Available at
www.ComputerScienceWeb.com
POWERED BY SCIENCE DIRECT®

Ad Hoc Networks 1 (2003) 3–11

www.elsevier.com/locate/adhoc

An Internet vision: the invisible global infrastructure

Leonard Kleinrock

Computer Science Department, UCLA, Los Angeles, CA 90024, USA

Abstract

This paper takes us on a short journey from traditional desktop computing to the three emerging dimensions of: *Nomadicity* (the system support needed to provide a rich set of computing and communication capabilities and services to nomads as they move from place to place in a way that is transparent, integrated, convenient and adaptive); *embeddedness* (small intelligent devices embedded in the physical world and connected to the Internet); and *ubiquity* (Internet service availability wherever the nomad travels on a global basis). These three dimensions give us a powerful system that supports global access for mobile users interacting with smart spaces. When we add intelligence distributed across this global infrastructure, we form a new "space", which is the basis of a vision I articulated for the Internet in 1969 and which has yet to be achieved. In that vision, I foresee that the Internet will essentially be an invisible global infrastructure serving as a global nervous system for the peoples and processes of this planet.

Keywords: Internet vision; Nomadicity; Embedded devices; Smart spaces; Ubiquitous computing

1. Introduction

This journal's specific goal is "complete coverage of all topics of interest to those involved in ad hoc and sensor networking areas". As the lead-in paper for this journal, I choose to take a longer-range vision with a broad view of what domain that vision should encompass and what that vision will look like. Specifically, I will not deal with the link-level and network-level issues so common in ad hoc and sensor network studies; these will be amply covered in this journal. Rather, I choose to discuss the implications of the wide availability of those and related emerging technologies. Indeed, as these technologies deploy broadly in the Internet, we will see the emergence of an invisible global infrastructure that serves as a pervasive global nervous system with enormous potential and capability.

Since the beginnings of telecommunication technology 100 years ago, we have witnessed a number of major shifts in the application of communications technologies to the needs of our society and industry. In that process, we have seen the marriage of wireline and wireless technologies, of analog and digital technologies, of voice, data, video, image, fax, streaming media and graphics to create a computer-communications infrastructure that spans the globe and serves billions of people. The Internet is the current manifestation of these developments and has penetrated every structure of our society.

We are now in the midst of an accelerating groundswell in this field of computer communications

E-mail address: lk@cs.ucla.edu (L. Kleinrock).

1570-8705/$ - see front matter © 2003 Published by Elsevier B.V.
doi:10.1016/S1570-8705(03)00012-X

in its most visible and useful sense—not simply the wires and networks, but also the infrastructure, middleware, applications, services, modes of use, and users of the technology. Let us begin with the vision and its history.

2. The vision

On July 3, 1969, UCLA put out a press release [8] announcing the forthcoming birth of the Internet (known originally as the ARPANET) which would take place two months later on September 2, 1969. The opening sentence of that press release begins, "UCLA will become the first station in a nationwide computer network ...". In the final paragraph of that press release, I am quoted as saying, "As of now, computer networks are still in their infancy. But as they grow up and become more sophisticated, we will probably see the spread of 'computer utilities' which, like present electric and telephone utilities, will service individual homes and offices across the country". The "computer utilities" comment foresaw the emergence of web-based IP services; the "electric and telephone utilities" comment foresaw the ability to plug in anywhere to an always on and "invisible" network; and the "individual homes and offices" comment predicted ubiquitous access. Basically, I articulated a vision of what the Internet would become. The part I did not include in my vision 34 years ago was that my 95 year-old mother would be on the Internet today (and indeed, she is).

That vision for the Internet can be broken down into five elements:

1. The Internet technology will be everywhere.
2. It will be always accessible.
3. It will be always on.
4. Anyone will be able to plug in from any location with any device at any time.
5. It will be invisible.

The Internet *almost* got it right. Indeed, the first three elements have already come about. However, the Internet as we know it today has not yet achieved the last two elements of the vision. Why have these last two been lagging?

