
Neuronet

A Distributed Real-Time System for Monitoring Neurophysiologic Function in the Medical Environment

Don Krieger, Gerald Burk, and Robert J. Sclabassi

University of Pittsburgh

This article focuses on a specific medical application at the Center for Clinical Neurophysiology at the University of Pittsburgh. The application is served by Neuronet, a large software package integrated with 80 computer workstations and instrumentation. It provides both off-line and real-time capabilities. The application addresses many problems that must be solved for computer-based systems that function in the medical environment. The application's real-time components can be life critical; the administrative support elements interact in a complex fashion with the life-critical ones.

The center provides diagnostic and intraoperative monitoring services throughout the university health center. These services assess patients' central and peripheral nervous system function using electrophysiological and vascular measurements. The numerous staff-level clinicians who participate in the service are academicians with individual research interests and styles in administering procedures. Thus, Neuronet has been configured to provide a host of readily accessible and continually expanding tools for running clinical neurophysiology studies.

System overview

The environment in a tertiary health care center such as the University of Pittsburgh

Neuronet provides immediate access to real-time life-critical data being acquired at multiple sites across the health center and allows one neurophysiologist to simultaneously monitor multiple surgical procedures.

medical center requires provision of service without delay or inconvenience to the referring physician at all times. Although it is expensive to provide enough equipment to cover peak times, the primary difficulty is in providing competent staff neurophysiologists to oversee multiple simultaneous procedures throughout the complex.

The university's health center encompasses seven hospitals. At any time, a dozen or more procedures may simultaneously

require a neurophysiologist's supervision. To enhance the observational ability of the staff neurophysiologist on call, a variety of remote capabilities were incorporated in Neuronet, including manual remote display, automatic remote display, and manual remote-message passing.

The sequence of events that occurs to initiate, provide, and support clinical neurophysiology services is as follows:

- (1) The physician orders a diagnostic study or intraoperative monitoring.
- (2) The study is scheduled with the patient.
- (3) Staff and equipment resources are committed and positioned at the designated time and place.
- (4) The patient is brought to a staffed diagnostic laboratory, or staff and equipment are brought to the operating room at the appropriate time.
- (5) The study is obtained.
- (6) A report is generated, and a note is entered in the chart.
- (7) The data are archived.
- (8) A bill is generated and dispatched.

The network of person-to-person transactions that enables this sequence is shown in Figure 1. This sequence of events is facilitated by conceptualizing the basis of the service as a local area network (LAN). The overall system is run on this network in a distributed environment. The file system

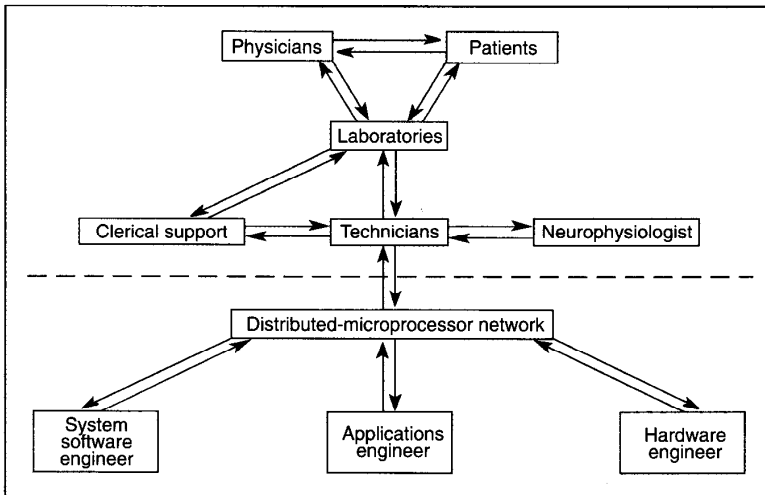


Figure 1. Summary of transactions in the neurophysiology group. The distributed-microprocessor network and its supporting personnel (below the dashed line) are transparent to the user community (above the dashed line).

is decentralized, the database is distributed, and servers for diskless nodes are placed in pairs to provide redundancy in case of hardware failure. The network has a 10-megabit-per-second Ethernet backbone (IEEE 802.3) with proprietary 12-megabit-per-second token-passing-ring subnetworks from Hewlett-Packard/Apollo. Figure 2 shows the currently operating network of computer systems in the health center.

Attached to the LAN are both fixed and portable computers, each associated with appropriate electronics for providing services for which that equipment rack will be used. This network of machines provides the following key capabilities:

- scheduling,
- applying stimuli with simultaneous data acquisition,
- computer processing for arithmetic data manipulation and display,
- archiving data,

