NeuroNet: Collaborative Intraoperative Guidance and Control

Robert J. Sclabassi, Donald Krieger, Robert Simon, Ray Lofink, and Greg Gross
University of Pittsburgh

Donald M. DeLauder
Computational Diagnostics, Inc.

It has long been appreciated that a patient's physiological status is dynamic and subject to rapid, life-threatening changes during surgery. Anesthesiologists routinely use extensive physiological monitoring to maintain the patient's homeostasis. This monitoring can be thought of as establishing a control loop between the anesthesiologist and patient for the purposes of life support. Some information from these monitoring procedures reflects stress on the central nervous system, for example, changes in heart rate related to both brainstem and vagal stimulation. However, evaluation of CNS functional status either by clinical means or by the physiological monitoring tools commonly available to anesthesiologists is relatively limited. Thus, considerable effort has gone into developing intraoperative neurophysiological monitoring to add another dimension to the assessment of patient status during surgery.

Intraoperative neurophysiological monitoring can also be thought of as establishing a real-time control loop—this time between the surgeon and patient (Figure 1). The primary goals are to reduce morbidity and to dynamically assess structure-function relationships of the patient's nervous system during surgical manipulation. This dynamic assessment can guide a surgeon by providing specific, sensitive measurements that reflect operative maneuvers and their impact on the patient's CNS functioning. These goals require real-time measurements of CNS functions that can be closely correlated to operative manipulations within a time frame valuable to the progress of the operation. Thus, multiple types of neurophysiological data must be acquired, processed, and displayed in real time. In many instances, the proper interpretation of these data requires immediate consultation between the surgeon and a remotely located neurophysiologist.

NeuroNet is a system designed to provide the tools for acquiring, processing, and displaying multiple types of neurophysiological data in real time and for facilitating communication, collaboration, and information sharing among members of a neurosurgical team (including neurosurgeons, neuroanesthesiologists, and neurophysiologists) in a real-time mode.[1-3] The system was developed to support intraoperative neurophysiological monitoring for a large health center and the surrounding medical community (which looks to the center for consultation), where the demand for this service far exceeded what the few trained individuals could provide.

NeuroNet has three primary objectives: (1) to acquire and process multimodal data; (2) to integrate these data into display formats suitable to various applications; and (3) to present the various data types in a way that allows medical personnel at distributed sites to consult meaningfully about the shared data.

Neurophysiological monitoring assesses CNS structure-function relationships during surgery. NeuroNet supports remote performance of this task through real-time multimodal data processing and multimedia network communication.

Intraoperative guidance and control

Neurophysiological monitoring warns the surgeon when damage is being inflicted on the patient's nervous system. Many apparently benign surgical manipulations can significantly affect the patient's neural responses and resultant clinical condition. For example, retraction of structures close to a neural pathway, noise and vibration from drilling, and heat diffusion from lasers all affect the underlying neural tissue and neurophysiolgical responses, and may significantly damage the neural tissue. Information obtained through intraoperative monitoring permits the surgeon to dynamically modify the operative approach and thereby minimize the injury.

The commonly accepted goal of intraoperative monitoring is to prevent morbidity. At a certain level, this is true. The more fundamental goal, however, is to give the surgical team information that lets them accomplish

0272-17-16/96/$5.00 © 1996 IEEE

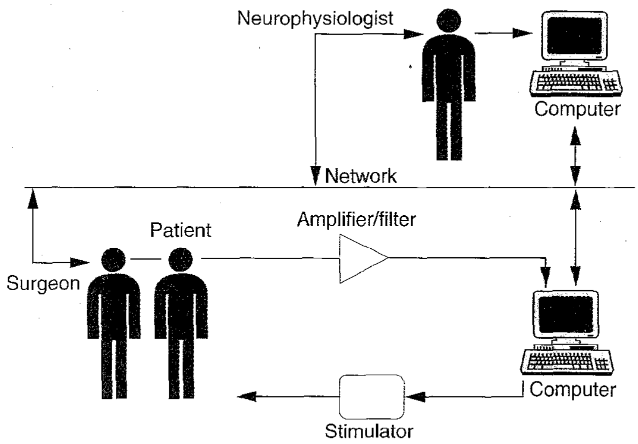

1 Neurophysiological monitoring establishes a control loop around the surgeon and the patient. The neurophysiologist must continuously interpret the monitoring data.