Basically, the mistake regarding element 4 above—any device plugged in at any location at any time—was that the Internet's TCP/IP protocol assumed that end users, their devices, and their IP addresses would all be found in the same location and would all be tightly coupled. These were correct assumptions at the time since the mentality then was that of a deskbound model for computing platforms. However, it is no longer the case that we stay at our desks. Rather, we are nomads and we travel constantly from our office to our home, airplane, hotel, automobile, coffee shop, branch office, conference room, bedroom, etc. The fact is that end users do not always access the Internet from their fixed-location offices, do not always use the same device, and the IP address they use may not be one familiar to every subnetwork they encounter in their travels (indeed they may use different IP addresses in their travels). That is, the users are nomads, and the issues associated with nomadic computing were not anticipated by the network protocols that grew up in the Internet. Indeed, we have now entered the era of *nomadic computing* wherein the mobile or nomadic user seeks to be provided with trouble-free Internet access and service from any device, any place, at any time.

The problem with element 5—invisibility—is that the Internet is anything but invisible in the sense of being easy to use in ways that do not assault our human senses with irritating input and output interfaces. (A similar consideration regarding the invisibility of computers was recently addressed in [6].) However, the rise of ad hoc networks, sensor networks, nomadic computing, embedded technologies, smart spaces, and ubiquitous access, enable cyberspace to move out into our physical world and open up new vistas. The concept of these technologies disappeared into the infrastructure (as has electricity) suggests some far-reaching capabilities in terms of how these disappeared technologies are organized into global systems that serve us and our information and decision-making needs in adaptive and dynamic ways. Mark Weiser had a vision for invisibility and disappearing technologies which he articulated in [9,10] and an entire issue of a new publication is devoted to his early ideas [7].

This paper will take us on a short journey from traditional desktop computing to the three emerging dimensions of: *nomadic computing and communications*, or *nomadicity* (the system support needed to provide a rich set of computing and communication capabilities and services to nomads as they move from place to place in a way that is transparent, integrated, convenient and adaptive); *embedded technology*, or *embeddedness* (small intelligent devices embedded in the physical world and connected to the Internet); and *ubiquitous access*, or *ubiquity* (Internet service availability wherever the nomad travels on a global basis). These three dimensions form a new "space", which is the basis of where I see our technology heading. I will elaborate on these three dimensions and then add distributed intelligence to provide the grand vision of an invisible global infrastructure that I foresee.

3. Nomadic computing and communications: nomadicity

The first "dimension" of the new space is that of *nomadicity* (nomadic computing and communications). Nomadicity refers to the system support needed to provide a rich set of capabilities and services to the nomad as he moves from place to place in a transparent and convenient form [4].

Until the mid-1990s, most computer users traditionally functioned in a world of desktop computing appliances that were connected through corporate or private networks to servers. But, in fact, most computer users are nomads, moving between office, home, airplane, hotel, automobile, coffee shop, branch office, conference room, bedroom, etc. In so doing, we often find ourselves with significant variation in the computing platform to which we have access, in the quality of the printers and displays that are available, in the communication device we use as well as in the communication bandwidth that is available to us (including the now-common case of complete disconnection). Moreover, we may choose to do computing and/or communication while on the move.

Today's variety of portable *computers* is impressive, including laptop computers, notebook computers, tablets, handheld PCs, PDAs, smart credit card devices, even wristwatch computers. In addition, the *communication* capability of these portable computers is advancing at a dramatic pace ranging from low-speed modems, to Ethernet, to email receivers on a card, to spread-spectrum hand-held radios and cell phones, to cellular digital packet data transceivers, to portable GPS receivers, to WiFi, to Ultrawideband, to gigabit satellite access, etc.

This combination of portable computing and portable communications is changing the way even the casual home computer user thinks about information processing. The business user, as well as every other kind of user, now recognizes that access to computing, communications, and services is necessary, not only from our "home" base, but also while we are in transit and when we reach our destinations.

Even without portable computers or communications, many people travel to numerous locations in their business and personal lives and thus require access to Internet services through equipment that is available when they arrive at their destination. Indeed, even moving from one's desk to a conference table in the same office constitutes a nomadic move, since the computing platforms and communications capability may be considerably different at these two locations. When a user arrives at a "foreign" destination, that user appears to be an alien in that environment. Our nomadic technology must provide the ability for the user to easily gain access as a "friendly" in this new environment.

