
CEWES MSRC/PET TR/99-21

# Experiences with Using TANGO Interactive in a Distributed Workshop

by

Troy Baer
David Ennis
James Giuliani
Leslie Southern
David E. Bernholdt

04h02099

**Work funded by the DoD High Performance Computing Modernization Program CEWES Major Shared Resource Center through**

Programming Environment and Training (PET)

Supported by Contract Number: DAHC94-96-C0002
Nichols Research Corporation

# Experiences with Using TANGO Interactive in a Distributed Workshop

Troy Baer, David Ennis, James Giuliani, and Leslie Southern
*Ohio Supercomputer Center*
*1224 Kinnear Road*
*Columbus, Ohio 43212*

David E. Bernholdt*
*Northeast Parallel Architectures Center*
*Syracuse University*
*111 College Place*
*Syracuse, NY 13244-4100*

** Author for correspondence. E-mail bernhold@npac.syr.edu*

## Abstract

Together, the Ohio Supercomputer Center (OSC) and the Northeast Parallel Applications Center (NPAC) at Syracuse University delivered high-performance computing (HPC) training courses to a geographically distributed Department of Defense (DoD) user community. In September of 1998, the Ohio Supercomputer Center (OSC) delivered a one day offering on the Fortran 90 programming language from the CEWES MSRC in Vicksburg, Mississippi. In January of 1999, the OSC delivered a two-day offering on OpenMP from the CEWES MSRC. The TANGO Interactive collaborative software was used to deliver these courses simultaneously to participants at DoD HPC Modernization Program Major Shared Resource Centers and Distributed Centers. This report describes these prototype distance HPC courses, our experiences, and provides instructor and student guidelines.

---

## Table of Contents

---

# Introduction and Background

One of the important missions of the DoD HPC Modernization Programming Environment and Training (PET) program is to provide user training to assist with the transition to scalable parallel computers and related leading-edge technologies. However the fact that the Modernization Program's user community is distributed over such a large geographical area relative to the four Major Shared Resource Centers at which the PET program is based poses a real challenge to the development of a training program that is both effective and convenient for users and instructors. "Distance learning" is one obvious response to this challenge. This paper describes a project aimed at transferring the tools and experience gained from a successful academic distance education effort into the more compressed, and in some ways more demanding environment of short-duration training, as traditionally used in the PET program.

Since the fall of 1997, the Northeast Parallel Architectures Center (NPAC) at Syracuse University has been working with the Computer Science Department at Jackson State University (JSU) to deliver semester-long academic credit courses to JSU students over the Internet, using the TANGO Interactive collaborative framework, developed at NPAC. This work, which has been described in more detail elsewhere [ 1 , 2 ] has been quite successful at several different levels. It has been used to successfully deliver four classes (as of Spring 1999) covering material in large part unavailable locally to the participating students. At a more fundamental level, it has served as an excellent venue to explore the social dynamics of distance education -- teaching styles, student and instructor interaction with the delivery tools -- and the technical demands. In addition to new and improved tools to facilitate distance learning, the distance education effort has also driven significant improvements in the stability and robustness of the core TANGO Interactive system. By the fall of 1998, we judged that TANGO, and our distance education experience, had reached a sufficient level that it made sense to begin the staged transfer of this technology into the PET training program.

The training environment differs from a regular academic environment in several important ways. Most importantly, the timescale is significantly compressed -- training typically runs one to three consecutive days six to eight hours a day rather than two to three days a week (typically three hours) for twelve or more weeks. This means that technical problems with the delivery tools or the network can be much more disruptive in a training, placing a greater demand on the tools and the support staff. Connected to this is the fact that attendees who experience technical problems with the class are less likely to return for another class using the same type of delivery. In the interest of having things run as smoothly as possible, we made the conscious decision that in this early stage we should deliver trainings only to training rooms with support staff with some knowledge of TANGO present locally. The ultimate goal of "direct to the desktop" delivery to individual researchers was deferred to future stages of the distance training deployment.

Also important to the use of TANGO in training was the fact that so far, the instructors for all of the distance education work to date had been NPAC staff who were thoroughly familiar with TANGO. In practical use in training, however, a great many different instructors, perhaps with little or no prior experience in distance learning, would have to develop and adapt courseware to this new medium and deliver it. It was primarily for this reason that the Ohio Supercomputer Center (OSC) was involved in this project. OSC has over ten years experience with high-performance computing, and as such, continues to build a repertoire of quality training offerings for the high-performance computing user. At the beginning of this project, one of the two OSC instructors involved attended a two-day training on the installation, operation, and support of TANGO, while the other had no formal training with TANGO at all.

In the course of this project, two distance trainings were delivered. The first, in September 1998 was a one-day, lecture-only class on a topic familiar to the instructor, Fortran 90, presented to 24 participants. i The second, delivered to 34 participants and presented in January 1999 was a two-day class on OpenMP on the SGI Origin2000, which included a hands-on laboratory component and was newly developed, so that the instructor had not taught it before either in normal or distance fashion. The first class was also more conservative in that the number of remote sites was intentionally limited to two -- the ARL MSRC and OSC's training room -- while the second class was offered to all four MSRC training rooms, as well as to the Naval Research Lab in DC, and OSC. Both were successful in delivering the class to a geographically distributed audience located at multiple sites. At a more fundamental level, they were also successful as a framework in which to examine the issues around a more routine use of TANGO Interactive for the delivery of remote training. The remainder of this paper describes in more detail the training conducted in this project, critically assess the results, and discusses prospects for expanded use of interactive distance training in the PET program.