the operative objective with as optimal a surgical strategy as possible while having a clear idea of the morbidity induced along the way. The latter goal is particularly important in cases where the degree of difficulty is high and the preclusion of morbidity is virtually impossible. In such cases, the guidance obtained from intraoperative monitoring not only reduces morbidity but also gives the surgeon feedback on how operative procedures affect the CNS. The result is reduced costs to both the patient and society by decreasing the time spent in the intensive care unit, the hospital, and rehabilitation.

Depending on the surgical procedure, electrical activity indicates the functioning of

- the brainstem (brainstem auditory-evoked potentials and brainstem somatosensory-evoked potentials),
- the cortex (electroencephalogram, somatosensory-evoked potentials, and visual-evoked potentials),
- the spinal cord (somatosensory- and motor-evoked potentials and electromyograms),
- the various cranial nerves (electromyograms), and
- peripheral nerves (compound-action potentials and electromyograms).

These signals tend to be small (on the order of a microvolt) and to occur rapidly after the application of a stimulus to the neural tissue (on the order of a few milliseconds). Signal averaging may be required to extract them from the background noise. The measures used must be specific to the neural tissue being manipulated and sensitive to changes produced by the surgical manipulations. Many of these measures must be obtained, displayed, and interpreted simultaneously to permit multidimensional assessment of the integrity of the neural structures at risk. In addition, many of the measures provide information not only about function itself, but also about variables that directly or indirectly affect function, such as blood flow, hypoxia, and hypotension.

Signal measurements are obtained continuously during operations lasting as long as 20 hours. All data require continuous, rapid interpretation by a trained observer. In addition, multiple operative procedures may occur simultaneously, each requiring the same level of monitoring. In our institution, as many as 15 operative procedures have required simultaneous monitoring. The monitoring task can be considered a surveillance task, in that all the data must be available and continuously examined, although significant events occur infrequently. When such events do occur, however, they require immediate response.

Thus, the goal of intraoperative monitoring is to provide specific, reliable, sensitive, and timely information to surgeons in a way that lets them modify their operative strategies in order to reduce CNS injury.

The NeuroNet system

We have been actively researching the development of distributed computer networks for the acquisition, integration, and assessment of neurophysiological data for several years.[1-3] This work has resulted in NeuroNet, a distributed computer system used extensively at the Center for Clinical Neurophysiology of the University of Pittsburgh Medical Center (UPMC) and a number of other teaching hospitals not associated with UPMC. The system is fully integrated, transparently combining the collection, processing, and presentation of real-time data sources, including all physiological monitoring functions, with non-real-time functions and extensive on-line database information.

Figure 2 presents a control-flow diagram of the system with its major components identified. The user station is any host that can support NeuroNet functions, ranging from workstations to personal computers. NeuroDisplay is a user interface screen and constitutes the front end to NeuroNet. It is based on X Windows and Motif. The real-time application software supports those tasks that must guarantee low-latency end-to-end performance, such as acquisition, processing, and display of neurophysiological data, while the non-real-time application software supports tasks requiring average throughput and maximum average delay, such as file transfers and text-based e-mail.

Workstations are mounted in instrumentation racks and configured with appropriate electronics to support various data acquisition tasks including electroencephalograms (EEGs), electromyograms (EMGs), and multimodality evoked potentials. Multiple racks can be used in parallel on the same case if the number of variables to be monitored exceeds the capacity of a single rack. The data acquired on these systems is transparently accessible, in real time, across the network for both review and analysis.

