Advanced Network Infrastructure for Distributed Learning and Collaborative Research

Parvati Dev, PI
Stanford University
•
W. LeRoy Heinrichs, co-PI
Stanford University

Steven Senger, Co-Investigator
University of Wisconsin, La Crosse
•
Brian Athey, Co-Investigator
University of Michigan, Ann Arbor

National Library of Medicine Contract No. N01-LM-3-3512
September 30, 2003 – September 29, 2007

Stanford University Parvati Dev • W. LeRoy Heinrichs • Patricia Youngblood • Ken Waldron • Sakti Srivastava • Larry Mathers • Robert Chase • Robert Cheng • Margaret Krebs • Craig Cornelius • Sean Shyh-Yuan Kung • Kingsley Willis • Kevin Montgomery • Rich Shavelson • Dale Harris • Camran Nezhat • Jeremy Durack • Juan Carlos Aragon • Chris Enedah • Aneesh Sharma • David Gutierrez • Brian Lukoff • Mari Kieft **University of Wisconsin, La Crosse** Steven Senger **University of Michigan, Ann Arbor** Brian Athey • Ameed Raoof • Alex Terzian • Ted Hanss **Indian Institute of Technology, Delhi** Sakti Srivastava **CSIRO, Canberra and Sydney, Australia** Chris Gunn • Patrick Cregan • Duncan Stevenson **Northern Ontario School Of Medicine** David Topps • Kevin Smith • Donna Newhouse • Mike Korolenko **Korea** Young Sung Lee • Min Suk Chung

http://havnet.stanford.edu/

Executive

We have developed a framework for an advanced network infrastructure for health education and medical research. We have shown that such a network infrastructure requires a Core Middleware System that monitors and reports network conditions to network-aware applications that can self-scale and self-optimize based on network "weather reports." The core system and applications have been developed within the context of two medical testbeds, a Clinical Anatomy Testbed and a Clinical Skills Testbed. Each testbed focuses on applications that challenge networks in unique ways.

The Clinical Anatomy applications build on our work creating, networking, and teaching with image databases, including stereo images. They are image-intensive and serve many clients. Thus, they provide insight into applications that require high bandwidth, low latency, and support for collaboration. The Clinical Anatomy Testbed activity included further development of the Remote Stereo Viewer multi-client interactive application to demonstrate the ability to self-adapt to network "weather" conditions, and to demonstrate scalability. It was used in an intensive series of tests of end-to-end system performance in which a class at the University of Michigan was instructed in anatomy by an instructor located at Stanford. This application was also used as an engine for displaying stereo images from our Bassett collection. The Clinical Anatomy applications were made openly available using our iAnatomy site. This site was set up to register learners into our InformationChannels system to allow interactive use of our applications. The Remote Stereo Viewer and other applications were used in collaborative activities with the University of Michigan, and with the Northern Ontario School of Medicine.

The Clinical Skills testbed is an experimental testbed that extends our previous work on simulators incorporating haptic devices to teach surgical skills. These devices are highly sen-

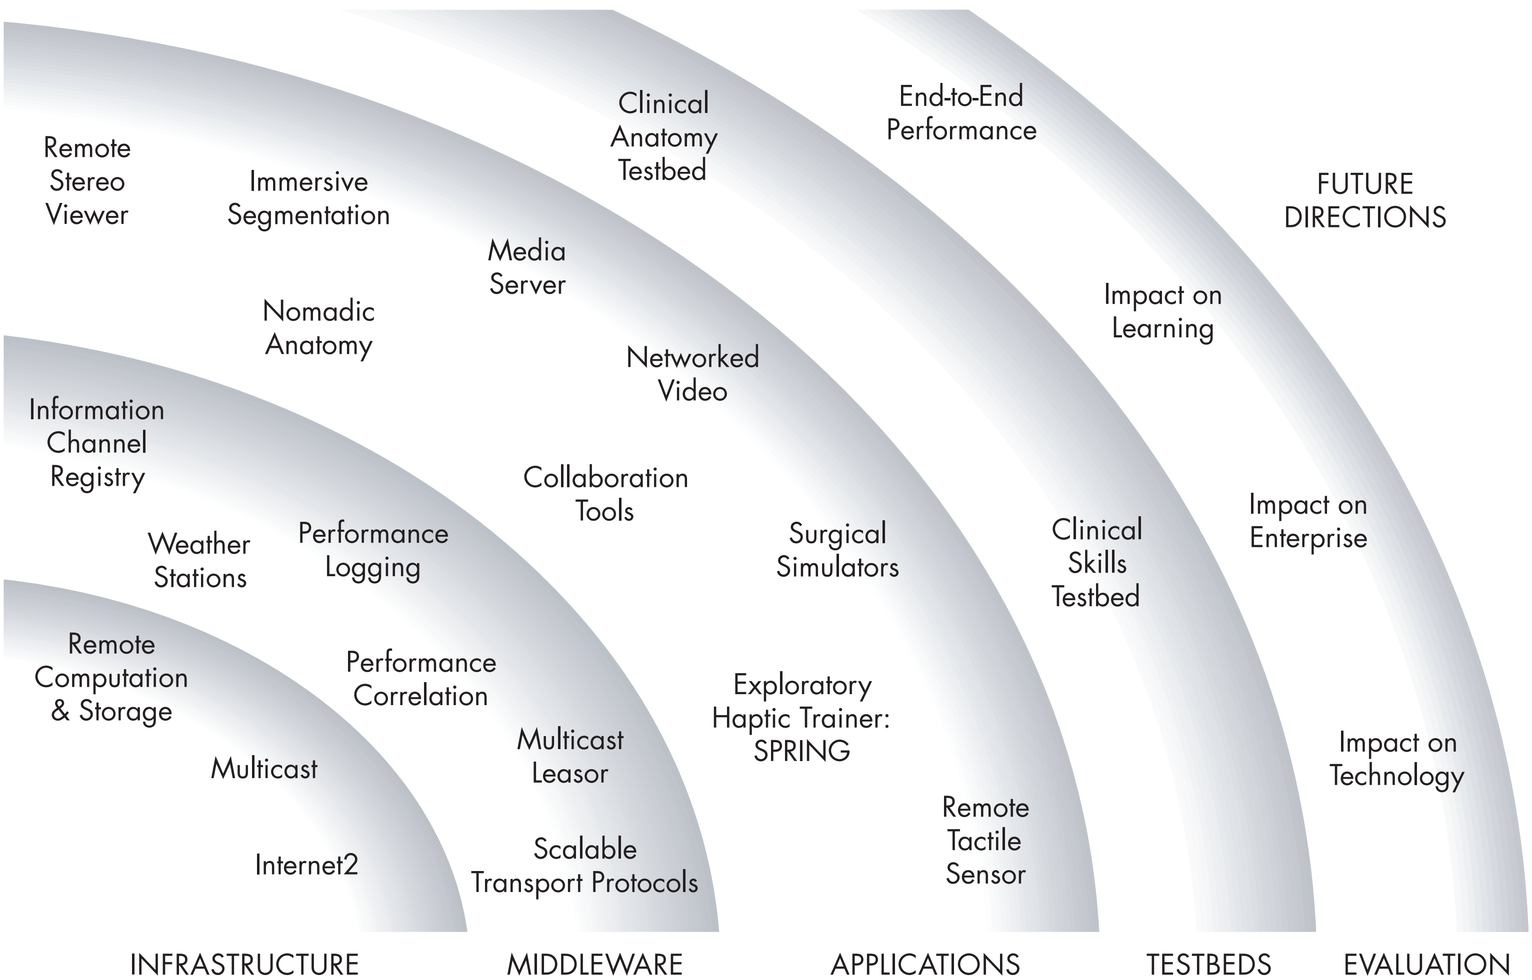

Summary

sitive to network latency and jitter. The Clinical Skills Testbed activity included further development of the SPRING surgical simulation application. This application provides a multi-learner, interactive, soft-tissue model of anatomy, together with learner-controllable models of surgical tools with haptic feedback. SPRING was made available open-source to the research community through the SourceForge website. It is being used as the core engine in a variety of simulators.

This testbed also includes new development of a remote tactile sensor capable of allowing a remote learner to feel simulated or actual lesions as though palpating with a fingertip. This testbed has generated insight into needed theories and research for systems that incorporate haptics, the uses of haptics within surgical simulation and surgical training, and the technological and network requirements of these systems. We have further investigated the abilities of subjects to evaluate tissue stiffness by probing, and have compared the abilities of subjects to distinguish surfaces of different stiffness using both a physical "durometer" and a virtual simulacrum.