# Course Materials

## Fortran 90

The first training class, presented in September 1998, was on the Fortran 90 programming language. In particular, the course is designed to present to users familiar with the Fortran 77 language standard the new features available in the Fortran 90 standard, such as dynamic memory allocation, derived data types, modules, and array syntax. The course has been taught several times by instructors at OSC, typically in an intensive one-day lecture format.

The notes used in this Fortran 90 course were originally written by Dr. David Ennis of OSC using Frame Maker. They were then converted to HTML by Victoria Sauber and placed on OSC's technical information server, and in this case they were presented by Mr. Troy Baer. These notes are split into sections dealing with specific Fortran 90 topics:

1. Preliminaries
2. Source, Types, and Control Structures
3. Procedures and Modules
4. Array Processing
5. Pointers
6. New I/O Features
7. Intrinsic Functions

The course is taught as a series of lectures on the topics listed above. In "local" presentations, the notes usually are projected on a screen at the front of the classroom, using either an overhead projector and transparencies made from the Frame Maker source, or a projector attached to an instructor computer station which runs a Web browser to display the HTML version of the notes. The TANGO-based delivery used a shared web browser to present the HTML notes, and as is normally done in OSC's local presentations of the class, attendees are given the option of following along using either a hard copy of the notes or the online HTML notes using a Web browser on the classroom workstations.

### Parallel Programming on the Origin 2000 Using OpenMP

In January 1999, the second training entitled *Parallel Programming on the Origin2000 Using OpenMP* was presented. This course targets the experienced programmer who is interested in the Origin 2000 architecture and the OpenMP application program interface (API). OpenMP is a portable, scalable model that gives shared-memory parallel programmers a simple and flexible interface for developing parallel applications.

The materials used in this course were developed using Microsoft PowerPoint by Dr. David Ennis of OSC and first delivered in conjunction with this project. They include the following sections:.

1. Introduction
2. MIPS R10000 Processor
3. Origin 2000 Architecture
4. Processes and Load Sharing
5. OpenMP Parallel Programming
6. Data Distribution Directives
7. Automatic Parallelization

After discussions of the processor, architecture, and creation and scheduling of parallel threads, the OpenMP directives were presented in detail along with examples of their use. The course concluded with an equally important topic of how to distribute the data used by parallelized OpenMP regions among the local memories on the Origin 2000.

The notes were projected on a screen at the front of the classroom, using a projector attached to an instructor computer station which runs a Web browser to display the HTML version of the notes. Local and remote attendees were given the option of following along using either a hard copy of the notes or the online HTML notes using the Tango Interactive shared web browser on the classroom workstations.

# TANGO Interactive

The software infrastructure used to deliver this course content simultaneously to several sites over the Internet is TANGO Interactive, from the Northeast Parallel Architectures Center (NPAC) at Syracuse University. This software consists of a plugin for the Netscape Navigator Web browser and several Java, Javascript, and native (i.e. compiled C++) collaborative applications. The TANGO suite allows applications to be shared between users in a master/slaves or peering collaboration.

To use TANGO, the user first connects to a web page containing a Javascript applet which brings up a TANGO startup window (shown in Figure 1). This allows the user to select the interface (the set of

shared applications available) as well as the TANGO collaboration server to which to connect. This TANGO server need not reside on the same system as a Web server, although it is often convenient to collocate the two. Once the user selects a collaboration server and interface, he or she is presented with a login window (shown in Figure 2) where they can enter an identifier for themselves. Once an identifier is entered, the user is presented with the TANGO control application (shown in Figure 3), which is used to initiate sessions of shared applications or to connect to already existing sessions.

There were four TANGO applications used in the two distance trainings: shared browser, Buena Vista, chat, and whiteboard. The first of these, the shared browser, is a normal Netscape Navigator window, except that it tracks the current URL in the shared browser window of the user who is the master of the shared browser session. When the master user's shared browser moves to a new URL, the shared browser of all the other users in that session also move to the new URL. The shared browser is designed so that only the URL is shared through the TANGO framework, and each browser loads the URL as if it had been typed in locally. This makes it possible for WWW caches at recipient sites to reduce the bandwidth consumed over the long-haul networks.

The second TANGO application used for these classes was the Buena Vista audio/video conferencing software. Buena Vista is a two-way streaming audio and video conferencing tool which can be run in either full or half duplex ("click to talk"), though in most cases so far, the half-duplex mode has been used based on practical site considerations, described in more detail below. Versions of Buena Vista are available for Intel based PCs as well as SGI workstations, and the two versions are completely interoperable. For the Fortran 90 workshop, Buena Vista was used to deliver a "talking head" consisting of audio and video streams of the instructor's lectures to the remote sites, as well as the facilitation of discussions between the instructor and remote attendees. During the OpenMP training, video was not used.

The third TANGO application was the chat client (shown in Figure 4). This is a fairly conventional chat program, similar to many IRC (Internet Relay Chat) clients. The chat tool provides a mechanism for students at remote sites to pose questions without worrying about audio configurations, and for support personnel to discuss and handle technical problems with minimal intrusion on the class itself (using a separate instance of *chat* from the ones the students and instructor use).