A fundamental way in which nomadic computing differs from conventional desktop operation is the huge variability in connectivity to the rest of the user's computing environment. That level of connectivity often includes extended periods of low bandwidth or no communication capacity at all. Since many users and programs alike make intermittent, but nevertheless essential, use of "off-machine" information and services, they are unable to operate effectively unless extraordinary steps (like reconfiguring their IP address, changing their netmask, removing their proxy, etc.) are taken by sophisticated users or their network administrators. The goal of nomadic computing is precisely to permit users and programs to be as efficient as possible and as unaffected as possible in this environment of uncertain connectivity and unfamiliar locations. That is, nomadicity makes the sometimes-connected computer operate in the same way and as effectively in a foreign location as when it is connected as a friendly to his organization's information network.

The researchers and developers creating the invisible universal infrastructure are deeply interested in those capabilities that must be put in place to support nomadicity. The characteristics needed for nomadicity include independence of: location (while at the same time providing location-aware services), motion, computing platform, communication device, and communication bandwidth. All of this must come along with widespread availability of customized access to remote files, systems and services. The notion of independence, as I define it, does not necessarily refer to the quality of service one receives, but rather to the perception of a computing environment that *automatically adjusts* to the processing, communications and access available at the moment and at that location, all customized according to that user's profile, preferences and privileges. For example, the bandwidth for moving data between a user and a remote server could vary from a few bits per second (in a noisy wireless environment) to hundreds of megabits per second (in an Ethernet environment). The computing platform available to each user could vary from a low-powered handheld PC while traveling, to a powerful supercomputer in a climate-modeling science laboratory.

Nomadicity exacerbates a number of the problems users face routinely every time they turn on their machines. The nomad experiences disconnectedness, variable connectivity due to voluntary changes (possibly traveling), or unpredictable changes (possibly a noisy wireless connection), variable routes through a network (possibly changing virtual circuits), variable requirements, resource replication (possibly copying files in multiple locations and devices), the need for the nomad to become aware of the changing environment, the need for the environment to become aware of the nomad's presence and location, the need for adaptivity to accept the nomad in "alien" environments, and a general need to manage distributed "stuff" including applications, files, services and numerous distributed resources. All of this networking complexity should be hidden from the user and managed by intelligent technology supporting that user.

Some key system parameters with which the user must be concerned include bandwidth, latency, reliability, error rate, delay, storage, processing power, component-to-component interface, interoperability, user interface, and cost. These are typical concerns for any computer-communication environment; what makes them of special interest for us in the context of nomadic computing and communications is that their values change *dramatically* (and sometimes suddenly) as the nomad moves from location to location anywhere on Earth. In addition, the nomad has some totally new and primary concerns such as weight, size, battery life, loss, theft, and damage to portable devices.

To address the changing broadband marketplace, service providers must deploy value-added web-based self-provisioning and service creation technologies in their networks. One way to provide this flexibility is to place intelligence at the edge of the network so the network adapts to the nomadic users appearing at that edge, instead of asking the users to adapt to the network. The edge may be defined in a number of ways, but perhaps the most effective is to recognize that the edge is that place in the network where the unmanaged collection of end user devices (such as laptops, palmtops, email pagers, and IP-enabled cell phones) first meets the managed infrastructure of the Internet.

In the office, users have three wonderful resources: (i) a high-performance workstation (e.g., a laptop computer), (ii) high-speed access to the Internet (e.g., 10/100/1000 Ethernet), and (iii) a network administrator who makes all of it work for them (installs new hardware and software, manages configuration settings, keeps the network running, etc.). These three collectively allow the user to gain access to Internet services. As soon as a user leaves the office, he/she leaves behind the last two of these resources and is unable to easily access Internet services. One might argue, correctly, that the Internet is often present at the nomad's destination, so why is there a problem? The answer is that the user's laptop appears as an unacceptable alien in it's new environment, and the third resource (the network administrator) is not present to configure one's laptop to make it acceptable. Indeed, a major nomadic computing issue that must be faced is how to provide the same set of Internet services to the nomad no matter where he or she may move.

