

Next generation networking: distributed multimedia information for healthcare

Michael J. Ackerman

Published online: 7 February 2007
© Springer Science + Business Media, LLC 2007

Abstract This paper is derived from a keynote address given at the DMS-03 meeting. It chronicles the need and development of next generation networks (NGN) in the United States. Specific organizational examples are derived from the Internet2-Abilene Network. The technical characteristics of a next generation network versus the Internet are discussed. Examples are given from the point of view of the need for a quality of service based network to deliver distributed multimedia healthcare information to the point of need. The concepts of network trust and of a network based scalable information infrastructure for the reliable delivery of distributed multimedia information is also introduced.

Keywords Quality of Service · Next generation networking · Distributed multimedia · Healthcare · Scalable information infrastructure · Internet2 · Abilene

1 In the beginning

The Internet started in the late 1960s as the ARPAnet, the Advanced Research Project Agency Network. ARPA had a need to link defense computers located at industrial and government sites with computers located at universities. A digital network was invented to do this, the original ARPAnet. Since scientists like to share, the scientists that were using the ARPAnet told their scientist friends and before long the ARPAnet was being used by the science community at large. By the 1980s ARPA decided that the ARPAnet was no longer serving as a DoD resource but rather a science community resource. Responsibility for the ARPAnet was transferred to the National Science Foundation (NSF) and the network became known as the NSFnet or NREN, the National Research and Education Network.

It's hard not to share a good thing. By the early 1990s, the university scientific community had introduced the humanities community to the NSFnet. The general academic community embraced the NSFnet. NSF decided that the NSFnet was no longer exclusively a

M. J. Ackerman (✉)
National Library of Medicine, Bethesda, MD 20894, USA
e-mail: ackerman@nlm.nih.gov

scientific resource but rather a public utility. The NSFnet was privatized into the public domain and became known as the Internet. At that time the Internet was not heavily used because in order to use it, you had to know how to use a text based UNIX like computer interlace.

In 1993 the World Wide Web was invented and public use of the Internet began to skyrocket. The Internet became "the Web" and with it came hackers, viruses, and congestion. The scientific community soon discovered that the Internet, which they had come to depend on for communications and collaboration with colleagues, was no longer useable for that purpose. The ARPAnet, NSFnet, Internet was no longer reliable enough for use by the scientific community. The American scientific community had lost an assumed infrastructure resource.

2 And then there were two

In 1996 two things happened. In September, representatives of about 30 research universities came together in Chicago to discuss what to do about the loss of the Internet. They decided to form a new corporation, the University Corporation for the Advancement of Internet Development, UCAID. The purpose of UCAID would be to rebuild a new, high-speed, reliable Internet for the academic community. The original Internet can be thought of as Internet1. The new academic Internet would therefore be known as Internet2.

News of the loss of the use of the Internet by the scientific community reached the White House. In October 1996 President Clinton announce a new government wide program called Next Generation Internet, NGI. The purpose of the NGI Program would be to "Connect universities and national labs with high-speed networks that are 100–1,000 times faster than today's Internet. These networks will connect at least 100 universities and national labs at speeds that are 100 times faster than today's Internet, and a smaller number of institutions at speeds that are 1,000 times faster" (<<http://www.ngi.gov/white-house/background.html>>).

3 Joint goals

The government NGI Program and the university UCAID program had very similar goals. They worked well together and complimented each other (<<http://government.internet2.edu/ngi.html>>). The first joint goal was to promote research into next generation network technologies. Without changing the technology, networks cannot be made to run faster because the underlying network technology is not scalable. The network will perform differently, often unpredictably, at high speed than it did at its design speed. In order to make a network run faster requires network research.

The second joint goal was to develop a next generation network testbed to connect universities and federal research institutions at rates that are sufficient to demonstrate new technologies not possible on slower networks and to support future research. UCAID was founded to create such a new Internet2 research network. That network has become known as the Abilene Network

The third joint goal was to demonstrate new applications that require next generation networking that meet important national goals and missions. This is where the National Library of Medicine (NLM) and several other government agencies are participating by

sponsoring the development of applications that demonstrate the utility and advantage of such networks.

