Byte July 1996

Teamwork on the Net

Adaptive multimedia applications will aid collaborative work.

by Bob Emmerson and David Greetham

Of the 44 million computers that will sit on Europe's corporate desktops in 1998, 71 percent will be connected to a LAN. Most network managers think of LANs as coherent topologies and regard WANs as links between islands of information. But what's really needed for computer-supported collaborative work is a seamless environment for multimedia applications, where terms like LANs and WANs are irrelevant.

Mobile professionals and corporate workers in far-flung locations need more bandwidth and less-expensive services to work effectively together. The public WAN infrastructure is circuit-switched and employs synchronous communications. LANs, meanwhile, employ a packet-switched infrastructure, so their communications are asynchronous. These services are all fine in their own way, but they don't work well together.

However, they all work with the Internet. Access is ubiquitous, easy, and affordable, so the Net has become an effective integrator--a common denominator that's ideal for e-mail and file transfers. Therefore, all the Net needs to become the ideal platform for collaborative work is a higher quality of service (QoS)--a guaranteed bandwidth for multimedia applications and minimal latency, or delay--which is critical for voice communications.

But the Net currently comes up short on bandwidth. And packet-switching latency is a major issue.

Shared Workspaces

The European research community, one of the early adopters of the Net, is currently pioneering many new technologies, applications, and systems that enhance collaboration. GMD, the German National Research Centre for Information Technology (Sankt Augustin, Germany), for example, has implemented the Basic Support for Cooperative Work (BSCW) shared workspace system on its Web server (located at http://bscw.gmd.de).

BSCW, which works with a regular Web browser, includes so-called awareness services, where icons attached to objects indicate events that have occurred with regard to the objects. Thus, team members immediately recognize if an object is new, if it's been read recently, or if a link has been renamed.

Each member of a workspace has a personal list of all events in that workspace, so everyone is always aware of oth-er members' activities. The BSCW site also offers public access for creating your own shared workspaces .

GMD says it will eventually support an e-mail interface that allows users to send e-mail messages whenever a certain event occurs. Thus, users will automatically get an update of the activities within the workspace even when they're not on-line.

Real-Time Collaboration

Packet switching is ideal for data communications where latency is not an issue. It's also efficient, since the communications link can handle different types of traffic coming from any number of parties. With packet switching, links can work to their maximum capacity, but the real-time behavior varies widely. Therefore, because of its packet-switched nature, the Net is currently a poor communications medium for voice traffic, since it wasn't designed to handle isochronous information.

Bandwidth is also a key issue for real-time collaboration on the Net. The replacement for the current IP protocol is being addressed by the Internet Engineering

Task Force (IETF). IP version 6 (IPv6), the next generation of Net protocols, has been designed to run on high-perfor mance networks and still be efficient on wireless and other low-bandwidth connections.

The benefits that IPv6 will bring include more efficient addressing, faster routing and delivery mechanisms, multimedia enhancements, and the potential to make the Net more secure. Deployment has already started, but about two more years are needed before IPv6 becomes the de facto stack for desktop machines.

IPv6 addresses the latency issue by enabling prioritization. For example, a video bit stream can be split into several channels that have progressively higher definition. When congestion occurs, the higher-definition channels are dropped, thereby decreasing the bandwidth requirement and allowing a lower-resolution image to get through to the receiver. This technique, which is known as hierarchical encoding , is particularly suitable for broadcasting.

The IETF is also working on the Reservation Setup Protocol, or RSVP, which operates above the IP layer. RSVP aims at receiver-initiated allocat ion of resources to reserve bandwidth over the link and thereby simulate a circuit-switched connection. This network-control protocol will allow Net applications to obtain a higher QoS for their data flows.

When an application in a host requests a specific QoS, RSVP is used to pass the request to each router along the paths of the data stream. The protocol enables multicasting (i.e., one-to-many links) and is receiver-oriented. To efficiently handle heterogeneous receivers and dynamic group memberships, RSVP makes receivers responsible for requesting resource reservations. Although it's a potentially expensive solution, it's an excellent way of minimizing the problem of latency, which is a particular problem for voice transmission.

Data Compression

Circuit-switched voice signals are sampled eight times per second, and each sample has 8 bits of data; thus, high-quality speech requires a bandwidth of 64 Kbps. Compression techniques can reduce the data-throughput rate and still give acceptable quality, but they all steal cycles from the CPU.

Video compression demands less timing accuracy than audio, since large portions of a picture, such as the background, don't change from frame to frame. However, MPEG is not a solution for the Net, since this codec requires fixed bandwidth (see "Toss your TV," February BYTE).

In 1992 and 1993, the French research organization Institut National de Recherche en Informatique et en Automatique, or INRIA, tested multimedia compression algorithms for the Net that are now incorporated into the IETF's Real-Time Protocol (RTP). RTP, which specifies how to adaptively encode audio and video information, is increasingly being used on the Net's multicast backbone, MBone (see the sidebar "The Internet's Multicast Backbone"). And INRIA's Rodeo Group implemented RTP in its Internet Videoconferencing System (IVS), which is also based on MBone.

Adaptive Applications

INRIA's research indicates t hat multimedia applications over the Net work best if they are able to adapt to the varying bandwidth of this heterogeneous network. Adaptive applications reduce the amount of data transmitted during periods of congestion; this results in either a lower-quality transmission (known as throughput adaptation) or a reduced frame frequency (known as delay adaptation). These applications detect delays and then establish a correct-size "playback" buffer on the recipient's system.

Throughput adaptation can be achieved using hierarchical encoding or variable compression ratios. The former is attractive for large broadcast applications, where closed loop control would be impractical. The latter has been implemented by INRIA in its IVS system.

Senders constantly poll the receivers being used in a multicast group to estimate the quality of reception and then compute the available bandwidth and set the compression ratio accordingly. As a result, a high-quality signal is transmitted when the network link is lightly loaded, and a poorer-quality signal is sent when it's crowded.

Maximum efficiency of adaptive applications (a relative parameter) does not depend on network resources (an absolute parameter). This adaptation concept was first implemented for video signals, and INRIA has now ported the technology to audio signals. The system also ensures good "cyberspace citizen" behavior, since feedback control prevents video applications from swamping the Net's resources.

The polling process is part of RTP. Therefore, any system that uses this protocol, such as the IVS system, becomes adaptive. This means that the application makes optimum use of network resources and no hardware upgrades are needed.

The quality of audio and video over the Net is quite restrictive; however, new protocols, the MBone, and adaptive applications point the way toward the QoS we're going to get in a few years. But don't expect the multimedia equivalent of a free lunch. Clients will require an ISDN conn ection in order to get optimal video quality.

Until the Net is upgraded to better handle multimedia applications and there is smoother integration among today's LANs and WANs, the optimum solution would seem to be a system that runs across all main media types (e.g., LANs, Public Switched Telephone Network [PSTN], ISDN, and TCP/IP). Cybertec's (Newton Abbot, U.K.) V-COM software meets these requirements.

V-COM has a software codec, so transmission speed depends on the PC, the modem, and, in the case of the Net, the time of day (because of the variable delay). However, Cybertec says that V-COM brings 24 color frames per second on an analog line using a 28.8-Kbps modem and approximately 4 to 5 fps over the Net. It also allows interactive, real-time application sharing and whiteboarding. Expect to see more all-in-one videoconferencing solutions of this kind hitting the market fairly soon.

So ignore the hype and accept the current limitations of the Net. If you need to collaborate with collea gues, clients, and the like, the Net is going to be the medium of choice for people who pay their own phone bills. Services will improve, and collaborative multimedia applications are coming.