
Technical Update

Networked Multimedia: Are We There Yet?

by Bill Wyman,
University of Colorado—Boulder

Once Upon a Time...

When Dave Bowman phoned home during his trip to Jupiter in *2001: A Space Odyssey* and got an image of his family wishing him happy birthday, a couple of generations of Americans began to wonder how long it would be until The Phone Company (remember them?) could make that trick available to ordinary natives of the planet. Granted it was just a recording; it took too long for radio waves to travel back and forth from Earth for Dave to carry on a conversation. Transmission rate was the bottleneck. That was 1968. Radio and television broadcasting and the telephone dominated communications for the next 20 or 25 years, and all the while, some continued to ask: Couldn't you make something so we can talk back and forth and see each other at the same time? Mentioning the "videophone" in 1996 will probably produce a rash of titters and wincing nostalgia.

By the end of the 80s, we could at least describe the dream technology that we wanted with superb technical precision—real-time interactive multimedia communications—even if we still weren't getting it. This concept takes in more than just videoconferencing, and as computer networking technology matures and flourishes, new products, like the Cheetah video server we'll be looking at here, are moving that dream technology not only to within our waking reach but also into the classroom. The stumbling blocks have been formidable, and removing them has both caused and reflected changes in our society. A quick retrospective on the technology leads to a checklist of issues and of the features you will want to consider when shopping for networked multimedia solutions.

Video on Demand

Once we had slain The Phone Company, we refocused our expectations on The Cable Company, and repackaged our synaesthetic aspirations as "video on demand." If we couldn't talk to and see each other across great distances whenever we liked (at least not until we could find someone to pay the stunning costs of the rewiring), maybe we could get a library of movies delivered to our television sets whenever we liked by hooking our home computers up to the cable, or to the Internet, or both. After all, the cable is already there, isn't it? It was only a short time until the new bottlenecks in this vision had names: compatibility and adequate bandwidth.

Videocassettes

While video on demand began taking on a life of its own in the public imagination, on the campus many foreign language teachers were ready for a more modestly scaled multimedia system that would allow students to learn and practice their skills intelligently, using their eyes as well as their ears. Cassette videotapes, convenient and inexpensive, were already a well-established system for delivering a stream of recorded sounds and images. VCRs allowed students to stop and start the stream, and to choose their own point of immersion. But the student remained a spectator, and attempts to heighten the instructional capabilities of videocassettes by marrying them to computer technology stumbled over the familiar analog addressing bottlenecks of queuing and pausing. The more often tapes stretched and broke, the more acutely we felt the need for a different medium.

Laser Discs

Instructional multimedia had evolved significantly by the time the first wave of computer-controlled laser disc technology and software arrived at the beginning of the 1990s. New ground was broken by the French instructional program *À la rencontre de Philippe*. The multimedia was there, as was a kind of interactivity; one or two students can collaborate on a *Philippe* session and come up with the game-like strategies it requires. Its multiplexed plot line did simulate some of the uncontrolled aspects of real communication, and thanks to its ingenious interface, *Philippe* can be used effectively by students with various degrees of competence in French. However, *Philippe* remains essentially a stand-alone learning tool, and its price tag informed us that, even when pressed into service for the worthy goal of instruction, filming a decent soap opera still costs a very big bundle. Filming it in multiple versions to boot pretty well guarantees an under-crowded market.

CD-ROM

When interest in the Internet was set ablaze by whatever is responsible for that having happened, some networking devotees (like me) were already looking askance at the emerging instructional multimediu(m) of CD-ROMs. While their storage capability was impressive and inexpensive, their limitations were all-too-evident. Throughput was halting and on the slow side. Teaching twenty students would probably require buying twenty copies of the software. Twenty students sitting in a room using an instructional CD-ROM program were still twenty students accessing a machine and a recording—a kind of group stand-alone. Getting multimedia CD-ROMs to work right on PCs was often less than fun (and still can be). Dealing with their nit-noid quirks has gotten easier with time, but like broadcast radio and television, cable television, cassette videotapes, and laser discs before them, CD-ROMs deliver information that flows in one direction only: at the student. Their proper place in an educational setting may be more the reference library than the laboratory or the classroom.

