Reflections on Three Years of Network-Based Distance Education

David E. Bernholdt, Geoffrey C. Fox, Nancy J. McCracken,
Roman Markowski, and Marek Podgorny

Northeast Parallel Architectures Center
Syracuse University
Syracuse, NY 13244-4100

May 26, 2000

1 Introduction

Distance education promises to be an important tool for universities in the not too distant future. It provides a means to reach students who perhaps would not otherwise have an opportunity to take classes they desire. In the context of DoD's High Performance Computing Modernization Program (HPCMP), this shows up in several ways. For example, DoD scientists and engineers may be interested in formal coursework, and even degree programs, but their work may constrain them from traveling to the university offering the courses, especially if it is not nearby. Another example is the preparation of (traditional) students at the undergraduate and graduate levels. These students, who constitute the next generation of DoD scientists and engineers, may benefit from certain courses not available at their home institution. It may be that the institution lacks faculty expertise in a fast-moving area, where teaching is closely tied to research. Or it may be that faculty teaching loads do not allow sufficient time to keep pace with a rapidly advancing area. Both of these situations are particularly evident today in computing related areas, which are, of course, at the heart of the HPCMP.

As part of an effort to pilot the use of distance education technologies to benefit the HPCMP, the Northeast Parallel Architectures Center at Syracuse University (NPAC) has delivered six semester-long academic credit courses over the last three years to Jackson State University (JSU) and other institutions. These courses, in the general area of computational science, have been offered at both undergraduate and graduate levels, both independent

of and together with local Syracuse courses, and by a number of different instructors. All classes have been taught synchronously (in real-time) in traditional lecture format using the Tango Interactive network collaboration tool. Previous reports?? have described the first four classes in detail, including the environment (computers, network, etc.), the course material, and various "lessons learned". In this report, we will describe more briefly the most recent courses and focus primarily on observations of our experience over the last three years. The intent of this report is to try to capture some of the more practical issues pertaining to the running of a program of synchronous distance education in the hope that our experience can benefit others contemplating distance education activities.

2 1999-2000 Courses

As with our past distance courses, lectures were presented on a regular twice weekly schedule over the course of a normal academic semester. Syracuse-based instructors and teaching assistants presented all course material and graded all homework assignments and projects. We have also encouraged receiving institutions to provide a local faculty "mentor" to monitor course and students, and to be available to help answer questions. NPAC also provided computer resources and tools (in particular our Virtual Programming Lab?? in CPS615) for students to complete most of their homework assignments. Students received credit for the class from their local institution.

2.1 CPS 640: Topics in Networking and Multimedia Applications

This course was offered in the Fall term, by Drs. Roman Markowski and Marek Podgorny.?? A total of 60 students (37 from Syracuse) took the course for credit, with additional students auditing the class.

The primary objective of the course is to provide a comprehensive panorama of networked multimedia applications and the underlying computer and networking technology. We will discuss how existing and emerging data communication technologies can meet multimedia application requirements. The course covers LAN and WAN Technologies, Bridging, Switching, Routing, Networking Protocols, Management, Design and Security as well as Multicast, Videoconferencing, Multimedia Collaboration Technologies and Audio/Video compression and coding. The course material is designed as an introduction to the field and a practical guide for designing and planning networks for multimedia applications. Note that the word "topics" in the title means that the course content will vary to reflect current or interesting topics in the field.

The course material is divided into 12 main modules and the structure reflects the following sequence of questions:

- What is the current and future networking infrastructure?

- What are existing and future multimedia applications?
- What requirements does multimedia place on the underlying networks?
- What solutions are available to satisfy the requirements?

Course Modules:

1. Introduction: Syllabus, Projects, Basic Definitions, Traditional Network Design, Wiring, Repeaters, Bridges, Switches, Routers, Gateways, Standards, Open System Interconnection Reference Mode; Tango Interactive
2. LAN Technologies: Ethernet, Fast Ethernet, Giga Ethernet, ISO Ethernet, Token Ring, Token Bus, 100VG Any LAN, FDDI/CDDI, Wireless LAN/Mobile computing, ATM
3. WAN Technologies: Modems, HSSI, CSC/CDPD, T1/T3, ISDN, Switched 56K, xDSL, SDLC, X.25, Frame Relay, SMDS, ATM
4. Bridging, Switching, Routing: Transparent Bridging, Source Route Bridging, Translational Bridging, Source Route Translational Bridging LAN Switching, Tag Switching, Data-Link Switching, Cell Switching, Implementing LAN switching Routing: RIP, IGRP, OSI, OSPF, E-IGRP, BGP, CIDR
5. Networking Protocols: TCP/IP, AppleTalk, NetWare, Banyan VINES, XNS, DECnet, SNA, NetBEUI, NFS, SLIP/PPP, OSI
6. Network Design and Management: SNMP, VPN, Unix/PC Integration, VLAN
7. Network Security: TCP/IP Security Issues and Tools, Security Recommendations, Public Key Cryptography, Secure Electronic Messaging, SSL, SET, firewalls
8. Lookup Services: DNS, LDAP, JINI
9. Network Requirements for Multimedia: Integrated Services Model Multicast Architecture: MBONE, DVMRP, MOSPF, PIM
10. Quality of Service and Streaming Protocols: RTP, RTCP, RSVP, 802.1p/Q HTTP/RTSP
11. Multimedia Applications: HTTP, Videoconferencing, Video on Demand, Multimedia Electronic Mail, Hypervideo Documents Introduction to XML/SMIL, Media Indexing Collaboration Systems: Tango Interactive Overview of H.323
12. Internet/Internet 2/Nysernet 2000