Another activity under the Clinical Skills Testbed was high quality transmission of live video. This was used in a demonstration of real-time high-resolution video transmission of a live surgical procedure, with expert commentary, from the Stanford Hospital, to a group of surgical residents in Sydney, Australia. Similar live multi-site events have been conducted with other sites.

Beginning with local testbeds, we have extended our testbeds to national and international scope, and have evaluated them for educational, technical and enterprise impact. Our research provides insight into new directions in networked access to complex data matched with powerful interfaces that support medical research and education.

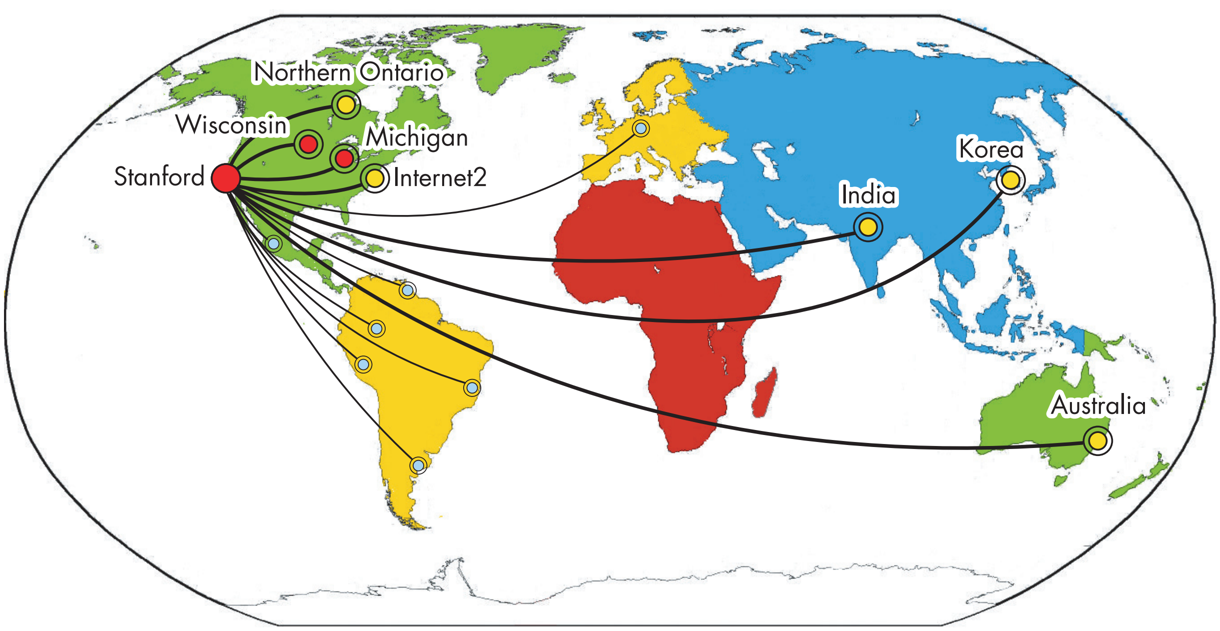

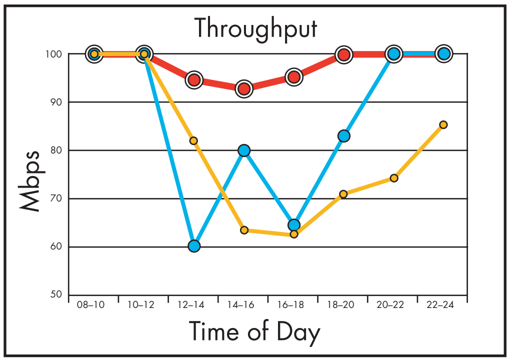

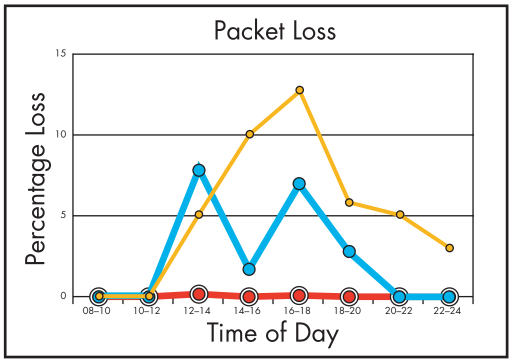

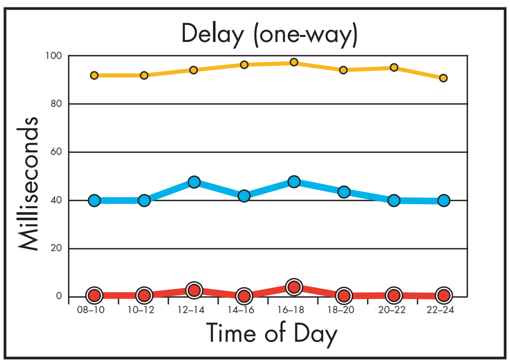

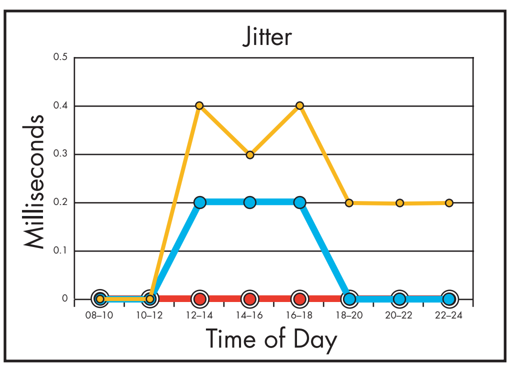

Remote Computation & Storage

The applications and infrastructure developed in this project are based on an architecture in which large databases like the Visible Human data set, and the core applications that work with them, reside on a small number of powerful servers. A large number of client workstations interact with the server applications to provide a collaborative working environment that can encompass multiple sites.

Multicast

Our client applications permit multiple views of the same data set. They also permit collaborative interactions to be led from any workstation. The leader can control the views displayed at all workstations and can position a pointer visible to all participants. The ability to do this depends on the Multicast protocol. Unfortunately, Multicast does not appear to have been uniformly implemented, and network support personnel appear to have little knowledge of its implementation. This became a major limitation on the use of our collaboration-supporting applications.

Internet 2

In all of our work we have assumed the availability of Internet 2 level service. Shortly after commencing the project we

infrastructure

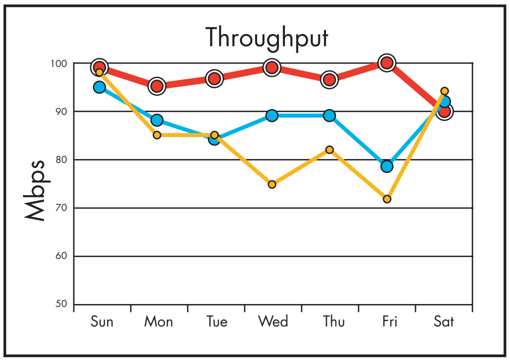

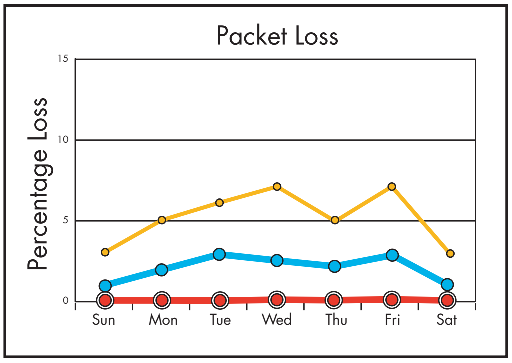

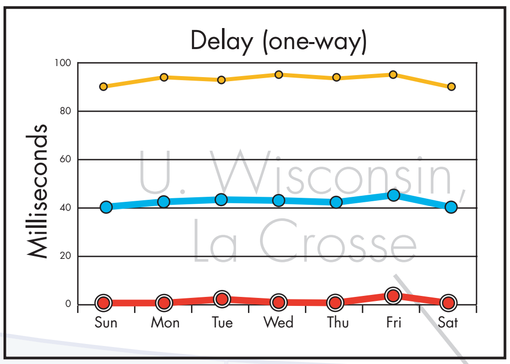

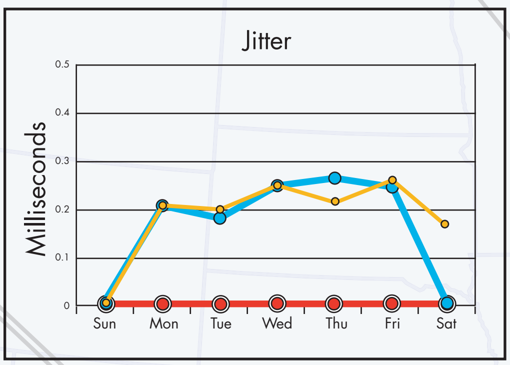

upgraded our connectivity from 100 Mb/s to 1 Gb/s. This involved a solution to the "last mile problem" upgrading the connectivity from the SUMMIT laboratory to the university's gateway, and thence to the CENIC backbone.