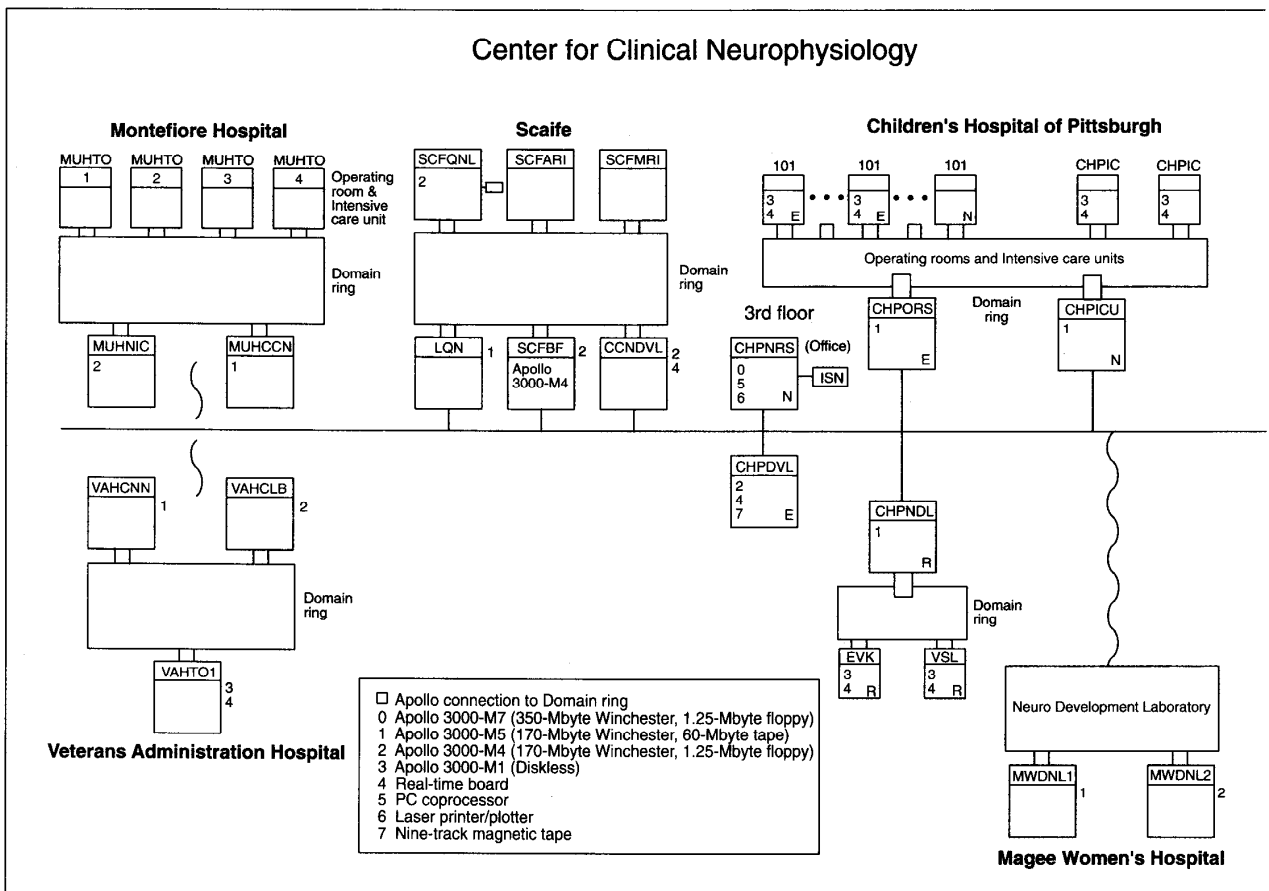


Figure 2. Current Neuronet configuration. The horizontal line spanning the figure represents the Ethernet backbone. Each token-passing ring is attached to the Ethernet through a gateway. The connections to Montefiore and the Oakland

- managing resources, and
- billing.

Neuronet applications span a wide spectrum of real-time performance requirements. In the operating room, evoked-potential acquisition, processing, and immediate display can be life critical. For other functions, brief delays are tolerable, for example, for database interactions while scheduling a patient over the telephone. Still other functions can be performed off line, for example, archiving data, managing resources, billing, and generating reports. In addition to time considerations in designing the system, reliability in the face of regular software and hardware upgrades is a key system requirement.

Hardware upgradeability was one of the primary requirements that led to the use of a LAN. Software upgradeability is maintained by conforming to the spiral life-cycle model¹ for software development with

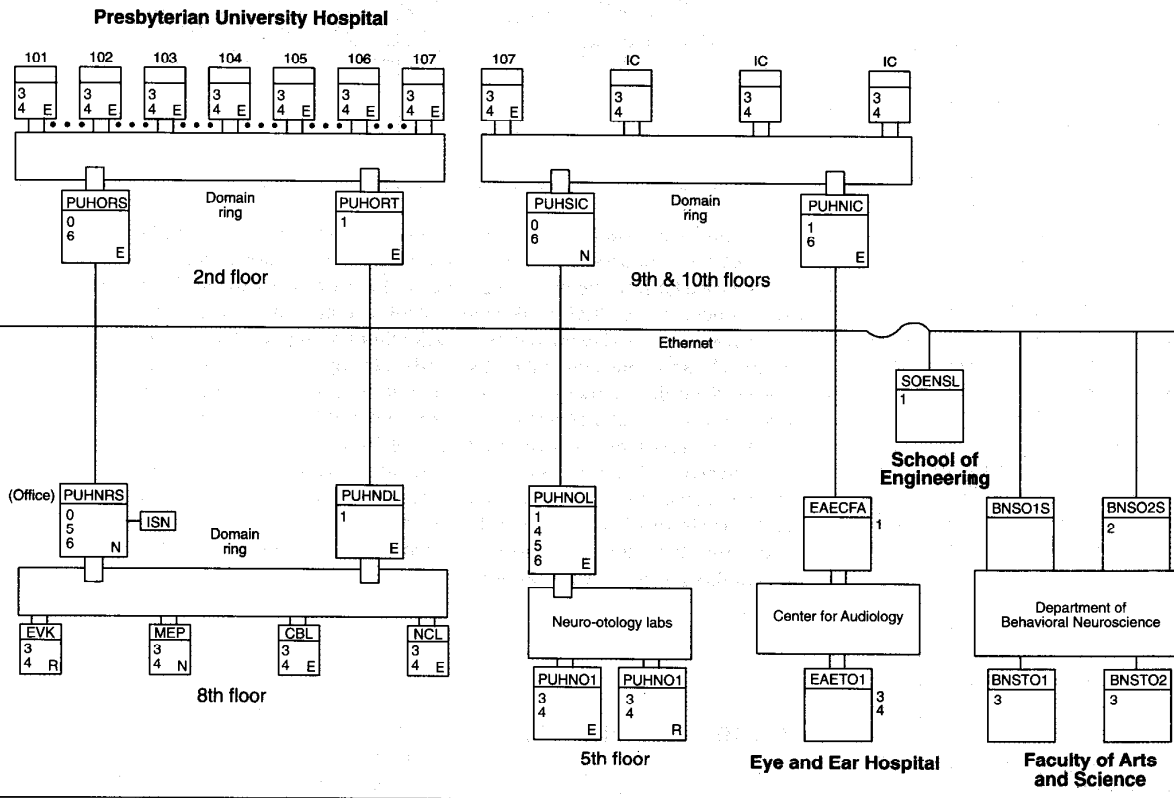
a sequence of more and more extensively exercised installations provided to the user.