2.2 CPS 615: Computational Science for Simulation Applications

This course was offered in the Spring term by Prof. Geoffrey Fox.?? A total of 14 students (6 from Syracuse) took the class for credit with additional students auditing the course.

This course was developed at Syracuse University in the College of Engineering and Computer Science as the graduate level introductory course in computational science for scientific simulations. It is designed to teach the basic tools from mathematics and computer science that are needed to give high performance computational solutions to scientific and engineering problems. The course outline is:

- Status of High Performance Computing and Computation HPCC nationally
- Application driving forces - Some Case Studies
  - Importance of algorithms, data and simulations
  - What is Computational Science Nationally and how does it with Information Technology
  - Systems approach: Problem Solving Environments
- Technology driving forces
  - Moore's law and exponentially increasing transistors
  - Dominance of Commodity Implementation
- Basic Principles of High Performance Systems
  - We have a good methodology but software that trails applications and technologies
  - Elementary discussion of Parallel Computing in Society and why it must obviously work for computers!
  - Overview of Sequential and Parallel Computer Architectures
    - Comparison of Architecture of World Wide Web and Parallel Systems (Clusters versus Integrated Systems)
  - What Features of Applications matter
    - Decomposition, Communication, Irregular, Dynamic ...
  - Issues of Scaling
  - What sort of software exists and Programming Paradigms
    - Data parallel, Message Passing
- Three Exemplars: Partial Differential Equations (PDE), Particle Dynamics, Matrix Problems

- Simple base version of first Example – Laplace's Equation
  - Illustrate parallel computing and lead to a
- General Discussion of Programming Models
  - SPMD and Message Passing (MPI) with Fortran, C and Java
  - Data Parallel (HPF, Fortran90) languages will be discussed later but NOT used
- Visualization is important but will not be discussed
- Return to First Example: Computational Fluid Dynamics and other PDE based Applications
  - Grid Generation
- Return to Parallel Architectures
  - Real Systems in more detail
  - Trends, Petaflops
- Second Exemplar: Particle Dynamics
  - Simple Example $O(N^2)$ Law – best possible parallel performance
  - Real Applications: Astrophysics and Green's functions for Earthquakes with Fast Multipole Solvers
- Third Exemplar: Matrix Problems
  - Matrix Multiplication is "too easy"
  - Linear solvers are demanding but use libraries
  - Sparse Solvers are most important in practice
- Advanced Software Systems
  - OpenMP and Threads
  - Parallel Compilers and Data Structures
  - Tools – Debuggers, Performance, Load Balancing
  - Parallel I/O
  - Problem Solving Environments
- Application Wrap-up: Integration of Ideas and the Future

3 Tango Interactive and the Technology Environment

As in the past, lectures were presented using the Tango Interactive network-based collaboration system??, developed at NPAC. It is implemented with standard Internet technologies and protocols, and runs inside an ordinary browser window. Tango is a general collaboration tool, which in this instance is used to synchronously deliver lectures and course material.

The primary Tango window is called the *Session Manager* (SM). From the SM, participants have access to many tools, including:

- Shared Browser, a special-purpose browser window that "pushes" web documents into remote client workstations;
- White Board, for interactive text and graphics display;
- BuenaVista, for audio-video conferencing;
- Chat tools for text-based interaction;
- VNC (Virtual Network Computer)??, for sharing displays.

At the beginning of each class, the instructor and the students at each recipient site all start up a copy of the Tango on their workstation by clicking a browser link. The instructor initiates a "session" of each tool that is to be used in the class and "joins" each student to the session. This starts up a window on the student's workstation. Course materials are generally presented as web pages using the shared browser, and can be produced directly in that form or (more often) are HTML exports of presentations written using Microsoft's PowerPoint. Java applets can be shared using the shared browser if they are made collaboration-aware or using tools like VNC if they are not. VNC can in fact be used to share arbitrary windows, so demonstrations need not be Java-based.