Network Traffic Data

Network traffic conditions change by time of day and day of week. They also change as the network architecture evolves within the institution, regionally, and on the national and international backbone. In Fall 2004, we monitored network traffic over many weeks. The resulting traffic patterns are shown on the accompanying graphs.

Local National International

The ability of our applications to adapt to network traffic conditions depends, in part, on the availability of data on current network conditions. We have made that data available by placing WeatherStation units, described later, at strategic network locations. We expect that this data will eventually be provided as a routine network function.

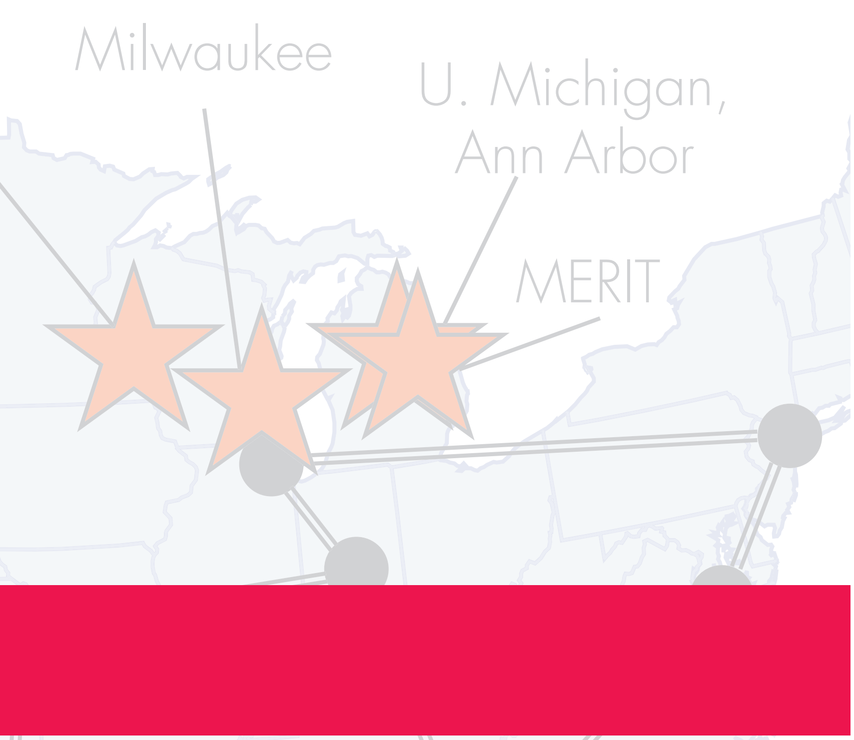

Middleware

The middleware developed under this project supports the goals of demonstrating applications with self-scaling and self-optimizing behaviors that support the end-to-end performance experience of the learner. The middleware builds upon work done under a prior project (the Next Generation Internet), by specifically extending the InformationChannels framework and the network monitoring WeatherStations. On top of these systems we have implemented a mechanism for correlating application performance experience to measured network metrics and using this correlated data to provide performance advice to applications. We have fully integrated this system for the Remote Stereo Viewer (RSV) application and partially integrated it into several other applications.

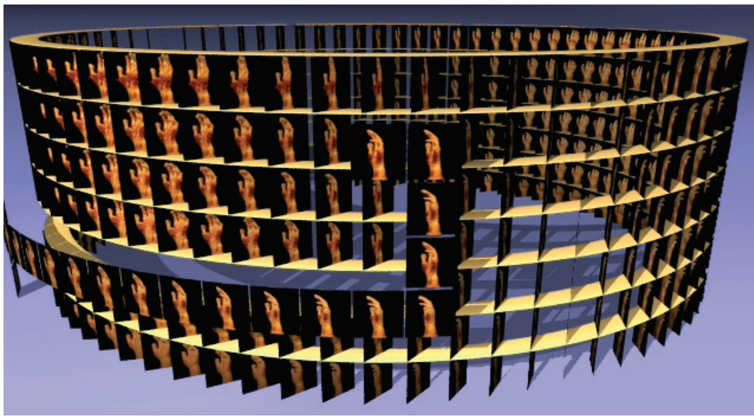

"we have implemented a mechanism for correlating application performance experience to measured network metrics"

The middleware provides mechanisms for logging performance from two sources. First, the WeatherStations installations and software provide the ability to periodically run network tests to measure key metrics. Second, applications log their actual performance experience as they operate. Both sources are placed into a common SQL database. These sources of data are then correlated to provide information on an application's performance experience when certain network conditions prevail. Applications can use this information to both adapt their behavior and provide learners with an expectation of performance as compared to previous experience.

InformationChannels and the Registration Server

At the base of the middleware is a framework initially developed for use in the prior NGI project. The InformationChannels framework supports the formation of complex networks of communicating application components (clients and servers) with minimal direct learner intervention. The framework accomplishes this by focusing on the channels of information exchanged between application components.

Applications announce their ability to provide or consume channels of information by periodically sending channel announcements on a multicast address. A channel is identified by the provider's name, the channel's name and the channel's type. A channel announcement specifies the time interval used for resending the announcement. The announcement contains the network parameters (IP address, port number etc.) required to establish the connection.

A registration server (possibly with multiple instances for redundancy) listens to the multicast address and maintains a list of all active channels. The registration server responds

middleware

http://peabrain.visu.uwlax.edu/WthrStns/v2/ http://peabrain.visu.uwlax.edu/cgi-bin/IC2 to queries about currently active channels. Application components can use this capability to discover the availability of remote services. An automatically-generated web based query page also allows learners to discover the existence of channels.

For example one application, the Remote Stereo Viewer (RSV), provides on-demand access to large, multi-dimensional image sets of anatomical subjects. Learners navigate through the image set in an unpredictable way and view only a small fraction of the total image set as a part of an interactive activity. Consequently, the application transports individual images in very high bandwidth short bursts (generally a fraction of the roundtrip time).

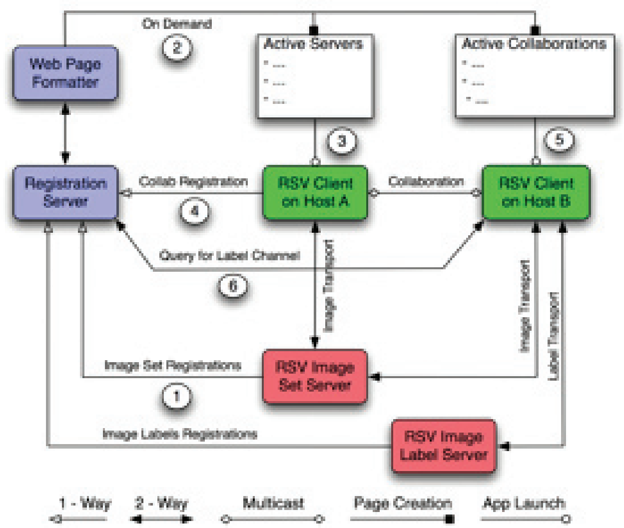

A Collaboration Framework

The RSV application supports the formation of a group of learners through the following process. (1) The RSV server announces its ability to provide access to an anatomy image set. The registration server receives this announcement. (2) A learner, through the web interface, queries for active RSV channels. (3) The learner selects the link for one of the advertised image sets. The RSV client is automatically launched and connected to the RSV server for the image set. The learner can now interact with the anatomy image set. (4) The learner indicates a desire to collaborate with other learners and the client begins announcing a collaboration channel. (5) A second learner finds the collaboration channel through the web interface and selects the link. This automatically launches the RSV client on the second learner's computer which now connects to the image set server as well as to the multicast address used for communication with the first client. Actions of the first learner are now seen by the second learner. Continuing this one step further, (6) either client can query the registration server and discover the existence of other services, such as, a label set server for the image set being viewed.

