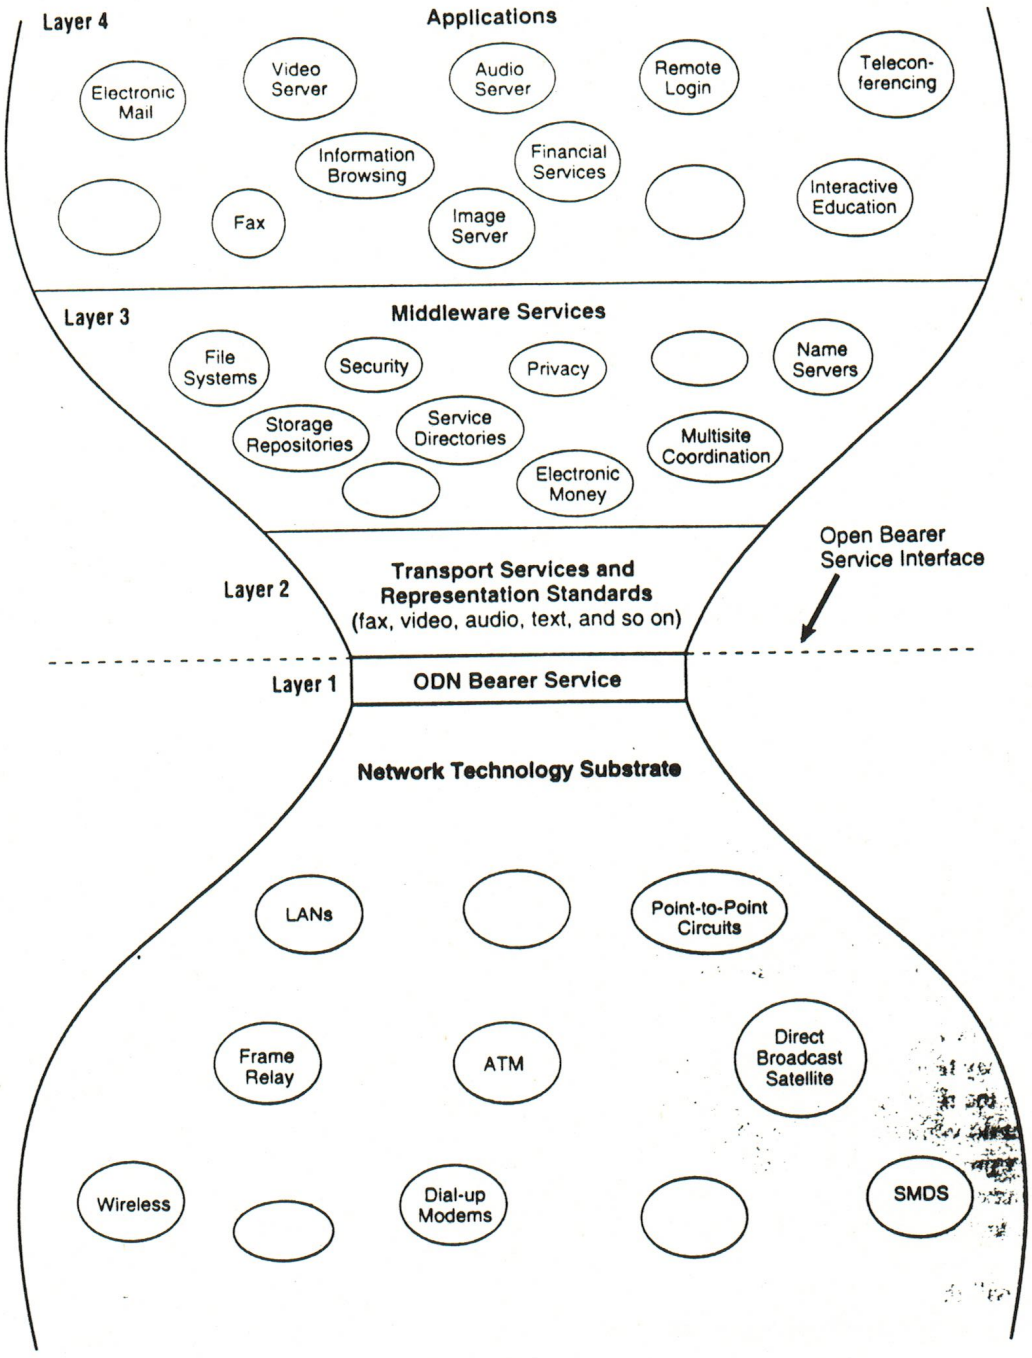

FIGURE 5.1 A four-layer model for the Open Data Network. Reprinted from Computer Science and Telecommunications Board, National Research Council. 1994. *Realizing the Information Future: The Internet and Beyond.* National Academy Press, Washington, D.C., Figure 2.1.