While previous classes had used Tango Interactive 1.4 (or earlier releases), for the Spring offering we used Tango 2.0. This release of Tango offers significant improvements in robustness and functionality, along with a significantly revised user interface for the SM application. Among other features, Tango 2 provides the possibility of using an interface tailored to educational applications, which launches a standard set of tools as a group rather than the instructor having to launch each individually. The size and placement of the resulting application windows in the education interface are setup with sensible defaults that minimize the need to drag and resize applications – sometimes a distraction for students.

Though we had previously had discussions with recipient sites about the platform required to receive Tango-cast courses, Spring 2000 was the first time we provided a detailed recommendation of a Tango platform??. Though it is not the only platform on which Tango will run, this has been in our experience the best:

- Name brand PC (do not by cheap clones)

  - 400 MHz Pentium III or faster
  - 128 MB RAM or more
  - 8 GB HDD or more (can be IDE or ATA)
- Good quality monitor (19" min) and appropriate graphics card capable of 1240x1024 True color recommended; 1152x864 True Color acceptable
- Video: Winnow PCI Videum Conference Pro $359
- Audio
  - Separate card unnecessary if Winnov card installed; otherwise SoundBlaster 64
  - External speakers are desirable in some circumstances and are quite inexpensive, so should be purchased
  - Headset + microphone: VXI Parrott 10-3 $59
- Ethernet 10/100
- CD-ROM drive
- Software
  - Windows NT 4.0 SP5 (upgradable to Windows2000 as soon as possible)
  - Microsoft Office
  - Textpad
  - WinZIP
  - FTP client
  - Netscape Communicator 4.7
  - Microsoft Internet Explorer 5
  - *Recommended:* FTP server to allow access to NPAC support staff. (MS FTP server from Web Peer Services is acceptable)

However one of the problems we have encountered in each and every course is that it is very hard to get platform (and other) information from recipient sites, so we cannot say if these guidelines were actually followed.

For the first time in our distance education work, lectures this year were also captured during delivery using the LecCorder system?? developed at NPAC. LecCorder combines the audio/video stream of the instructor delivering the lecture with the "clickstream" of changing slides and allows such presentations to be published easily on the web. Students can then review the lecture, viewing the slides synchronized with the audio and video streams.

Table 1: Recipient sites for NPAC distance courses, and their network connectivity. Some sites received the full course synchronously, while others, depending on student interest and local logistical considerations, chose to support only asynchronous use of the LecCorder recordings of lectures. Mississippi State University did not participate in the Spring 2000 class, "Commercial" connectivity indicates that traffic between Syracuse and the remote site traversed the "commercial internet" through multiple providers rather than a specialized network such as DREN or vBNS.

	Mode of Reception		Connectivity	
Site	Fall 1999	Spring 2000	Pre-vBNS	Post-vBNS
ARL MSRC, Aberdeen, MD	Asynch.	Asynch.	DREN	DREN
ASC MSRC, Dayton, OH	Asynch.	Asynch.	DREN	DREN
ERDC MSRC, Vicksburg, MS	Synch.	Synch.	DREN	DREN
Jackson State Univ., Jackson, MS	Synch.	Synch.	DREN	DREN
Mississippi State Univ., Starkville, MS	Synch.	–	Commercial	vBNS
Morgan State Univ., Baltimore, MD	Synch.	Synch.	Commercial	Commercial
NAVO MSRC, Stennis Space Center, MS	Synch.	Synch.	DREN	DREN
Naval Research Lab, Washington, DC	Synch.	Synch.	DREN	DREN

It is important to recognize that Tango is a network-based collaboration tool, and that as used in our courses, not only the courseware, but also the audio-video stream of the lecturer. Given the lack of quality of service (QoS) infrastructure on current networks, this can pose a particular challenge to this kind of distance education. Moreover, because users (students) are generally not familiar with the subtleties of networks and collaborations tools, such problems may be (incorrectly) attributed to the tools (i.e. Tango), or to the mode of use (i.e. distance education), possibly resulting in negative perceptions about the wrong aspects of the experiment.