WeatherStations

Network performance is monitored by WeatherStations that make use of the InformationChannels framework. To monitor end-to-end performance across the national testbeds used in the HAVnet project we have operated weather stations at Stanford, CENIC – Sunnyvale CA, WISCNet – Milwaukee, UW – La Crosse, the Michigan Center for Biological Information and the University of Michigan Medical School.

Each weather station announces a channel indicating its ability to conduct various network measurements. By querying the registration server each weather station discovers the existence of peer stations and schedules tests. The tests consist of:

- Ping loss, roundtrip time
- TCP throughput
- UDP max send rate with less than 1% loss
- Multicast group setup time and loss rate

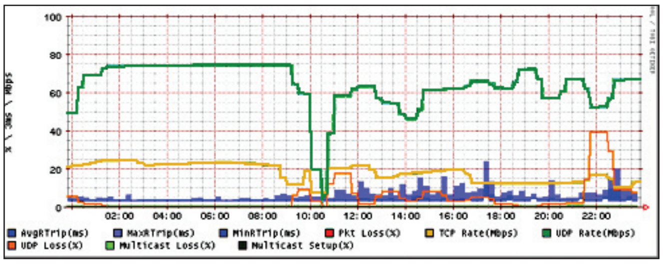

Weather stations are periodically polled for the results, which are placed into the performance database as well as being used to update performance graphs.

MultiServ: Multicast Lease Server

This server provides a mechanism for "leasing" multicast addresses from an available pool. This allows RSV to obtain a multicast address to use with collaborating peers and allows WeatherStations to obtain multicast addresses for use in conducting network tests. The server utilizes the InformationChannels facility to create the lease mechanism. When an application is given an address to use the application begins to announce an "MLease" channel where the name of the channel is the multicast address. As long as the channel continues to be reasserted by the application the multicast lease server marks the address as being unavailable.

LogServ: Logging Performance Data

This server advertises a channel that applications can use to relay performance data, which this server then logs into the performance database. Applications communicate with the LogServ through a client side API that queues log messages, sends them as single UDP packets to the LogServ and continues to resend them until acknowledged by the LogServ. If the application quits before all log messages are acknowledged the remaining queue of messages is written to a local file. The next time the application is run it will attempt to resend these messages.

Both RSV and the WeatherStations query for the presence of a log server and if found use it to log all performance data.

COR & CORQ: Correlating Performance Data

This server correlates the weather station measured metrics and application data. Since specific performance data are unique to each application the COR server creates correlation tables unique to each application. The general principle is to aggregate application performance information based upon the prevailing weather conditions at the time the application was running. As a secondary correlation, application data is also aggregated by the time of day and day of week. The correlation mechanism allows individual computers to be treated separately or for the performance data from a group of machines (e.g. a teaching lab) to be combined.

The correlation query server, CORQ, allows applications to ask for their aggregate performance experience for the weather conditions currently prevailing between a specific client and server locations. Clients can ask for the performance experience from a specific computer or for the combined experience of all machines within their local zone.

Scalable Transport Protocols

Each application will have unique ways in which it can adapt its behavior to changing network performance conditions as a result of design choices. RSV is designed to support unpredictable navigation through large image sets by transporting images on demand. This implies a transport protocol that attempts to send image packets at as high of a rate as possible while balancing this against the goal of zero-packet loss since lost packets incur a round-trip time penalty to recover.

A Network-Aware Scalable Application: End-to-end Performance

The middleware components described are integrated into an adaptive system where the application's features are selected based on network weather conditions.

A learner selects an anatomy image resource, such as "Lung and Pleura" on the course's iAnatomy web page. The click on the data link automatically launches the Remote Stereo Viewer client program on the learner's computer, and links to the corresponding program on the server. The lung images begin displaying on the learner's screen.

In the background, the application sends its network performance data to the LogServ program, and queries (CorQ) the performance data base about any changes it needs to make to its transmission settings. The current network weather (bandwidth, delay, packet loss, jitter, multicast setup) is checked against the information stored in the LogServ database.

If CorQ suggests that current network weather requires changes to the application settings, the scalable application selects some of the many possibilities in a predetermined priority order. The possible changes we have tested include a change in image resolution, frame transmission rate, and in the rate of interleaving of prefetch data according to the application's speed in decoding compressed image data.

Experiments conducted by us suggest that responsiveness of the application is one of the most important characteristics sought by the user. Other scalable features, such as image resolution, can be modified as long as they remain adequate for the learning task.

Applications

Applications were selected and customized for the HAVnet project to support collaborative and distance learning. Selected applications were enhanced significantly so as to support our research in creating network-aware, scalable applications. They are described below and in the highlight boxes.

Immersive Segmentation

Immersive Segmentation supports the interactive visualization and segmentation of large volumetric data sets such as that of the Visible Human Project. The application has a client/server design and uses a stereoscopic, haptically enabled, learner interface. The application is both visually immersive and computationally immersive since the learner is placed at the center of a fine-grained computational loop where learner choices continuously affect the course of the segmentation algorithms. The application utilizes the InformationChannels framework and can log performance data.

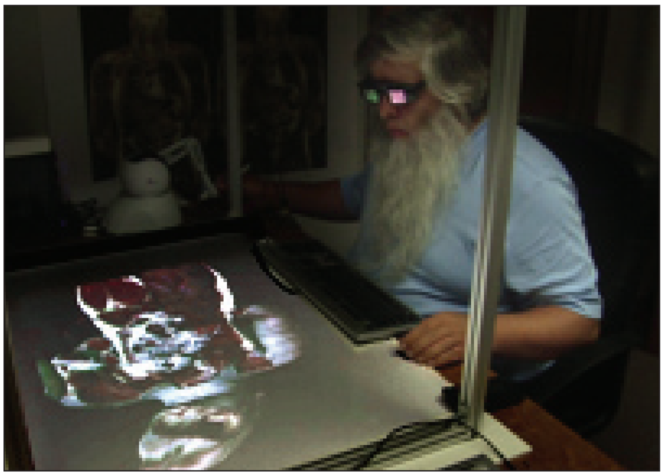

Nomadic Anatomy Viewer (NAV)

This application puts the Visible Human Project (and similar volume data sets) in your hand via a wireless handheld computer. The interface is pen-based. A small collection of on-screen controls allow the learner to arbitrarily position a slice plane relative to the volume. The application utilizes the InformationChannels framework and supports collaboration between clients.

Remote Stereo Viewer

Description and Past Work: The Remote Stereo Viewer (RSV 2.0) is a client/server application that provides immediate access to large sets of high-resolution stereo images. The image sets are organized as multi-dimensional grids. Each image set consists of a series of images forming a 360° revolution of viewpoint around an anatomical subject. Consecutive images around this rotation axis are used to construct a stereo view of the object. Horizontal mouse motion moves the image in this dimension through the image set. Additional dimensions, mapped to vertical mouse motion, may show the structure at various levels of dissection or magnification. The application assumes that the learner will view only a "small" selection of the images and consequently does not warrant downloading the entire image set. RSV 2.0 uses the InformationChannels framework and includes the ability to form collaborative groups of learners. The content includes both hand and knee image sets, labeled for use in anatomical study.

Network Characteristics: Image sets are stored on the server as JPEG images and transported to the client on demand using a UDP layer protocol. The transport rate is configured to be equal to the JPEG decompression rate of the client, allowing the client to interleave receipt and decompression of data. Between Stanford and LaCrosse, transport produces bursts of between 30 and 40Mbps.

End-to-end Performance Support: RSV 3.0 incorporates the new collaboration and end-to-end performance support of the core in-

applications

Media Server

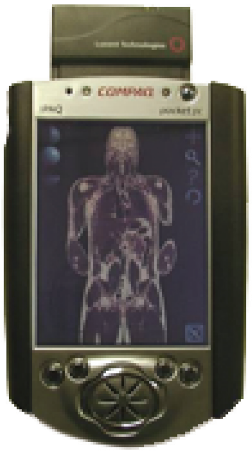

The Stanford MediaServer is a centralized web-based management system for the collection, organization, authoring, and distribution of media collections in the biomedical research and education communities at Stanford. A specialized version, the Bassett MediaServer, hosts a digitized version of the original images from The David L. Bassett Atlas of Human Anatomy. Stereo images of these detailed cadaveric dissections are now available through the Bassett MediaServer for use in our testbed, through the iAnatomy web site.