The final TANGO application used in these classes was the whiteboard (shown in Figure 5). The whiteboard is a shared vector drawing package, which was used to display diagrams and example code fragments.

For the OpenMP class, an additional TANGO application was used experimentally. This was a shared telnet session, designed to allow instructors to share a live telnet session with students for demonstration purposes during a lecture, or during a hands-on laboratory session to allow students to show the instructor problems they are having when asking for assistance. The tool was a prototype, created by interfacing an existing Java telnet/terminal emulator to the TANGO system. For security purposes, it was designed so that keyboard input was accepted only from the master instance of the tool, with the others displaying only the characters returned from the remote host. In a hands-on lab setting, the idea was that students would each have their own separate (not shared) telnet session, and when they needed to consult with the instructor they would use TANGO to launch a copy of their session to the instructor's workstation. The tool was used briefly by the instructor during a demonstration; due to time constraints, there was insufficient time to try to use it in the hands-on laboratory environment.

Except for Buena Vista, the TANGO applications used require minimal bandwidth, on the order of 10 kB/s or less. Buena Vista, because it does synchronous audio and video streaming, requires considerably more bandwidth depending on the quality of the audio and video streams generated. The TANGO developers recommend that for best results, networks should be able to sustain 100 kB/s with minimal latency or jitter. As will be seen below, this can be difficult to achieve over the "commercial" Internet backbone, but is generally available on the DREN (Defense Research and Engineering Network), to which most would-be PET training recipients are connected.

# Classroom and Network Configuration

Both classes were presented by OSC instructors in the Training and Education Facility (TEF), located in the Information Technology Laboratory (ITL) at the Army Corps of Engineers Waterworks Experiment Station (CEWES) Major Shared Resource Center (MSRC) in Vicksburg, Mississippi. This classroom is outfitted with SGI Indy workstations at each of the student desks and an SGI O2 workstation at the instructor desk. The O2 is connected to a projection system which projects an image of the O2's display onto a screen at the front of the classroom; it is also equipped with a digital video camera and a headset microphone, as well as a PA system and analog video cameras for capturing classes on videotape.

The second classroom for these classes was the "Fishbowl", located at the Ohio Supercomputer Center (OSC) in Columbus, Ohio. This glass-walled classroom is configured similarly to the TEF at CEWES, except that the student workstations are mostly SGI O2s rather than Indies. There is also no analog video camera in the Fishbowl.

During the first training, a Windows95 workstation at the Army Research Laboratory (ARL) MSRC served as the other remote recipient. For the second training, ARL made use of their training room, located at the HEAT Center, which is equipped with SGI O2 workstations.

Other participants in the second training were the Aeronautical Systems Center (ASC) MSRC at Wright-Patterson Air Force Base (WPAFB) in Dayton, Ohio, the Naval Oceanographic Office (NAVO) MSRC at the Stennis Space Center in Mississippi, and the Naval Research Laboratory (NRL) in the District of Columbia. ASC and NAVO used SGI O2 workstations in their training rooms, while NRL used two WindowsNT systems. Both classes were also monitored at NPAC for support purposes.

All four MSRCs and NRL are connected to the Defense Research and Engineering Network (DREN) backbone, shown in Figure 6. Network traffic between NPAC or OSC and the DREN sites was carried partly over the commercial Internet. Through its work with Jackson State University, NPAC has been able to work with Syracuse University's Internet service provider and DREN staff to obtain an efficient routing of network traffic from NPAC to the DREN (JSU's distance education classroom is connected by a T-1 link to the CEWES MSRC, and is part of the DREN). The distance training made use of the same routing, so network performance to NPAC was not a problem. However OSC, which did not have any historical special arrangements, suffered from some degradation of audio performance during the late afternoon, when the Internet is most heavily loaded.

## Preparation and Support

Part of the original design of this project was the use of centralized training facilities with local support staff knowledgeable in the computer hardware and software configurations available. Our concern was

two-fold: first, proper setup of TANGO, particularly the audio/video conferencing system (Buena Vista), while not terribly complicated does require some attention; second, we felt it was important to have someone at each site with at least a basic familiarity with TANGO and the distance training setup to assist students and handle problems that might crop up during course delivery.

Familiarity with TANGO was developed in two ways. Although it was not required, some of the local support staff (from ARL, CEWES, NAVO, and OSC) were able to attend training offered by NPAC staff on the installation, operation, and support of TANGO. Before each of the two trainings, we also undertook a number of test sessions involving site support staff and the instructor with the dual purpose of increasing the experience and familiarity of everyone involved with the distance delivery setup and the particular tools and techniques that would be used to deliver the class, and to try to identify in advance any potential problems with the hardware and software configurations at each of the sites. Two or three advance tests were held before each training, each test lasting from 30 to about 90 minutes. Typically the first test was designed to get all of the sites connected with at least one workstation and to run through the specific tools that the instructor planned to use during the class to insure that they were operational at all sites (we experienced several problems due to the use of different versions of Netscape Navigator at different sites which were identified during these tests, and handled by reverting to a simpler version of the tool in question). This gave all sites a working installation, which they could then replicate on the rest of their computers. The second test was originally intended to allow sites to double check the configuration of the remainder of their systems, though in fact it was not really used in this way. Primarily, it served as an opportunity to develop more experience among the sites involved.

