

TELELAB: Secure Software for Remote Health Monitoring on the Internet

Rex E. Gantenbein

Department of Computer Science

University of Wyoming

Laramie, Wyoming 82071-3682

Phone: (307) 766-4226

Fax: (307) 766-4036

EMAIL: rex@uwyo.edu

WWW: <http://www.cs.uwyo.edu/~rex/>

Introduction

NASA has always been a pioneer in *telemedicine* -- the use of telecommunications and computer technology to improve health care -- particularly in the context of space flight. From monitoring astronauts' heart rates with biotelemetry to collaborating with Russian physicians in the "Spacebridge to Russia" project, NASA researchers have used state-of-the-art telecommunications technology to support health care in remote or extreme environments. NASA's leadership in this field has produced many advances in medical care technology for space.

The nature of NASA's role in remote medical care is changing, however. A new model of human space flight is emerging, with longer and more complex missions involving crews, physicians, and technicians from different agencies, nations, and cultures. New health-care technology integrated into flight programs such as the International Space Station must provide appropriate support to missions under this model.

One critical issue in health care for future space missions is the collecting and sharing of medical information over long distances. There are many problems inherent in remotely observing the health of crews on a space flight, including the collection of health-related data in space and transmitting it to a remote observation site on the ground. Even when this problem is solved, it is not always feasible for a physician monitoring particular aspects of a crew member's health to travel to a central observation site to view the results of the data collection.

Furthermore, the nature of medical monitoring for long-duration space flights often requires observation at relatively infrequent periods, over several weeks or even months of a mission. Most medical personnel cannot relocate to a central observation site for several months, nor would they find it possible to travel to the observation site for every

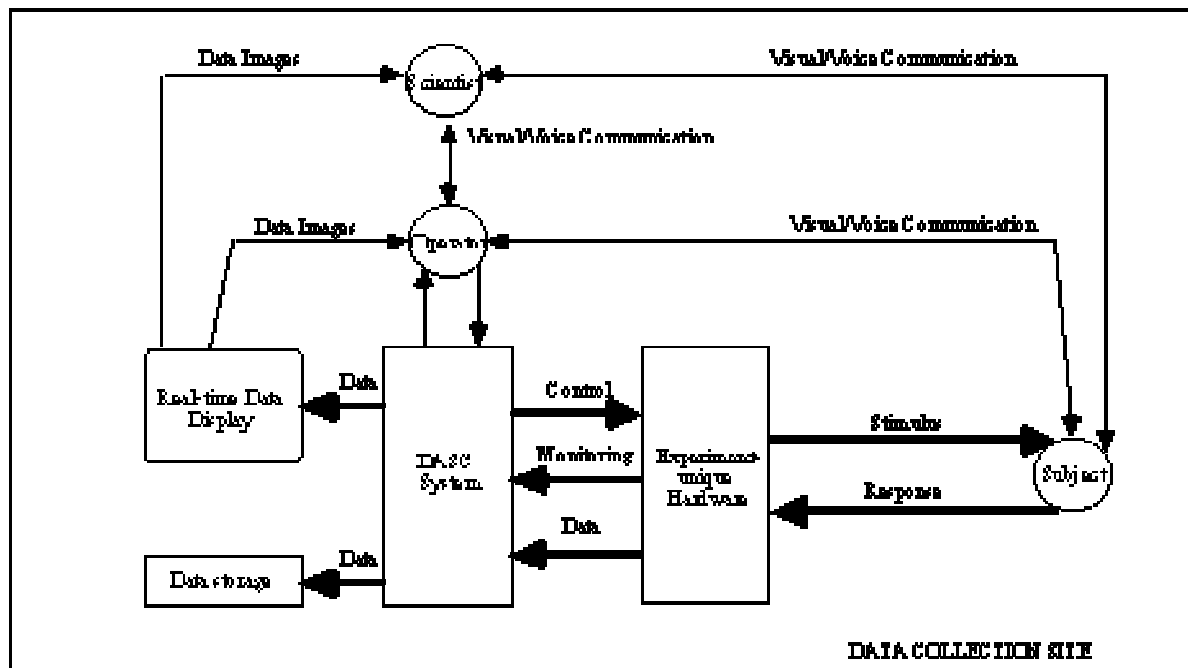
individual data collection. In addition, evaluation of health data is typically not performed by an individual, but rather by teams of collaborating specialists; if all members of the team cannot be present at the observation site, then the data must be able to be quickly and easily distributed for analysis.