Exploratory Haptics Trainer: SPRING

SPRING is a software platform for developing surgical simulation applications, initially developed by Dr. Kevin Montgomery of the National Biocomputation Center in 1999. The set of networked, cooperating applications in SPRING offer an interactive 3-D world of tissue and tool models, using standard data formats, and with support for both rigid and deformable models with customized physical tissue parameters. SPRING also includes 3-D surgical tools with programmable behaviors (Probe, Grasp, Cut, etc.), and 3D interfaces with both position information and haptic feedback. (see highlight box)

Video Collaboration

The Collaboration Room opened in 2005 and provided the HAVnet team with greatly improved ability to participate in and host video collaborative networked events and research. It supports numerous styles and quality of video collaboration, shared applications, and very flexible setup for educational activities. (see highlight box)

frastructure. Specifically, we modified the client to report actual image transport performance information to the WeatherHistory mechanism. The client also queries current weather conditions and uses this information to optimize its behavior accordingly.

All clients in a collaboration group currently have independent connections to the image server and submit separate image requests. Thus, the server repeatedly sends the same image to each client in a collaborating group. The new version incorporates a multicast address group to transport image data to collaborating clients, significantly improving the application's scalability.

Scalable Transport in RSV: RSV is designed to support unpredictable navigation through large image sets. Typically, images are sent in high bandwidth bursts where the burst duration is a fraction of the total round-trip time. Since the transport link would be idle for significant periods of time, the client prefetches images in the neighborhood of the last request to utilize the available capacity of the transport link as long as this does not negatively impact transport efficiency. The client can also adjust the image resolution requested to better balance use of the transport link. The performance data logged by RSV consists of its transport success experience (packet loss percent, percent of images transported with zero loss) at different send rates.

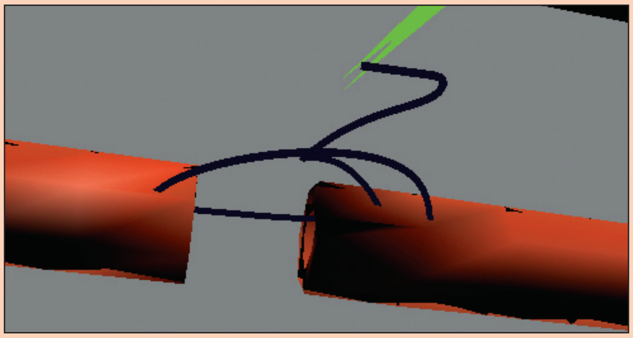

Exploratory Haptics Trainer: SPRING

SPRING is a software platform for developing surgical simulation applications, initially developed by Dr. Kevin Montgomery of the National Biocomputation Center in 1999. The set of networked, cooperating applications in SPRING offer an interactive 3-D world of tissue and tool models, using standard data formats, and with support for both rigid and deformable models with customized physical tissue parameters. We have enhanced the software of SPRING while making it more accessible to potential learners and developers. Our major activities have included updating the code base and development environments for use with modern software engineering tools, improved algorithms, and resolution of many software problems. The software is networked, cross-platform, and open source, with a web site, availability on SourceForge, and extensive documentation. SPRING has been presented at workshops and conferences as a development platform on Open Source Surgical Simulation. Partly as a result of these efforts, SPRING is now in use as a research tool and simulator development platform at multiple university sites around the world. We intend that improved availability and support will allow computer science groups and medical schools to design and develop simulation applications for medical education and training.

The Collab Room

The *Collaboration Room* opened in 2005 and provided the HAVnet team with greatly improved ability to participate in and host video collaborative networked events and research. It is a 400 sq.ft. experimental space to study the impact of new collaboration technologies and new teaching methods on medical education of the future. This space is used for planning and staging interactive high bandwidth educational events, virtual reality team training sessions, and pilot studies for simulator validation. The room is equipped with a high resolution videoconferencing system, Access Grid, capable of connecting teams simultaneously from multiple locations worldwide. Rich media data streams, such as large projected stereo images from remote simulators are part of the learning experience. The infrastructure supports easy switching to other collaboration tools, such as standard videoconferencing, desktop-sharing and videoconferencing from many vendors (Microsoft, Marratech, Skype, and others) or other web-based communication tools such as Adobe Connect.

Around the room are mobile carts with state-of-the-art surgical simulators where residents and fellows evaluate modules on clipping, dissecting, grasping and other basic surgical skills. Key to the room's design, therefore, is flexibility so that the room can be set up and torn down for the variety of activities required for their educational research.

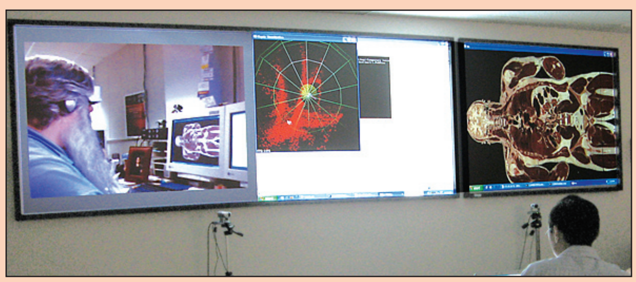

applications

 http://spring.stanford.edu http://havnet.stanford.edu/about/collab.html

Remote Tactile Sensor

The dermatologist, conducting an examination of a remote patient, needs both high quality video,and high quality haptics for diagnostic purposes. We have designed and constructed a system for remote diagnosis of dermatological disorders employing both visual examination and palpation. It employs a haptic master-slave robot system. This sub-project turned out to be more challenging than was originally envisaged, and the full system is still being tested. However, we did make significant progress in many different, and relevant, areas.

We developed and tested a tactile probe suitable for the detection of skin texture, and the evaluation of skin profile. The unit is compact and light in weight suitable for mounting on a slave robotic arm. The sensing is done by an array of piezoelectric sensors formed on a piece of PZT film. As part of the design of this sensor we did extensive work on mathematical modeling of tactile sensors interacting with compliant substrates, such as skin for the purpose of evaluating texture. The tactile probe has been tested on human subjects and has been successful at discriminating between the textures of skin on the back of the arm and on the front of the arm and hand. We developed separate means of sensing of skin profile: that is the palpation of lumps in order to evaluate their geometry, location and mechanical properties. This requires a combination of the tactile probe, and the native haptics of the slave manipulator. We have successfully operated the tactile probe mounted on an Sensable Technologies Omni slave robot and programmed to maintain constant pressure on the skin, in order to sense the skin profile.

The network parameter of particular relevance to the Remote Tactile Sensor is latency. Closing a control loop around a delay of the order of that generated by speed of light and switching delays (10's or 100's of milliseconds) results in dynamic instability. Control algorithms, such as the wave variable algorithm, cause a softening of the haptic sensation. We are investigating calibration methods based on the average latency for each specific session. We are also investigating a model-driven approach to collecting haptic data where the learner then interacts with that model, rather than the actual patient.

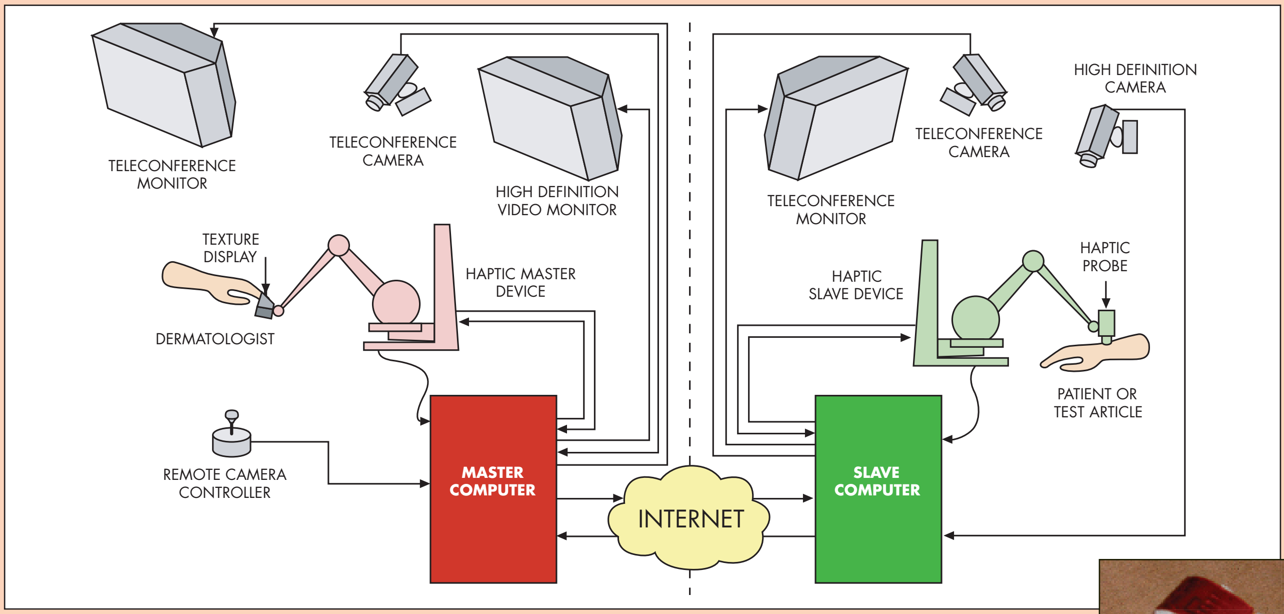

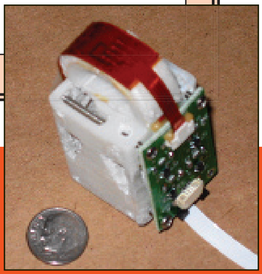

Clinical Anatomy Testbed

We designed and implemented local, national, and global testbeds to link learners at distributed sites through Internet2, and provided access to remote media resources through richly interactive interfaces. We used these testbeds to demonstrate, investigate and evaluate technical and pedagogic requirements for collaborative learning among students at distributed sites. Lessons learnt in each use of a testbed were used to improve the configuration and process for subsequent testbed experiments.