Essentially, the goal of nomadicity is to provide the illusion of connectivity even when the nomad is disconnected and to provide access to Internet services seamlessly wherever the nomad travels. To achieve this, not only must the infrastructure be enhanced to provide these capabilities, but also it is necessary for applications to become nomadically enabled. Currently, most applications fail in the face of dramatic changes in connectivity, latency, delay, or operational immediacy. Nomadically enabled applications should recognize such eventualities as "the normal case" and be able to provide Internet services anywhere, from any device at any time.

These ideas form the essence of a major shift to nomadicity.

4. Embedded technology or "smart spaces": embeddedness

The second "dimension" of the new space is that of *embedded technology*, or *embeddedness*. This refers to small intelligent devices ubiquitously embedded in the physical world and connected to the Internet. Presently, users see cyberspace as trapped behind the screen of their workstations. But most users have no idea what is going on behind that screen, hence they see cyberspace as trapped in the Netherworld. We are fast approaching a time when the Netherworld of cyberspace will move out into the physical world. Most things in our physical real-world environment will be Internet-enabled via embedded technology.

A number of technologies support the growth of embeddedness. These include: MEMS technology, nanotechnology, digital electronics, sensor technology, small battery technology, and much more. These embedded devices will interact with each other to provide a "smart space" that adds intelligence to the environment [3].

Indeed, our environment will be alive with this embedded technology. It will appear in the walls, in the floors, in our desks, in our lamps, in our clothes, in our eyeglasses, in our refrigerators, in our automobiles, in our hotel rooms, in our wristwatches, in our belts, in our fingernails and in other places throughout our bodies. We will likely have a "bodynet" connecting all the devices we are carrying and which will act as our surrogate in communicating with the body nets of others, as well as the rest of the "smart space" in which we will be immersed. This embedded technology will be made up of sensors, actuators, logic, memory, processors, communicators, cameras, microphones, speakers, displays, RFID tags, etc. When I walk into a room enabled with this embedded technology (an intelligent room), the room will know I just entered. I will be able to converse with the room in natural language asking for information on a given subject, and perhaps four books will reply with their table of contents (and possibly one will inform me that it is located in my colleague's office down the hall). The web will present me with links and information via natural language speech, video, images, eyeglass displays, holograms, or other human-centered intuitive interface technologies.

We are already seeing the beginning of some interesting applications of smart spaces. For example, eSuds is a program [2] using a smart-chip system that manages laundry rooms at college campuses. eSuds lets students swipe smart cards, which could possibly also be their student ID card, to bill their accounts for washing or drying clothes, and also connects washers and driers to the Internet, letting students monitor the status of their wash load or find an empty machine via the network. When the washer finishes or clothes are dry, the student can receive an e-mail or pager alert.

The time is soon coming when computing devices connected in a wireless web will permeate our entire physical environment.

5. Ubiquitous computing: ubiquity

The third "dimension" of the new space is that of *ubiquitous computing*, or *ubiquity*. This refers to

Internet service availability wherever the nomad travels on a global basis. It implies the widespread deployment of network access across the infrastructure, making it available to the nomad everywhere. This has been a goal of the industry and a desire of the user population for many years, and is still in its very early stages.

The first appearance of the move to ubiquity was evident more than a decade ago when Internet access became widely available from corporate networks, i.e., local area networks connecting corporate computing devices to the Internet via high-speed links to the backbone network. The appearance of laptop computers enabled road warriors to travel to their homes and to hotels, conference rooms, etc., leading to a remote access problem. The industry responded by providing dial-up access (at relatively low speeds) using analog modems attached to the telephone system. This low speed, but reasonably ubiquitous, access exposed the existence of the "last mile" problem.

Basically, what was lacking was high-speed access from our homes, airports, hotels and those other places to which we traveled. Thus, when the technologies of copper-based DSL, of cable-based cable modem, and of satellite access appeared, there was great anticipation that broadband would be delivered to our homes, small offices, and our other destinations. However, a number of impediments prevented a timely rollout of these services. Among these impediments were the immaturity of the technologies, the high cost of deployment, the complexity placed on the end-user, the failure of CLECs who were deploying the technology, the foot-dragging by the telephone companies, and much more. At present, true ubiquitous access has still not arrived.

