

# Multimedia MedNet

## A MEDICAL COLLABORATION AND CONSULTATION SYSTEM

**Robert Simon**  
**Donald Krieger**  
**Taieb Znati**  
**Raymond Lofink**  
**Robert J. Sclabassi**

*University of Pittsburgh*

**MedNet and its predecessor, NeuroNet, have been in use since 1985. The current system provides real-time monitoring and multiparty consultation during 1,600 brain surgeries per year.**

**T**he field of medicine is a natural application for distributed, collaborative multimedia systems. The data types, manipulations, and user interactions are broad and demanding, since the medical environment involves potentially life-critical decisions and integration of a rich data set.

One area where a collaborative system can substantially improve patient care is clinical neurophysiology, a consultative service of diagnostic and monitoring techniques used to assess nervous system function. Diagnostic techniques include neuroelectric measurements from muscles, peripheral nerves, and the brain, and blood flow measurements from intracranial arteries. During brain surgery, monitoring techniques help prevent damage to nervous system structures by continuously measuring activity recorded directly from the brain in real time.

This article discusses an ongoing, long-term, distributed multimedia project developed at the University of Pittsburgh Medical Center and used on a daily basis at seven hospitals and multiple diagnostic and research laboratories. Multimedia MedNet functions as a real-time medical command, control, and communication system running in a general-purpose hospital computing environment. MedNet provides real-time monitoring and multiparty consultation and collaboration during brain surgery for approximately 1,600 cases per year. This intraoperative monitoring places a real-time control loop around the patient and surgeon to warn the surgeon when the patient's nervous system is being damaged. These real-time monitoring and collaborative functions distinguish MedNet from other distributed medical multimedia systems, which are used primarily for picture archiving and teleconferencing.<sup>1</sup>

MedNet design and implementation began in 1985 and has included several phases.<sup>2-4</sup> A key goal was to build a production-quality system by combining off-the-shelf components, such as an unmodified Unix operating system (with no real-time scheduler), with hardware and software developed specifically for MedNet and its earlier version, NeuroNet.

A major challenge in implementing MedNet is to develop techniques for the appropriate processing and display of real-time multimodal medical information. The phase currently under implementation, Phase II MedNet, uses a purely digital system. Phase I experiences generated two major design criteria that are crucial to the success of Phase II:

- Develop effective communication control strategies to handle the wide variety of real-time data types that arise in a medical environment, including audio, neurophysiological, and video.
- Design effective system-level support for users' collaborative activities within MedNet.

In addition, the communication strategy for Phase II must support a large distributed system and should be interoperable with future communication infrastructures.

Phase II incorporates a new group-communication protocol for distributed multimedia systems called DIPCS (Distributed InterProcess Communication System).<sup>5</sup> The basic design tenet of DIPCS is that communication management and support are determined from human activities during a particular communication session, not from the types of media involved in that session. Phase II MedNet represents the final integration of all sources of information into a unified digital system.

## MEDNET AS A MULTIMEDIA COMMAND AND CONTROL SYSTEM

NeuroNet,<sup>2,3</sup> the precursor to MedNet, was designed to support real-time neurophysiological intraoperative monitoring, intensive care unit monitoring, and diagnostic evaluation. It also provided distributed consultative support. The Apollo Domain file system was NeuroNet's basic communication system.

The current implementation, Phase I MedNet,<sup>4</sup> includes over 100 computer nodes in service at the University of

### Intraoperative monitoring within MedNet

The following case shows the time-critical nature of MedNet data and the value of remote collaboration. A patient was being operated on for a right carotid endarterectomy, which is associated with decreased blood flow to the brain. The consulting neurophysiologist was remotely monitoring the progress of this case, along with several other cases at different hospitals, via data obtained from the right and left hemispheres of the patient's brain.

The accompanying screen shows the data file for this case. There are two channels of neurophysiological information, corresponding to the output of two sensors:


- Channel 1 measures activity from the right side of the brain, where the operation is taking place.

- Channel 2, corresponding to the left side of the brain, is used as control data.

Both channels have distinct troughs and peaks in the waveform. These are the areas of interest to the neurophysiologist. The negative trough running down both channels is called N20 and corresponds to electrical activity projecting to the cortex, while the positive peaks next to it are called P30 and correspond to excitation of the cortex itself.

Shortly after the surgeon clamped the right carotid artery, the response from the right hemisphere appeared reduced in amplitude. This observation was made remotely, across the network, by the neurophysiologist. This can be seen in waveform A (in the figure) along with the corresponding

annotation from the technician. The neurophysiologist immediately informed the surgeon of the change. Upon closer inspection of the data in the operating room, it was hypothesized that the patient may have had a thrombus, a blockage to the right middle cerebral artery, at the time the right carotid artery was clamped due to the formation of clotted blood particles. The surgeon placed a shunt to rule out that possibility, but the shunt did not correct the observation (waveform B). So the patient was started on barbiturates to reduce the metabolic demand, and an arteriogram, an intraoperative evaluation of blood flow through the arteries, was performed. Prior to performing the arteriogram, the surgeon completed the endarterectomy. The arteriogram detected a blockage at the mouth of the middle cerebral artery, and this blockage was then dissolved intraoperatively. The barbiturates were discontinued. Upon completion of the case the EPs (evoked potentials) had returned to 50 percent of baseline (waveform C) amplitude, and the patient awoke with a mild weakness, which eventually resolved.



