CCR-4/95

NOMADIC COMPUTING - AN OPPORTUNITY

by

Leonard Kleinrock
Chairman, Computer Science Department
UCLA

ABSTRACT

We are in the midst of some truly revolutionary changes in the field of computer-communications, and these offer opportunities and challenges to the research community. One of these changes has to do with 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. This new paradigm is already manifesting itself as users travel to many different locations with laptops, PDA's, cellular telephones, pagers, etc. In this paper we discuss some of the open issues that must be addressed as we bring about the system support necessary for nomadicity. In addition, we present some of the considerations with which one must be concerned in the area of wireless communications, which forms one (and only one) component of nomadicity.

1. INTRODUCTION[1]

There are few things in one's professional life as gratifying as finding a powerful new analytical result that provides insight into a class of interesting computer-communications problems. Results of this kind change the way we think about such a class of problems. Many of us, as former winners of the ACM SIGCOMM Award, have been recognized for just that kind of work.

However, among those activities that do compete with such "results with insight" for excitement and gratification, is the

[1] This work was supported by the Advanced Research Projects Agency, ARPA/CSTO, under Contract J-FBI-93-112 "Computer Aided Design of High Performance Network Wireless Networked Systems and under Contract DABT-63-C-0080 "Transparent Virtual Mobile Environment".

exploration of emerging new technologies that represent a major shift in the way we do things, as opposed to how we think about things. It is just such a new technology, namely, Nomadic Computing, that we discuss in the present paper.

The fact is that the field of computer-communications in its largest sense (i.e., not simply the wires and networks, but also the infrastructure, the middleware, the applications, the uses and users of the technology) is in the midst of an accelerating groundswell. Witness the fact that the Internet is now a household word (just ask your neighbors). The use of the Worldwide Web (WWW) is growing faster than any other application we have ever witnessed in 25 years of networking (from the day the ARPANET was born at UCLA in September 1969 up to the present); and the WWW is still in its infancy!

Three things have converged in the last three years to bring our field into center stage of technology, science, and society:

- The focus by the present Federal Administration on the National Information Infrastructure (NII).
- The explosive growth of the Internet.
- The recognition by the commercial and entertainment world that networking has an enormous market potential.

Taking center stage brings with it opportunities and responsibilities. For example, do we as technologists have a responsibility to evaluate the likely scenarios that follow from the differing visions of each of the three events above? Is there a need to find a vision and an architecture for the NII which permits an integrated view of these differing visions? How do we address the needs of certain public interest communities (e.g., those of research, education, library) who do not have a natural platform for expressing their needs and who depend, in large part, upon government support for their activities? What is the role of government in these matters? What are the opportunities and responsibilities of our profession? We choose not to dwell any further on this here, but refer the reader for a discussion of these and many related issues to a report released last year by the National Research Council [1].

What is noteworthy is that we are in the midst of some very significant changes in many aspects of information technology, and these changes reach far beyond the purely technical aspects of our profession. It is at times like this that new innovations in thinking are called for, that major shifts in the use of technology must be recognized and anticipated, and that future, unknown, killer applications of the technology must not be precluded by the imposition of shortsighted architectural constraints. Yes, indeed, this is a time of great opportunity for our technology and its applications.

2. A MAJOR SHIFT

Currently, most users of think of computers as associated with their desktop appliances or with a server located in a dungeon in some mysterious basement. However, many of those same users may be considered to be *nomads*, in that they own computers and communication devices that they carry about with them in their travels as they move between office, home, airplane, hotel, automobile, branch office, etc. Moreover, even without portable computers or communications, there are many who travel to numerous locations in their business and personal lives, and who require access to computers and communications when they arrive at their destinations. Indeed, a move from one's desk to a conference table in one's office constitutes a nomadic move since the computing platforms and communications capability may be considerably different at the two locations. The variety of portable *computers* is impressive, ranging from laptop computers, to notebook computers, to personal digital assistants (or personal information managers), to smart credit card devices, to wrist watch computers, etc. In addition, the *communication* capability of these portable computers is advancing at a dramatic pace from high speed modems, to PCMCIA modems, to email receivers on a card, to spread-spectrum hand-held radios, to CDPD transceivers, to portable GPS receivers, to gigabit satellite access, etc.