4 Abilene and its design

The Abilene network is accessed by Internet2 member institutions through a GigaPoP. A GigaPoP is not a breakfast cereal. A GigaPoP is a gigabit point of presence. If you're building an interstate highway and you put an interchange every mile along the highway the traffic on the highway will not be able to flow at speed. The interchanges need to be spread out and placed at least some minimum distance apart. This is how Internet2's Abilene Network is designed. UCAID required the universities in an area to get together and build one point of presence, the GigaPoP, which would then be connected to the Abilene Network in order to minimize the number of network backbone connections. There are 31 GigaPoPs distributed very much like the interstate highway system. Where there are more universities, there are more GigaPoPs. Where there are fewer universities, there are fewer GigaPoPs.

Membership in UCAID requires the payment of dues. But it also requires the payment of "in kind" dues. That means that the member university has to make a pledge to spend a certain amount of money on advanced networking technology. For example, when a university rewires the dorms from a 10baseT to a 100baseT Ethernet, that cost accrues towards this pledge. A graduate student's scholarship, assuming the graduate student is working on advanced networking algorithms, accrues towards the pledge. Members not only have to pay the dues but also have to be actively engaged in the field by making good on the "in kind" pledge (<<http://members.internet2.edu/membershipfaq.html>>).

Why the Internet2 network is called Abilene? One of the UCAID member companies is Quest Communications. Quest is in the process of putting fiber all over the country. They currently have a lot of dark fiber. So Quest, as its "in kind" pledge, offered to allow UCAID to use some of its dark fiber for the Abilene Network. An interesting and relevant story is the method by which Quest lays its fiber. Quest owns a train. An engine pulls three cars. The first car digs a trench on the side of the track. The second car puts the fiber in the trench. The third car covers it over. They say the train runs at 4 mi/h laying fiber. Quest believes that sooner or later the fiber will be needed and Quest will have fiber lines close to every destination. Quest's method of laying fiber associates it with the railroad. The Abilene Network is a coast to coast network, it's transcontinental. One of the terminus points on the eastern side of the transcontinental railroad was the town of Abilene, Kansas. And so in order to honor Quest, its railroad train and the transcontinental nature of the network, the network was named the Abilene Network (<<http://abilene.internet2.edu/about/>>).

5 Next generation networks and healthcare

The NLM joined UCAID as part of Next Generation Internet program. The job of the NLM, at least with respect to advanced networking, has been to promote the use of advanced network technology in healthcare. NLM specifically wants to discover needed health care applications that cannot be realized because of inadequacies in current network technology. NLM encourages the use of the Internet2 Abilene Network as the testbed. What are the candidate application areas that might benefit from advanced networking technology?

One potential area includes those applications that are based on the use of high-resolution images, either big pictures or databases containing big pictures. Large images take a lot of time to transmit. If the picture is going to reach its destination in a timely way, in time to help somebody out of a life threatening situation, the transmission must happen rapidly, or as they say in medicine, "stat." So high-resolution imaging is part of NGI.

Another potential application area includes real-time telemedicine. Today it is very common to call your physician over the telephone and seek advice by describing your condition that the doctor cannot see. An advanced network infrastructure would allow a video call from the computer in your home to the computer in your doctor's office. In addition the appropriate transducer appliances could be attached to your home computer and the doctor could make virtual house calls via the network using the appropriate home based transducers.

NLM is the library of record for the world's medical literature. As we move into the twenty-first century, the world's literature includes images and therefore the NLM should also serve as a medical image reference library. If a pathologist had a specimen slide and wasn't exactly sure of its significance the pathologist could scan the slide into the computer and request that the computer retrieve a slide from the NLM image database that looks like this one. Along with the retrieved slide would come the gold standard medical diagnosis associated with that slide. But we soon came to understand that this vision was very shortsighted. What we are really talking about is retrieval from the multimedia electronic medical reference library that each of us leaves behind us as we move through life.

Each of us leaves a multimedia library at different doctor's offices and at various hospitals which needs to come together at the point of need. A good part of a diagnosis is based on the change in a person's condition since their last examination, rather than the condition at the moment of the current examination. Mammography serves as an easily understood example. A current mammogram has clinical significance. But its true significance can only be understood in relation to the mammogram that was taken the year before. If the current mammogram shows something that is suspicious but that has not changed in a year it has much less clinical significance than same suspicious finding which has changed radically in the last 6 months. The problem is that the previous mammogram is often not available because we did not visit the same facility as we did the last time. The case can be easily made that each of us is leaving a distributed multimedia library behind and the technology and the networks are needed to bring all the information together to the point of need.