**Networking:
A History**

As CD-ROMs began their ascendancy, local area networks (LANs) and intranets (a.k.a. wide area networks or WANs) with all of their versatility and special capabilities, as well as the unprecedented communications power of the Internet, were largely still standing outside the door of the classroom, waiting their turn. While it was common even five years ago to find students exchanging Internet email in a laboratory, networked multimedia faced significant old and new technical challenges. How do we make systems that are affordable, compatible, usable, and with sufficient capacity and bandwidth for instruction? We have come to expect sounds and images to be of the level of quality set by the analog recording industry and by conventional television and radio broadcasting. Not cheap.

Digital capture, storage and transmission of multimedia are still “young technologies;” i.e., unfamiliar, costly, and unstandardized. Affordability requires mass production, and for profitability mass production presupposes an established or, even better, a growing market of ready consumers. Compatibility implies standards, treaties or contracts that determine how independent manufacturers of equipment existing in competitive relationships will, nonetheless, agree to work together. Government regulation doesn’t go over well in the 1990’s. The most expensive issues have proven to be capacity and bandwidth, and they are interrelated. The greater the storage capacity of any device, the more data that will exist in

one place, and the greater the bandwidth that will be required to move it to another place in an efficient manner. Likewise, the greater the available bandwidth, the more the data thus can be moved and stored in an efficient manner. Big reservoirs beget big pipes and vice-versa.

Terminal/Host Networking

Under the old “dumb terminal” model, all the data and the programming to manipulate them were stored on a main-frame host that we logged into and then talked to over telephone-like lines. These terminal displays were monochromatic (green), incapable of displaying sophisticated graphics—much less full-motion video—and mute. Terminal/host transmission rates, too, were a problem and set limits on both the degree and kind of user interactivity that were possible.

Microcomputers had their own built-in processors, software programs, internal hard drives for storage, color monitors, and increasingly clever operating systems and capable graphic interfaces. Stereo sound was only a few years away. The storage and processing demands of multimedia, however, easily outstripped even their capabilities, and for interactivity with other users, some way would have to be found to make microcomputers compatible, efficient communication devices. New data compression technology—a counterintuitive process that allows more data to be stored in less space and transmitted using less bandwidth—stretched network carrying capacity, and the task of connecting microcomputers and servers on high-speed, high-bandwidth networks got the full-court press.

Client/Server Networking

Anyone who has sent or received data over the Internet or surfed the Web regularly in recent months cannot have missed the crucial role bandwidth plays making a communications system efficient and usable. The same holds true for any network of any size. In a client/server networked environment (be it local, an intranet, or the Internet itself) the data being disseminated—whether text, graphics, audio, video or any combination of those—is stored and distributed by servers. The processing power and software to use that data reside in the user’s client computer. Result: synergy. Servers consist of a network operating system (NOS), a CPU with one or more processors, input/control devices (keyboard and mouse), data storage devices (disk drives and tape), and a network interface card (NIC) and/or a modem. Stored data are retrieved and utilized by clients. A client consists of an operating system (OS), a CPU generally with a single processor, input/control devices, data storage devices, network client software,

and a NIC or modem. The physical network, then, is a system of copper wires and fiber optic cables that connects servers and clients.

A key feature of client/server topology is that a user working or interacting with data retrieved from the server can modify that data and send it back to server where it can be stored anew, and, through a series of permissions, shared with other users even if its originator is not on-line at the moment. This is how email, for example, works in a client/server environment, and it's basically what we watched Dave Bowman do enroute to Jupiter. If data is shared between two or more network users while they are on-line, then the communication is said to occur in real-time. Internet Relay Chat (IRC) permits text messages to be exchanged this way, while videoconferencing schemes like CU-SeeMe and PictureTel add the dimensions of sight and sound, albeit at differing levels of transmission quality, speed, and expense.

Intraoperability and Interoperability

Networks in the broad sense can consist of similar and dissimilar devices. When networked devices are sufficiently similar, they can communicate with each other by relying on the built-in compatibility of their components and operating system (ex., Appleshare, Appletalk, Windows for Workgroups, and Windows '95). When all the hardware and software devices within a network are compatible with one another, that network is said to possess "intraoperability," and that's a good thing. In reality networked devices are often dissimilar. Such is life in free-market America and the Wide World Out There. As a result, an intraoperable network may or may not be able to communicate with functionally comparable networks or devices that happen to be made by other manufacturers. This lack of standards is called "a familiar nuisance."