The combination of portable computing with portable communications is changing the way we think about information processing [8]. We now recognize that access to computing and communications is necessary not only from one's "home base", but also while one is in transit and when one reaches one's destination.[2]

These ideas form the essence of the "major shift" to nomadic computing and communications that we choose to address in this paper. The focus is on 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.

3. NOMADIC COMPUTING[3]

We are interested in those capabilities that must be put in place to support nomadicity. The desirable characteristics for nomadicity include independence of location, of motion, of platform and with widespread presence of access to remote files, systems and services. The notion of independence here does not refer to the quality of service one sees, but rather to the perception of a computing environment that automatically *adjusts* to the processing, communications and access available

[2] Moreover, one may have more than a single "home base"; in fact, there may be no well-defined "home base" at all.

[3] Some of the ideas presented in this section were developed with two groups with which the author has collaborated in work on nomadic computing and communications. One of these is the Nomadic Working Team (NWT) of the Cross Industrial Working Team (XIWT); the author is the chairman of the NWT [4]. The second group is a set of his colleagues at the UCLA Computer Science Department who are working on an ARPA supported effort known as TRAVLER, of which he is Principal Investigator.

at the moment. For example, the bandwidth for moving data between a user and a remote server could easily vary from a few bits per second (in a noisy wireless environment) to hundreds of megabits per second (in a hard-wired ATM environment); or the computing platform available to the user could vary from a low-powered Personal Digital Assistant while in travel to a powerful supercomputer in a science laboratory. Indeed, today's systems treat radically changing connectivity or bandwidth/latency values as exceptions or failures; in the nomadic environment, these must be treated as the usual case. Moreover, the ability to accept partial or incomplete results is an option that must be made available due to the uncertainties of the informatics infrastructure.

The ability to automatically adjust all aspects of the user's computing, communication and storage functionality in a transparent and integrated fashion is the essence of a nomadic environment.

Some of the key system parameters with which one must be concerned include: bandwidth; latency; reliability; error rate; delay; storage; processing power; interference; interoperability; user interface; etc. These are the usual concerns for any computer-communication environment, but what makes them of special interest for us is that the values of these parameters change dramatically as the nomad moves from location to location. In addition, some totally new and primary concerns arise for the nomad such as weight, size and battery life of his portable devices. And the bottom line consideration in many nomadic applications is *cost.*

As researchers and users, there are a number of enchanting reasons why nomadicity should interest you. For example, nomadicity is clearly a <u>newly emerging technology</u> that users are already surrounded with. Indeed, this author judges it to be a <u>paradigm shift</u> in the way computing will be done in the future, so why not begin working in the field now? Information technology trends <u>are moving in this direction</u>. Nomadic computing and communications is a <u>multidisciplinary</u> and <u>multiinstitutional</u> effort. It has a huge <u>potential for improved capability</u> and convenience for the user. At the same time, it presents at least as huge a <u>problem in interoperability</u> at many levels. The contributions from any investigation of nomadicity will be mainly at the middleware level. The products that are beginning to roll out have a short term focus; however, there is an enormous level of interest among vendors (from the computer manufacturers, the networking manufacturers, the carriers, etc.) for long range development and product planning, much of which is now underway. Whatever work is accomplished now will certainly be of immediate practical use.

There are fundamental new research problems that arise in the development of a nomadic architecture and system. Let us consider a sampling of such problems. Below, we break these into Systems Issues and Wireless Networking Issues.

3.1 Systems Issues

One key problem is to develop a full *System Architecture and Set of Protocols* for nomadicity. These should provide for a transparent view of the user's dynamically changing computing and communications environment. The protocols must satisfy the following kinds of requirements:

- Interoperation among many kinds of infrastructures (e.g., wireline and wireless)
- Ability to deal with unpredictability of: user behavior, network capability, computing platform
- Provide for graceful degradation
- Scale with respect to: heterogeneity, address space, quality of service (QoS), bandwidth, geographical dimensions, number of users, etc.
- Integrated access to services
- Ad-hoc access to services
- Maximum independence between the network and the applications from both the user's viewpoint as well as from the development viewpoint
- Ability to match the nature of what is transmitted to the bandwidth availability (i.e., compression, approximation, partial information, etc.)
- Cooperation among system elements such as sensors, actuators, devices, network, operating system, file system, middleware, services, applications, etc.