Below this layer are the system support modules. They include an operating-system layer, a database/file system level, a communication support layer, and a network control layer. The NeuroNet Communication Support (NNCP) layer encapsulates the communication control structure, while network control functions provide the lowest level access to the underlying network. The databases act as the integrating agent for all the data acquired and manipulated through the system.

Data acquisition and processing features

NeuroNet permits simultaneous data collection and on-screen viewing of multiple modalities, each having user-determined observation intervals and stimulus rates that can be independently displayed and processed in real time on any other system on the network.

All NeuroNet software is integrated—there is no concept of separate packages for each type of data collection. This provides maximum capability for collecting and analyzing combinations of different data types. The EEG capabilities include compressed spectral arrays on all available channels, digital EEG filtering, and real-time spectral computations on the incoming data with arbitrary-length spectral averages. The system provides

- real-time remote viewing of all acquired data;
- multiway communication across the network, either digital audio or text; and
- unified user interfaces for local and remote systems, thus requiring familiarity with only one user interface.

Instrumentation racks perform a number of functions, including stimulus control and generation, data acquisition, signal processing, and data display. Each workstation has a high-resolution (1,024 × 1,280 pixels) 24-bit color monitor. Data are acquired through a custom-designed unit that includes a 12-bit analog-to-digital converter with a 16-channel multiplexer. This unit can be expanded to 64 channels and is used for acquiring physiological data from the anesthesiology monitors simultaneously with neurophysiological data. All data manipulations are handled by calls to the Neuro Data File (NDF) library (see the section below, "Data structures").

NeuroNet has an extensive package for collecting evoked-potential data and presenting it to the user. All modalities can be collected individually or mixed simultaneously. Data trending over time is flexible, allowing each channel of each modality to be independently displayed and controlled. Routines available for the analysis of all signals include digital filtering, standard averaging, odd/even averaging, noise estimation, and peak marking (both time and amplitude).

Users can enter comments at any time during data collection, and they can store and retrieve predefined comments through a pop-up window to annotate an open record quickly. The system can display baseline data for any waveform (both in the real-time displays and in trended displays). Users can retrieve the baselines from any channel of any data file and thus include baselines from preoperative studies.

Users also fully control artifact rejection. They can define two time windows per channel for artifact rejection. Furthermore, they can set the amplitude-rejection criterion, along with a "spike allowance" parameter that permits a percentage of the data to exceed the artifact-rejection limits without throwing away the trial.

The data-acquisition nodes can be operated independently of the network. Whether networked or standing alone, the acquisition nodes give users

- complete control over every acquisition parameter;
- multiple stimulus patterns including uniform, jittered, or bursting trains;
- total flexibility in setting artifact rejection;
- digital filters with unlimited combinations (low-pass, high-pass, median, and user-defined);
- baseline data from either current data or any other data previously gathered from the patient;
- feature marking;
- simultaneous spectral analysis;
- filtered and spectral waveforms displayed with or without underlying data;
- waterfall displays for tracking sequential operative data;
- compressed spectral array displays;
- filtered data on separate traces or overlaid;
- odd/even averaging;
- noise estimation; and
- scrolled EEG.

Network structure

A fundamental and unique feature of NeuroNet is its ability to support multiple mobile instrumentation carts with remote viewing capabilities. All data acquired at any one of the instrumentation carts can be viewed at any other cart connected to the network. Thus, one neurophysiologist can monitor several procedures at the same time. The system has built-in remote monitoring to provide shared data as well as instant typed or audio