In previous years, we were primarily delivering course to Jackson State University, to which we have had rather good network connectivity. With help from Syracuse's Internet Service Provider (ISP), we were able to engineer a routing which reached the high-speed Defense Research and Engineering Network (DREN) in a few hops on the "commercial internet". The classroom at JSU is serviced by a special T-1 line provided by the Engineer Research and Development Center (ERDC) which is accessible from DREN. This year represented a significant extension of delivery to include multiple sites, not all accessible via DREN. In addition, in the middle of the Spring 2000 term, Syracuse connected to VBNS, the academic ATM-based national network, which provided another high-speed path to some recipient institutions. Once Syracuse's vBNS connection was running, traffic destined for DREN sites was able to make use of vBNS instead of the commercial internet to reach the DREN interconnect. A summary of the recipient sites and their network connectivity to NPAC is presented in Table 1. Our experience has been that network performance and

quality has been improving, but does still occasionally cause problems.

4 Observations and Discussion

4.1 Organizational Issues

Each course we delivered required a concerted organizational effort beginning well in advance of the term. When our offerings were opened to institutions besides Jackson State University, we began by "recruiting" institutions to participate in the class. This required providing of information about both the course content and the other requirements for participation, and setting deadline for responses which would allow enough time to complete installation and testing at all sites. We have found that two months lead time is a practical minimum – one month to organize the participants and one month for everyone to install and test the necessary tools.

Most institutions we contacted did not have specific funding to support their participation in our classes, so they would be trading their local support costs (personnel, equipment, etc.) for a "free" course they might not otherwise have in their curriculum and the opportunity to participate in a technological experiment. Given that faculty and staff at most institutions are already quite busy, and especially so at minority serving institutions (MSIs), which were a major target of our recruiting efforts, it is not surprising that the ultimate response rate was fairly low. In some cases, we also encountered situations where interested institutions ended up not participating because of insurmountable technical obstacles (i.e. firewalls with uncooperative administrators), or lack of student interest in the particular course offering. A common problem at the MSRCs was the geographically distributed nature of the user base, and concerns about being able to provide proper support for Tango Interactive to remote users. This combined with low numbers of people expressing interest lead both ARL and ASC MSRCs to support only asynchronous access to the LecCorder recordings of the lectures.

In the classes we have offered, students have received course credit from their local institutions. This usually requires some effort in advance of the class, as the instructor (located in this case at Syracuse) and the course content must be approved by the recipient institution, and perhaps appointed to an adjunct faculty position. This does not normally pose a problem for qualified instructors, but does take time. A related problem involves institutional constraints. For example, we investigated the possibility of delivering academic courses to the HEAT Center, which is affiliated with the ARL MSRC and the primary location for its "on-site" training activities. The HEAT Center is also affiliated with the State of Maryland's university system and is prevented by its charter from offering academic classes from other institutions. It is possible that increasing use of distance education will lead to simpler methods of recognizing and certifying courses provided by other institutions, but for the time being it is merely another, possibly time-consuming hurdle that has to be taken into

consideration.

Scheduling is another organizational issue that can be challenging. Academic institutions normally have their own calendar and daily course schedule, and there is no guarantee that they will match across many institutions. The only practical response to this is a flexibility among the participants, allowing a mutually agreeable solution. It is worth noting that the typical week-long spring break together with variations in start dates can make the Spring term a particular challenge. Since our Spring 2000 offering was primarily a graduate-level course, we had the advantage that the majority of students at all sites were resident year-round at their institutions, and less likely to have problems attending classes on a schedule which did not precisely match their university's.

Finally, we have found that the involvement of a local faculty "mentor" at each site to be quite helpful. Such a person takes an academic responsibility for the class, and may be the instructor of record for the course at their institution. They provide the students with a familiar face, helpful in dealing with local issues and available to respond to student questions. The faculty mentor also provides a conduit to help the the remote instructor gauge student response to the material, and deal with issues pertaining to specific students (i.e. why a registered student might not be turning in homework assignments). We have found that faculty mentors can also serve their institutions by helping to bring the course they are mentoring into the local curriculum. For example, serving as faculty mentor to a Syracuse-delivered class helped a Jackson State professor prepare to deliver the same class locally and via distance education.

4.2 Delivery and Support Issues

Problems with delivery of a class have, in our experience, been fairly rare. But it is important to recognize the impact they can have on *perceptions* of a distance-delivered class even if their actual impact is minimal. Both for setup of the class, and routinely during delivery, we have found that having knowledgeable, well trained, and responsive technical staff supporting the class at all sites. We have found this to be a continual challenge, undoubtedly because synchronous distance education is a time- and resource-intensive activity and must compete with other demands on the staff member's time. Fortunately, we can point to few cases where support problems actually resulted in delivery problems with the class, but it has caused a significant level of anxiety and frustration here at NPAC, where we perhaps have more of a vested interest in the overall success of the endeavor.