For scheduling, managing resources, generating reports, billing, and eventual direct interfacing to the main hospital information system, a Codasyl (Conference on Data Systems Languages) database was chosen and integrated with the software package. The real-time applications include a wide and constantly expanding set of signal processing capabilities. In addition to local processing, a sequence of innovations has been introduced that enables remote real-time data display. Data can be acquired at one location on the network and displayed simultaneously in many other locations, both across the network and over voice-quality telephone lines. A significant portion of the program development effort has focused on the development of an effective and well-organized human interface that is used consistently throughout the software package.

The practice of clinical neurophysiology includes a wide range of diagnostic and monitoring techniques.^{2,3} Diagnostic techniques include neuroelectric measurements from muscle, the scalp, vascular measurements, and others. A neurophysiological measure provides information about nervous system function, distinguishing it from anatomical techniques that provide information about structure — for example, computed tomography, also known as CT.

The intraoperative neurophysiological techniques are primarily based on electrical measures, although the direct dependence of nervous system function on blood supply calls for blood-flow measurements in some cases. The electrical measures of interest have relatively short time constants, on the order of 1-10 milliseconds. However, the electrical environment in the operating room is noisy, and there is background electrical activity produced by the

Center for System Neuroscience



Veterans Administration hospitals are via voice-quality phone lines. The backbone is attached to the Internet through the University of Pittsburgh's gateway.

central nervous system. Therefore, most neuroelectric measures taken in the operating room are collected over many seconds and arithmetically processed to produce interpretable information. This processing includes averaging, digital filtering, and estimation and background noise removal for single trials.⁴

Case example

The aorta, which supplies blood to most of the body, including the brain and spinal cord, is abnormally narrowed in some children. Repair usually requires cross clamping — that is, clamping the vessel on either side of the narrowed area, removing the narrowed area, and rejoining the severed vessel. In a small percentage of coarctation repairs, about 1 in 300, the spinal cord does not tolerate the loss of blood supply that occurs with an aortic cross clamp. To prevent spinal cord injury, cord function can be monitored by repeatedly applying electrical stimuli to sensory nerves in the lower limbs and recording the response of the central nervous system at the scalp. If the neural pathways that carry this evoked response — including the spinal cord — are compromised, a change in the response will be detected.

The evoked response at the scalp is about 5 microvolts and appears on a background of normal brain activity of 20 to 100 microvolts. To extract this small signal from the background and electrical noise in the operating room, it is necessary to present many stimuli, typically 70 to 120, and process the response to each stimulus to produce an interpretable response. At a stimulus rate of 3 to 5 hertz, an interpretable evoked response is obtained every 20 to 35 seconds, providing relatively rapid feedback to the surgeon regarding spinal cord function.

Figure 3 shows sample results from a repeat coarctation repair in a 12-year-old. As indicated by the arrows in Figure 3a, responses from the beginning of the case until the moment of aortic cross-clamp application were consistent. As soon as the artery was clamped, the waveform lost its consistency (traces 86 to 89). The surgeon immediately removed the clamp, and the responses returned. To maintain blood flow, the surgeon then placed the vessels that supply the spinal cord on a pump that bypassed the aortic cross clamp. Figure 3b shows that the response remained stable following bypass placement, even with the aortic cross clamp.





Figure 3. Case example. Shown are two intraoperative monitoring displays of peroneal nerve responses to twin-pulse stimulation. These displays are available during the case both in the operating room and at any node attached to the network. Responses were obtained during critical moments in an aortic coarctation repair procedure in a 12-year-old girl. Each stimulus consisted of a shock to the left peroneal nerve followed 100 milliseconds later by a shock to the right peroneal nerve. Stimulus artifacts are prominent at the 0 and 100 milliseconds marks in each trace. A total of 128 responses were averaged for each trace. The recording electrodes were located at International 10/20 System placements Pz (active) and Fz (reference). Positive voltage is represented upward. Time and voltage scaling are shown at the bottom. The number of each trace and the clock time at which it was obtained are shown at the left. Note that traces proceed in temporal sequence from the top of the figure to the bottom. Comments entered by the technician during the case are shown at the right. Two superimposed baseline

Remote capabilities

A coarctation repair typically lasts three to five hours, with responses obtained about every 45 seconds. Since there is significant

spinal cord compromise due to aortic cross clamping in only 1 out of 300 cases, the intraoperative monitoring, although time critical, is a vigilance task. There are other cases in which the contribution is more immediate and occurs with much higher



traces obtained early in the case are shown at the top with the cortical evoked response indicated by arrows.

Trace 86 (part a) was obtained during the first minute following aortic cross clamp. Note the disintegration of this waveform compared with those that preceded it. The clamp was removed during the course of trace 87. The trace showed some variability for several minutes due to dramatic blood-pressure changes following aortic unclamping.

Traces 149 through 159 (part b) were obtained following aortic cross clamp with a circulatory-assist bypass pump in operation. There was no significant change compared with preceding traces other than what could be accounted for by blood-pressure changes. The bypass pump was turned on at trace 148 as indicated by the increased noise level.

probability. For example, in reconstructive hip surgery, the incidence of sciatic-nerve palsy is approximately 1 out of 10 cases. In surgery for large acoustic neuromas, the likelihood of facial-nerve palsy and/or hearing loss is higher still. In those

cases, monitoring is more efficient in terms of the number of cases for which a contribution is made, although the seriousness of many of the potential complications is considerably less. The value of multiple observers in a vigilance task, in addition to

the personnel resource considerations, has stimulated the development of a set of remote display capabilities. Although the need for these capabilities has been recognized, there are few reports available and limited development of such capabilities to date.⁵⁻⁷

All remote capabilities are implemented through files that can be accessed from anywhere on the network. Although this approach is slightly slower than one in which messages are passed directly from processor to processor, it is more conservative of critical system resources: CPU time, network time, and memory. This can be seen by examining a remote display for evoked-potential data that is being acquired in the operating room. The computer that acquires and processes the data for local display is a diskless workstation sitting in an instrumentation rack. Diskless machines are used for several reasons:

- (1) The instrumentation racks are wheeled from room to room. Using diskless machines avoids the danger of mechanical damage to a Winchester disk.
- (2) They are less expensive.
- (3) Using disks in fixed machines assures that all disk space will be accessible at all times.
- (4) The machines that run in the operating room are relieved of the CPU and bus loading associated with disk services.