**Figure A.** Evoked potentials collected using MedNet during surgery. There are two channels of data, from the left side and the right side of the brain. Channel 1 is where the operation is taking place; Channel 2 is used for control purposes. The troughs and peaks in both channels represent the strength of nervous system response and are used to judge nervous system integrity during an operation.

Pittsburgh's seven hospitals. Besides neurophysiological data, Phase I MedNet incorporates real-time data streams of digitized audio as well as extensive physiological and computer-generated data. Phase I also incorporates video streams transported by a parallel, broadband, analog network into MedNet workstations. Phase I notably enhances the diagnostic and consultative functions of the original NeuroNet by more comprehensively portraying a patient's condition and internally providing effective multiway consultation among remote health care providers.

MedNet allows continuous surveillance of the patient's condition during brain surgery; neurophysiologists can provide immediate feedback to neurosurgeons. For example, the "Intraoperative monitoring with MedNet" sidebar presents a case where the use of MedNet's remote monitoring prevented serious damage to a patient's nervous system during an operation.

Neurophysiological data enables assessment of a patient's nervous system during neurosurgery. However, an intraoperative assessment is best made when based on neurophysiological information in conjunction with accompanying real-time video and audio data from the operating room. The reason is that correct interpretation of intraoperative data highly depends on surgical manipulations occurring in the operative field. Without remote access to the video images, accompanying audio information, and the neurophysiological data, the neurophysiologist must be in the operating room more often, limiting simultaneous case monitoring by one neurophysiologist.

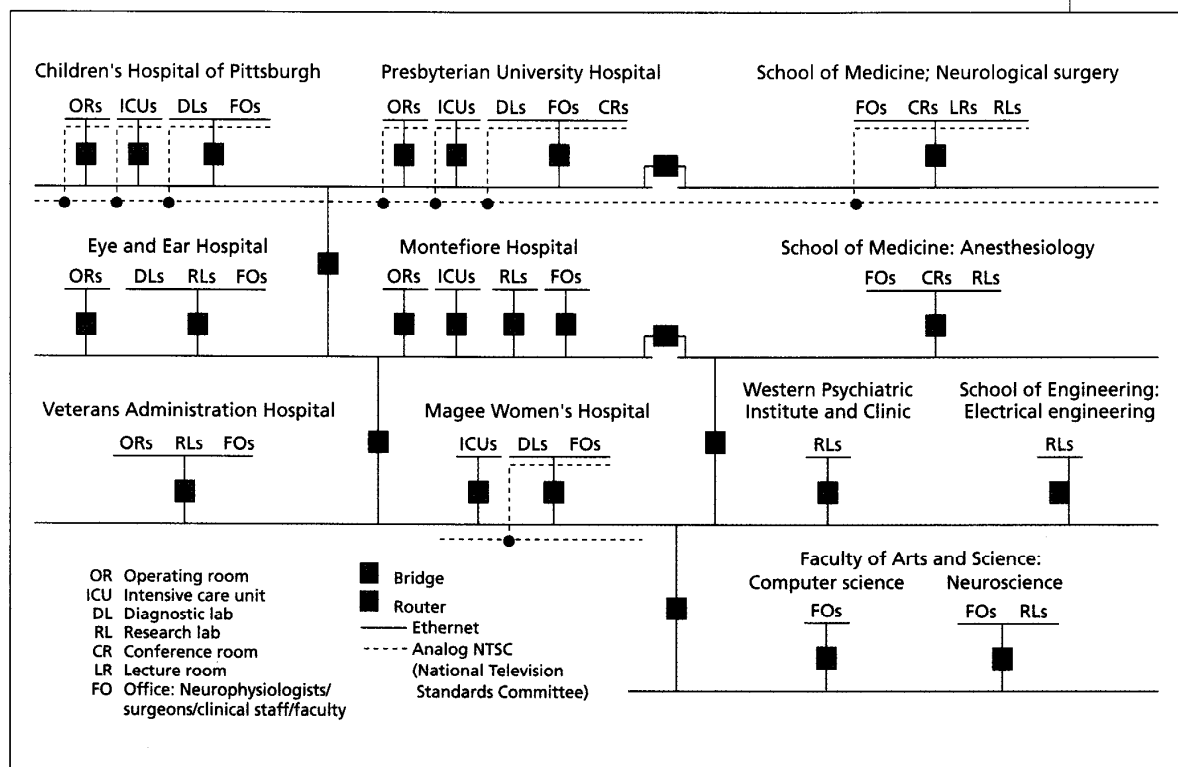
In practice, a neurotechnician operates a mobile MedNet node connected to the network during an entire operative or diagnostic procedure. Both the technician and the mobile node are at the procedure site. The technician runs various tests, which can be displayed at any MedNet node on the network, and often textually annotates such tests. Neurophysiologists can monitor these tests remotely via a MedNet node that may include video images of the operation and audio monitoring and multiway communication between operating room personnel and neurophysiologists. All diagnostic tests and procedures are stored for subsequent retrieval, analysis, and report preparation.