In an ideal world, it would be possible to setup a distance education classroom at the beginning of a term and use it, exclusively for that purpose and without change to the hardware or software for the entire duration of the course. In this case, problems could only arise due to equipment failure, which is usually easy to diagnose, or external (i.e. network) trouble. Recognizing that this is impossible in practice, it is important to protect such installations from change to the maximum extent possible (for example, by bored students

fiddling with a computer when the room is used for some other class) and to coordinate systems hardware and software changes with the organizers of the class to help insure that (a) changes will not conflict with the tools being used for the class, and (b) appropriate testing can be done on the modified systems prior to use in class.

Because installations do change, and because external networking and other factors can effect course delivery, we also strongly recommend that preparations for each lecture begin thirty minutes in advance. This gives time for all participants to setup the collaborative session and verify that everything is functioning normally. If problems are discovered, it also provides time to troubleshoot them and (in the case of network problems) try to contact the appropriate carrier to report the problem. If the same machines and configurations are used routinely for a class, we've found that thirty minutes is often sufficient to fix or work around minor problems that do crop up.

From the instructor's viewpoint, actual delivery of the lectures can be rather more "busy" than delivery of a traditional face-to-face lecture. In addition to actually lecturing, the collaborative system must be monitored for problems, and the tools by which remote students can provide feedback also need to be monitored (chat, "raise hand" tool, two-way audio if available). During the 1999–2000 courses, we also captured lectures using LecCorder, which is not well integrated with the Tango Interactive framework, and so added to the complexity. From the beginning, we have found it quite helpful for the instructor to have an assistant who can attend to some of these details. We have not always operated this way, but with the addition of the LecCorder we found it to be a necessity. Because of the numerous tools being used at once, space on the computer screen can become precious. The Tango Interactive 2.0 education-specific interface helps manage screen real estate, but there are still advantages to using multiple displays (which may run on separate computers).

There are also a number of technical aspects of actual class delivery that bear mention. First, and perhaps most important, is the importance of the audio channel this kind of collaborative application. The audio setup requires the most attention, and can be the greatest single point of difference between a face-to-face lecture and a distance lecture. Specifically, we are used to two-way audio channel, allowing students to ask the instructor questions. In our distance education work so far, we have had to limit it to a one-way channel from the instructor to the student, using other means (chat tool) for the "back channel". In principle, a "push to talk" arrangement is possible for an audio back channel, but the practical implementation can be quite problematic because of issues about mic location (feedback concerns), switching controls, etc. Providing headsets (with boom mics) for each student would also allow for two-way audio but unless used in "push-to-talk" mode would likely distract the instructor with background noise. The other problem with this solution is that without using IP multicast (network support for which is often lacking) or other techniques, this requires one audio stream per student instead of one per site, which can be a drastic increase in bandwidth requirements. Full "open room" two-way audio should be possible with the addition of echo cancellation hardware at all participating sites, but this is a demanding ap-

plication (due to delays common in network-based audio transmission) and quality hardware is expensive (each unit costs several thousand dollars). This is not a solution we've had an opportunity to experiment with.

Secondly, depending on the size of the class, downloading of course materials (lecture slides) can be taxing on the network and the web server that provides them. The shared browser used to display the slides is usually run on a per-computer basis (typically one or two students per computer), and since all browsers try to access the new material at essentially the same instant, the result is a burst of requests, of perhaps 20–40 kB each (for GIFs exported from PowerPoint). In a couple of cases, where inadequate web servers (software or hardware) were used this has posed problems, but by and large the growing capacity of networks and computers have minimized this problem. Nevertheless, it bears consideration, especially for larger classes. Another way to address this problem would be for participant sites to run local web caches, so that one request from each site would be transmitted to the actual web server, while the others would be answered from the local cache. It is even possible to pre-load the cache by stepping through the lecture slides in advance of the class. We have experimented with this approach at JSU and found it quite effective, but we have not had enough problems so far to warrant the effort required to put it into routine use.

One interesting feature of the Spring 2000 class is that it was the first course to be delivered from multiple sites. In addition to Syracuse, a number of lectures were delivered from Florida State University and one from ERDC MSRC. Apart from making sure the infrastructure (Tango Interactive, etc.) were properly setup and tested (more of a concern when a site is used only occasionally), these classes went smoothly. This is a particular advantage of distance education compared to the traditional classroom. With the growth of personal digital devices and wireless network connectivity, it is possible to imagine that in the near future, the instructor can become even more mobile. Using a cellphone for the audio channel, an electronic book, palmtop or laptop computer to hold a local copy of the courseware (assuming wireless bandwidth is insufficient to support direct access to the content on the web server used by the students), and a wireless device interfaced to the collaboration system to control the shared browser, lectures could in principle be delivered from any location with wireless connectivity.