6 Quality of Service (QoS)

In 1998, NLM commissioned a study the US National Research Council [1]. The purpose of the study was first to identify the medical applications that require Next Generation Internet technology and then to identify the technical factors that are missing from today's Internet which would be needed for the identified applications to become practical. The technical factors can be categorized as "Quality of Service" (QoS) features. The Internet is designed to get information from here to there as quickly as possible. In that sense it's like the post office. You put your envelope into a mailbox and the Post Office will get it to its destination as fast as they can. Sometimes it takes a day, it usually takes 3 days, and every once in while we hear about an envelope showing up 50 years later. Delivery is guaranteed, but they cannot guarantee when. It's exactly the same for the Internet. In health care we

have to guarantee the time of delivery. We need the overnight delivery service version of the Internet.

The needed Quality of Service features are dependent on the needs of your application. Some very smart network programming often obscures these needs. For example, the first time one goes to a new Web site, it often takes a long time to download. The second time it appears to download instantly, not because the network was not busy but because a copy of that Web site was made in the cache on your computer the first time you accessed that Web site. The second access is from your cache, not from the network. The way to prove this is to go to a Web site that has a clock. Go back a few minutes later and look to see what time it is—has the clock changed? That's important to know because I may look at a Web site displaying my patient's current cardiogram. I come back a few minutes later to take another look and the cardiogram still looks normal. But my patient has actually just suffered a crisis. I don't see the problem cardiogram because the Web site is being displayed from my computer's cache, not from my patient.

Under Quality of Service are certain specific features that are important in a networked health care environment: bandwidth, low latency, low jitter, variable priority, data integrity, selectable loss rate, and security. Variable priority means that everything doesn't have to go at the same time, some things can wait. Data integrity means that not only must all the data properly arrive at its intended destination, but it also must be presented in an accurate way. This includes the transducer. What good is it to send full high fidelity stereo across a network if the playback is going to be on a 1-in. speaker built into a laptop. All that work of network data integrity was negated because it wasn't played back with the appropriate transducer.

As surprising as it may seem, there are times that you do not have to guarantee data integrity. Selectable loss rate is acceptable provided you do it in a smart way. For example, the diagnostic information in an echocardiogram is in the motion of the image, not the detail. If there is a decrease in available bandwidth you can cut back the detail, maintain the frame rate, and still get the clinical information through.

7 Lessons learned

With this information in hand, the NLM funded test-bed projects to demonstrate the use of NGI capabilities by the health community. These demonstrable capabilities included quality of service, security and medical data privacy, nomadic computing, network management, and infrastructure technology as a means for collaboration. The demonstrations were designed to improve our understanding of the impact of NGI capabilities on the nation's healthcare, health education, and health research systems in such areas as cost, quality, usability, efficacy and security. The details concerning the applications that were funded by the NLM and the lessons learned can be found on NLM's web page (<<http://www.nlm.nih.gov/research/ngisumphase2.html>>).

In summary, we learned that only a few applications require high bandwidth, most do not, and that Quality of Service is really the key. This was our hypothesis, if Quality of Service could be provided, the network would be usable by healthcare applications. But it turns out that although Quality of Service is necessary, it is not sufficient. You also need smart applications that are capable of interacting with these advanced network features. There needs to be an interactive partnership between the computer on your desk and the network.