It is clear that there is a particular need for telemedical systems that will support remote collection and distribution of health data over wide geographical areas. Fortunately, the growth of the Internet, coupled with the increased availability of powerful desktop computers, has made the long-distance sharing of computer-acquired health data much more feasible. Through advances in software and communications technology, it is now possible for a computer in the United States to access a computer in Moscow, and vice versa. Data and files can be quickly and easily shared among computers anywhere in the world that the Internet reaches. It seems, therefore, that a system integrating computer-based data acquisition with efficient and practical distribution of data over the Internet would provide the foundation needed for this new model of international cooperation in space medicine.

As described in this paper, we propose to create such a foundation through the development and evaluation of TELELAB, a computer and communications system that will provide remote access to computer-based data acquisition systems through the Internet. The main objective of TELELAB is to free medical personnel from the need to be physically present at an ground-based observation site to monitor the health of humans in space. This "virtual laboratory" will be accessible to not just one observer but several, and will support both real-time and off-line observation of health-related data collection, thus enhancing the participants' ability to analyze the health of a crew member and diagnose potential problems.

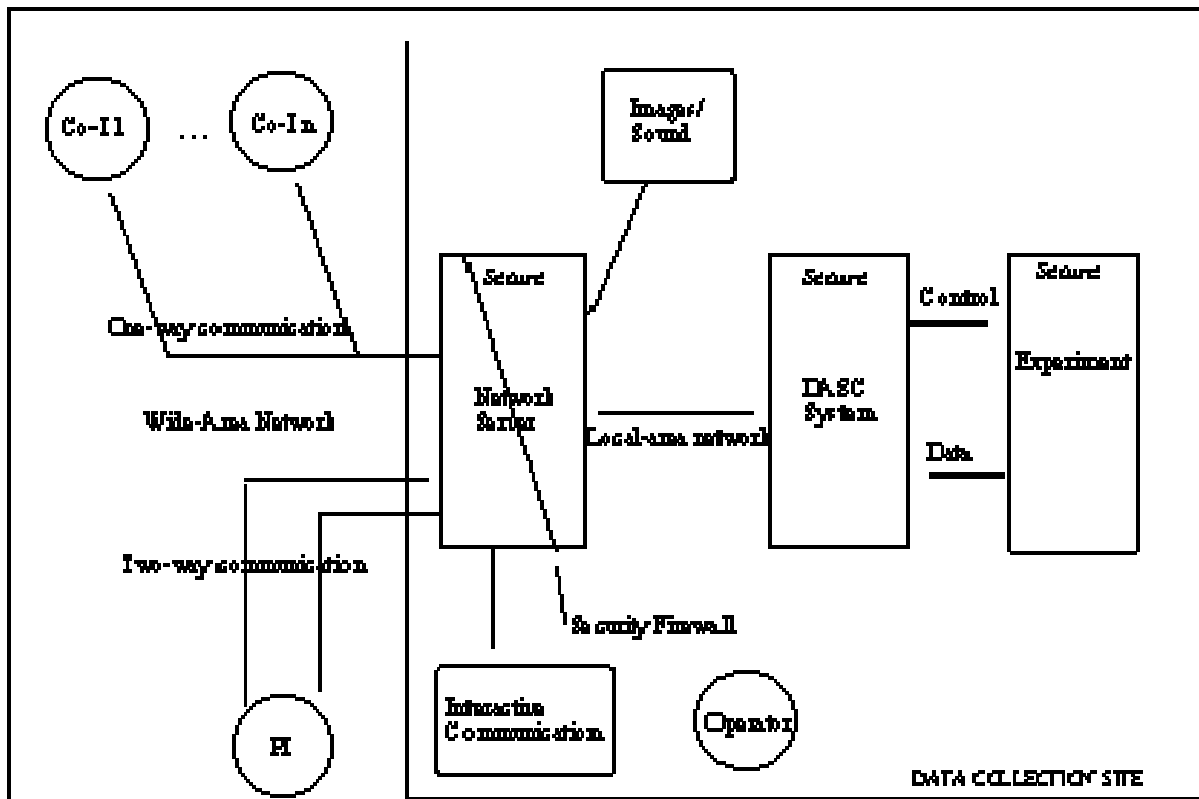
Although intended for use in space medicine, TELELAB may have other applications in monitoring human health in remote or extreme environments on Earth. There is growing interest in the use of telemedicine systems in applications such as rural health care delivery, remote consultations with specialists, cardiography, and radiology. The work we propose would provide the framework for many health-related uses.