Meanwhile, another very significant technology was being deployed, namely that of digital wireless access. The first manifestation of this access technology was that of digital cellular technology. The wonderful thing about digital cellular is that it is a highly ubiquitous technology. The downside of cellular is that it is of relatively low bandwidth, similar in many ways to wireline dialup access. In addition, the wireless spectrum was acquired by the cellular carriers at extremely high prices, thus damaging the economics in many ways.

The second and very significant manifestation of digital wireless appeared in the form of broadband wireless access, namely WiFi. WiFi is a broadband technology that operates in the unlicensed spectrum (i.e., there is no charge for the spectrum). This technology currently offers bandwidths in excess of 50 Mbps over hundreds of meters and with newer technologies, involving clever antenna designs, will offer greater bandwidths and greater distances. WiFi has become enormously popular and truly qualifies as a grass roots technology that is rapidly deploying and gaining widespread use. Many homes have personal WiFi networks. We are already seeing many thousands of public "hotspots" that offer WiFi service. The appeal of wireless LANs based on WiFi is that it has no license fees, the equipment and deployment costs are dropping, the installation is easy, it provides high-speed connectivity, it is a widely supported global standard, there is considerable interoperability among WiFi equipment, it supports many operating systems across many platforms, and it has a large and enthusiastic customer base. Moreover, there is a certain level of maturation occurring at the time of this writing whereby we are seeing alliances and settlement brokers appearing who are beginning to offer global roaming to the subscriber. With this service, just as in cellular telephony, the subscriber need sign up with only one provider and still be able to roam into other providers' networks in a seamless fashion; the stage is now set for a very deep penetration of global WiFi access. Along with these many positive attributes, WiFi currently presents a number of challenges and potential shortcomings. Among these are: support for real-time traffic, truly secure operation, serious privacy issues, interference in the channel, relatively short range which results in "spotty" coverage, compatibility with digital cellular systems, etc. It is clear that WiFi is a disruptive technology with great promise and great risks to the providers. In some sense, we have let loose a "monster", and it is now our challenge to tame it into a stable and functional technology.

An important extension of these one-hop base-station oriented digital wireless systems is that of multi-hop mesh networks. In these networks, it is possible to deploy wireless technologies in environments devoid of base stations by using the wireless nodes as relays. An entire technology is being developed based on these ideas [1,5].

In spite of the fact that these technologies are in their early stages as of this writing, they are adding considerable ubiquity to our infrastructure.

6. The triangle of nomadicity, embeddedness and ubiquity

Let us now picture the dimensions of the space we have been discussing. In Fig. 1, we see a representation of a three-dimensional space depicted as a triangle. The axes are: *nomadicity*, *embeddedness*, and *ubiquity*. The interpretation of an axis is that if one is at a point on the triangle which is near the axis, then that point has a large amount of the attribute of the axis. For example, the point A on the axis of embeddedness is far from the two other axes, and so this point indicates that there is a considerable degree of embedded technology that the user has available (i.e., a smart space), but little ability for the user to travel from place to place seamlessly (i.e., little nomadicity), and relatively few locations where smart spaces are deployed. Alternatively, point B on the axis of ubiquity indicates that there are many places where users have access to the Internet, but these locations are devoid of embedded technology (i.e., a dumb space) and recognize only users known to the location (i.e., static users). Lastly, point C on the axis of nomadicity indicates that there is a considerable amount of mobility for the user (he may travel anywhere there is access and be recognized as an acceptable user who is given access to services), but there are relatively few locations that provide access and those locations are dumb spaces.

The three corners of the space are especially interesting:

- At the lower left-hand corner we possess a high degree of nomadicity as well as smart spaces, but find that there are relatively few locations that provide access. Indeed, we may then interpret the lower left-hand corner as a system which supports mobile users, lots of embedded devices, and essentially no ubiquity—(i.e., the embedded devices are all in one small locale offering only local access).
- At the lower right-hand corner, we possess nomadicity at numerous locations, but find that these locations are devoid of embedded technology. Thus this corner may be classified as supporting lots of nomadicity, a high degree of ubiquity, and no embedded technology—hence supporting mobile users in many locations (global access), but with no embedded intelligence in the physical space at these locations (a dumb space).
- At the top corner, we have lots of embedded technology, ubiquitously deployed in many locations, but no ability to accept alien (nomadic) users. That is, it provides global access: a smart embedded environment, but only for those users known in their locale.