8 A scalable information infrastructure


The need for a computer–network partnership was highlighted in the findings of a workshop, “New Visions for Large-Scale Networks: Research and Applications,” held in June 2001 (<<http://www.itrd.gov/iwg/pca/lsn/lsn-workshop-12mar01/>>). The workshop identified the following networking opportunities, all of which NLM feels are needed in the truly networked healthcare environment:

Adaptive, dynamic, and smart networking: Automated discovery of resources
Measurement, simulation, modeling, and scalability: End to end performance measurement and metrics
Trust (security, privacy, and reliability): Quality of Service
Middleware
Collaboration environments
Revisiting network fundamentals: Wireless technology

Based on our Next Generation Internet findings and the results of this workshop, NLM entered into a new research effort “Applications of Advanced Network Infrastructure Technology in Health and Disaster Management.” NLM is funding health related applications which demonstrate *self-scaling* technology, utilize *self-optimizing end-to-end network aware* real-time technology and/or middleware, may be dependent on *wireless* technology, involve *advanced authentication* methodologies, e.g. biometrics or smartcards, or demonstrate nomadic technology applications and/or applications using *geographic information systems* (GIS) techniques (<<http://www.nlm.nih.gov/research/siiawards.html>>).

References

1. Networking Health: Prescriptions for the Internet. National Academy, Washington, DC. 2000. <<http://books.nap.edu/catalog/9750.html>>



Michael J. Ackerman, Ph.D. Michael J. Ackerman received his Ph.D. from the University of North Carolina, Chapel Hill, in Biomedical Engineering. After graduation he served as a research physiologist in the Hyperbaric Medicine and Physiology Department, Naval Medical Research Institute, where he studied the effects of the hyperbaric environment on neurophysiology and behavior. He later became head of the Institute's Biomedical Engineering and Computing Branch responsible for the application of computers to real-time medical data analysis and the control and monitoring of diving systems. Dr. Ackerman came to the National Library of Medicine in 1987. He served as the Chief of the Educational Technology Branch of the Lister Hill National Center for Biomedical Communications, applying interactive technology to medical education, and as the Associate Director for Specialized Information Services responsible for the Library's non-bibliographic databases. He is currently NLM's Assistant Director for High Performance Computing and

Communications, providing guidance for NLM's telemedicine, distance collaboratory, advanced networking and imaging interests. He holds academic appointments as an Associate Professor in the Department of Computer Medicine at George Washington University and as an Assistant Professor in the Department of Medical Informatics at the Uniformed Services University of the Health Sciences, and has published over 140 papers and book chapters.

Dr. Ackerman is active in the field of medical informatics. He was a charter member and served as Treasurer of the American Association for Medical Systems and Informatics (AAMSI). He had been a member of the Board of Directors of the Symposium for Computer Applications in Medical Care (SCAMC) from 1976 to 1988 and served as its President. He was the Program Chair for the 9th SCAMC and Finance Chair for Medinfo '86. He is a founding member of the American Medical Informatics Association (AMIA), and served as Treasurer, Secretary, chair of the Meetings Committee, and as a member of the Board of Directors. He was co-chair of the 1992 Health Science Communications Association (HeSCA) Annual Meeting, a consultant to the Radiological Society of North America's Electronic Communications Committee, a member of the American Physical Therapy Association's Advisory Panel on the Resource Center for Research and the United States Pharmacopeial Convention's Advisory Panel on Consumer Interest and Health Education. He was elected a Founding Fellow of the American Institute of Medical and Biological Engineering (AIMBE) in 1992 and a Fellow of the American College of Medical Informatics (ACMI) in 1985. He currently serves on the Editorial Boards of the *Telemedicine Journal & e-Health*, the *Journal of the American Medical Informatics Association*, the *IEEE Transactions on Information Technology in Biomedicine*, and *Medicine on the Net* and as a member of the Board of Directors of the American Telemedicine Association (ATA) and of AIMBE.

Dr. Ackerman's work has been recognized through numerous awards including the 1998 Johns Hopkins University Ranice W. Crosby Distinguished Achievement Award, 1997 Government Technology Leadership Award, 1996 National Institutes of Health Director's Award, the 1996 Friends of the National Library of Medicine Public Service Award, the 1996 Satava Award for Medical Applications of Virtual Reality, the 1995 Public Health Service Special Recognition Award, the 1994 American Institute for Medical and Biological Engineering Dedicated Service Award, the 1993 and 2003 American Medical Informatics Association President's Awards, the 1993 Health Sciences Communications Association Special Achievement Award, and the 1992 National Institutes of Health Award of Merit. His work on the Visible Human Project was nominated as a finalist for a 1995 Discover Magazine Award for Technological Innovation in Software and a 1996 Smithsonian Award for Information Technology.