Hence, MedNet's functional requirements include

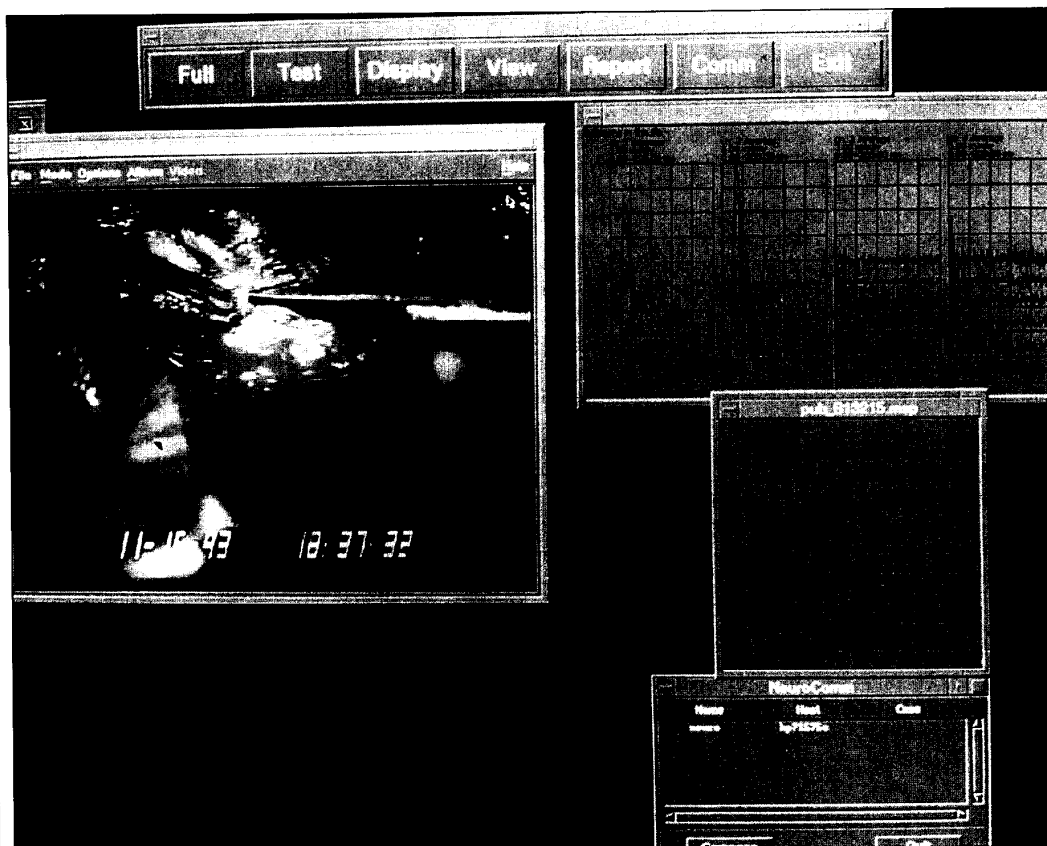
- simultaneous acquisition and distribution of video, audio, neurophysiological, and autonomic data;
- real-time display and analysis of digital data acquired locally and remotely;
- a reliable, appropriate platform for distributed monitoring and collaboration activities within the application; and
- dependable, convenient archival storage of all data.

## IMPLEMENTATION OF PHASE I HARDWARE

Phase I MedNet was built with stimulators, recording amplifiers, and off-the-shelf computer and communication hardware.



**Figure 1. Diagram of medical center digital and analog networks, showing operating rooms, intensive care units, labs, offices, and the connections between them.**



**Figure 2. Phase I MedNet screen. The video picture is from a video camera aligned with the operating microscope. The two windows showing waveforms are from intraoperative neurophysiological information that is acquired and digitally processed in real time. Each waveform window represents a different case. The NeuroComm window shows other users on the network; they communicate through a multiway audio program or an X-Windows based program.**

## Network

The communication infrastructure contains analog and digital networks, as shown in Figure 1. The system is designed around an Ethernet backbone for digital information (including audio) and a parallel CATV system for video and audio (broadband, 450 MHz).

Two video sources from the operating room are integrated in the system. Many neurosurgical operations use microscopes equipped with video cameras, and in some operating rooms, a ceiling-mounted video camera focuses on the operative field. Both video sources have an accompanying audio channel. These video images, along with the audio channel, are transported by the analog broadband network and displayed directly into MedNet workstations via Hewlett-Packard VideoLive software and hardware. Sequences are captured, stored on disk, and retrieved for later use with a dedicated video card. An additional source of audio data from hardware within the MedNet nodes is transported through the digital network, which has connections in 44 operating rooms, two interventional neuroradiology rooms, 20 neuro-intensive care beds, over 180 other ICU beds, 10 diagnostic laboratories, and several research laboratories.