Background: Distributed Data Acquisition and Security



The layout of a typical computer-based data acquisition system is shown in Figure 1. A data acquisition system computer (DASC) at an experiment site collects subject data in response to stimuli; the data is either analyzed in real time through observation of a display of the data as it is collected, or is stored as a file and processed at some later time. Although a trained operator can carry out the actual data collection, a medical or scientific observer must be present to carry out the analysis and evaluation of the data.

This model of data acquisition can be integrated with telemedicine technology, making the DASC and the data it collects accessible from remote computers, through *distributed computing*. Distributed computing architectures use networks -- both local-area networks, which cover small geographic areas such as a single building or laboratory, and wide-area networks like the Internet that link distant local-area networks together. Through these interconnected networks, multiple workstations, or *nodes*, performing individual computational tasks can be linked into a cooperative, coordinated system.



A distributed data acquisition system, as shown in Figure 2, would allow the DASC at a collection site to transfer data acquired from a subject there to a remote computer, where an observer could view digital images of the acquired data -- exactly like those at the collection site -- in *near-real time* (that is, as the data is being collected, delayed only by the few seconds required to transmit the data across the network). Through two-way communication with the site, a "primary" observer could remotely control and configure the DASC. Other observers with network connections to the data acquisition site could view the collection simultaneously on their computers.

This architecture would also support the distribution of acquired data in *store-and-forward* mode (that is, the acquired data is saved as a file on the DASC system and then transferred at some later time to an observer). Instead of simply being transferred between computers, however, the file would be "replayed" by the software used for viewing of the data, simulating real-time observation of the data collection. This mode of data distribution is particularly useful when observers need not or cannot be available at the time of a data collection, as when several time zones separate them from the observation site or data collections occur at long intervals over periods of weeks or months.

An obvious problem in such a system, however, is security. As has often been demonstrated recently, computer networks, particularly public networks like the Internet, are notoriously susceptible to security problems. These problems typically fall into two categories: unauthorized access to computers connected to the network, and breaches of confidentiality for data stored on or transmitted between those computers. Both of these

present significant challenges to distributed data acquisition, especially in medical research, where information about a specific individual's health is normally released only with the express consent of the patient and his/her physician. Therefore, in a medically oriented distributed data acquisition system, precautions must be taken to preserve data confidentiality at all times.

Private, dedicated networks that are inaccessible to all but authorized users are one solution, but using such networks for international scientific collaboration is difficult due to cost, problems arising from international agreements, and limited access to private networks in many countries. Fortunately, security mechanisms for public networks also are available, ranging from encryption on transmitted data to user authentication and authorization schemes for computers connected to the network. Most security breaches on public networks are caused by lack of a comprehensive security policy or by failure to enforce existing policies rather than failures in the mechanisms used to implement security. Nonetheless, the confidentiality of human health data in a distributed data acquisition system can be achieved only if a strong defense against unauthorized access to both internal and transmitted data exists within the system. However, it is not yet known how secure such systems can be made to be, or how security mechanisms affect their collection and delivery of scientific data to remote observers.