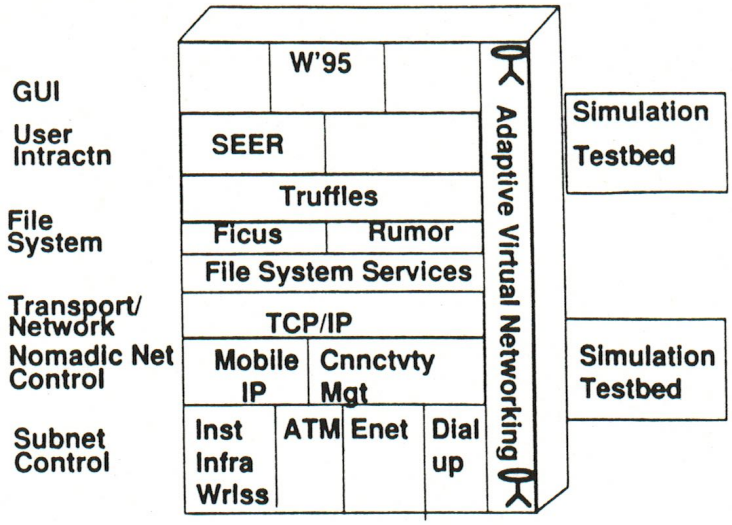

FIGURE 5.2 Architectural model for integrated nomadic support developed at the University of California at Los Angeles. Reprinted with permission from Leonard Kleinrock. Copyright 1996 by Leonard Kleinrock.

sloppy, ugly, and expensive (building the system and measuring how it behaves are certain to bankrupt you). The right answer is probably some hybrid mix to do the part that works best in each environment.

Consider adaptive agents (Figure 5.3). The classic assumption is that there is a big fat network between client and server. However, it is not client-server, it is client-network-server. The assumption is that you put a lot of stuff here because a big network will deliver whatever you need. In the nomadic environment, though, the network may be thin or zero. As the network skinnies down, you may want to move some functionality around in anticipation of a need for tools and data that you do not have with you. Adaptive agents at the middleware level are also called surrogates, proxies, helpers, and knowbots. We need a kind of theory or formalism, an architecture, a language for agents. They should help the nomads, the applications, the network, the servers, the communications devices, and the computing devices. They can sit everywhere. In a peer-to-peer application, or maybe a client-server application, agents may help to decide what goes on. We certainly need some of these adaptive agents inside the network as well to do the thinking for us, the compression, the connectivity management, and so on. This is the challenge.

So what do you observe from all of this? Nomadicity is here, you cannot escape it, and the needs are real. It makes every problem you ever thought about much harder. Each one gets an order of magnitude more difficult in this environment. It is a fascinating area (I think). The payoff can be huge. There is a severe lack of any integration, and there is chaos and confusion right now. You simply cannot afford to ignore the challenge of nomadicity.

Where will the next innovation come from? Ask your children—they will tell you, we won't.

John Major

I had a professor years ago who taught me that it is key to watch the moments in life when big things—new things—happen, because important things will follow. Words matter. When was the first time you heard the word *picosecond* referring to the speed of operation and *gigabyte* referring to storage capacity, or other words like those?