The analog broadband network is installed in four operating rooms at Presbyterian University Hospital and three operating rooms at Children's Hospital of Pittsburgh. These parallel networks extend to all the neurosurgery and neurophysiology faculty offices at the health center, as well as several conference and lecture rooms. This enables these facilities to support not only more effective clinical care but also research, teaching, and general consultation tasks.

## Nodes

Three types of MedNet nodes can be directly attached to the network. The first is a mobile rack-mounted node for patient monitoring and testing in operating rooms and intensive care units. If necessary, these nodes can function in a stand-alone mode. A second type of data acquisition node is stationary and is used to test patients in diagnostic or research laboratories. The third type is a workstation for remote monitoring. It contains a MedNet node with audio hardware and an HP VideoLive card that can display video images on the workstation screen. Video input is obtained through a local tuner in a VCR attached to the CATV network. Primary users of these nodes include

neurophysiologists (performing remote monitoring), computer support personnel (developing and monitoring network functions), and administrative and clinical staff. These nodes support both distributed consultative tools and general-purpose computing.

All patient testing and monitoring is performed on workstations equipped with a data acquisition controller that provides an interface between the workstation and the electronics attached to the patient. The data acquisition controller can acquire up to 64 input channels of physiological and neurophysiological data, with an aggregate sampling rate of 160 Kbytes/second. It provides digital and analog output for the generation of stimuli used to evoke activity in the patient's nervous system.

Extensive signal-processing algorithms are implemented on the local acquisition nodes. Signal processing routines at all MedNet nodes let the remote user control data analysis and display. Phase II will incorporate additional physiological and other types of monitors, such as a transcranial doppler. This instrument monitors blood velocity through intracranial arteries. Its primary output is a video data stream.

## IMPLEMENTATION OF PHASE I SOFTWARE

Conceptually, Phase I software has three parts: the interface, the communication protocol, and the file system.

### Interface

Medical multimedia systems face a major challenge in integrating extensive real-time, multidimensional data before its use in the decision-making process. The MedNet interface is divided into four distinct X-Windows based programs: Test, Display, View, and Report.

The Test program establishes an interface to the data acquisition hardware that lets the operator select one of a series of procedures or signal-acquisition modes. Test feeds the resulting raw data to the Display program on the local node, and passes either processed or raw data to the transport layer.

The neurotechnician selects signal acquisition procedures and display routines through mouse-driven, push-button menus. Display provides both an oscilloscope-style window and a scrolling window to present the raw neurophysiological data as Test acquires it. The neurotechnician can textually annotate the wave forms, specify baseline traces for comparison, and pass the data through various signal-processing routines. The View program receives the data from Test and presents a variety of trending displays, each independently controlled. When an operative procedure is started, View creates an independent window for each data component or other case being acquired. Figure 2 shows two View windows. Each window displays four data channels. View receives data from the transport layer and updates the appropriate window.

Users can review any component of a case by clicking on the appropriate window and requesting a particular subset of the data. Because View maintains a waveform history for a current case within one window and can easily display a different test or case within a different window, a neurophysiologist can rapidly assess multiple operations' progress and communicate with any appro-

priate party. Either an audio communication program, called the Talker, or an X-windows based program, called NeuroComm, supports this communication. Remote users are shown in the NeuroComm window in Figure 2. This interface allows one neurophysiologist to oversee all the clinical activity for which he or she is responsible.

Report is generally used to review and print waveforms after a procedure has been completed. It combines the annotation capabilities of Display, the trending facility of View, and a capability to generate a standardized report showing selected waveforms and the annotations.

The Talker program permits multiway audio communication—typically between neurophysiologists, neurotechnicians, and MedNet support staff—using the hardware available on MedNet nodes. Audio packets are transmitted through UDP (user datagram protocol) sockets.

### Communication protocol

The MedNet communication protocol (MNCP) manages network support for distributing neurophysiological information. MNCP is a distributed-group process-management system that provides a global naming and location service, general communication primitives, and group membership functions. It also establishes a transport service for real-time neurophysiological data. This implementation uses Network Computing System remote procedure calls and Unix domain sockets. MNCP services are available to MedNet clients, such as View and Display, via library calls.

The interface to MNCP on each MedNet node is through a local Information Service Daemon (ISD). MNCP uses a Data Location Daemon (DLD) for a central naming and location server. The DLD contains directories and locations of cases and announcements about new cases or cases no longer available.

MNCP clients send and receive data through their local ISDs. Every neurophysiological data transfer, via the network, is directly from one ISD to another. Each time an ISD starts, it registers itself with the DLD. MNCP clients then register themselves with their local ISDs. Clients can obtain case directories by making calls to their ISDs, open or close cases, and perform data reads and writes. On behalf of a client, an ISD can also request information about cases that are available remotely and ask the DLD to inform it when a new case becomes available. The DLD maintains this knowledge through information it receives from ISDs.

### File system

Data acquired by MedNet is stored on and retrieved from disk files in the Neuro Data Structure file format. Built on top of the Network File System, NDS files consolidate all data related to a given case. Different data streams for a case are stored in individual files. A channel manager structure identifies and manages data streams. This structure contains information for each data type in a header and handles variable-length records. The file header contains information for synchronizing the data streams. Data types are currently defined for classes of neurophysiological, physiological, anesthesiological, audio, video, and imaging data.