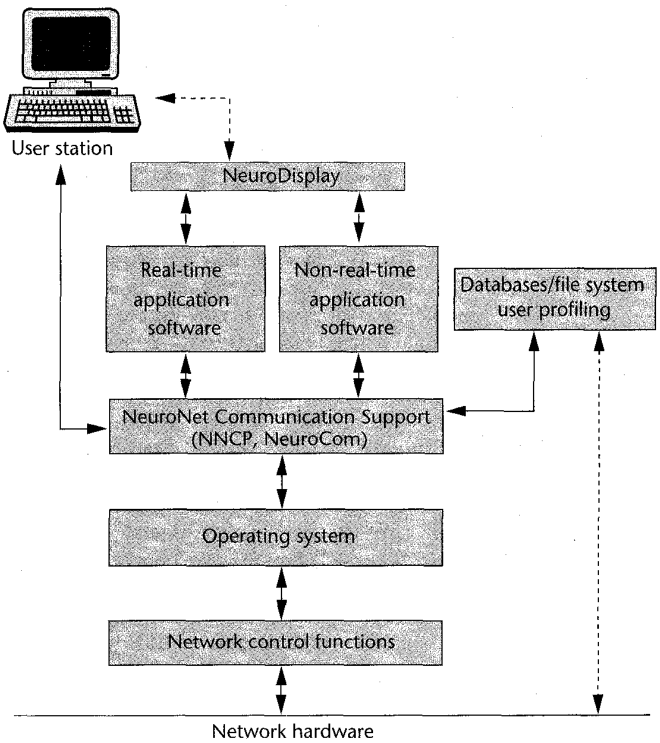

2 NeuroNet control-flow overview shows the layers of the system's organization.

communications between users. In addition, when on the network, NeuroNet can automatically back up the data collected by the instrument carts. NeuroNet can be integrated with existing heterogeneous computer systems that support TCP/IP and X Windows, including all workstations and IBM-compatible or Macintosh personal computers. The "network" can be either Ethernet (thick, thin, or twisted-pair), Token Ring, ISDN, FDDI, ATM, or combinations of these. Therefore, virtually any site can be connected in some form to another site.

The current UPMC network uses an Ethernet backbone with connections to 44 operating rooms, two interventional neuroradiology rooms, 20 neuro-intensive care beds, over 180 other intensive-care-unit beds, and 10 diagnostic laboratories, as well as a number of research laboratories. This network extends into all the neurosurgical and neurophysiology faculty offices at the health center, as well as several conference and lecture rooms, allowing these facilities to support not only more effective clinical care but also research, teaching, and more general consultation tasks. Data can be displayed and analyzed anywhere within the network and can also be observed off-site via modem connections to personal computers if a high-speed network is not available.

Communication protocol

NeuroNet Communication Support (NNCP) encapsulates the communication control structure, providing global naming and location information. NNCP makes extensive use of a distributed software package called Parallel Virtual Machine. PVM, developed at Oak Ridge National Laboratory and Emory University, enables a collection of heterogeneous computers to be used as a coherent and flexible concurrent computational resource.[4] The individual computers can be shared- or local-memory multiprocessors, vector supercomputers, specialized graphics engines, or workstations. They can be interconnected by a variety of networks, such as Ethernet and FDDI. User programs written in C, C++, or Fortran access PVM through library routines. PVM provides global naming services, dynamic process groups, message passing, multicasting, and global synchronization functions.

NNCP is organized in two layers supported by distributed daemons: a PVM daemon (PVMd3) and an information services daemon (ISD). An instance of a PVMd3 runs on each NeuroNet machine. In addition, there is a single instance of the PVM group server daemon on the network.

An ISD also runs on each NeuroNet machine. This daemon uses the message transport services provided by PVM to receive and service requests for data or for lists of active cases. Thus, the ISD is the server process in a client-server architecture, where the applications are the clients. Each ISD maintains a list of both historical and active cases on the machine where it is running. Active cases are defined as those that have collected and saved data within the past hour. The ISD is a polling process in the literal sense. It polls the directory where case data are preserved and examines the time and date each file was last modified. It also polls the message queue maintained by PVM for requests from other processes, and it services these requests whenever they are detected.