The instrumentation rack workstations perform the following functions:

- data acquisition and stimulus presentation;
- artifact rejection via several algorithms, including a primary clipping algorithm;
- signal averaging;
- digital filtering, when selected;
- plus/minus averaging, when selected (this is a noise estimation procedure);
- initiation of file storage of completed evoked responses (completion is determined when a set number of nonartifact-contaminated trials are averaged);
- initiation of ongoing disk storage of the partially completed evoked response to a special place in the file every three seconds;
- constant update and display of the ongoing partially completed evoked response; and
- constant update and display of the completed evoked responses in waterfall form, compressed-spectral form, or density-spectral form.

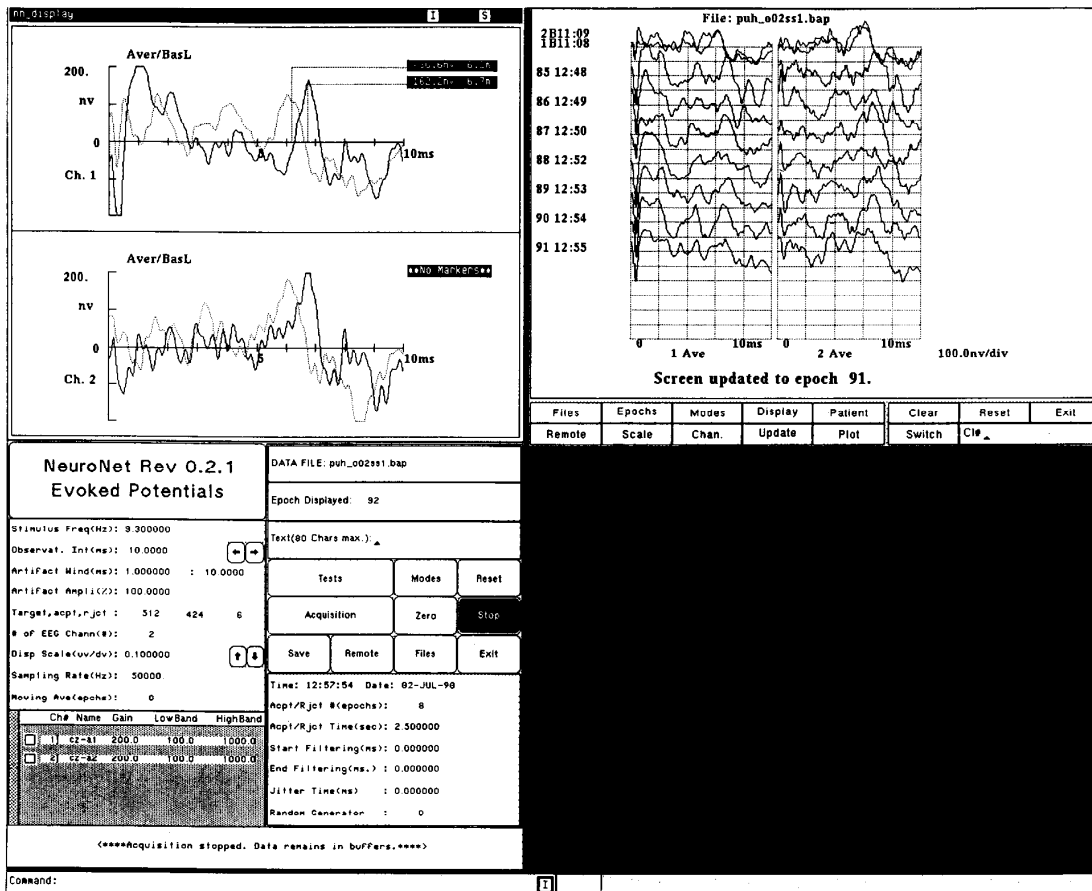


Figure 4. Typical operating room display during evoked-response monitoring. The upper left panel shows the current, partially completed, evoked response (solid line) superimposed on a baseline evoked response from the beginning of the case (dotted line). Note that the current response is delayed by 0.5 milliseconds compared with the baseline. The lower left panel shows all the parameters and controls for acquiring evoked responses. Note that the stimulus frequency is 9.3 hertz, the observation interval is 10 milliseconds, etc. These are typical parameters for acquiring brain-stem, auditory evoked responses as shown here. The upper right panel shows a waterfall display of evoked responses acquired over the past several minutes. This window is similar to those shown in Figure 3.

Figure 4 shows the display from an instrumentation rack workstation that acquires and displays evoked-response data. The server node — from which the diskless node is booted — performs all data-storage functions. If a remote node elsewhere on the network requests data from the file, the server handles the request, thus relieving the diskless node from further computational or other burdens. The file can be opened for read-only access by multiple nodes across the network, or even across the Internet using FTP (file-transfer protocol). The file is opened by a process running on a node elsewhere on the LAN from

which a user wishes to view the data acquired in the operating room; the data is read and displayed as the remote observer wishes. The remote observer can also access the ongoing, partially completed, evoked response.