In addition, the components that can help in providing these requirements are:

- An integrated software framework which presents a common virtual network layer
- Appropriate replication services at various levels
- File synchronization
- Predictive caching
- Consistency services
- Adaptive database management
- Location services (to keep track of people and devices)
- Discovery of resources
- Discovery of profile

A second research problem is to develop a *Reference Model* for Nomadicity which will allow a discussion of its attributes, features and structure in a consistent fashion. This should be done in a way that characterizes the view of the system as seen by the user, and the view of the user as seen by the system. The dimensions of this reference model might include:

- System state consistency (i.e., is the system consistent at the level of email, files, database, applications, etc.)
- Functionality (this could include the bandwidth of communications, the nature of the communication infrastructure, the quality of service provided, etc.)
- Locality, or Awareness (i.e., how aware is the user of the local environment and its resources, and how aware is the environment of the users and their profiles)

A third research problem is to develop *Mathematical Models* of the nomadic environment. These models will allow one to study the performance of the system under various workloads and system configurations as well as to develop design procedures.

As mentioned above, the area of nomadic computing and communications is multidisciplinary. A list of the disciplines which contribute to this area are (in top-down order):

- Advanced applications, such as multimedia or visualization
- Database systems
- File systems
- Operating systems
- Network systems
- Wireless communications

- Low power, low cost radio technology
- Micro-electro-mechanical systems (MEMS) sensor technology
- MEMS actuator technology
- Nanotechnology

The reason that the last three items in this list are included is that we intend that the nomadic environment include the concept of an *intelligent room*. Such a room has imbedded in its walls, furniture, floor, etc., all manner of sensors (to detect who and what is in the room), actuators, communicators, logic, cameras, etc. Indeed, one would hope to be able to speak to the room and say, for example, "I need some books on the subject of spread spectrum radios." and perhaps three books would reply. The replies would also offer to present the table of contents of each book, as well, perhaps, as the full text and graphics. Moreover, the books would identify where they are in the room, and, if such were the case, might add that one of the books is three doors down the hall in a colleague's office!

There are numerous other systems issues of interest that we have not addressed here. One of the primary issues is that of security, which involves privacy as well as authentication. Such matters are especially difficult in a nomadic environment since the nomad often finds that his computing and communication devices are outside the careful security walls of his home organization. This basic lack of physical security exacerbates the problem of nomadicity.

We have only touched upon some of the systems issues relevant to nomadicity. Let us now discuss some of the wireless networking issues of nomadicity.

3.2 Wireless Networking Issues

It is clear that a great many issues regarding nomadicity arise whether or not one has access to wireless communications. However, with such access, a number of interesting considerations arise which we discuss in this section.

Access to wireless communications provides two capabilities to the nomad. First, it allows him to communicate from various (fixed) locations without being connected into the wireline network. Second, it allows him to communicate while traveling. Although the bandwidth offered by wireless communication media varies over an enormous range as does the wireline network bandwidth, the nature of the error rate, fading behavior, interference level, mobility issues etc., for wireless are considerably different so that the algorithms and protocols require some new and different forms from that of wireline networks [3].

The cellular radio networks that are so prevalent today have an architecture that assumes the existence of a cell base station for each cell of the array; the base station controls the activity of its cell. The design considerations of such cellular networks are reasonably well understood and are being addressed by an entire industry [5]. We discuss these no further in this paper.[4]

There is however, another wireless networking architecture of interest which assumes no base stations [2,6]. Such wireless networks are useful for applications that require "instant" infrastructure, among others. For example, disaster relief, emergency operations, special military operations, clandestine operations, etc., are all cases where no base station infrastructure can be assumed. In the case of no base stations, maintaining communications is considerably more difficult. For example, it may be the case that the destination for a given reception is not within range of the transmitter, in which case some form of relaying is required; this is known as *multi-hop* communications. Moreover, since there are no fixed location base stations, then the connectivity of the network is subject to considerable change as devices move around and/or as the medium changes its characteristics. A number of new considerations arise in these situations, and new kinds of network algorithms are needed to deal with them.