When a manufacturer supplies a network system whose hardware and software components are compatible with one another, but are also intentionally or unintentionally incompatible with those of other manufacturers, the system is said to be "proprietary." Our stuff works with our systems, but it may or may not work with anybody else's. We experience this situation daily because Macs and PCs, for example, are regularly unable to run the same software code. Platform wars. Older Macs and PCs sometimes cannot run newer versions of platform-appropriate software. Version-itis. However, when manufacturers of dissimilar devices and operating systems agree on a lingua franca, that is known as a "protocol," and by using it very many different devices, operating systems, and whole networks achieve "intraoperability." This is the

principle championed by TCP/IP, the Internet protocol, and a genuine reason for users to hope for a brighter future.

Multimedia Networking

On a network, multimedia are just another kind of data file. The best approach to networking multimedia successfully is to begin with a network that conforms to current industry standards and has generally accepted operating characteristics. The checklist below summarizes those standards and characteristics. No one can guarantee that you will be able to address all the needs of all users all the time, but you can get pretty close to that goal for now and for some years out. Consider these good talking-points when you are dealing with vendors or network designers:

Bandwidth—Throughput should be in the range from 10-100 megabits/second. A current standard for micro-computer video servers calls for them to deliver sound and images in a stream at 1.2 megabits/second. Keep in mind that each user gets their own stream, and there may be other kinds of concurrent traffic, such as IRC, email, and Web browsing. When in doubt, consider buying the bigger pipe.

Scaleability—Networks tend to grow over time. The growth in the number of network service users shows no sign of slowing down. The physical size of your facility is finite; you can only accommodate a limited number of users at a given time. But networking means that not all of your users need to be present at the same time—or ever for that matter—to use your facility. Be wary of proprietary hardware solutions and software packages that seem to restrict the number of users unduly, as well as pricing schemes that add a large premium for the next increment of user growth.

Intraoperability—The three leading local/Intranet network hardware systems at this writing are Ethernet/Fast Ethernet, FDDI (Fiber Distributed Data Interface), and ATM (Asynchronous Transfer Mode). You should be hearing about one or more of these technologies from your network vendor or designer. Building your network on these technologies is the best way to ensure that the network devices (servers, clients, printers, network interface cards, etc.) and software you select (1) will be compatible with one other and work together correctly, (2) will be fast enough and have enough bandwidth to meet your needs, and (3) can scale up and be expanded in number and functionality conveniently and relatively inexpensively.

Interoperability— The more protocols your network can support, the greater the number of hardware and software devices that will be compatible. With its roots in the UNIX operating system, the Internet protocol TCP/IP (Transmission Control Protocol/Internet Protocol) has become the global standard for communication. There is no substitute. At the local/Intranet level, popular protocols include IPX (Novell), Ethertalk (Apple Macintosh), and NetBeui (Microsoft). Networks can support many protocols at the same time, and there are other protocols to consider. Keep in mind that we live and work increasingly in a cross-platform world. Even if you “only use PCs” in your facility today, an engaging new network product for Macintosh may beckon tomorrow. Think inclusively.

Cheetah Multimedia Network Server

Experiments and prototypes are springing up in the arena of broadcasting or multicasting full-motion video (interactive and non-interactive) over the global Internet. Conventionalized commodity solutions are not here yet, and for that reason, the kind of multimedia network servers that FL teachers are considering work on a local area network, or perhaps the intranet level. The latter group is dominated by the video-on-demand entertainment systems found in large hotels and on some airlines, as well as networked training materials used by some major corporations. Lots of niftiness, to be sure, but at a price guaranteed to make your dean and mine blanch.

TNCi (The Network Connection), with their Cheetah and Hi-Per Cheetah family of video-capable file servers, has been making significant inroads into the market that many foreign language teachers shop in. The Cheetah is impressive both for its intelligent design and its varied capabilities. Rather than a proprietary hardware design, the Cheetah uses an ordinary Intel Pentium processor. You can customize its configuration to include one or multiple Pentium processors when ordered, or add additional processors as your needs expand. To keep throughput up under heavy user load, TNCi adds an APIC (Advanced Programmable Interrupt Controller) chipset, and a RAID 5 disk array that also provides lots of disk storage, up to 112 gigabytes. You will need that disk space when you start capturing and storing full-motion video at the customary 30 frames/second. Happily, Cheetah also supports an array of video compression technologies, including MPEG1, MPEG2, JPEG, MJPEG, AVI and DVI. That means you can stretch that 112 gigabytes of storage to hold 224 hours of on-line video.