Project Description

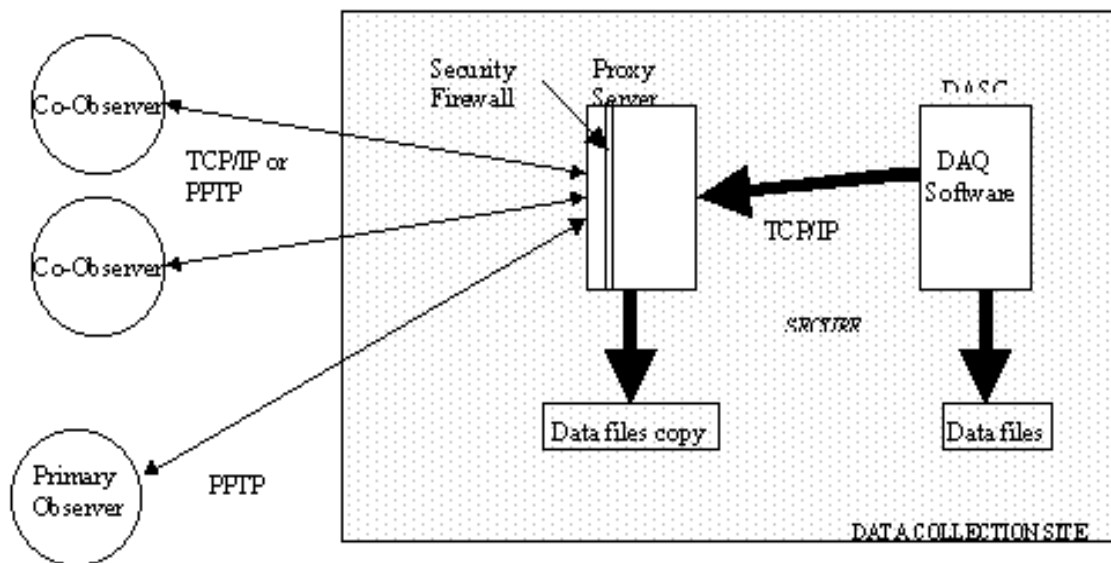
The work we are proposing here will evaluate the ability of computer-driven data acquisition systems to securely distribute human health data, with the emphasis on remote monitoring and control, confidentiality, and use of existing networking protocols. From this work, we hope to show how distributed data acquisition systems can (1) increase the accessibility and quality of health data, (2) reduce the costs of sharing data among multiple observers, and (3) preserve the confidentiality of all data collected by the system.

We will do this through the development and evaluation of TELELAB, a prototype "virtual laboratory" providing secure Internet-based sharing of data between a data acquisition computer and remote computers. Phase 1 of TELELAB, which supports distributed data acquisition, but does not address security issues, is nearly complete. The software for this system is written in Labview, a commercial programming language widely used for developing data acquisition systems. It consists of three components:

1. A *data acquisition* program, which collects real-world data, converts it to digital form, and saves it on the hard drive of a data acquisition system computer (a Power Macintosh with analog/digital converter hardware). This program is adapted from one developed by the author during a sabbatical at NASA JSC in 1994-95 for the Neuroscience Laboratory of the Space Life Science Laboratories at JSC.
2. A *proxy server* program, which uses the TCP/IP networking protocols to support interaction between the DASC and remote computers, making data from the data acquisition program available for network access.

3. A *client program*, which can be run on any computer that supports Labview and has an Internet connection. This program interacts with the proxy server to access data collected by the DASC and display it in a continuous stream, similar to the display on the DASC itself. It can also interact with the data acquisition program to remotely start, stop, and configure an experiment. In addition, the client program can be set up to "replay" data transferred as a file to the remote computer.

The proposed extension of the TELELAB project will involve developing and evaluating mechanisms for ensuring the confidentiality and the reliability of the data being distributed. As mentioned above, public networks like the Internet are notoriously insecure, so security protocols must be part of the data distribution. We will use the next phase of TELELAB's development to study how secure distributed data acquisition can be, as well as evaluate the effects of security on the speed and fidelity of the data distribution.



The Phase 2 version of TELELAB will incorporate commercial proxy server software (available from a number of vendors, including Microsoft and Sun Microsystems) that interacts with the data acquisition software in a manner similar to the current Labview proxy server, but with several extensions for security. This software, which will run on an independent network server connected to the DASC via TCP/IP over a local-area network as shown in Figure 3, will create a secure "firewall" between the DASC and remote observers. The firewall will control the distribution of data files and provide user authentication services to keep the DASC "private" (that is, invisible to Internet users not connected to the local network). Furthermore, the proxy server software will restrict access to the network server itself to specific Internet addresses or domains; coupled with its user authentication service, the proxy server will protect the entire local network from access without explicit authorization. This authorization will be reinforced by positive

authentication protocols, which use digital signatures to verify the identity of all sites that can access the data.

The Phase 2 TELELAB software will enhance the client program to make the remote observer interface simpler and more reliable. The existing client interface closely resembles that of the original data acquisition component, which requires some training to use. We will replace this client program with one based on World-Wide Web technology, making remote access to a data collection available to any observer with a Web browser. This can be accomplished through new Labview features that will support standard Internet connections into the proxy server component.