7. Putting it together: the invisible global infrastructure

From Fig. 1, we recognize that it would be useful to possess the attributes of all three dimensions simultaneously. We can achieve this graphically by redrawing the figure to create a tetrahedron as shown in Fig. 2. In this figure, the three dimensions which were represented by lines in a triangle of Fig. 1 are now represented by triangular faces in the tetrahedron. The left front triangular face of the tetrahedron represents the

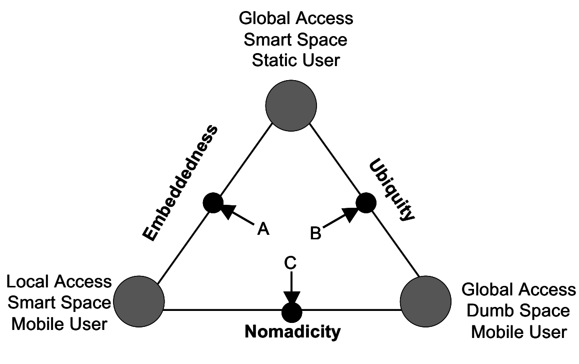

Fig. 1. The triangle of nomadicity, embeddedness and ubiquity.

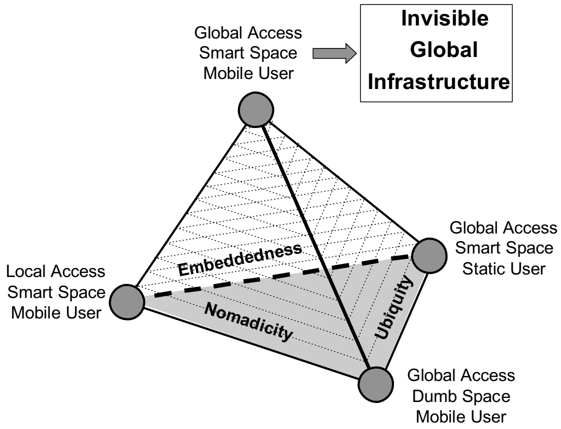

Fig. 2. The invisible global infrastructure.

dimension of *nomadicity*. The triangular face at the rear of the tetrahedron represents the dimension of *embeddedness*. The right front triangular face of the tetrahedron represents the dimension of *ubiquity*. We now see that the three dimensions (i.e., attributes) merge at the top of the tetrahedron, giving us a powerful system that supports *mobile users, smart spaces and global access*.

However, we see that this tetrahedron has only provided us with an infrastructure. To this infrastructure we must also add applications and services. Once we do so, we recognize that a missing component is the presence of *distributed intelligence* across this global infrastructure. This intelligence will support the applications and services, for example: roaming capability, location-based services, services based on a user's privileges, profiles and preferences, services offered by proxies. This suggests that there will exist intelligent software agents deployed in the infrastructure to help generate, store and deliver these applications and services. In the next five years, I expect to see the deployment of intelligent adaptive agents (also known as surrogates or proxies). Their purpose is to carry out tasks on behalf of the nomads as they wander around the smart space enabled Internet. For example, an adaptive agent may decide to send a low-resolution black and white picture or perhaps an outline of a document to other nomads who are poorly connected, instead of a full-resolution full-color picture or full-document text or hologram. An adaptive agent can act as an "impedance match" between the network and the things attached to it. In general, they would be there to support nomads, along with their applications, the network, servers, communication devices and computing devices.

Indeed, with the recognition that intelligence will exist in the entire infrastructure, and that the infrastructure is widely and flexibly deployed, we envision that the technology and intelligence will toil behind the scenes, vanishing into the infrastructure, hence resulting in an invisible global infrastructure.