A software library called Neuro Data Format (NDF) sup-

ports NDS file input and output. NDF lets a calling application access all NDS data objects through the same function calls, regardless of file type. Furthermore, the calling application does not need to know where the data resides. In this manner, NDF functions as a location-transparent file system. Depending on the data type and application choice, data can be returned in integer, real, scaled, or a completely unstructured format appropriate for multimedia playback.

## **EXPERIENCES WITH PHASE I MULTIMEDIA MEDNET**

NeuroNet has been in continuous, active use since 1985, and Phase I MedNet enhancements have been in use since early 1992. MedNet can be classified as a soft real-time system, even though it doesn't use real-time scheduling for any of its components. Real-time scheduling is not needed, partly due to the specialization of MedNet components, such as digital signal processing (DSP) boards and video and audio boards, which tend to function independently and therefore do not compete for resources.

We have learned by using Phase I MedNet that the range of end-to-end and synchronization performance requirements for a real-time call have more to do with the *context* in which a call is placed than the media types involved in the call. (A multimedia call is a single session of logically related data streams, potentially between multiple senders and receivers.) Collaboration in MedNet includes two broad types of activities.<sup>6</sup> One communication type is essentially multiparty interactive, where several participants discuss a particular situation or problem. This requires that the communication latency be short enough not to affect the collaboration. Another type of activity is long-term monitoring of a situation with only an occasional need for multiway interaction. This communication type can tolerate longer communication latencies and still perform adequately. Similar requirements exist for the synchronization of media types. Some applications (such as teleconferencing) require tight synchronization, while others (such as coordination of video data with the digital display of neurophysiological data) can tolerate looser synchronization. Finally, privacy and security are important in this environment. Participants in a collaborative setting, as well as those whose actions are being monitored, should be aware of one another, both to avoid needless effort in communicating and to maintain assumed standards of privacy.

These application requirements suggest that call management should let the performance requirements of the media types within the call and of the call itself (collaborative, monitoring, and so on) determine the system support level. Furthermore, call management semantics should appropriately address privacy, mobility, and a heterogeneous equipment base. We have developed a communication management system, the Distributed InterProcess Communication System (DIPCS), that meets the requirements of MedNet and other similar multimedia systems.

## **DESIGN OF PHASE II MEDNET**

Phase II MedNet provides a command and control system for surgeons and neurophysiologists that handles all

MedNet information through a unified digital system. Phase II MedNet will run over a high-speed digital backbone, such as an ATM (asynchronous transfer mode), that can offer real-time performance.

Since the functional requirements for the current MedNet implementation have proven to be satisfactory and need not be changed, the major challenge in implementing Phase II is to develop effective communication control. A unified communication control system is required, since all information sources within the system will be digitized. The control mechanism provides a simple programming abstraction to simplify management of multiparty collaboration and interaction patterns. It also specifies and controls real-time connections in the anticipated BISDN-type (Broadband Integrated Services Digital Network) environment.

## **Distributed InterProcess Communication System**

Phase II MedNet uses the DIPCS model for connection and communication management.<sup>5</sup> This model has three basic objects: a multimedia device abstraction model, stream model for the devices' communication characteristics, and a process group model for controlling communication management and multiparty collaboration.

The DIPCS model is one of several distributed multimedia system models that represent hardware and software multimedia devices in an abstract way, and enable separation between general device control and a particular communication instance of the device.<sup>7</sup> However, unlike other models, DIPCS is designed as a general-purpose, end-to-end, connection-level management tool, targeted toward a real-time network that can offer performance guarantees via resource allocation. Furthermore, DIPCS uses a distributed multimedia process group model for all call management. This simplifies controlling collaboration patterns within MedNet. This is in contrast to some real-time connection management tools that neither specifically target real-time networks nor provide group operations for control of multiple multimedia communication streams (for example, see Schulzrinne et al.<sup>8</sup>).

**DIPCS DEVICES AND STREAMS.** Multimedia processing devices are often highly specialized (for example, digitizing video cameras, speakers, and so on). As a result, multimedia devices tend to be unidirectional; they can be senders or receivers, but not both. This also applies to multimedia capture and display equipment, which may be embedded in a conventional workstation. We refer to multimedia devices as either media sources or media sinks. (This includes multimedia information produced entirely in software, such as computer-generated animation.)

From a communication standpoint, different types of sources and sinks are distinguishable by the data stream on and the speed at which they operate. DIPCS represents sources and sinks of multimedia information through a device abstraction. Each device represents a communication endpoint for a source or sink. The actual endpoint may be a buffer in computer memory or a hardware device directly connected to the network. Each instance of a device's communication requirements is represented to the

network by the use of a stream. Each stream contains a message rate specification  $\lambda$  and a reliability parameter  $\alpha$ .

A multimedia site is a collection of streams along with their associated devices. A site represents a work area, such as a workstation or diagnostic lab, that can be under one user's control. Sites associated with users become members of AlphaDeltaPhi (ADP) groups.