Automatic remote viewing is accomplished by a shell script. Data from intraoperative monitoring and diagnostic studies are stored in a specific subdirectory on each server. The shell script searches this subdirectory continuously on all network servers. Any time it finds a file that has been updated in the last 10 minutes, it brings up a copy of the remote display

process with all display parameters automatically set to appropriately display the data. A window is created on the screen of the automatic remote monitoring node that continuously updates the display as new data is saved by the diskless acquisition node in the operating room or diagnostic laboratory. If a file goes more than 10 minutes out of date — that is, the case or diagnostic study is over — the shell script terminates the remote display process and clears the screen space for another display process. Thus, a dynamically changing picture of all the studies across the network is maintained automatically and continu-

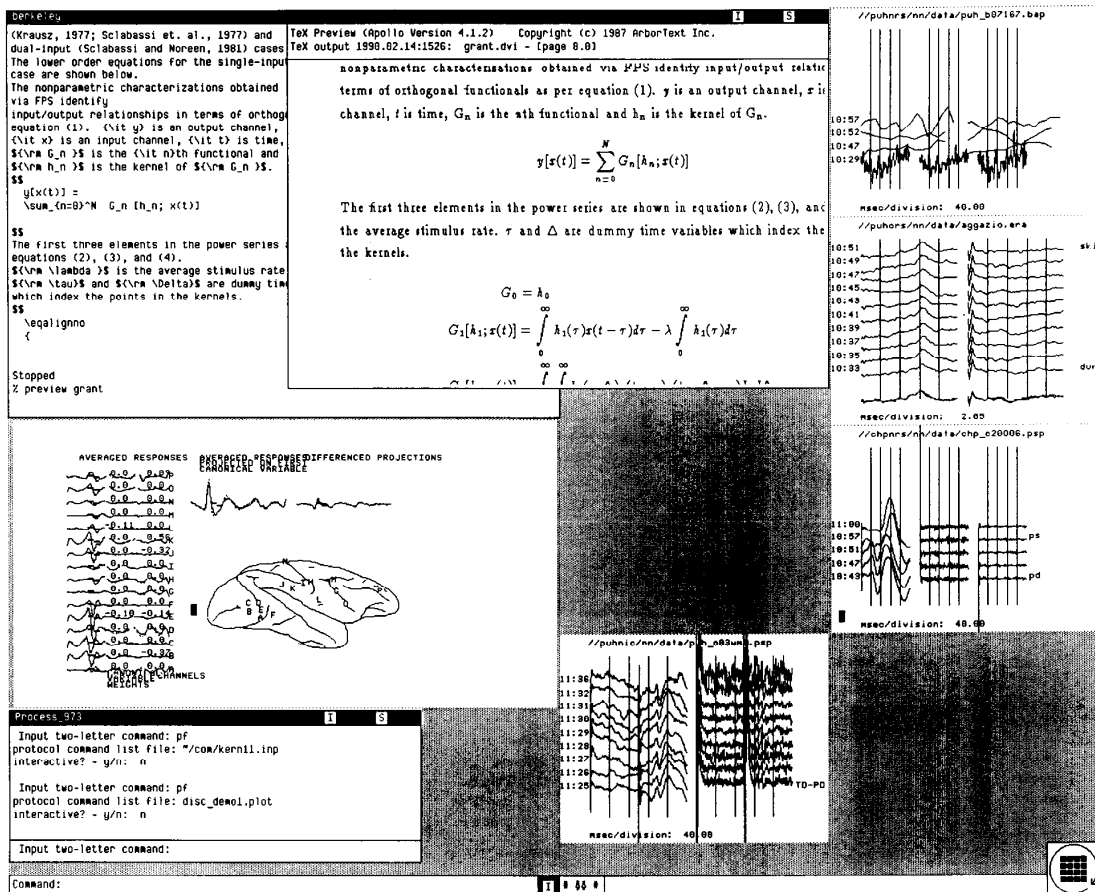


Figure 5. Typical remote monitoring display. This figure shows a display obtainable from anywhere on the network. The small box at the lower right corner contains a process that is automatically and continuously searching for active operating room or diagnostic studies. The four small windows containing waterfall displays were spawned by this process and display one ongoing diagnostic study (upper right corner) and three operating room cases. The two text windows at the top show a manuscript in preparation. The figure in the center was produced using the process shown in the lower left window.

ously by this automatic remote display function. From any node on the network, a single neurophysiologist can oversee all the clinical activity in the health center for which he is responsible, while at the same time maintaining other activities. A typical display from a remote monitoring node is shown in Figure 5.

Remote message passing is carried out through the remote display processes and can be initiated from anywhere on the network. Messages can be sent in both directions and appear along with a bell that sounds to alert the user. In this way, the staff neurophysiologist can communicate

quickly (one to two seconds) and at will with any operating room and with multiple technicians.

These remote capabilities constitute partial surveillance of events that occur in the operating room and diagnostic laboratory. To extend these functions, we plan to add a private phone system and a local cable TV network controlled through the same mouse-driven button. Extending the neurophysiological data transmission functions to the multimedia ones that are under development⁸ will significantly improve a single neurophysiologist's ability to effectively monitor multiple simultaneous procedures.

Signal processing

In each acquisition and display routine, there are multiple hooks provided for signal-processing capabilities. In most programs, the user can access these capabilities through mouse-driven push-button menus. New signal-processing schemes are added by adding subroutines to the code. The user invokes these schemes by clicking on new lines that appear in the menus. In display routines, whenever signal-processing functions are invoked, the raw signal is also displayed. This assists in recog-