While Leonard Kleinrock and I were getting acquainted to do this presentation, he pointed out that he had just obtained a new laptop with two 1.2-gigabyte hard drives in order to have enough information when he is

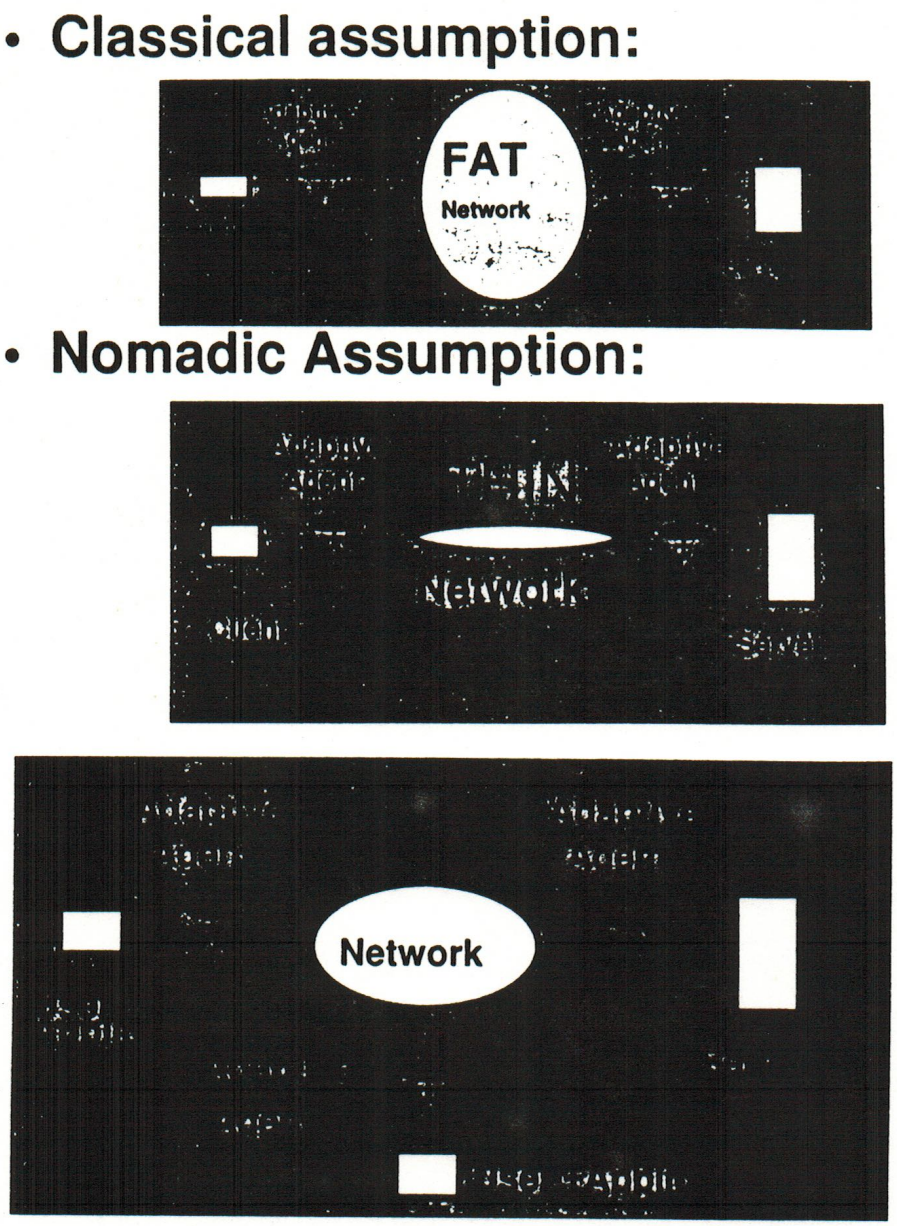

FIGURE 5.3 Adaptive agents. Reprinted with permission from Leonard Kleinrock. Copyright 1996 by Leonard Kleinrock.

"nomading" around. This was the first I had heard of two 1.2-gigabyte hard drives in a laptop, which is serious mobile computing. So I started to think that the vision I have of the virtual world will be well received.

I told Leonard about my concept of teleconferencing—that people at some point will travel virtually and no longer physically. In fact, that is the title of this talk, "Creating the Virtual World." With all sincerity he said, "That will never work!" So if I could not convince someone with two 1.2-gigabyte drives in his laptop that we are going to have a virtual world, perhaps I do not have much chance with the rest of you. The point is that this is a difficult concept. There is a lot to gain, but also a lot needs to be developed to make it possible.

Does anyone remember the 1987 movie *Planes, Trains, and Automobiles*, with John Candy and Steve Martin?