To ensure secure communication between the network server and remote clients, the Phase 2 TELELAB will employ a secure network protocol such as Microsoft's Point-to-Point Tunneling Protocol, or PPTP. PPTP and similar protocols support *virtual private networking* over public networks, extending the normal TCP/IP connections by encapsulating the network data packets into encrypted datagrams that cannot be intercepted and read by unauthorized receivers. Remote clients equipped with a PPTP connection can be securely connected to the network server, and thus to the DASC through the proxy server. Other clients could be connected through normal TCP/IP networking, but their access to data would be strictly limited to files on the network server.

Once these security mechanisms have been developed, integrated into the Phase 1 TELELAB prototype, and tested, we intend to evaluate their effect on the performance and quality of the data distribution. We will develop tests and tools for measuring performance and overhead for the distribution of data. The specific data to be gathered from the components of the system include:

- throughput of the collected data for the various components, including encryption and password protection;
- response times of the requests for data from the remote sites;
- locality of memory references in the local-area network;
- communication bandwidth achieved in both the local- and wide-area networks;
- processor speed available in each of the nodes of the system; and
- transmission latency over the wide-area network.

This information will be gathered through the use of monitoring software designed to avoid biasing the performance results as little as possible. The performance and overhead data gathered from these experiments will be analyzed to determine the effectiveness of acquisition, distribution, communication, and security mechanisms in the prototype.

Clearly, the success of any distribution project such as this depends on the quality of the transmission link between the sending and receiving sites. This "quality" is difficult to measure in a quantitative way. Performance measures such as throughput, response time, and transmission latency can provide some indication of the quality of the transmission, but many such measures are subjective: visual impression of transmitted images, location of information on the workstation screen, even ease of use of the interfaces.

For local distribution of the data, quality of the data can be practically measured by comparing the original data input to the DASC against the data as received by the network server. A simple side-by-side comparison program can check the fidelity of the acquired and transferred data. The transmitted fidelity of encrypted data can be similarly evaluated by comparing it to an encryption of the original data.

The quality of the wide-area distribution of the data and the reliability of the communication links can be measured through the use of checklists to be developed during the course of the Phase 2 extensions to TELELAB. As we work with the network transmission mechanisms, problem areas should become apparent. Once these have been identified, both the operator at the "collection" site and the observer at the "remote" site will use a checklist during tests. Features perceived by these users as less than adequate will be reviewed and improved where possible. Repeated tests will ensure the validity of the responses as well as isolate any transient events.

Project Objectives

Our goal in creating TELELAB has not been to create a new set of communications protocols for real-time distributed data acquisition, but rather to make as much use as possible of off-the-shelf (OTS) technology for supporting telemedicine. The feasibility of transmitting human health data among geographically separated sites using telecommunications has already been demonstrated in our Phase 1 prototype. The continuation of this research will allow us to examine the ability of existing network technology to support secure distribution of health data in both near-real time and store-and-forward modes.

In this kind of work, "results" are difficult to predict in the usual experimental sense. As rapidly as networks are growing and changing, the technology available now may be obsolete or superseded by the time this project is completed. We are proposing to develop our prototype with the newest communication and security technologies and integrate them with existing data acquisition technology. Through this research, we hope to discover, first, how secure distribution of human health data can be accomplished, and, second, how security measures affect the distribution of the data. By evaluating the strengths and limitations of the available technology, we will create a framework in which the "next generation" of such systems can be effectively designed.

The results of this study should prove valuable to NASA researchers considering development of new data acquisition systems for remote monitoring of the health of humans in space, as well as to the medical community at large. While certain costs can be

reduced by the use of distributed systems, and benefits such as convenience and concurrent data analysis accrue, other costs will increase; trade-off studies will be necessary to determine what price must be paid for a given level of service. This is particularly true with respect to security, which is difficult to quantify in any meaningful way, but can be related to the associated costs and trade-offs required to achieve it at varying levels.

The growth of distributed systems, and particularly those capable of operating over wide geographic separation, has not coincidentally accompanied unprecedented connectivity among networks and the proliferation of tools to take advantage of these internetworks. It is clear that over the next few years the ability to transfer information between remote sites will increase dramatically. As distribution technology advances, the use of off-the-shelf components for this type of work should become more common and the costs decrease. We are strategically placed to take advantage of these advances and utilize them over the duration of the project.