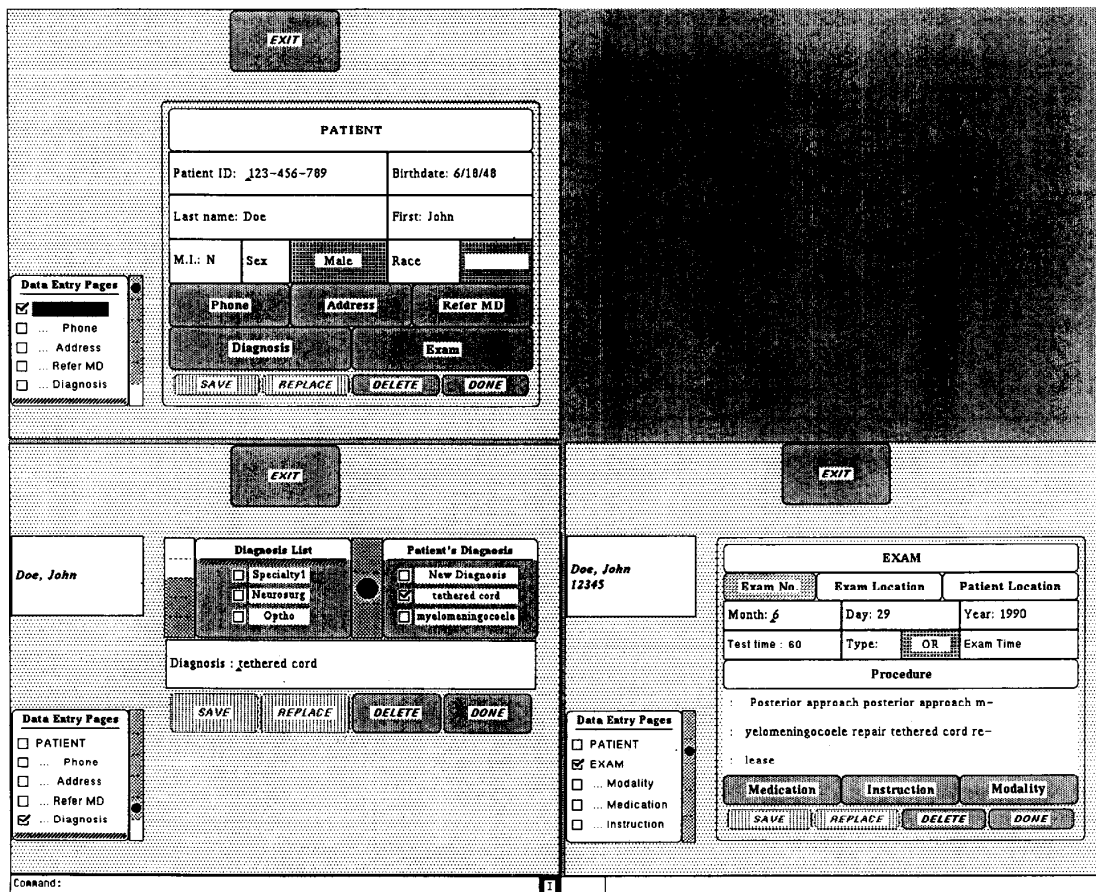


Figure 6. Typical demographics display. Three copies of the demographics database routine are shown to demonstrate program features; in normal use, only one copy is used by one user. The three windows represent three records in the top two layers of the data structure shown in Figure 7. The upper left window shows the menu of items used to enter or display a patient. The lower left window shows the menu of items used to enter or display the patient's diagnoses. The lower right window shows the menu of items used to enter or display an exam. Note that only one record can be displayed at a time, but that the parent and offspring records of the displayed exams are shown in the lower left panel. This panel can be used at any time to transit the data structure. The upper left panel contains one key item from each parent and grandparent record.

nition of any distortion that may be introduced by the signal-processing techniques. It also facilitates the recognition of the signal features for whose enhancement the signal-processing technique has been applied.

An arbitrary digital filter can be applied to any signal under study. Digital filters are stored in files that can be selected from a menu. A new filter can be defined by invoking the function for that purpose. Additional signal-processing capabilities that have been included in one or more of the acquisition or analysis routines include frequency analysis using Fourier trans-

form or discrete Wigner Ville time-frequency representation, notch filtering using regression, plus/minus averaging on alternative epochs to provide a measure of background noise, and others. A variety of other techniques are being studied and developed in prototype routines. In addition to signal processing, several different display modalities are provided, depending on the program. For example, for spontaneous brain activity, a user can select raw display, compressed-spectral array, density-spectral array, fast Fourier transform averaging, or time-frequency density spectra.

User interface

The user interface was designed to be easy to learn and use and to respond rapidly. In most cases, it is a mouse-driven interface. Parameters are set by clicking on arrows that raise or lower them. When a parameter is set to a nonstandard value, the user types in the desired value. All mouse buttons are used in a standard way across all programs. Clicking the left button enables a selection. Clicking the right button obtains a short help message for the item selected.

For prototype program development, much work was done with a terminal interface. This assisted in rapid conversion of programs that were developed on main-frame computers with terminal access only. More recently, prototyping with the mouse-driven interface has been used because the programming tool for this interface, Domain Dialogue, allows rapid development of prototype software. An additional advantage is that using a mouse-driven button interface from the beginning provides a first cut at the final user interface; thus, the final software version is one step closer to completion.

User-interface development was an important motivating factor in the use of an evolving prototype software model. This model has enabled numerous revisions of the interface in response to input from a wide variety of users around the health center.

User interface examples are shown in Figures 4 and 6. Figure 4 shows the interfaces for the evoked-response acquisition and waterfall-display programs. Figure 6 shows the interface for the D3M-based program to enter and access demographic information (see **Database**). Figure 4 shows that critical real-time functions, such as starting and stopping the acquisition and saving an epoch, are large and easily identified. Additional functions are presented in other buttons. For example, the user can select a baseline epoch from an arbitrary file anywhere on the network, a variety of digital-filtering techniques, numerous data-display options, etc.

Database

This package required the following functions of a database management system integrated with real-time acquisition and processing:

- storing demographic information,
- documenting and tracking real-time data, and
- storing secondary variables defined during data scoring as part of the effort to produce a report.

The secondary variables include peak heights and latencies from evoked potentials, a variety of blood-pressure and velocity measurements from trans-cranial Doppler, and English-language comments and interpretations that can be integrated as part of a report.

The database management system that

was chosen was bundled with the Hewlett-Packard/Apollo workstation software, D3M. D3M is a Codasyl database with a number of special features that make it attractive for development.

D3M is distributed; access to it is provided by server processes running on nodes all over the network. Many users can simultaneously access a particular database at any time. Any collisions — attempts to change the same record by multiple users at the same time — are adjudicated by the server processes. Furthermore, a database can be divided among many files existing on separate disks across the network. This provides considerable latitude in constructing redundant records for robust operation in case of disk failure.

The system software provides access to D3M via calls from a variety of programming languages, including Fortran, C, and

Pascal. In addition to this access, all calls that involve transactions to D3M contain English-language statements that describe what is to be done. Thus, each call that includes a transaction is documented in the call itself.

The typical Codasyl data structure shown in Figure 7 was developed to support these functions. This data structure provides tremendous flexibility for storing data in an organized fashion while providing rapid access through a hash table. It is integrated with other software components to enable multiple access at multiple levels for different functions.