I travel a lot. I find it impossible to watch this movie all the way through because it hurts too much. Everything bad that has happened to me on one trip or another—not having money, the rental car not working, and the airplane schedules being unreliable—my whole life is being played out before my eyes in 120 minutes. It is a great movie, but it encapsulates every bad travel experience, and as such, it humorously captures part of the need for a virtual world.

You would think we would be highly motivated to change this. In fact, when I committed to this concept of describing the virtual world, I figured I would just go out, grab a few books, and catch up on the emerging theories. Then I would explain it to the audience, leave with loud applause, and feel very comfortable with myself. I discovered that, today, there is no agenda to create a virtual world. In fact, on some level, most of us believe that what we go through when we travel is important, part of what we do, and necessary to our society. In other words, we feel that there is no other way.

Yet, we know that travel is inefficient and that it significantly lowers productivity. It also costs a lot of money. The last time I looked, the federal government was spending about $40 billion a year to maintain the highway system. Office buildings cost money and isolate workers from their customers. Meanwhile, we are competing in a global world. If I believe it is important to be able to walk in to see my Motorola team and rally their spirits and convey information to them, and I do this for the team in Schaumburg, what about the team in Bangalore? What have I done about our lab in Australia?

If I think that being physically there matters, I should never go home—simply fly forever. Well, you cannot fly forever. With a truly global company, you cannot be in enough places. All of your time would be consumed in airplanes. So there is something wrong if you really want to run a global corporation or have a global team be as productive as it should, unless you are able to create, in some way, this virtual world.

How do you do it? How do you eliminate physical travel and commuting? How do you enable virtual travel and commuting? I will offer that it is not going to be via e-mail, fax, teleconferencing, or videoconferencing as we know them today. These represent a beginning; they are important steps, but they will not solve the problem. They may not even be reflective of the form of the eventual solution. To illustrate this point, imagine if low-cost robots had been invented by 1940 and were ubiquitous. Then someone said, "I am going to invent the dish-washer." Can you imagine what that dishwasher would look like? It would look like a robot doing the human tasks for washing dishes. As it was, there were no low-cost robots, so someone had to be very innovative. Today, as a result, dishes are washed in the dishwasher very efficiently, very conveniently, and completely differently from the way people wash dishes. This is the kind of fresh thinking that may well be required for virtual travel.

We have already seen some examples in communications of how new technology and uses change old paradigms. Today we have e-mail, which is in no way an analogue of what we saw before in the old teletype and telegram era. It is a completely new embodiment of messaging—transitional, to be sure, but very different. Remember centralized fax? One of my favorite memories is the time I suggested that I should have a fax in my office so I could communicate better with my colleagues around the world. I got a stern note from corporate, pointing out that fax was something that we had to be very careful about. There was an office you could go to if you wanted to send a fax. It was a couple of hundred yards from where I was. This centralized fax analogue has nothing to do with how we use fax today. So we have been able to step into new paradigms, but they are always very difficult to envision. In 1976, Jack Nilles said, "Telecommuting is the substitution of telecommunications for travel." This was the first time that the term was used. In 20 years, not much has changed, but there has been some progress.

There is some work being done in this area. There is some hope. There is some evolution. BellSouth has made a prediction that by the year 2000, 25 percent of corporate employees will, in fact, be working outside their traditional offices. By then, we will have evolved to the point at which a substantial amount of the work product will be separated from the physical office.

There are signs of progress. The growing use of e-mail is one. We are, in fact, learning to work as a global team with the primitive tools we have today. At Motorola, we have 140,000 employees and 120,000 computers. Since I manage this activity, it begs another question I am frequently asked by the chief executive officer, How can Motorola have almost one computer per person when well over half of our people work on assembly lines and have nothing to do with computers at all? Part of the answer is that we use these computers to run our factories, but the

countries around the world. So I do not think we should shrink from the magnitude of the problem. We must understand that interaction may be a critical piece of it. There is surprisingly little work at understanding what the issues are.

ROBERT BONOMETTI: A few weeks ago, the Federal Communications Commission had what I think is going to become a historical notice of proposed rulemaking whereby it proposes to take 350 megahertz of spectrum and allocate it for a so-called national information infrastructure band or a supernet band or a hyperband, depending on whose buzzword you like. What are your thoughts as far as how this may be a critically important enabler for nomadic types of computing?