# Problems Encountered

As might be expected, a number of problems occurred during the delivery of the two classes. These problems were annoying and/or temporarily disruptive but were fixed or worked around so that they did not jeopardize the overall class. We describe some of the problems here in order to provide a feeling for what instructors and support staff might be called upon to address on short notice.

## Fortran 90

The first problem, seen about an hour before the workshop started, was that NPAC's web server (`trurl.npac.syr.edu`) was down during part of the morning of the workshop. This caused problems with OSC's TANGO installation, which expected to be able to download the TANGO startup web page from the NPAC web server. Luckily, this necessary web page was replicated on the CEWES TANGO server (`tango.wes.hpc.mil`), and the OSC installation of TANGO was temporarily modified to use the CEWES version of the page instead of the NPAC version.

An ongoing problem throughout the day was the network congestion seen between OSC and CEWES. This became progressively worse as the day wore on and the public Internet backbones became saturated from cross-continental network traffic. This caused the downloading of applets and web pages to slow to as low as 3 kB/s. It also created problems with Buena Vista such that no one at OSC who joined the workshop's Buena Vista session after 10 a. m. (EDT) was able to get both audio and video streams to work consistently. OSC had multiple Buena Vista clients running, so in this instance, transmissions in a multicast manner would have been preferred. To alleviate this problem, only one Buena Vista client was used during the second course.

Another Buena Vista problem, unrelated to network performance, was that the SGI version of Buena Vista initially overrides the current audio settings of the workstation and sets the output/speaker level to between 90 and 100 percent of maximum. This can cause discomfort or potentially even hearing damage for users wearing headphones. It can also lead to acoustic feedback problems on systems where a microphone is set to transmit.

## Parallel Programming on the Origin 2000 using OpenMP

One problem encountered during the testing phase in advance of the second training was that different sites were running different versions of Netscape Navigator. Certain of the recently developed TANGO applications ("interactive chat" and "object-oriented whiteboard") have problems due to bugs in some releases of Navigator. Since it was not practical for all sites to upgrade to the latest version of Navigator, we fell back to earlier versions of the tools in question which did not suffer the same problems ("simple chat", "paint"). This was a minor but noticeable inconvenience to the instructor, who was expecting to use the newer versions of the tools.

Several problems with TANGO and the network infrastructure were experienced throughout the workshop. Some of these problems were caused by high latencies, while others were caused by application bugs or in one case a server being down. Though not the fault of TANGO per se, these problems made it difficult to get consistent performance out of the TANGO applications, especially Buena Vista.

Due to an equipment problem at OSC immediately prior to the start of the OpenMP training, there was a problem with audio echoing due to switched cables through the sound mixer at OSC. The problem was diagnosed and resolved during the first hour, but it was annoying to the participants until it was fixed.

Another audio related problem that occurred only during the second training was a series of "drop-outs", where ARL lost the audio feed. These drop-outs either ended without intervention after a period of 20-60 seconds or the ARL Buena Vista client was restarted to restore audio. While obviously disruptive for the ARL students, and generally undesirable, we have been unable to reproduce, or even adequately explain the problem. In several years of using Buena Vista, we have never experienced such a problem before, and in this case it was only ARL that experienced the problem. This rules out problems with the transmitting client (at CEWES), which should have randomly affected all receiving clients. Nor is there any evidence that there were network problems between CEWES and ARL. The possibility that there was a problem with the receiving workstation at ARL was not investigated.

Finally, it is worth noting that twice during the class, the Netscape Fasttrack HTTP server being used to provide the slides froze and had to be restarted. These outages were believed to be due to the relatively high demand placed upon it by the size of the class. This type of behavior has been experienced before with this server (which is commonly available because it is part of the standard software installation on recent SGI systems), and as part of the guidelines we have developed for future training (see Appendix), we recommend avoiding this HTTP server.

# Student Reactions

## Fortran 90

Much of the confusion surrounding the TANGO format centered on getting the software loaded and configured. When students entered the classroom, they were given a logon account, password and a instruction sheet on how to load TANGO. The local instructor went through the TANGO procedure before class, but students who arrived late missed out on the TANGO walkthrough.

Once connected and observing the lecture, students did not seem to have any problems. They did seem to benefit from the video image of the instructor as his hand movements did add clarification. Most seemed comfortable with the headsets and audio settings, once they were shown how to bring up the audio panel that contained the volume setting. The audio quality was good, while the video image was of marginal quality. The quality of image did not seem to bother the students a great deal. Some asked if it could be improved, but seemed satisfied when they found out that it could not.

A few students lost their Buena Vista session midway through the course and were unable to reconnect. As discussed above, full Buena Vista sessions could not be established after approximately 11 a.m. (EDT). This did create some problems as once students became used to the video image, they had difficulty adjusting to it's absence.

Students also appeared uncomfortable asking questions via the TANGO tools. To avoid background noise, the mics were turned off when students were asking questions. Since there was no mute switch on the headsets, they had to stop and start Buena Vista audio or adjust the mic input lever in the audio panel to activate their mic to talk. Usually by the time they got ready to ask the question, the instructor moved on to new material. The natural network delay was also a detriment. The students did not feel comfortable interrupting the instructor. If the instructor paused and asked for questions, they seemed more willing to speak up. Posting questions to the chat window and having the instructor address them at his convenience might be a better approach.