Multiple applications can run on each node. The applications use case listings and case data, fetched by the ISDs and transported by the PVMd3's, to generate data displays. NNCP uses the dynamic group services provided by PVM to identify server processes (ISDs) and the PVM message transport services to send requests and receive responses.

Data structures

Data acquired by NeuroNet are stored on and retrieved from disk according to a format we developed called the Neuro Data Structure. The NDS is designed to support the concept of a *case abstraction*, that is, a logical grouping of all data pertaining to a single patient. A case abstraction is the fundamental object used by NeuroNet to access information, and the NDS supports it by allowing the logical grouping of heterogeneous data types. Different data streams are identified and managed by a "channel manager" structure. In NeuroNet, data types are defined for classes of neurophysiological, physiological, and anesthesiological data. To support these different data types, the channel manager contains all pertinent information for each type in its header portion, and handles variable-length records.

NDS file input and output is provided through an applications programming interface called NeuroData Format. NDF allows a calling application to access all NDS data objects through the same function call, regardless of NDS type. In addition, the calling application does not need to know where the data reside. 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.

NDF calls are designed to access parameters in a file one element at a time. NDF works internally by using a functional programming style. This allows the simple introduction of new parameters and new data types in a structured fashion, without affecting previous programs using NDF calls. The NDF data access library is currently supported on HP 9000s (HP-UX), PCs with DOS/Windows or Unix, and Macintoshes (PowerPC).

User interfaces

NeuroNet user interfaces are based on X Windows and Motif and allow for the manipulation and presentation of all data types organized in the system. NeuroDisplay is an oscilloscope-style display, which allows current physiological data to be reviewed on a local machine. All the graphics programming for NeuroDisplay has been coded using Xt Intrinsics, XLib, and Motif convenience functions. NeuroDisplay consists of a display area for display waveforms and a menu (Figure 3). The options available from the menu let the user control how the display area presents data.

NeuroView is used to view data being acquired across the network. As with NeuroDisplay, all the graphics programming was coded using Xt Intrinsics, XLib, and Motif convenience functions. NeuroView relies on NNCP to support network transport and on NDF to support data access.

NeuroView has five distinct components: Application Shell, Spawner, Help File Reader, Application Thread, and Display Shell. All these components, except the Help File Reader, rely on a context structure, which is modeled in NeuroNet as a process. The structure contains all the information needed to characterize the state of a display shell.

The Application Shell consists of the user interface, callback functions to set elements of the context structure, and interfaces to the Spawner and the Help File Reader. The Display Shells are the NeuroView components that the user interacts with most frequently. A shell is composed of a drawing area, widgets that let users select a style of information display, and all the functions actually required to perform the display.

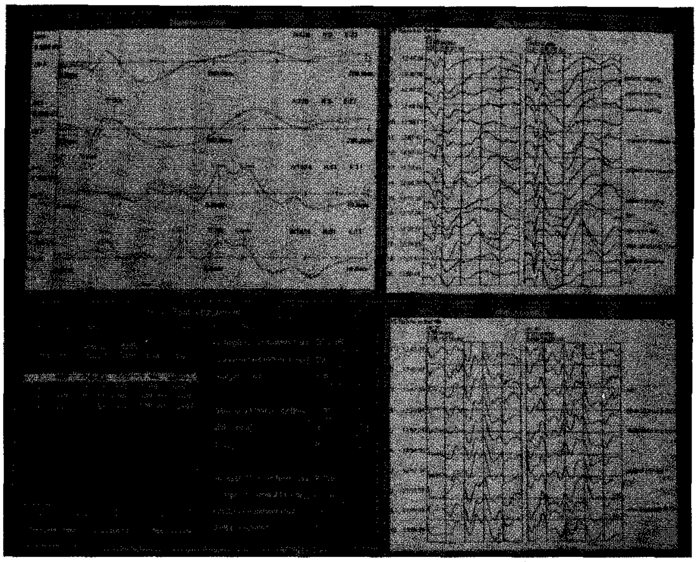

3 NeuroDisplay window (upper left-hand quadrant) shows two responses being simultaneously obtained and viewed. NeuroView windows (right half) show a waterfall display of the two data types being acquired.