LEONARD KLEINROCK: Well, the problem of getting at the nomad when he is out of harm's way—for example, near fiber or wire—is a major one. Certainly the wireless spectrum provides a solution, be it satellite or local wireless. I think this is very important. I am anxious to see how the technology buys into it and what the devices are. This kind of freeing up of spectrum is important, but I do not think it is the most important problem. I think the systems issues are far more difficult and have to be developed first, but this is a component.

DAVID FARBER: A couple of months ago the dean at the University of Pennsylvania described me as his first virtual professor. He said it with pride in front of the board of trustees. With respect to John's comment about there needing to be some people supporting change, I think you are slowly beginning to see this when a largely liberal arts institution can absorb the concept without collapsing.

Let me make a couple of comments and then ask a question. When I travel (I am probably one of the most gadget-prone people around, courtesy of John and others), I tend to have a lot of gadgets. Half of my luggage consists of power supplies, batteries, everything incompatible with each other. So I would argue that there is a real need for standardization. It is pretty ridiculous when I have to carry five power sources so I can keep running.

I think that you alluded to a serious problem, Lenny, but let me see if I can draw you out a little bit. One of the big problems is the change of context every time I move from device to device, every time I move from one bandwidth package to another. I am sitting in an office, allegedly in the future, watching you via video link. I walk outside to my car and get on the road; I am on the way home, and I arrive at home. Each time that I pass through different bandwidth domains, everything seems to want to change. I think it is the discontinuity that I, at least, find very, very difficult to handle. I have to carry too many models in my head. Do you think this is going to change?

KLEINROCK: I think it will improve. It will never be as ideal as you would like. There will be ghostly images of me following you one way or another, be it by voice as you are driving or via a still picture of my unshaven face as you are shaving. There will be components like that. I think you are quite right. The discontinuities in bandwidth should be smoothed out somehow. One can anticipate this. One can, in fact, cache some of it. Predictive caching is a big component here, and it is very hard to do. This is why I say that there are a lot of good research problems, but these issues are key.

MAJOR: I think Len's concept of the intelligent agent that rounds off some of the edges and, in fact, modifies the communications environment automatically in response (where you simply get the headline and then can probe for the rest of the article) is the kind of concept that could be really contributory. To me, it is a fresh thought that the middleware community has not touched so far.

PATRICE LYONS: I was interested in your reflection on trying to liberate a few billion dollars from local and perhaps national highway programs to support the development of infrastructure. Well, if you had a bag of money and were a local government official, what minimal elements would be at the top of your priority list to encourage the kind of things you are thinking about?

What comes to mind from some of my legal debates in the last couple of years is the question of what is the proper role for government in a hierarchy of certifying authorities? For example, if you are a lawyer working in a suburb, your main office is downtown, and you are discussing things of an extremely confidential nature, you would want to make sure that there was some minimal integrity in the communications pathway. Perhaps this would encourage more people to do so. What would be the top one or two things that you would think about?

MAJOR: I will suggest that I do not think there is a quick answer here. I think the first pile of money has to go into basic research to understand the nature of the problem and free up the solution set from what would appear to be the direct path to what I described—to what might be a successful path to solving the real problem.

JERRY MECHLING: One of the problems we faced with the London project was about 200 to 300 people

around the world who were not paid by us. They were like a network of different governments, universities, different entities—a complete mix. We were going to standardize on some collaborative work. We were going to do Lotus Notes. We decided not to, because the training requirements became too difficult and you had to fly them all someplace to get trained anyway. If you used video, you had to have all this language capability to translate, which was a problem because you had to explain the translation of the words. It became such a quagmire to figure out whether there was a way to standardize that we ended up with a network that was a complete hodgepodge of everything from PCs to fax machines to the most advanced systems.

Is there anything in the works that is going to address what I think is a problem people are not looking at: how do you conduct the collaboration when the players are not in the same organization? If it involves just one company, you would make a corporate decision, but you cannot do it in an interagency deal.

KLEINROCK: From a technology point of view, I cannot give you any answers to that question. I mentioned foreign language, the one you selected as a problem. Sure, translators will help, but they are a long way off. I am not sure what technology will overcome that problem. I think it is too early to standardize on a lot of this. We are just probing now; I would be frightened to do it at this point.

DEFINING A DECADE

Envisioning CSTB's Second 10 Years

Computer Science and Telecommunications Board
NATIONAL RESEARCH COUNCIL