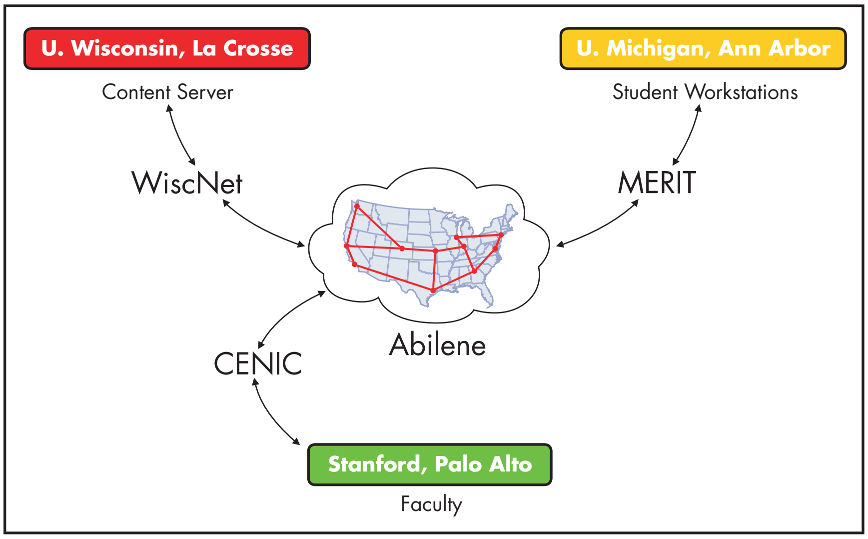

Local Testbed: The local testbed was set up at Stanford University, between the Anatomy division's classrooms and our laboratory. Faculty in the laboratory guided students in the classrooms through a series of stereo views of anatomy. The actual teaching sessions were preceded by numerous test sessions that were needed to set up the technology.

The primary application used in the testbed was the Remote Stereo Viewer (RSV). Using RSV, students accessed stereo images of the dissected hand and the knee, to get a three-dimensional view of dissected anatomy. The viewpoint could be moved interactively, allowing rotation of the anatomy and dissection through multiple layers. RSV supports collaboration by allowing students to form groups to observe others' interactions with the images.

The local testbed was critical in identifying the technical infrastructure essential for remote collaboration in learning. Setting up the collaborative group, simultaneous viewing of an image with its cursor, and transfer of leadership, all required use of a key feature of the network – Multicast. Our local testbed showed that, even within Stanford, the multicast protocol was not implemented uniformly over the network, leading to failure of multicast between some pairs of learners. A second requirement for collaborative learning was the use of video collaboration. The teacher found it essential to see the tacit cues that were indicative of failure to understand or follow the lesson. A third requirement was good audio, with no echo. This was achieved either through the use of headsets or through echo-cancelling microphones. A fourth requirement was the configuration of identical equipment and versions of software at each site, in order to reduce the chances of system incompatibility for complex applications. A major outcome of the experiments with the local testbed was the specification of a new experimental learning space, the Collab Room, as a flexibly configurable space, suitable for testing many types of distributed learning.

National Testbed: For our national testbed, we linked faculty at Stanford University with students at University of Michigan, Ann Arbor, using Internet2 at 1 Gbps. Both accessed image resources located at Wisconsin and at Stanford. The goal of this testbed was to determine the issues that would arise as we scaled up the testbed.

clinical anatomy

Creating a stable network infrastructure that would support our application, Remote Stereo Viewer, proved even more difficult in the national testbed than in the local testbed. Each site worked with its institutional network provider to obtain adequate bandwidth. However, obtaining multicast functionality required testing of every point-to-point connection, with the involvement of the regional networks, CENIC, WiscNet and MERIT, as well as the network engineers at each institution. Intervention of Internet2 network engineering leadership was required for multicast problems to be addressed. This heroic effort made it clear that the network infrastructure for collaboration is inadequate, and that, until these services are available at the network level, collaboration may have to be a function of each application.

International Testbed: For our international testbed, we partnered with ORION, Ontario's high performance network group, and the Northern Ontario School of Medicine (NOSM), a new medical school, with an inherently distributed student body. Classes are held simultaneously at their two campuses, Sudbury and Thunder Bay, 500 miles apart.

Faculty at NOSM accessed the Bassett collection of stereo images of dissection remotely, to prepare image sets that were uploaded to the Remote Stereo Viewer server at Stanford. The images were labeled in RSV and then used in interactive small group learning by both faculty and students. Since NOSM provides only prosected dissection specimens, and no actual dissection experience, the three-dimensional stereo images proved attractive to the students. NOSM continues to use this collection for teaching even though the testbed experiment is complete.

In the evaluation, all students agreed or strongly agreed that the use of the images in RSV was worthwhile. 88% agreed or strongly agreed that 3-D virtual reality and the Bassett collection should be incorporated into their curriculum resources. The students had some excellent practical suggestions on how the technology and the material could be incorporated into their curriculum.

As an extension to our testbed teaching experiments, we undertook numerous demonstrations to institutions around the world. Demonstrations to most parts of the world were over Internet connections of 2Mbps or less, and with significant delay. In these situations, slide presentations using videoconferencing applications were the most effective. In some demonstrations, Internet2 quality bandwidth was available. The remote site could access our resources, such as RSV images, interactively, though only in mono- viewing, not stereo. However, true multi-site collaboration was not available. Screen-sharing applications, such as vnc, supported collaboration but interaction was slow. Some sites had stereo-viewing capability, enhancing the viewing experience. Others, such as the KISTI and MEDRIC sites in Korea, had high-definition videoconferencing capability, thus providing an extremely strong sense of presence between our laboratory and theirs. However, only those sites, such as NOSM and a few others, who had installed the RSV client and had multicast capability, experienced the collaborative interaction that our application and content was able to provide.

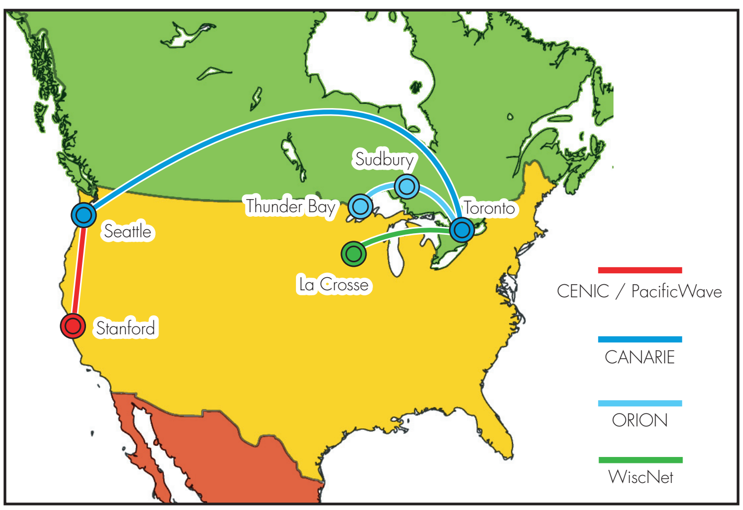

testbed

Clinical Skills Testbed

The clinical skills testbed encompassed many areas, most including a component of haptics perception.