The Spawner is an entry point where the ISD can notify NeuroView that a remote process is active. The Spawner is also responsible for starting Display Shells on user requests. The Application Thread is a routine that waits for data from an acquisition process or a file and then updates the associated Display Shell as necessary. The Help Reader gives the user on-line help about NeuroView. It consists of an index of topics and a scrollable area for displaying the help file.

The context data structure contains all the information necessary to update the Display Shell with information. This information is stored in a "stack" of structures, one structure per display shell. A shell's context is similar to the context of a process. When one shell has an operation affecting it, the other shells store their states and do not change. By accessing various elements on the stack, NeuroView uses the same routines to quickly and efficiently update up to 16 displays, each with different parameters. NeuroView data can also be accessed via phone dial-up for remote viewing on PCs. For example, all members of UPMC's Center for Clinical Neurophysiology can access any activity on NeuroNet from home, allowing them to consult on cases late at night.

NeuroNet database

The NeuroNet DataBase (NNDB) is integrated with the real-time acquisition and processing software and is accessible in a distributed fashion across the network. The two main purposes of NNDB are to provide an organizing structure for the data based on the concept of a case abstraction and to support the creation of a specific case abstraction when a subject is entered into the system.

NNDB is the first layer for user interaction with any component of NeuroNet. The database front end, based on X Windows and Motif, is designed to isolate both the user and the applications interface from the underlying databases.

NNDB is organized around a system of case abstraction parameters and pointers, where the parameters define the attributes of an individual case and the pointers identify where the data for that case actually reside. Each work step in the following sequence is mediated by the database:

- patient appointment for diagnostic and/or intraoperative procedures;
- scheduling staff and equipment resources to perform the procedures;
- specification of preparatory instructions by the staff neurophysiologist, which include protocols for controlling the computerized data acquisition and locations of electrodes on the patient for stimulation and recording;
- execution of the prepared procedures in the diagnostic laboratory and/or operating room;
- initial interpretation of the neurophysiological data by a junior staff member or training fellow in preparation for generating a written report;
- final interpretation of the data by a senior staff member and report generation;
- patient billing;
- data review at any time following the study; and
- review and analysis of classes of studies for research purposes.

Each of these functions is performed at a different time and possibly from a different workstation on the computer network. In each case, the user accesses the appropriate computer capabilities through NNDB. Thus, the database interface is what the user sees and interacts with to perform the function. Furthermore, the database itself is a common pathway through which all commands must pass, enabling thorough tracking of all work done and data collected by NeuroNet.

Discussion

NeuroNet is in daily use as a production system at UPMC. Its computationally powerful features for multi-

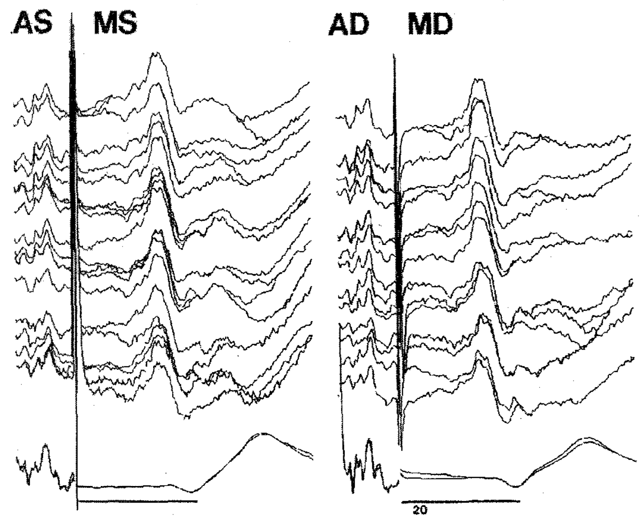

4 Simultaneous acquisition of brainstem and cortical data during removal of a tumor at the base of the skull. The data labeled AS and AD are auditory brainstem-evoked potentials; the data labeled MS and MD are somatosensory-evoked potentials.