4.3 Courseware and Tools Issues

The course materials for our computational science classes have been developed and refined in the course of use in Syracuse classes or the instructors' external presentations. These materials have been produced in Microsoft PowerPoint, and HTML exports are used with the Tango shared web browser as the primary tool to deliver the courseware in the distance classes. The shared browser represents a particularly powerful tool because of its generality and flexibility. Many different document tools (wordprocessors, presentation tools, etc.) can export web-compatible versions of the documents. Documents may be static or dynamic.

Dynamic document adds a new level of complexity to the collaboration tool designed to share them, but we have demonstrated the ability to share documents with embedded Java applets and also forms and similar HTML-based content.

The sharing of Java applets for purposes of demonstration raises some interesting issues regarding their interactions with the underlying collaboration framework used to share them. Collaboration systems typically use some combination of "shared event" and "shared display" models. The shared display approach involves actually sharing the rendered screen image (or a portion thereof), while in the shared event case GUI events (button clicks, widget updates, etc.) are intercepted within the application being shared and transmitted to the remote instances. Shared event systems tend to give the greatest network efficiency and give the application tremendous flexibility in "how" it collaborates, but at the expense of the application having to be "collaboration aware" in order to trap events and send and receive them. Shared display systems tend to be more resource intensive, but have the advantage of simplicity since the application is oblivious to the fact that it is being shared.

Tango Interactive provides a shared event model of collaboration which is in large part responsible for the power and flexibility of the system. Thus, Tango clients must be collaboration aware, and share events. Beca created TangoBeans?? to make it easier to produce collaborative Java applets following the JavaBeans model. However few of the existing Tango clients, can be described as "compute intensive", whereas one might expect that many of the codes one might want to use as teaching tools in computational science fall into this category. The sharing of compute intensive applications poses additional problems with respect to synchronization in the shared event model. The tremendous range of performance in Java implementations across different platforms means that the pace of a compute-intensive application can vary widely and must be governed by the slowest client. Late-comers pose a similar problem: to synchronize with the others, the new client must either step through the simulation until it reaches the same timestep as the existing clients, or it must obtain the complete current state of the simulation (which may be quite large) from one of the other clients. In adapting an existing computational fluid dynamics demonstration applet for use with Tango, Elmohamed, Kim and Fox?? developed a general algorithm for the event-based synchronization of time stepped simulations which insures that all clients remain synchronized to within both a give temporal interval and a given number of simulation steps, on the assumption that it is rarely necessary to run the simulations in lock step, but it is important that the instructor be able to stop and start them reliably with all clients at the same point. Late-comers are handled by having the new client step through the computational part of the simulation without displaying the results until it reaches the appropriate time step. This is a simple solution which takes advantage of the fact that in this case the display is more compute-intensive than the simulation itself. It is possible that very late entrants will never catch up, but it does handle the majority of cases.

Our experience with the CFD applet lead us to look for a simpler solution to develop other demonstration applets. We realized that in the typical educational setting, such demonstra-

tions are controlled exclusively by the instructor, while students may explore the demonstration in more detail independently outside of the class session. In this situation, the applet does not really need to be collaboration aware, and a shared display model suffices. Therefore, when developing new applets to illustrate the parallel implementation of a Laplace Equation solver?? and a matrix multiply algorithm??, we chose not to make them collaboration aware. We were able to share them using the VNC shared-display tool, recently integrated as a Tango client, however VNC itself was designed for this particular application and scalability problems within the package prevented routine use within our CPS615 class. We also experimented with variant of the shared display approach which might be termed "shared canvas". A display toolkit, such as the Java SWING classes are made collaborative in a shared-event fashion. Then an arbitrary applet which uses the SWING classes for display can easily be adapted to use the collaboratized version, changing at most a few lines of the original applet??. In essence, only the display portion of the applet, the "canvas", knows that it is being used in a collaborative environment. This is a promising approach which merits further investigation. In general it seems clear that both shared event and shared display models have their place in collaboration tools for use in education and in more general technical collaboration.

One might easily imagine that an electronic whiteboard would be an extremely useful tool in a distance education situation, but it is interesting to note that over the course of our distance education work, this tool has been used very rarely. The Tango Interactive whiteboard client includes the ability to load "background images", and so could in principle be used to display lecture slides and provide the ability for the instructor to annotate them in real time. It could also be used separately from the usual shared browser display of the lecture slides as a place to make special notes or expand on certain points or questions. Why this capability has not been exploited deserves deeper study with a larger group of instructors (NPAC's instructors have been in the habit of presenting lectures "online" directly in PowerPoint for some time, so it may be that others with different habits would take a different approach). It may also be that the whiteboard is too clumsy for practical use in class. As anyone who has tried freehand drawing or writing with a computer mouse knows, the results are often far from satisfactory. One way to investigate this possibility would be replace the mouse on the instructor's workstation with an input tablet, which allows the use of a pen-like object as the pointer, and which are now available for PCs in various sizes at reasonable prices.