Access to the upper layers is via the user interface shown in Figure 6 or, more commonly, through the real-time acquisition program and scoring programs. For example, the computer can sense the amplifier gain and filter settings during evoked re-

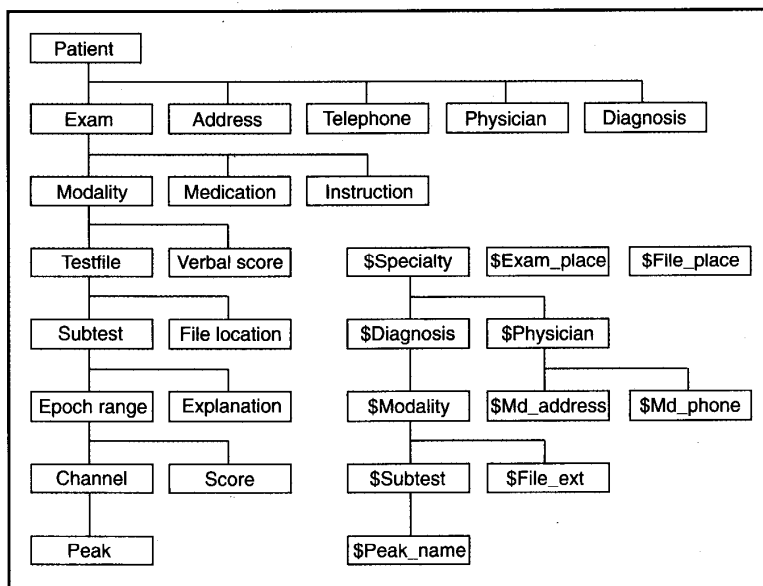


Figure 7. The Codasyl data structure used in Neuronet. Each element in the diagram is called a record and is composed of one or more items of information. Records are owned by records above them in the diagram. There can be multiple versions of a particular kind of record attached to a single record above. For example, each patient can have multiple exams, physicians, etc.; each exam can have multiple modalities, etc. The items in the records in the upper two layers of the structure on the left contain demographic information. These layers are primarily accessed through the user interface shown in Figure 6. The lower layers are for storing waveforms or other scores, amplifier settings, and other parameters used for specific tests. The records on the right whose names begin with a dollar sign contain both site-specific items (for example, names of physicians) and generally useful information (for example, common interpretation phrases). These items can be changed only by the system administrator.

sponse and electroencephalogram studies. Once a study is done, these values are set in the database for the permanent record. These parameters can be changed from one epoch (complete averaged evoked response) to the next, and the new values will also be stored. Conversely, previously stored parameters can be used to automatically set up a study. Parameters that can be controlled through the database include amplifier gain and bandpass, acquisition rate, stimulus rate, and others.

Reliability and upgradeability

The software package was originally developed on LSI-11 (Digital Equipment Corporation) computers attached to an Ethernet backbone.⁹ The software was then ported to a workstation environment. At the same time, the Ethernet backbone was extended with token-passing-ring subnetworks.¹⁰ Development continued in this multiprocess programming environment to the form described here. To carry development forward efficiently and dependably, while maintaining the functionality of the computer network on a day-to-day basis, the spiral life-cycle model has been used.¹

Each software unit is treated as an evolving prototype. New functions are developed as stand-alone routines or as subroutines that are used in a variety of ways. However, new routines are always backed up by the tried-and-tested routines that form the service's bread-and-butter functions. As new functions are developed and tested for their functionality and user interface, a professional applications programmer incorporates them in the standard package and further tests them in the development laboratory.

A sequence of four complete software installations is maintained at all times on the network. If a previously unrecognized catastrophic fault is present in a recent software revision, a user can rapidly retreat to a previous, better tested version. This retreat is accomplished by logging out and then logging in to the desired environment through a different account name. A user logging in to a particular environment is automatically and transparently pointed to the software revision corresponding to that account name through a symbolic link. This link is "soft" by virtue of a logical environmental variable included as part of the link name. This use of soft links enables rapid installation and change of the

software of a specific account by simply changing the environmental variable. Although this approach is cruder and slower than the *N*-version approach to fault tolerance,¹¹ it is functionally equivalent for this purpose.

The four installations with their functions are listed below:

(1) Betaneuro. Only one version of Betaneuro exists on the network. Betaneuro is the most recently revised software version and is used under close programmer and neurophysiologist supervision to test new functions for utility, dependability, etc.

(2) Newneuro. Once Betaneuro is tested by the neurophysiologists and deemed useful for the operating room without close technical supervision, Betaneuro is installed on every server in the network under Newneuro. The installation of this and all other versions of the software on the network provides maximal robustness in the face of potential network failure. If the Ethernet backbone goes down, or if one server goes down in an operating room ring, the software is still available on a local server to the diskless nodes running in the operating rooms.

(3) Neuro. When Newneuro is installed, the previous Newneuro is moved to Neuro. Thus, what is installed in Neuro is a version of the software that has been extensively tested — not only in the laboratory but also in the full network environment in the operating rooms and diagnostic laboratories under Newneuro. This environment, Neuro, is the one that is routinely used to support the clinical service.

(4) Oldneuro. Whenever Newneuro is moved to Neuro, Neuro is moved to Oldneuro, providing double redundancy in the software.

In general, this revision scheme is one in which modest changes are made for new revisions. Big changes are made in only one program at a time so that it is known which program must be exercised most thoroughly to validate the software revision. This development scheme has been very effective in allowing software revisions on a monthly basis. With this rapid revision schedule, the software development group can be very responsive to user needs and, at the same time, maintain high system reliability.

The present revision scheme has been in operation over the past three years, since moving to the workstation environment. It has been used continuously and without failure during monitoring of 4,000 sequen-

tial cases and 12,000 diagnostic studies. There have been several network-wide slowdowns associated with the reception of large electronic mail messages from China following the Tienamen Square incident. Although these slowdowns caused no system failures, as might be expected from a chain reaction,¹² they were dramatic enough to require temporary shutdown of the electronic mail capability.