**DIPCS ADP GROUPS.** The use of process groups and group-level communication offers a powerful operating-system-level abstraction for developing and controlling cooperating processes in a distributed system.<sup>9</sup> Basic communication and connection level management in DIPCS is accomplished through an ADP group. The parameter  $\alpha$  refers to the percentage between 0 and 1 of a stream's messages that must be delivered over each synchronization interval with a delay of no more than  $\Delta$ . The parameter  $\Delta$  is the maximum end-to-end delay that messages can suffer in the network, and  $\phi$  is the synchronization interval for a group connection. The synchronization interval is the time required to present a multimedia playback unit, such as a video frame. Different values for  $\alpha$  imply different performance requirements from the underlying network.

ADP groups are related sets of cooperating processes whose communication is supported by a spectrum of quality-of-service and message-delivery options coexisting in one group. Connection management at the group level is useful for several reasons:

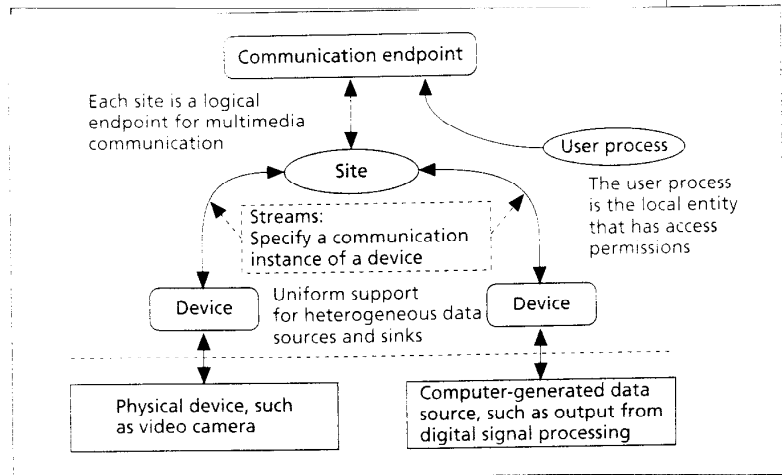
- Group communication is a well-understood distributed programming concept.
- Naming and routing information for group communication is supported in a general distributed system environment.
- ADP groups can be added with relatively little overhead to an existing group management system.
- Based on four quality of service (QoS) parameters ( $\alpha$ ,  $\Delta$ ,  $\phi$ ,  $\lambda$ ), a network request specification can be immediately produced.

There are two message classes, control and data, that can be sent to an ADP group. Control messages involve basic group management functions. Data messages in ADP group communication have  $\alpha$ -reliable semantics, where  $\alpha$  reflects the degree of on-time delivery guarantees required by the underlying application. For the network to guarantee an  $\alpha$  greater than 0, the ADP group communication network model uses resource reservation. Communication that is  $\alpha$  reliable implies that all participating sites can achieve the same quality of service even though all active members do not receive the same messages. It also means different multiple sinks of the same source data stream can receive a different QoS level. Because it is not realistic for the source to send multiple streams to each sink, different QoS levels can be supported at the network level. This enables better use of network

resources, since routes to sinks with lower QoS levels require fewer network resources.

ADP groups contain a collection of entities, called *comends* (communication endpoints). Each comend is a pair  $\langle \text{user, site} \rangle$ . A user may correspond to a human user or to some process that controls a system activity (for example, a database server). The user process has the required-system level of security privilege that allows it to use the site's computer and communication resources. The user is the entity at which all group membership actions are resolved (such as Join, Quit, Forward, and so on). Figure 3 shows the relationship of devices and streams to a comend.

An ADP group  $G$  with  $k$  members is defined as a set  $G = \langle \Delta, \phi, \text{comend}_1, \dots, \text{comend}_k \rangle$ , where  $\Delta$  is the group's maximum end-to-end delay and the value of  $\phi$  is the synchronization interval, which depends on the application's



**Figure 3. A member of an AlphaDeltaPhi (ADP) group is called a comend (communication endpoint). Each comend is a binding of a site, which represents a work area, to a user process. Sites contain devices, which are general representations of multimedia equipment, and streams, which are communication instances of devices.**

synchronization method. The network ensures that end-to-end delay of messages is less than  $\Delta$  and jitter is limited; therefore, synchronization can be guaranteed. We have developed a network support scheme to provide guaranteed performance for an ADP group connection.<sup>5</sup> The initiator of the group sets  $\Delta$  and  $\phi$  for the group.

To send or receive data messages from an ADP group, a comend must be part of that group. ADP group membership, which is granted only if all current group members explicitly approve the new member, is *closed*. The closed group requirement stems from the nature of collaborative work, common assumptions about working environments and privacy, and our experiences with MedNet. The principle is, "If I say something, I should know who I am saying it to," and "If I hear something, I should know who else is hearing it."

After a group has been formed, an application sends a request to join a group by calling Join (group\_id). Based

on the encapsulated stream types, the appropriate source/sink connections in the network are made. These connections are between streams of opposite direction (for example, sink to source) but of the same type (for example, video). Multiple source streams can be connected to a single sink stream (for example, multiple audio source streams can be connected to a single audio sink stream). This mixing mode is device specific. The group connection causes the allocation of the appropriate network resources. DIPCS provides other group membership primitives, such as Quit, Hold, Resume, and Forward.