As noted above, NPAC provided students with the computing facilities they needed to complete their coursework. The Virtual Programming Lab (VPL)?? has been used in several classes, both local to Syracuse and distance offerings to help provide students a simpler web-based environment in which to complete their programming assignments. VPL might be seen as an early example of a "computing portal"?? and provides for the editing and compilation of source code, and parallel execution on a workstation cluster. By constraining the execution environment, VPL can actually help prevent problems which would impact other users of the

system, or at least make them easier to identify and fix. It also makes it easier for the student to complete their work without worrying about the details of the computer and operating system they are using. It has proven very effective so far, and with the rise of the "portal" concept in computing, may well foreshadow the future of both distance and traditional computer education. One can also imagine incorporating collaborative capabilities into such a portal, so that it could be used as part of demonstrations or between an instructor or teaching assistant and a student during remote "office hours".

4.4 Human Factors

Since distance education uses a computer (or other electronics) to mediate the interactions of teachers and students, the human element, and human-computer interactions, clearly play an important role in the acceptance and perception of distance education. We have seen this manifest in many ways over the past three years, and raises many interesting questions which we have not been able to examine in the detail we would like due to constraints on time and resources.

One general issue is the question of how students perceive the course, and how they compare it to the equivalent face-to-face experience, both consciously and unconsciously. In formal course evaluations, students tend to rank the distance aspects of the class as slightly worse than a traditional classroom setting. This is not surprising, given that this work has been at the forefront of synchronous network-based distance education. On the other hand, students are often getting course material which would otherwise not be available at the institution, which can be quite valuable. For example, JSU students who took our early courses focusing on current web technologies (Java, databases, HTML/XML, etc.) reported much greater interest from prospective employers compared to their classmates who had not had this course.

A less direct measure of student reaction to the distance environment is the level of interactivity in the class. Interestingly, NPAC instructors have seen the level of interactivity decline over the duration of our experiments, until the most recent class, in which many lectures would pass with no questions from the remote audience, or even the local audience. The reasons for this trend are unclear. While we have not been able to provide two-way audio, as students would be accustomed to in a face-to-face setting, Tango has always provided a chat tool, which allows students to type a question, and which in the past, and in some of our distance training experiments, has been quite effective. Discussions with faculty mentors at recipient institutions yield several interesting insights for consideration. At one location, the faculty mentor reported that he was answering many student questions rather than conveying them to the instructor. At another site, we were told that the students were "intimidated" by the chat tool, thinking that it was more of an interruption to the instructor than a raised hand or verbal interruption would be in a traditional classroom. This is particularly interesting because as instructors and tool developers, we have generally considered the chat

tool to be *less* intrusive than a two-way audio exchange – the instructor can see the question and respond to it at the most appropriate point in their lecture with less break in the flow of the lecture than if the question had to be asked verbally.

This decline in interactivity as occurred at the same time we have tried to *increase* the variety of tools available to students to provide feedback to the instructor. The Tango "raise hand" applet is intended to serve as the electronic analog of a student raising their hand to get the instructor's attention. This tool has has been available since the beginning of our distance education work, and yet has been essentially unused by students. More recently a more sophisticated feedback tool has been introduced, which allows students to indicate if the instructor is going too slow or too fast, if there are technical problems, etc., with the results aggregated onto a compact display for the instructor. Once again, this tool is essentially unused during actual classes.

Another interesting comment we received from the faculty mentors pertains to this year's use of LecCorder to capture lectures for later asynchronous review. Students discovered the value of this capability during the Fall 1999 class, and when considering whether to participate in the Spring 2000 class, many cited the continued use of LecCorder as a deciding factor. Feedback from others who used *only* the LecCorder recordings (at locations where participation in the synchronous deliver was not supported) was also very positive. It is easy to believe that students would find the ability to review a lecture (or portions thereof) in their own time valuable, but we were surprised that it would be a deciding factor for some. Obviously this capability is rarely available in the traditional face-to-face context, so it is hard to make comparisons. It is possible that the availability of LecCorder recordings makes up for other (perceived or real) disadvantages of the distance presentation, but this also may suggest that the use of lecture capture tools in traditional classrooms may have significant value to students.