This article describes an extensive computer-based instrumentation system used in a large tertiary health care center. Many of the problems whose solutions are presented in the context of this specialized system are generally encountered in computer-based medical systems. Among these are expandability, robust operation in the face of potential hardware failure, regular revisions without reduction in dependability and with zero down time, real-time response, and integration with an extensive data-base management system.

The solutions that are presented are generally applicable and demonstrably effective and economical. The system has been used in its present configuration for more than three years, not only in time-insensitive applications — for example, data archival and report generation — but also in life-critical applications in multiple simultaneous procedures. ■

Acknowledgments

This work was supported by the Children's Hospital of Pittsburgh, the Presbyterian-University Hospital of Pittsburgh, the Ben Franklin Partnership for Western Pennsylvania, the National Institutes of Health (Grant No. NS43-NS-7-2393), Western Psychiatric Institute and Clinic (Grant No. 1096), Computational Engineering, and the University of Pittsburgh.

References

1. M. Deutsch and R. Willis, *Software Quality Engineering: A Total Technical and Management Approach*, Prentice Hall, Englewood Cliffs, N.J., 1988.
2. K. Chiappa, *Evoked Potentials in Medicine*, Raven Press, New York, 1983.
3. R. Cracco and I. Bodis-Wollner, eds., *Evoked Potentials*, Alan R. Liss, Inc., New York, 1986.
4. D. Krieger and W. Larimore, "Automatic Enhancement of Single Evoked Potentials," *Electroencephalography and Clinical Neurophysiology*, Vol. 64, Dec. 1986, pp. 568-575.
5. G. Dumermuth and R. Dinkelman, "EEG Data Acquisition and Preprocessing by a Microcomputer Satellite System," *Computer Programs in Biomedicine*, Oct., 1979, pp. 197-208.
6. J. Steben, L. Streletz, and R. Fariello, "Multiprocessing Computer System for Sensory Evoked Potentials and EEG Spectral Analysis for Clinical Neurophysiology Laboratory," *J. Medical Systems*, Vol. 9, May/June 1985, pp. 347-363.
7. S. Nystrom, I. Sulg, and E. Hokkannen, "Multichannel Telephone Transmission and Quantitative On-Line Analysis of EEG in Brain Surgery," *Acta Neurochirurgica*, Vol. 41, 1978, pp. 255-268.
8. R. J. Sclabassi et al., "The Multimedia Medical Monitoring, Diagnosis, and Consultation Project," *Proc. 24th IEEE Conf. System Sciences*, CS Press, Los Alamitos, Calif., Order No. #ICSS-24 (Four-volume set), 1991.
9. 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.
10. D. Krieger et al., "Neuronet: Implementation of an Integrated Clinical Neurophysiology System," *Medical Instrumentation*, Vol. 21, June 1988, pp. 296-303.
11. A. Avizienis, "The N-Version Approach to Fault-Tolerant Software," *IEEE Trans. Software Eng.*, Vol. SE-11, No. 12, Dec. 1985.
12. U. Manber, "Chain Reactions in Networks," *Computer*, Vol. 23, No. 10, Oct. 1990, pp. 57-63.

Readers can contact Don Krieger at the Department of Neurological Surgery, 3470 Children's Hospital, 3705 Fifth Ave., Pittsburgh, PA 15213, e-mail don@po.neuronet.pitt.edu.




Don Krieger is an assistant professor of neurological surgery at the University of Pittsburgh and was a clinical fellow in neurophysiology at the Children's Hospital of Pittsburgh. His research interests include functional brain mapping using bioelectric and biomagnetic signals, management and multivariate statistical signal processing of very large data sets, nonlinear systems analysis, and software quality engineering.

He received the BS degree from the Massachusetts Institute of Technology and the PhD degree from Maharishi International University, Fairfield, Iowa. He is a member of the IEEE Computer Society.



Gerald Burk is a senior programmer on the Neuronet project at the Children's Hospital of Pittsburgh. His research interests are real-time applications in digital signal processing and control.

He received the BS degree in electrical engineering from the University of Pittsburgh, where he is currently a graduate student.



Robert J. Sclabassi is a professor of neurological surgery, electrical engineering, mechanical engineering, psychiatry, and behavioral neuroscience at the University of Pittsburgh.

He received the BSE degree from Loyola University of Los Angeles, and the MSEE and PhD degrees in electrical engineering from the University of Southern California. Sclabassi received the MD degree from the University of Pittsburgh and has been a postdoctoral fellow at the Brain Research Institute at UCLA. He is a registered professional engineer and a member of the IEEE Computer Society.

BOOLE

Version 1.4

The software tool of choice for development and delivery of approximate reasoning models.

Arbitrary boolean formulas and a complete set of coefficients of relation allow precise descriptions of uncertainty.

Solver based on the Binomial Polytopes algorithms for unmatched performance and reliability.

Requires hardware capable of running PC/MS-DOS.

Manual, software: \$385.00
Visa MC Check



URSIC Computing

5210 Trafalger Place
Madison, WI 53714
(608) 241-0651
Mon. - Fri. 9am - 4pm

Reader Service Number 2



"BrainMaker is the most fascinating computer software I've ever seen."

John Dvorak
PC Magazine

See Jan. '91 Computer product review

By using the same architecture as your brain, **BrainMaker** Neural Networks give your PC human abilities such as finding patterns and hidden relationships in data, and learning from the past. **BrainMaker Professional** accepts larger data sets, provides more flexible modelling, graphics, and includes a run time license.

- Read in data from Lotus, dBase, Excel, binary, ASCII.
- Menus, dialog boxes, color, mouse support, great manuals.
- Stock, bond, crop forecasting, gene sequencing, cancer cell analysis, solar flare prediction, image recognition, and more.
- Educational discounts available.

30 day money back guarantee!

PC Magazine Best of 1990

Free Catalog:
800/284-8112 or 916/477-7481
Order Today! **BrainMaker is \$195**
Professional is \$795

California Scientific Software
10141 Evening Star Dr #6, Grass Valley,
CA 95945