## Parallel Programming on the Origin 2000 using OpenMP

For the second course, support staff prepared student workstations by preloading TANGO prior to each day of the course. Instead of having the students responsible for connecting to the appropriate applications, connections were the responsibility of the instructor. For students entering late, support staff provided assistance. Also, only one Buena Vista application ran from the instructor workstation alleviating some of the a/v problems students encountered during the first course.

All remote questions were received through the chat application. In this instance, students asked a multitude of questions. Video showing the remote classrooms may have enhanced this process as some students at other locations would have preferred less emphasis on certain topics. Through video, the instructor could visually survey the majority of students and focus on their reactions.

# Instructor Reactions

## Fortan 90

The instructor's reactions to this teaching environment were mixed. While the TANGO tools were generally effective in delivering the course content, answering remote questions required more time than expected. Simultaneously repeating questions so all the students could hear and formulating responses extended the class time and placed additional demands on the instructor.

The TEF classroom environment was such that it was difficult for the instructor to pay the same attention to the local users as the remote users. The O2 instructor workstation in the TEF was at a right angle to the main axis of the classroom. The instructor had to choose between facing the workstation and its digital camera and "ignoring" the local class, or facing the local class and "ignoring" the remote students. This was especially problematic when questions were asked from remote sites via the chat window.

## Parallel Programming on the Origin 2000 using OpenMP

The instructor had some problems with timing of the class, and was forced to rush through the final chapters of the planned material and essentially skip the hands-on lab exercises. This is likely due to a combination of factors. Some experienced TANGO instructors do find that the natural pace of a TANGO class is slightly slower (though they generally digress from the planned material less), however it is also true that this was the first presentation of the course material in any context, so it seems likely. It does seem likely that the minor technical problems, which can cut into lecture time, have a more pronounced impact in the training environment than in the academic environment because in the latter there is more opportunity to adjust both pace and content.

The instructor felt that having the students use the chat tool to ask questions was cumbersome. The instructor would hear the incoming signal, a bell, read the question, repeat the question out loud, and answer it. The general preference would have been for two-way audio. This was not feasible because of concerns about more audio/video clients possibly triggering known stability problems in the particular version of Buena Vista being used for this class. This problem is eliminated in Buena Vista 3.0, which is part of the TANGO Interactive 1.4 release distributed in March 1999, so future events should be able to experiment with two-way audio more readily.

The other main issues during this class from the instructor's viewpoint had to do with features and capabilities of the tools used to deliver the class which the instructor would have liked to have had:

- A shared pointer in conjunction with the shared browser used to present the slides, allowing the students to "see" the instructor pointing at a portion of the slide (as is natural in a face-to-face class) without having to say "If you look at the fifth bullet..."
- A way for the instructor to add annotations to slides during the presentation. -- combining the capabilities of the shared browser with the whiteboard.
- The small fixed font size for the prototype shared telnet application posed readability problems which limited its use in demonstrations. A more general shared display might also have been useful.

All of these features are in fact under development in various forms, but were not integrated in usable form in time for this class. As a more general form of shared display, we are investigating VNC (Virtual Network Computing) from AT&T Cambridge Research Labs (formerly Olivetti Research Lab), which is similar to X11, widely used on unix platforms, but can be used cross-platform with PCs and other display devices.

# Lessons Learned

Overall, the training was successful, both in terms of reaching the geographically distributed DoD training audience and with respect to helping understand the the differences and similarities between distance training and distance education. As might be expected in an early effort of this type, a combination of minor technical problems and instructor unfamiliarity combined to make the prototype distance training somewhat less effective than a face-to-face training might have been. There was also a non-trivial amount of effort required in advance of the class to arrange and carry out the necessary testing, however this experience proved to be quite valuable in the long run. We believe that these are "startup costs" associated with a new an unfamiliar tool/environment, which will ultimately fade away as instructors and support personnel gain more experience with this type of distance training.

The differences between training and education show up in several ways. Perhaps most significant are the tighter time constraints on training, which put more pressure on the instructor, the delivery tools, and the support staff. Lack of student familiarity with the tools, though generally a minor issue, also can take on more significance in the time-constrained training environment. From the instructor's viewpoint, it seems especially important to become very familiar with the TANGO Interactive tools, so that these do not become an issue during delivery. Practice with the tools will also help the instructor understand the differences in pacing compared to a face-to-face class, and thus plan the material better.

Adequate infrastructure remains an important issue, as in education, and it is important for everyone involved to realize that some factors (especially long-haul networks) are outside the control of the people organizing and supporting the class. The occasional nature of training classes, and the likelihood of new sites being interested suggests that extra care must be taken to insure that each and every site is appropriately configured, has adequate network links, etc. prior to each class. However much of this (especially workstation hardware and software configurations) can often be handled proactively by local support staff familiar with the issues around distance training.

Through the lessons learned from a combination of these prototype distance trainings and four semesters of academic course delivery, we have developed a set of guidelines designed to help with all aspects of planning and delivering distance training and education courses. These guidelines, given in the Appendix, will become a separate, living document, which we will update to reflect changes in our experience, changes and enhancements in the TANGO Interactive system, and comments and questions from readers of this report.