Each time a comend is instantiated, it is associated with a Connection Manager and a Group Process Manager (GPM). Each CM handles group membership actions and maps and controls streams and devices. Phase I Information Service Daemons (ISDs) will be merged into Phase II CMs to provide information about local cases. The GPM handles the network control portion of ADP group support. After a connection is accepted by the CMs, the GPM inserts the group name into the network naming system and establishes a communication tree linking comends together. The GPMs use the global naming service that the Phase I Data Location Daemons (DLDs) provide. Based on the encapsulated characteristics of the media sink streams, the GPM calculates the required network support level. Figure 4 shows the relationships between the various communication managers.

## EVALUATION

Phase I MedNet is an example of a working collaborative multimedia system that integrates a variety of real-time data streams. Extensive system use has clearly demon-

strated the medical benefits of remote access to sophisticated medical diagnostic procedures:

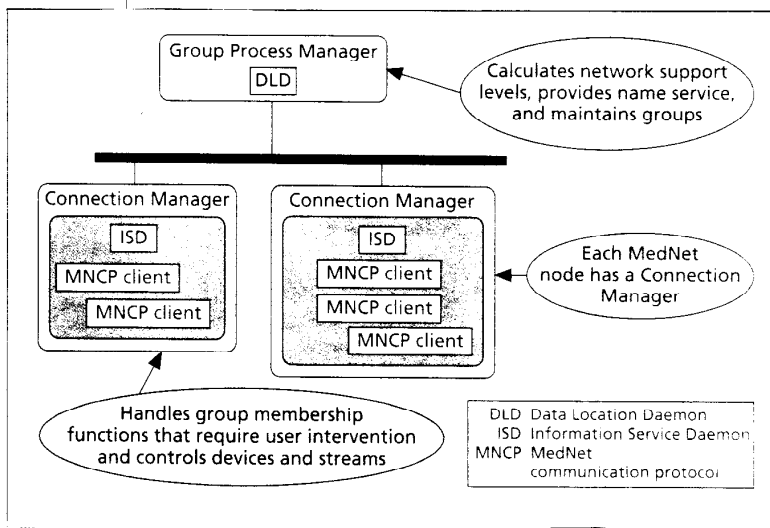
- Remote simultaneous monitoring of many cases makes optimal use of the staff neurophysiologist, who is the limiting resource in providing clinical services.
- Simultaneous display of all cases at multiple sites lets multiple observers monitor each case. This routinely enables continuous observation by both the technician in the operating room and the staff neurophysiologist, who may be at any location where the voice, data, and video are available. Simultaneous monitoring by multiple individuals minimizes oversights or delays in recognizing a change in a patient's nervous system.
- Remote simultaneous display promotes collaboration between staff neurophysiologists. This feature is especially valuable when data interpretation is problematic or when the case load is too great for a single staff neurophysiologist.

Phase II MedNet will continue to run on the fastest available workstation using conventional Unix operating systems. As system designers, we will not limit the number of applications running at once on a workstation to maintain a certain quality-of-service level, but will leave this decision to individual users. However, this philosophy does not apply to the underlying network, where multiple users compete for resources and where congestion and overload can rapidly degrade performance for all system users.

The major challenge in implementing Phase II is communication control. The DIPCS process group abstraction and programming model for distributed multimedia systems provides appropriate connection-level management for MedNet, and is being used for communication control. DIPCS supports Phase II MedNet communication semantics by

- letting the performance characteristics be set on a per call basis,
- having a closed group membership protocol,
- providing primitives appropriate for collaboration management, and
- reducing the complexity of distributed application development.

Another important problem which must be resolved for Phase II MedNet is how to manage collaboration patterns and the associated information glut. MedNet continues to increase its functionality in terms of neurophysiological monitoring capabilities. Users change location or are added to the system, and system connectivity increases. This connectivity is complicated by the fact that not all nodes have the same workstation and network capabilities. Some mechanism must be employed to automatically manage information overload and mediate the system's collaborative aspect. This must include the capability



**Figure 4. The Distributed InterProcess Communication System (DIPCS) communication management, showing the Connection Manager and the Group Process Manager. The GPM handles the network control portion of ADP group support and relies on the naming service provided by the Phase I MedNet DLD. Each CM handles group activities such as group membership actions, and maps and controls streams and devices; it uses the Phase I MedNet ISD to handle local cases.**



to distinguish between critical and noncritical collaborations and discussions.