The use of video of the instructor and of classrooms is directly related to the comfort level of the participants. The audio channel and the courseware presented in the shared browser essentially provide the entirety of the educational content to the class, so that video of the instructor (a "talking head") adds very little. We have tried on a number of occasions to transmit video of the instructor to remote sites, though for technical reasons, we usually have not done so (on the theory that we're likely to get better quality of service for the audio stream if we put less of a load on the network links). We've gotten mixed response to these unscientific experiments. Most people agree the the instructor video is not essential, but at least for some students, it does seem to increase the sense of connection and human contact with the instructor. On the other hand, some report that the jerky motion due to the low frame rate, and the generally low quality of the image due to size, lighting, and other factors, actually make is a distraction. One approach to this dichotomy might be to deliver video to each student computer instead of a central screen (we have recommended a one-instance-per-site video configuration, just like audio, in order to minimize network load) and let each student decide whether they wanted to display the instructor video.

Conveying video of the remote classrooms back to the instructor is something of a different matter, which we have been unable to experiment with to any real extent. To be useful, this requires mounting of one or more cameras in relatively inaccessible places in a classroom and both hardware and software setup on the computer(s) hosting the camera(s). Given the challenges we faced in obtaining the basic support needed to run the class at remote sites, it seemed to us inadvisable to undertake this kind of experiment. It is easy to imagine the possible utility of giving the instructor a view of the classroom. It would allow them to better gauge the reaction of the students, to facilitate "normal" communications by identifying a raised hand, etc. On the other hand the size and quality of typical computer-based video is typically not high, and it raises questions of how much information can actually be conveyed without going to the complexity and expense of multiple cameras. There is also a minor issue of screen real estate for the instructor. One can easily imagine an entire screen full of two or three video images from each of six or eight participating sites. While this is not technologically impossible, it does increase the complexity of the instructor's environment and would undoubtedly require more of their attention to make use of than looking out on a local classroom with an equivalent number of people in it.

5 Summary

NPAC has now delivered six semester-length academic credit classes to a total of eight different organizations (academic and DoD) affiliated with the DoD's High Performance Computing Modernization Program. These classes used the traditional lecture format, and were shared with participating institutions in real time using computer-based collaboration tools. In this period, we have demonstrated success in using state of the art technology to bring classes that would otherwise not be available to students at the undergraduate and graduate levels, as well as to professionals in the DoD workforce seeking to continue their education. These same techniques are also being used within the HPCMP's PET Training program.

Our experience clearly shows that at the present time, synchronous distance education requires a significant investment of time and effort on the part of both the provider and recipient sites. (Outfitting a distance education classroom with the necessary computer equipment is a non-trivial cost, but at most universities and ever-increasing number of suitable rooms are being setup to support on-campus activities.) Depending on the instructor there may be some additional time commitment required to put lecture materials into a web-accessible form, but the bulk of the extra time is associated with the "distance" rather than the "education" aspect of distance education. It is important the participants in distance education allow for both the organizational and support costs. As hardware, software, and network support for multimedia applications (i.e. collaboration tools) becomes more robust and more common place, we can expect that support costs will eventually decrease. It is possible that

organizational matters will also become easier as distance education increases (i.e. university consortia synchronizing their schedules and calendars to simplify coordination), but it may be quite some time before anything like this happens.

Collaboration tools such as Tango Interactive provide a rich environment in which to provide distance education, though in fact only a portion of Tango's capabilities were actually used in these classes. It is clear that collaboration technology is more than sufficient for distance education and has improved steadily over the course of our work (however the external environment – networking, audio and video hardware and software, etc. can still be troublesome). A rich environment like Tango allows one to imagine many ways in which it might be possible to add value over and above the simple analog of a lecture delivery, though it should be done carefully, as we have also observed that students do not avail themselves of all technological tools available to them (i.e. the Tango "raise hand" and feedback applets). Alternatively, from a reductionist point of view, our experience to date indicates that one can provide effective lecture-style distance education with a relatively simple tool set (basically audio, shared browser, and chat) if desired. One interesting new observation is that although the shared event model of Tango provides it with its richness of capability, shared display capabilities (which are not mutually exclusive) are valuable for the simplicity they bring to the development of certain kinds of educational courseware.

Perhaps one of the most important aspects of distance education is to understand better how the computer-mediated interaction effects the students and instructors. We have made a number of observations in this area based on our experience over the last three years. Having demonstrated the technological ability to deliver distance education, These observations, and other "human factors" issues will perhaps be the most important matters to pursue in the next phase of the development of distance education.

6 Acknowledgments

This work would not have been possible without the many people in NPAC and our collaborating institutions who worked very hard to insure the success of these courses. NPAC's Tango Group deserves particular mention in this respect.

This work has been funded wholly or in part by the DoD High Performance Computing Modernization Program ERDC and NAVO Major Shared Resource Centers through Programming Environment and Training (PET). Supported by Contract Numbers: DAHC 94-96-C-0002 with Nichols Research Corp. and DAHC 94-96-C-0008 with Logicon Inc.