# Future Directions

Based on the results of this project, we plan to make distance training delivered with TANGO a more or less routine part of the CEWES PET training program. Initially, this will primarily involve expanding the group of instructors and course offerings. Organized training will be offered to help acquaint new instructors with the tools and techniques, and they will be required to participate in preparatory trial

deliveries before their first actual distance training. To facilitate an appropriate level of support for these trainings, we expect to continue to emphasize delivery to training rooms. We do, however, plan to begin some experiments with direct to the desktop delivery with carefully selected individuals.

Based on these experiences, a number of enhancements to TANGO are in progress or planned which address many of the issues raised here. (Some were in fact under development before these trainings, but had not reached a deployable state in time for use here.) One of the most significant advances coming is WebWisdomNT, a sophisticated database-backed tool for the management and presentation of courseware. This tool will provide the shared pointer and annotation capabilities noted by the instructors (students will also be able to store their own annotations using WebWisdomNT). Also, an "educational" version of Buena Vista is under development which will give the instructor access to all student controls, which will make it easier to utilize audio for interaction (on appropriately setup workstations) with less demand on the local support staff or students to be familiar with the details of the audio system's operation.

# Acknowledgments

The authors wish to thank the support staff at the recipient sites for their very important help in making this project a success and for their critical readings of the manuscript: Brian Cladwell, Mason Colbert, John Eberle, Derek Moses, Tracey Smith, Brian Tabor, and also Geoffrey Fox, Nancy McCracken, Marek Podgorny and the TANGO group.

The major delivery technology, TANGO Interactive, was initially developed with funding from Rome Laboratory.

# References

1. Tom Scavo, David E. Bernholdt, Geoffrey C. Fox, Roman Markowski, Nancy J. McCracken, Marek Podgorny, and Debasis Mitra, Synchronous Learning at a Distance: Experiences with TANGO, Corps of Engineers Waterways Experiment Station Major Shared Resource Center (CEWES MSRC)/Programming Environment and Training (PET) Technical Report 98-29, http://www.wes.hpc.mil/CEWES/reports/.
2. David E. Bernholdt, Geoffrey C. Fox, Roman Markowski, Nancy J. McCracken, Marek Podgorny, Thomas R. Scavo, Qutaibah Malluhi, and Debasis Mitra, Synchronous Learning at a Distance: Experiences with TANGO, in SC98 Conference, Institute of Electrical and Electronics Engineers and Association for Computing Machinery, 1998.