8. Conclusion

So what is the vision? It is precisely the five elements I articulated in 1969:

1. The Internet technology will be everywhere.
2. It will be always accessible.
3. It will be always on.

The Internet has so far taken us those three steps. As described above, as we introduce nomadicity, embeddedness, ubiquity and distributed intelligence, we are presently moving toward the fourth and fifth steps of that early vision, namely,

4. Anyone will be able to plug in from any location with any device at any time.
5. It will be invisible.

In this vision of the future, I see users moving more into a mode of mobility wherein they access the net not only from their corporate desktop environment, but also ubiquitously at any time from wherever they happen to be with whatever device they have, in a seamless, secure, broadband fashion. I see small pervasive devices ubiquitously embedded in the physical world, providing the capabilities of actuators, sensors, logic, memory, processing, communicators, cameras, microphones, speakers, displays, RFID tags, etc. I see intelligent software agents deployed across the network whose function it is to mine data, act on that data, observe trends, carry out tasks dynamically and adapt to their environment. I see considerably more network traffic generated not so much by humans, but by these embedded devices and these intelligent software agents. I see large collections of self-organizing systems controlling vast fast networks. I see huge amounts of information flashing across networks instantaneously with this information undergoing enormous processing and informing the sophisticated decision support and control systems of our society. I see all these things and more as we move headlong into the 21st century. Indeed, I foresee that the Internet will essentially be an invisible global infrastructure serving as a global nervous system for the peoples and processes of this planet.

References

[1] Ian F. Akyildiz, David J. Goodman, Leonard Kleinrock, Guest editorial: Mobility and resource management in next generation wireless systems, IEEE Journal on Selected Areas in Communications 19 (2001) 1825–1830.

[2] Ian Austin, Automated system to do student laundry, The New York Times, September 2, 2002. http://www.usatech.com/pdfs/nytimes_020902.pdf.

[3] Deborah Estrin, David Culler, Kris Pister, Guarav S. Sukhatme, Connecting the physical world with pervasive networks, IEEE Pervasive Computing 1 (1) (2002) 59–69.

[4] Leonard Kleinrock, Breaking loose, Communications of the ACM 44 (9) (2001) 41–45.

[5] Leonard Kleinrock, On some principles of nomadic computing and multi-access communications, IEEE Communications Magazine 38 (2000) 46–50.

[6] Donald A. Norman, The Invisible Computer, MIT Press, Cambridge, MA, 1998.

[7] M. Satyanarayanan (Ed.), Reaching for Weiser's vision, IEEE Pervasive Computing 1 (1) (2002).

[8] Thomas Tugend, UCLA to be first station in nationwide computer network, UCLA Press Release, 3 July, 1969. http://www.lk.cs.ucla.edu/LK/Bib/REPORT/press.html.

[9] Mark Weiser, The computer for the twenty-first century, Scientific American (1991) 10–94.

[10] Mark Weiser, John Seely Brown, Designing calm technology, PowerGrid Journal 1 (1) (1996).

Leonard Kleinrock received his Ph.D. from MIT in 1963 and has served as a professor of computer science at the University of California, Los Angeles since then, serving as chairman of the department from 1991 to 1995. He received his BEE degree from CCNY in 1957 (also an Honorary Doctor of Science from CCNY in 1997, and an Honorary Doctor of Science from the University of Massachusetts, Amherst in 2000). He was first President and Co-founder of Linkabit. He is also Founder and Chairman of Nomadix, Inc., a high-technology firm located in Southern California. He is also Founder and Chairman of TTI/Vanguard, an advanced technology global forum organization based in Santa Monica, California. He has published more than 225 papers and authored six books on a wide array of subjects including packet switching networks, packet radio networks, local area networks, broadband networks and gigabit networks. Additionally, Dr. Kleinrock has recently launched the field of nomadic computing, the emerging technology to support users as soon as they leave their desktop environments; nomadic computing may well be the next major wave of the Internet.

Dr. Kleinrock is a member of the National Academy of Engineering, a member of the American Academy of Arts and Sciences, an IEEE fellow, an ACM fellow, an INFORMS Fellow, and a founding member of the Computer Science and Telecommunications Board of the National Research Council. Among his many honors, he is the recipient of the C.C.N.Y. Townsend Harris Medal, the CCNY Electrical Engineering Award, the Marconi Award, the L.M. Ericsson Prize, the NAE Charles Stark Draper Prize, the Okawa Prize, the IEEE Internet Millennium Award, the UCLA Outstanding Teacher Award, the Lanchester Prize, the ACM SIGCOMM Award, the Sigma Xi Monie Ferst Award, the INFORMS Presidents Award, and the IEEE Harry Goode Award.