We are working on an information-filtering and user-profiling scheme for Phase II. This will transparently support varying classes of users working in different settings under different constraints. A class of users might be neurophysiologists or graduate students. The settings are the physical resources available to a user, such as a workstation or a videophone hookup. The constraints refer to the time requirements and priorities of the tasks users undertake—for example, monitoring active cases. The user profile tracks not only user types such as neurophysiologists but also individual user sophistication, to generate appropriate interfaces. |

### Acknowledgments

We gratefully acknowledge the support of Computational Diagnostics Inc., The Benjamin Franklin Fund for the Commonwealth of Pennsylvania, and Hewlett-Packard for their generous support of the project. We wish to thank our colleagues for many valuable discussions, particularly Don DeLauder, Andre Srinivasan, Matthew Henry, Bill Madden, Scott Cline, Bob Leichner, Allan Kuchinsky, Walt Hill, Frank Prince, Bonnie Nardi, and Earl Doyle. The editorial assistance of Lawrence Rowe and the comments of the anonymous reviewers greatly helped improve the final version.

### References

1. L. Orozco-Barosa et al., "A Multimedia Interhospital Communications System for Medical Consultations," *IEEE J. on Selected Areas in Communications*, Vol. 10, No. 7, Sept. 1992, pp. 1145-1156.
2. D.N. Krieger, G. Burk, and R.J. Scabassi, "NeuroNet: A Distributed Real-Time System for Monitoring Neurophysiological Function in the Medical Environment," *Computer*, Vol. 24, No. 3, Mar. 1991, pp. 45-55.
3. R.J. Scabassi, R. Lofink, and E. Doyle, "NeuroNet: A Distributed Microprocessor Network for Clinical Neurophysiology," in *Microcomputers in Medicine*, M. Geisow and A.N. Barrett, eds., Elsevier Science, New York, 1987, pp. 283-303.
4. R.J. Scabassi et al., "The Multimedia Medical Monitoring, Diagnosis, and Consultation Project," *Proc. Hawaii Int'l Conf. System Sciences (HICSS-24)* Vol. 3: *DSS and Knowledge-Based Systems*, IEEE CS Press, Los Alamitos, Calif., Order No. 2121, 1991, pp. 717-728.
5. R. Simon, T. Znati, and R.J. Scabassi, "DIPCS: An Interprocess Communication Architecture for Distributed Multimedia Systems," *Multimedia Tools and Applications*, to be published in 1995.
6. B. Nardi et al., "Turning Away from Talking Heads: The Use of Video as Data in Neurosurgery," *Proc. Interchi 93*, ACM, New York, 1994, pp. 327-334.
7. *Multimedia System Services Version 1.0*, Interactive Multimedia Association (Hewlett-Packard, IBM, SunSoft), June 1993.
8. H. Schulzrinne et al., *Internet Draft draft-ietf-avt-rtp-06.txt*, work in progress, Nov. 1994.
9. L. Liang, S.T. Chanson, and G.W. Neufeld, "Process Groups and Group Communication: Classifications and Requirements," *Computer*, Vol. 23, No. 2, Feb. 1990, pp. 56-67.

**Robert Simon** is a PhD candidate in computer science at the University of Pittsburgh. His research interests include distributed multimedia systems, high-speed networks, real-time systems, and computer-supported cooperative work. He received a BA degree in history and political science from the University of Rochester, and an MS degree in computer science from New York University. He is a member of the ACM and the IEEE Computer Society.

**Donald Krieger** is an assistant professor of neurological surgery at the University of Pittsburgh. His research interests include functional brain mapping, management and multivariate statistical signal processing of very large data sets, nonlinear systems analysis, and software quality engineering. He received an SB degree from the Massachusetts Institute of Technology in science and humanities, a PhD in Neuroscience from Maharishi International University in Fairfield, Iowa, and a fellowship in clinical neurophysiology at Children's Hospital of Pittsburgh. He is a member of the IEEE Computer Society and a founding fellow of the American Society of Neural Monitoring.

**Taieb Znati** is an associate professor in the Department of Computer Science, with a joint appointment in the Telecommunications Graduate Program, at the University of Pittsburgh. His research interests include distributed systems, real-time communication networks to support multimedia environments, and local area networks, with a particular emphasis on design and analysis of communication protocols for distributed real-time systems. He obtained a PhD in computer science from Michigan State University and an MS degree in computer science from Purdue University. He is a member of the IEEE Computer Society and the ACM.

**Raymond Lofink** is the systems manager for the Center for Clinical Neurophysiology in the Department of Neurosurgery at the University of Pittsburgh's School of Medicine. His primary research focus is on designing maintainable software for continuously on-line production networks. His interests include configuration management and distributed systems management.

**Robert J. Scabassi** is a professor of neurological surgery, electrical engineering, mechanical engineering, psychiatry, and neuroscience at the University of Pittsburgh. He received a BSE degree from Loyola University, Los Angeles, an MSEE degree, an engineer in electrical engineering degree, and a PhD in electrical engineering, all from the University of Southern California, and an MD from the University of Pittsburgh. He was also a postdoctoral fellow at the Brain Research Institute at the University of California, Los Angeles. He is a registered professional engineer, a senior member of the IEEE, a fellow of the American Institute of Medical and Biomedical Engineering, and a member of the IEEE Computer Society.

Readers can contact the authors at the Department of Neurological Surgery, University of Pittsburgh, Pittsburgh, PA 15260; e-mail {simon, znati}@cs.pitt.edu or {krieger, lofink, bobs}@neuronet.pitt.edu.