The designers of Cheetah wisely steered clear of proprietary video-enabling software such as the schemes marketed by IBM, Oracle, and Starlight's *StarWorks*. Instead, they chose to base Cheetah on Windows NT, thus providing support for TCP/IP, IPX, NetBeui, and Ethertalk protocols. This means that you are not constrained to use "their kind" of video capture or storage, and you are free to choose clients and peripheral devices from a score of manufacturers, all the while confident that they will be compatible with your network server. Likewise, you are not forced to use "their software." In fact, Cheetah comes with no installed software as such. You can choose any video software products compatible with the Windows NT platform. Because it is a non-dedicated server, Cheetah can provide ordinary local area network services like file and printer sharing for Mac and PCs, as well as such Internet services as email and FTP hosting, and Web serving. In a word, Cheetah gets high marks for intra- and interoperability.

All Cheetah servers can support 120 concurrent users, and the Hi-Per Cheetah superserver scales up to handle 300 concurrent users. Its networking topology is open and compatible with Ethernet/Fast Ethernet, FDDI, and ATM, as well as the less common CDDI and 16 MB Token Ring. This attention to scalability should pay off for its makers as well as its users.

Frequently-Asked Questions

I chatted with Jim Reiner, TNCi's Vice-President for Engineering, about the questions he is frequently asked by prospective Cheetah owners, and asked him a few myself. Jim said that if your client machines are connected to your network using ordinary 10BaseT Ethernet, you will find the throughput provided by that arrangement adequate for videoserving, too. This may spare you rewiring your lab.

He emphasized that Cheetahs, like other local area network servers, are meant to be unattended systems. Acquiring a Cheetah videosever should not require you to hire another person just to keep it running, thanks to its non-exotic and non-proprietary engineering.

Because Cheetah comes with no installed software as such, customers often ask where they can get software information and advice. TNCi staff, Jim assured me, keep a sharp and self-interested eye on the software that's out there, and they are willing to consult with customers and refer them to known workable solutions.

The base price for a Cheetah server is about \$20,000, a sizable chunk of which is the cost of disk storage. The good news is that the price of commodity hard drives has been cascad-

ing downward as their capacity has doubled about twice a year. Keep in mind, too, that an Intel-based local area network server for file and printer sharing could cost you one-third or more of that \$20K price tag anyway. For more information, I recommend the TNCi Web site at <http://tnc.www.com/home/catalog/cserver.html>.

Demonstration Site

While I did not have the opportunity to meet and greet a Cheetah for this article myself, I did talk with Vic Svec of Phillips Academy, whose Andover, Massachusetts' Language Learning Center (LLC) doubles as a demonstration site for a Cheetah system. Vic, who is also Chair of the Russian Department, explained that with the help of a Cheetah server, the LLC has become a state of the art 100% digital foreign language learning facility in which all the out-of-class components of foreign language learning (audio, video, and computer work such as email and word-processing) have been put into one room and onto a single machine. At the LLC, that single machine is an Apple Macintosh PowerPC 5260, and there are 40 such network clients available for student use, each capable of independent operation. Each of those 40 students can concurrently access the same digitized video material at any point each student chooses and without affecting the other students. One or more students can be watching the beginning, while others are watching the middle, and still others the end of the same digitized video clip stored on the Cheetah server. At the same time, students can listen to digitized pre-recorded audio from the server, record their own voices, use a word-processor, browse the Web, or send and receive Internet email. For more information, visit Vic's Web site at <http://www.tiac.net/users/svecv/>.

Incidentally, Apple Computing has been busy cultivating ideas of their own in the testbed of Phillips Academy's LLC. Keep an eye out this fall for Apple Network Assistant, a new software product that combines an intercom system with features of Apple Network Installer and the capabilities of Timbuktu and Carbon Copy. The intercom component allows students and instructors to speak with one another directly through headsets, and page the instructor for assistance.

Just as this article headed for the printers, we also got wind of another recent Cheetah installation in progress at the Language Lab at George Washington University, directed by Brigitte Charlotteaux, the IALL Mid-Atlantic Regional Group Leader.

Conclusion

Are we there yet? True affordable commodity real-time global interactive multimedia communications technology is still a ways off. And in the interim, here's hoping we can find a shorter term for it. How far off? On the global scale, it looks as if it's well within the orbit of Jupiter, and on the scale of your average laboratory or classroom, we are, very nearly, there. ■

Bill Wyman heads the Computing Support Division of Academic Media Services at the University of Colorado—Boulder.