dimensional data extraction and display together with its distributed system design have provided a natural testbed for real-time consultation on complex multimodal data. Intraoperative neurophysiological monitoring imposes far more stringent time constraints on system performance than other collaborative applications do—it does no good to inform the surgeon 10 minutes after a significant CNS event occurs. This application has therefore inspired the development of methods for extracting and analyzing data rapidly and efficiently.

Successful neurophysical monitoring requires the simultaneous acquisition of as many appropriate neurophysiological variables as possible. The approach described here offers immediate assessment of operative effects on the nervous system with implications for adapting the surgical approach based on the feedback of functional information. However, it also requires a highly trained neurophysiologist to rapidly interpret complex data recorded in less than optimal conditions (for example, Figure 4 presents an example of data obtained sequentially during an operation that used the neurophysiological monitoring application). NeuroNet enhances the correct interpretation of neurophysiological measures by enabling the display of all acquired data in a way that facilitates comparison. These data include the baselines acquired both at the beginning of the case and from preoperative studies.

Additionally, in our institution, as in many others, the demand exists for monitoring many cases simultaneously. This means requiring a highly trained individual to be present in each operating room all the time or providing access to the data remotely along with multidirectional communication facilities. NeuroNet allows a single trained individual to consult on multiple cases simultaneously.

This research and development effort is significant in many dimensions. First, distributed networked computer systems hold great promise for decreasing the cost of health care. The NeuroNet implementation at the University of Pittsburgh has significantly reduced the costs of consultants participating in health-care delivery. It has also reduced the length of time patients are hospitalized by reducing the morbidity associated with major surgery. In addition, this technology has enhanced the productivity of people using it.

Second, several important, unresolved engineering issues emerge naturally from our formulation of this system. These issues include

- the properties of acquisition, manipulation, and communication of multiple data types including audio, video, and images as well as physiological data;
- the properties of high-speed network testbeds required to support the collaborative use of these extensive data types;
- alternatives for supporting real-time data acquisition and distribution;
- the development of user interfaces that support the user's natural level of understanding;
- the integration of heterogeneous computers sized for different tasks from data acquisition to image processing to data review;
- the organization and implementation of large-scale databases for multiple types; and
- the strategies for presenting all these data in a coherent and intelligible fashion.

Solutions to these issues address fundamental problems in the development of real-time distributed applications that facilitate consultation on complex multimodal data.

We investigated the impact of these multimedia features on the delivery of a neurophysiology consultation service with a team of cultural anthropologists.[5] The results described above encouraged us to see a real role for additional data types in a time-constrained collaborative system focused on real-time decision making. Over the past several years, we have expanded NeuroNet to support the integration of neurophysiological, audio, and video data into a unified consultative system.[6] This work has included the development and implementation of a distributed communication architecture for multimedia systems.[7] The system's multimedia design results from our assessment of how it would be used.[8] The initial implementation, called MedNet, was based on a parallel cable television system providing 450 MHz of bandwidth to support the integration of full audio and video.

The primary need for computer platforms is in the collection, storage, processing, display, and retrieval of data and for collaborative activities. Thus, we configured NeuroNet to run on heterogeneous platforms. For full functionality, each node will either have to run control software on its own or become an X station for control of the display. We plan to explore the use of supercomputing resources as compute servers to the system with particular emphasis on near-real-time image rotations and database manipulations on large (terabyte), complex databases (containing multimodality data). ■

Acknowledgments

We acknowledge the support of the Benjamin Franklin Fund for the Commonwealth of Pennsylvania and Hewlett-Packard, Inc., for their generous support of this project. We wish to thank our colleagues for many valuable discussions, particularly Matthew Henry, Bill Madden, Scott Cline, and Earl Doyle.