Skills Training with Simulators: Procedure simulators support practice of surgical tasks, from camera control, to grasping and clipping, to conducting a procedure such as excision of an ectopic pregnancy. We conducted training sessions with ob-gyn and other residents to understand the process of training with simulators, and the technical and pedagogic issues that may arise.

Performance Score in Simulation Tasks: Surgical simulators have embedded metrics such as time taken and path length of tool, many of which do not capture the overall performance of the task. We developed a statistical tool to evaluate performance scores for criterion-based training with simulators. (see highlight box)

Video-Supported Learning Events: Video telecast of real-time surgery is a common distance training method. To create a greater sense of presence and immediacy, we experimented with high-resolution two-way video (DVTS) and stereo video

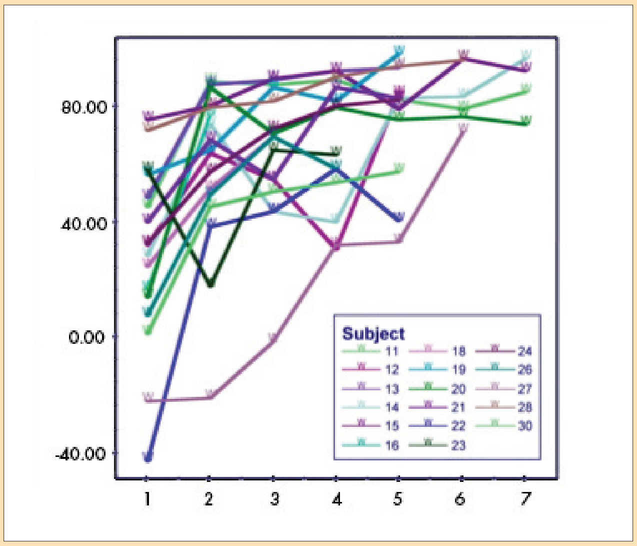

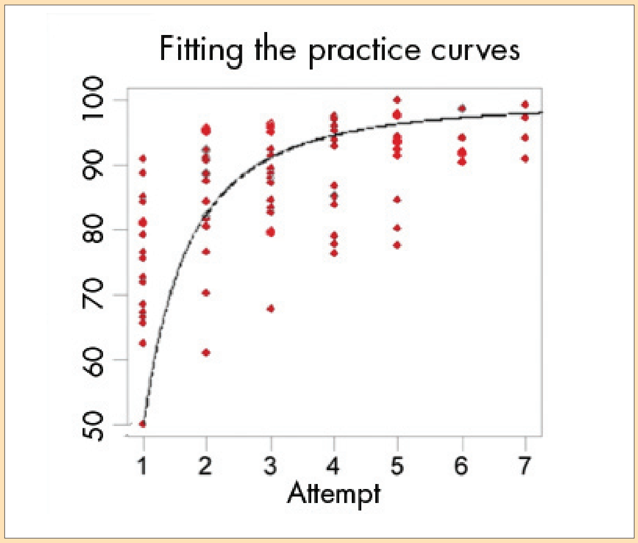

Performance Score in Simulation Tasks

We, in collaboration with the Society for Laparoendoscopic Surgeons, developed a collaborative, formal study for acquiring benchmark data for criterion-based training of basic technical surgical skills with five widely-used training simulators. Seventeen laparoscopic surgeons spent three half-days learning to use the simulators, and collecting data from which Performance Scores (on 4th-attempt data) could be rigorously calculated. Scores are a linear weighted sum of the results from metrics embedded in the simulator, and are of the form b0 + b1.X1 + b2.X2 + . . . + bk.Xk. The coefficients, b0, b1,. . . bk, are adjusted for each task, using a linear regression process, to generate a performance score for each task that improves from a value of 50 on the first attempt and asymptotes to 100. Fourth attempt data, using this score, was determined to be an accurate predictor of a subject's performance. This statistical method will be a valuable tool for evaluating performance scores for subsequent studies of networked simulators.

from the laparoscope camera. A laparoscopic appendectomy was conducted in an operating room at Stanford University, and was observed by residents in Sydney Australia. The project was a technical success and routine transmission of similar sessions is possible on Internet2.

SPRING: Exploratory Haptics Trainer SPRING software is inherently network-based, supporting geographically distributed surgical manipulators connected to a compute-intensive simulation server. Consequently, network emulation software, with controllable delay and other performance parameters, could be inserted between any device and the server. We used SPRING as the basis for our networked haptics perception experiments. We have also built a suturing trainer that will be used for experiments on suturing performance over the network.

Haptics perception experiments: To make effective use of touch in networked simulators, it is essential to understand how the network communicates this sense for individual and collaborative work in surgical simulation. In this work we focused on network latency as the primary network parameter to adjust. (see highlight box)

Experiments in Haptic Perception

We performed a series of experiments to measure haptic perception using both real and simulated forces. We first created a durometer containing latex sheets of various thickness. Learners probed these using laparoscopic tools, but without visual feedback. Subjects are able to feel differences in elasticity and can correctly rank the sheets in order of stiffness.

Virtual membranes with haptic output were modeled using the SPRING platform, comparing objects with four different constant stiffnesses to an unknown that matched one of the other four. The task is to pick the closest match to the center object by probing and sensing force differences. Network latency was also introduced, adding a time delay from from 0 to 100 milliseconds between probing and the force feedback. This apparently simple task proved to be difficult for most subjects.

We tested several profiles of force feedback as a function of penetration depth. Although learners judged the profile that increases force as the square of the depth (quadratic) to be most realistic, subjects performed at the same level with both linear and quadratic force profiles.

In an apparently simpler task, comparing two virtual surfaces using a low-friction, finger held haptic interface, subjects can usually pick the difference in the force levels reliably when one force is twice the magnitude of the other. A significant number of errors are made with smaller force differences.

When both membranes resist with similar force, but a network delay greater than 50 ms is introduced in sensing force on one membrane, subjects usually interpreted the two forces as being different. This is compatible with our earlier studies that indicated 100 ms as the maximum loop delay before a position-following task became unstable.

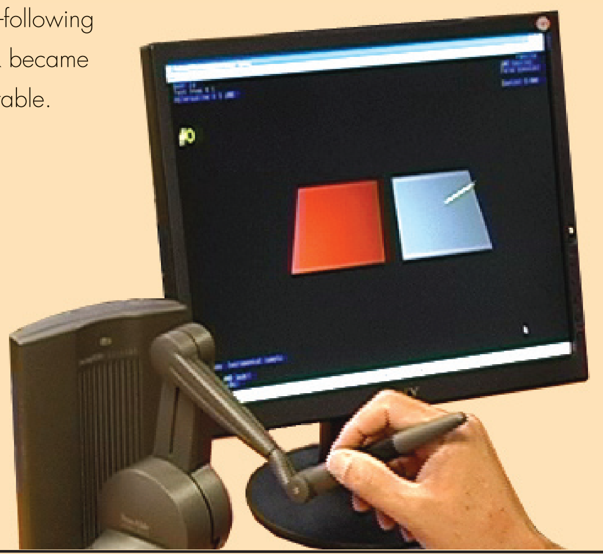

ed

Impact on Technology

Infrastructure and Middleware The work performed under this project in transport protocols, performance monitoring and end-to-end performance has had several impacts on the research community. Within the Internet2 community, the network requirements of our applications led, influenced the development of an API design for a new transport protocol in the Bulk Transport working group of Internet2. This transport protocol will have better performance characteristics than TCP by attempting to differentiate congestive loss from non-congestive loss by monitoring changes in connection latency.

Our extensive use of multicast has demonstrated that monitoring multicast performance should become a standard part of network monitoring. The project's use of multicast also resulted in multicast being made an operational priority within the Wiscnet network. The unique performance needs of the applications developed by this project have been used in planning and soliciting for a new state network backbone.

User Response to Network Performance Degradation: We conducted a study of student perception of RSV images under degraded network conditions and under degraded image quality. The degradations included reduction of the available bandwidth for the transport of the images through the network links, and two different latencies between server and client.

An important finding is that students rated with scores around 4 (considered very good) the scenario in which the RSV application works with a network connection set at 5Mb/s and with an image set at 50% resolution. This is a striking finding considering that in the usual non-degraded scenario this application works with a bandwidth above 70 Mb/s and 100% image resolution. This opens the opportunity to extend the use of RSV on IEEE 802.11 Wireless LAN environments, where bandwidth is scarce. We also investigated the impact of responsiveness and image quality on the overall learner perception. We found that the responsiveness component has a bigger influence on the overall perceived quality of the RSV application.