In order to elaborate on some of the issues with which one must be concerned in the case of no base stations, we decompose the possible scenarios into the following three:

<u>Static Topology with One-Hop Communications</u>: In this case, there is no motion among the system elements, and all transmitters can reach their destinations without any relays. The issues of concern, along with the needed network algorithms (shown in italics), are as follows:

---

[4] Wireless LANs come in a variety of forms. Some of them are centrally controlled, and therefore have some of the same control issues as cellular systems with base stations, while others have distributed control in which case they behave more like the no-base-station systems we discuss in this section.

- Can you reach your destination: *Power Control*
- What access method should you use: *Network Access Control*
- Which channel (or code) should you use: *Channel Assignment Control*
- Will you interfere with another transmission: *Power and Medium Access Control*
- When do you allow a new "call" into the system: *Admission Control*
- For different multiplexed streams, can you achieve the required QoS (e.g., bandwidth, loss, delay, delay jitter, higher order statistics, etc.): *Multimedia Control*
- What packet size should you use: *System Design*
- How are errors to be handled: *Error Control*
- How do you handle congestion: *Congestion Control*
- How do you adapt to failures: *Degradation Control*

Static Topology with Multi-Hop Communications: Here the topology is static again, but transmitters may not be able to reach their destinations in one hop, and so multi-hop relay communications is necessary in some cases. The issues of concern, along with the needed network algorithms (shown in italics) ARE ALL OF THE ABOVE PLUS:

- Is there a path to your destination: *Path Control*
- Does giant stepping [7] help: *Power Control*
- What routing procedure should you use: *Routing Control*
- When should you reroute existing calls: *Reconfiguration Control*
- How do you assign bandwidth and QoS along the path: *Admission Control and Channel Assignment*

Dynamic Topology with Multi-Hop: In this case, the devices (radios, users, etc.) are allowed to move which causes the network connectivity to change dynamically. The issues of concern, along with the needed network algorithms (shown in italics) ARE ALL OF THE ABOVE PLUS:

- Do you track or search for your destination: *Location Control*
- What network reconfiguration strategy should you use: *Adaptive Topology Control*

- How should you use reconfigurable and adaptive base stations: *Adaptive Base Station Control*

These lists of considerations are not complete, but are only illustrative of the many interesting research problems that present themselves in this environment. Indeed, in this section we have addressed only the network algorithm issues, and have not presented the many other issues involved with radio design, hardware design, tools for CAD, system drivers, etc.

4. CONCLUSION

In this paper we have presented nomadicity as a new paradigm in the use of computer and communications technology and have laid down a number of challenging problems. The field is current, exciting, draws from many disciplines, and offers a variety of kinds of problems whose solutions are of immediate importance. You can contribute to nomadicity in a number of ways, limited only by your imagination. We have thrown down the gauntlet - now it's your move.

REFERENCES

[1] Computer Science and Telecommunications Board, "Realizing the Information Future: The Internet and Beyond", National Academy Press, Washington, DC, 1994.

[2] Jain, R., J. Short, L. Kleinrock and J. Villasenor, "PC-notebook Based Mobile Networking: Algorithms, Architectures and Implementations", ICC '95, June 1995.

[3] Katz, R. H., "Adaptation and Mobility in Wireless Information Systems, " IEEE Personal Communications Magazine, Vol. 1, No. 1, (First Quarter, 1995), pp. 6-17.

[4] "Nomadicity: Characteristics, Issues, and Applications", Nomadic Working Team of the Cross Industrial Working Team, 1995.

[5] Padgett, J.E., C.G. Gunther, and T. Hattori, "Overview of Wireless Personal Communications", IEEE Communications Magazine, January 1995, Vol. 33, No. 1, pp.28-41.

[6] Short, J., R. Bagrodia, L. Kleinrock, “Mobile Wireless Network System Simulation”, UCLA Technical Report, CSD-950015.

[7] Takagi, H. and L. Kleinrock, "Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals", IEEE Transactions on Communications, Vol. COM-32, No. 3, pp. 246-257, March 1984.

[8] Weiser, M., "The Computer for the 21st Century", Scientific American September 1991, pp. 94-104.