References

1. R.J. Sclabassi, R. Lofink, and E. Doyle, "NeuroNet: A Distributed Microprocessor Network for Clinical Neurophysiology," in *Microcomputers in Medicine*, M.J. Geisow and A.N. Barret, eds., Elsevier Science Publishers, New York, 1987.
2. D.N. Krieger et al., "NeuroNet: Implementation of an Integrated Clinical Neurophysiology System," *Medical Instrumentation,* Vol. 21, No. 6, 1988, pp. 296-303.
3. D.N. Krieger, G. Burk, and R.J. Sclabassi, "NeuroNet: A Distributed Real-Time System for Monitoring Neurophysiological Function in the Medical Environment," *Computer,* Vol. 24, No. 3, Mar. 1991, pp. 45-55.
4. A. Geist et al., *PVM 3 User's Guide and Reference Manual,* available from netlib@ornl.gov.
5. R. Simon, T. Znati, and R.J. Sclabassi, "DIPCS: An Interprocess Communication Architecture for Distributed Multimedia Systems," *Multimedia Tools and Applications*, Vol. 1, No. 3, 1995, pp. 263-293.
6. R.J. Sclabassi et al., "The MultiMedia Medical Monitoring, Diagnosis, and Consultation Project," *Proc. Hawaii Int'l Conf. System Sciences, Vol. III,* IEEE Computer Society Press, Los Alamitos, Calif., 1991, pp. 717-723.
7. B. Nardi et al., "Turning Away from Talking Heads: The Use of Video-as-Data in Neurosurgery," *Proc. InterCHI* (ACM), Addison-Wesley, Reading, N.H., 1993.
8. R. Simon et al., "MultiMedia MedNet: A Medical Diagnosis and Consultation System," *Computer*, Vol. 28, No. 5, May 1995, pp. 65-73.

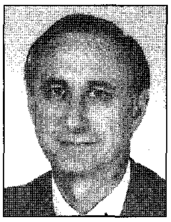

Robert J. Sclabassi *is a professor of neurological surgery, electrical engineering, mechanical engineering, psychiatry, and neuroscience at the University of Pittsburgh. He received a BSE from Loyola University, Los Angeles, an MSEE and PhD in electrical engineering from the University of Southern California, 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.*

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 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.*

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 in history and political science from the University of Rochester, and an MS in computer science from New York University. He is a member of the ACM and the IEEE Computer Society.*

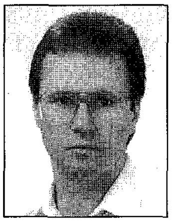

Ray Lofink *is the systems manager for the Center for Clinical Neurophysiology in the Department of Neurosurgery at the University of Pittsburgh's School of Medicine. His primary research focus is on designing maintainable software for continuously on-line production networks. His interests include configuration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

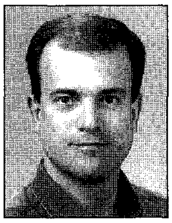

Greg Gross *is a systems analyst for the Center for Clinical Neurophysiology in the Department of Neurosurgery at the University of Pittsburgh, School of Medicine. His research interests are database and user-interface design for real-time and medical informatics systems. He received a BS in biomedical engineering from Case Western Reserve University.*

Donald M. DeLauder *is the director of engineering for Computational Diagnostics, Inc., in Pittsburgh, Pennsylvania, where he focuses primarily on the development of networked neurophysiology monitoring systems. He received a BSEE in 1979 and an MSEE in 1982 from West Virginia University.*

Readers may contact Sclabassi at the Department of Neurological Surgery, School of Medicine, University of Pittsburgh, PA 15213, e-mail bobs@neuronet.pitt.edu.