Open Source Surgical Simulation Software: SPRING: In December 2005, we released the SPRING package as a robust, publicly usable suite of software applications, with freely available development code, on SourceForge.net. We have developed and maintained a website for SPRING learners that describes surgical simulation, applications, and the use of SPRING for developing practical simulation. Documentation on the architecture and details of SPRING and its deployment are available, as well as links to the SourceForge project. We have hosted a 3 day workshop, and presented at conferences, to build awareness of this powerful software.

SPRING as an open source software has opened a technology platform to students, research groups, and developers around the world. In addition to raising awareness of surgical simulation, SPRING's availability has lowered the barrier to incorporating 3D modeling, deformable surface physics, and haptic interaction into projects around the world. We have initiated and facilitated growth of a community of learners, developers, and researchers in open source surgical simulation through this project.

Impact on Learning

The use of advanced technologies in medical and surgical education has a significant impact on learning. The impact is most easily recognized in terms of access. In many schools and many parts of the world, dissection is no longer a part of the medical school curriculum, so the opportunity to do a "virtual" dissection becomes essential for these students. Richly interactive, distributed learning technologies, such as the RSV application, give learners access to complex datasets that use powerful learning resources, 7 days a week, from any location—on campus, across town, or from another city!

evaluation

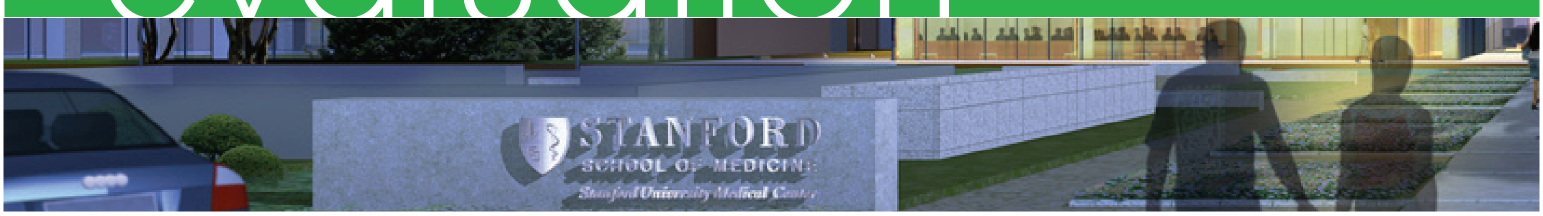

Surgical education, at the residency level, is being transformed by the introduction of simulators—both "stand alone" part-task trainers and networked simulations which give trainees the opportunity to work "one on one" with a master surgeon—manipulating virtual tissue and organs of the same simulated patient—even when trainee and surgeon are in different geographic locations. These new learning technologies also make it possible for surgeons to watch a live surgical demonstration in three dimensional stereo, as if they had a front row seat in the operating theatre, regardless of their actual location. The addition of force feedback, or haptics, to surgical simulators makes it possible for trainees to "feel" what the instructor feels, receive "hand on hand" guidance and respond to individualized instruction from a master surgeon in a distributed learning environment.

Evaluation of these technologies shows considerable acceptance in spite of technical difficulties. After a surgical learning event (SimTech, 2004) between Stanford University and surgical residents at Canberra, Australia, 100% of the residents rated the learning value of the stereo images as acceptable or better, and 97% felt that the performance of the networked application, over Internet2 and its global equivalent, was acceptable or better. Medical students at Northern Ontario School of Medicine assessed the RSV application, the images and the labeling capability as useful or better. Many suggested that their use in small group learning and teaching between students would be very useful.

Impact on Enterprise

Through the HAVnet project, SUMMIT has impacted the local enterprise at Stanford University, as well as institutions and organizations globally.

Learning and Knowledge Center (LKC): A newly funded building program, scheduled for opening in early 2010, will incorporate state-of-the-art learning technologies in classroom configurations that take advantage of networked facilities. SUMMIT's leadership, through HAVnet and other projects, has contributed the vision for learning spaces that enable networked visualization of media-rich anatomy, of simulations with standardized and virtual patients, surgical simulation for individual skills development, and virtual worlds for team training.

Center for Immersive and Simulation-based Learning (CISL): The strong focus on simulation at SUMMIT, as well as at centers at the VA, Childrens Hospital, and in the surgery department, has led to the formation of the new Center for Immersive and Simulation-based Learning. The center's mission is to improve patient safety, patient care, education, and research through innovations in immersive and simulation-based learning techniques and tools.

Institutions Outside Stanford: As a direct result of the Clinical Anatomy Testbed activities, The University of Michigan's medical school has developed a new learning space for collaborative and visualization-based learning. Network performance has been a key requirement for this space, including the requirement for multicast. This room is now in routine use at Michigan. Similarly, the Northern Ontario School of Medicine has pressed for 10 Gbps connectivity to both their campuses, with a requirement that this support simultaneous teaching at both facilities. The stereo images from Stanford will be the backbone of the visualization-based teaching component at NOSM.

Collaborative Learning Spaces: The Collab Room at SUMMIT was based on an innovative collaborative learning facility at Stanford in Wallenberg Hall. We customized the concept to create a room suitable for medical teaching and learning. Since the space opened in 2005, we have hosted about 100 groups each year, interested in the space and the applications it houses. Sites at Stanford and elsewhere have modified the design and implemented their own collaborative learning spaces.

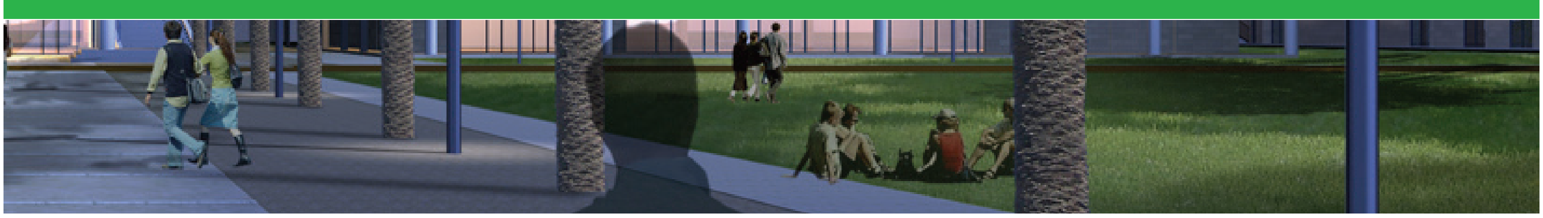

Future Directions

Specifically, we see several directions for our future re search. The first is the technical infrastructure and the user interfaces to support the increasingly complex data and algorithms that will be accessed and used by future healthcare workers. Collaborative learning, access to remote experts, and transparent access to computing, storage and collaboration resources will be essential to future research. Our past and current work provides an excellent foundation for future research in this direction.

New technologies that promise to be relevant to medical education continue to be developed. The consequence is a continually increasing demand on Internet performance. When we started the current project we had 100 Mb/s internet connectivity. That was expanded to 1 Gb/s for the purposes of this project. Several of our competitors now boast of 10 Gb/s connectivity, and for purposes of planning only a few years ahead we need to think in terms of 100 Gb/s. The implications in terms of the ways in which we will run interactive applications, and the ways in which software and hardware infrastructure will be affected remain to be explored. There will certainly be issues like the multicast protocol of the second generation internet that has not been consistently implemented in the way that was originally envisaged, and consequently remains a major limitation on interactive networked applications.

New visualization systems, high-resolution display walls, and high-definition video will become essential. Wave-front displays are becoming commercially available and offer a very viable alternative to stereo video. There are issues to be explored as to how one interacts with a wave-front model, and what are the implications for display of the very complex data sets characteristic of anatomical and physiological medical models. Functional MRI presents many of the same issues, but also those of real-time data segmentation, anatomical registration, and identification with a physiological model.

Hardware and software developed for gaming continues to be a fruitful source of potential adaptations for educational purposes, and particularly for medical education. We have been active in other projects in developing models of virtual environments suitable for training team interactions in operating and emergency room environments. Much work remains to be done before we can fully exploit the potential of these methods.

Our work with the Remote Tactile Sensor has demonstrated the diagnostic importance of synthesizing very diverse types of data obtained by different means, and with very different transmission and display characteristics. This work can obviously be expanded to many different domains.

Interpretation of complex medical systems in which function is important, as well as structure, as is, in fact, always true drives research in the direction of multiscale representations from the whole organism down to individual organs, to substructures and on down to the cellular and molecular levels. Understanding the relationships between these functional scales is, of course, crucial to modern medical education.

future directions