---

# Figures

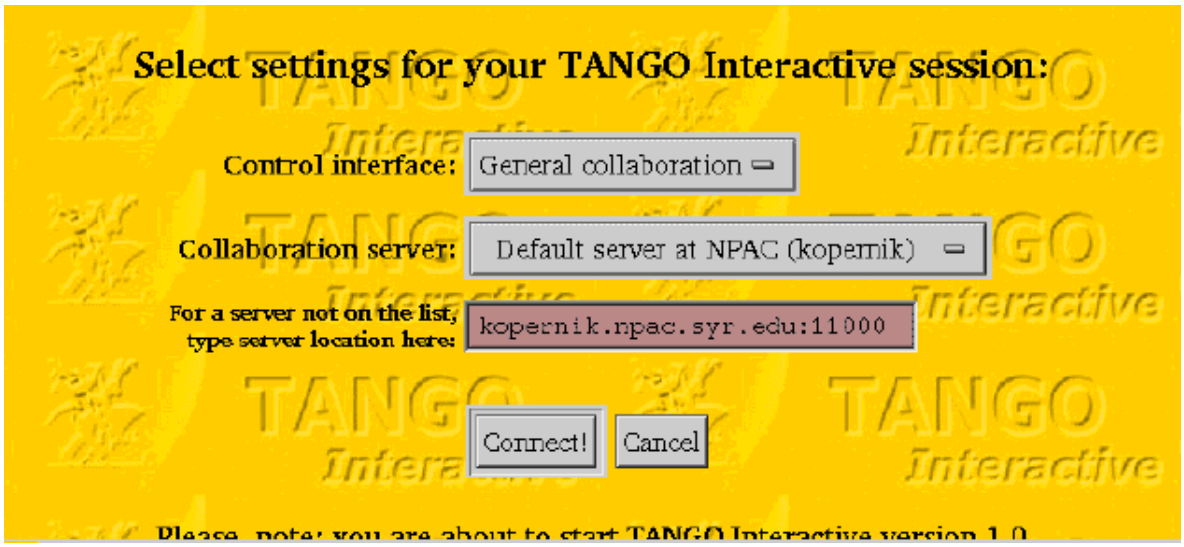


**Figure 1: TANGO Startup Window**


**Figure 2: TANGO Login Window**

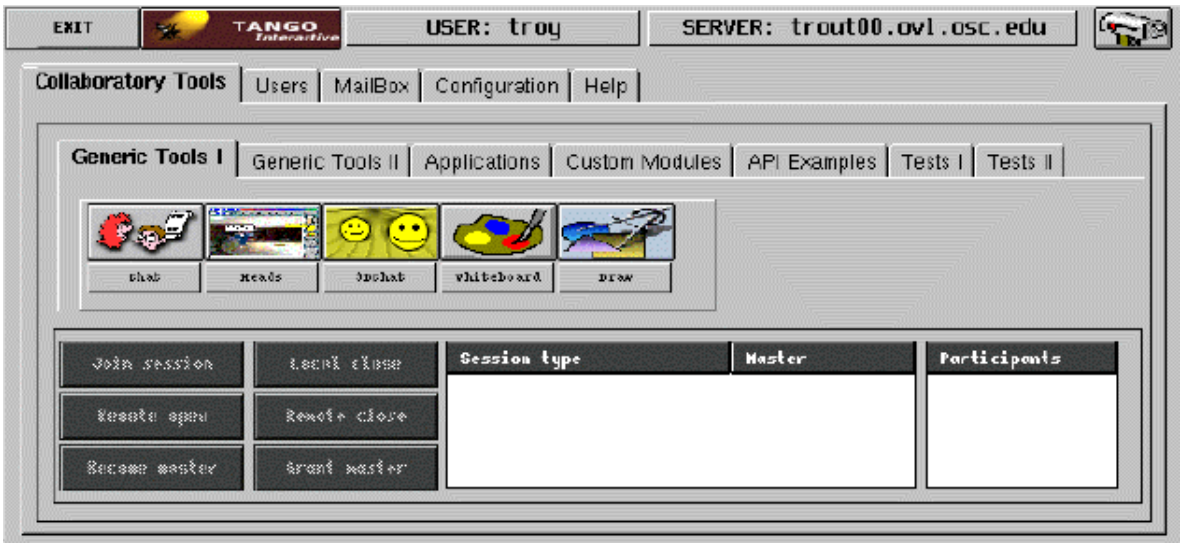


**Figure 3: TANGO Control Application**


**Figure 4: TANGO Chat Client**

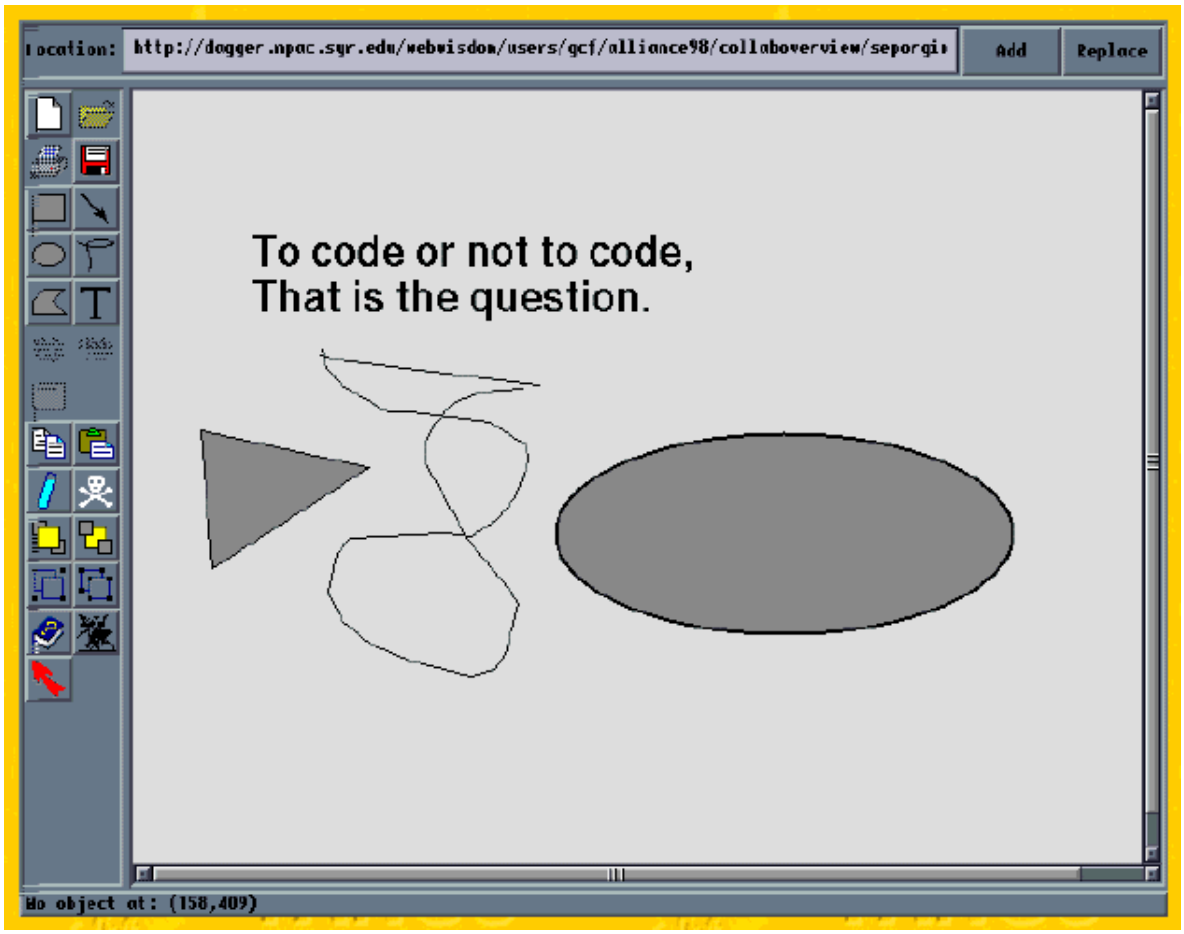


**Figure 5: TANGO Object Whiteboard**

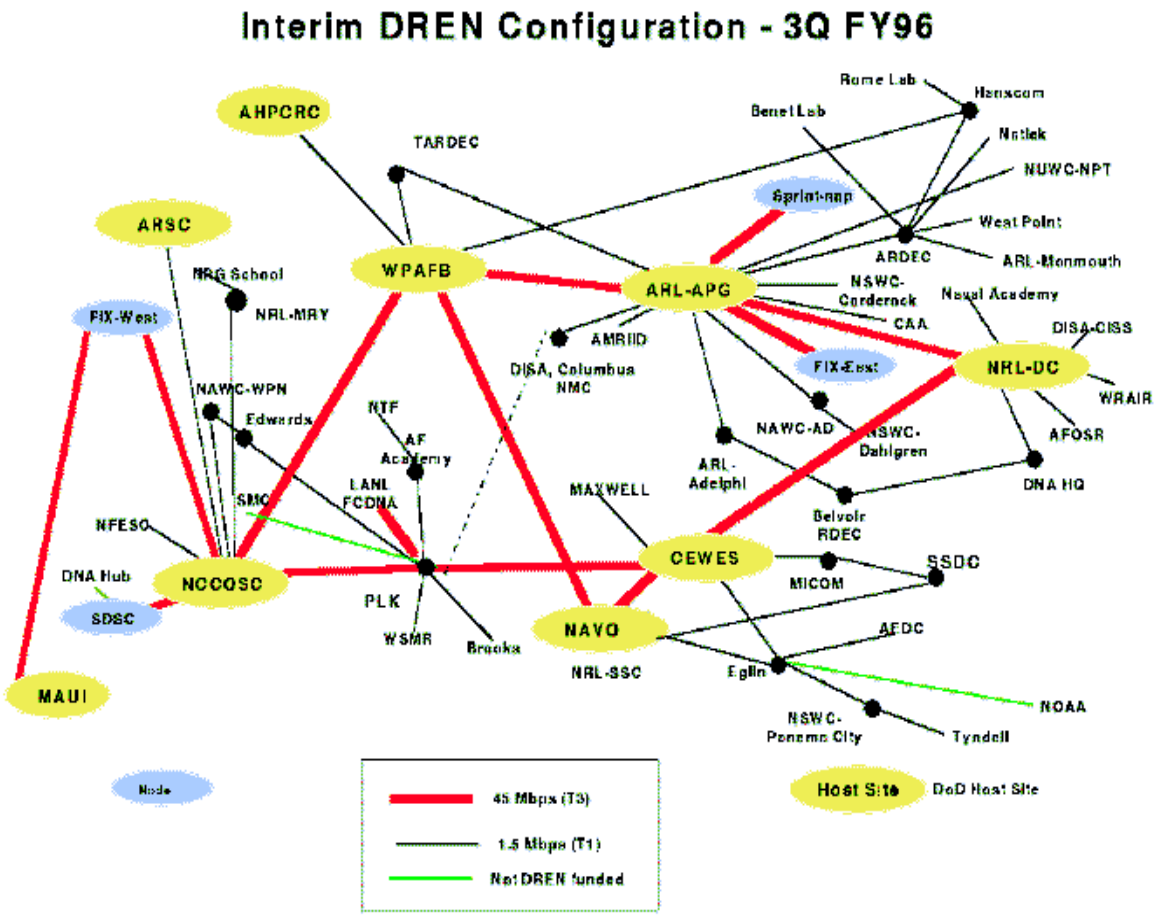


**Figure 6: DREN Connectivity Map**

# Appendices

## Guidelines for Distance Training with TANGO Interactive

These guidelines have been developed based on the experience gained from both academic and training course delivery over the past several years, including the training experiments described in this paper. They are intended to help would-be instructors, students, delivery sites and recipient sites better understand some of the important requirements and factors which, based on our experience, will help produce successful distance training delivered with TANGO Interactive.

### For Students

- TANGO-delivered classes are subject to technical problems with the network and other outside factors as well as the occasional glitch in TANGO itself. TANGO glitches can usually be remedied fairly quickly and easily with the help of experienced support staff. Outside factors, on the other hand, may leave the instructor and support staff as frustrated as you.
- We plan to begin controlled experiments with "direct to the desktop" delivery (instead of to a local training room) in the near future. Interested students should be aware that this will require some preparatory work with support staff (local or remote). The goal is not to discourage direct to the desktop delivery, but to do it in a controlled environment in order to help insure its success, and to better understand the issues that will come with wider deployment.
- While many people express an interest in "direct to the desktop" delivery instead of having to go to a centralized training room, even local, it is worth noting that others have expressed a preference to get away from the distractions of one's office in order to really concentrate on the training.

## For Instructors

- Although you may not be accustomed to practicing delivery of regular classes, it is very important to practice the delivery of TANGO-based classes until you are thoroughly comfortable with the tools. Some of your practice sessions should include at least one TANGO recipient to help point out common mistakes, such as pointing with your hand instead of using the mouse (for tools with a shared pointer display) or describing where you want to point (in general).
- Expect TANGO-based classes to move a little slower than local classes -- words must replace gestures and body language that would suffice in a local class, and things like student questions can be somewhat more cumbersome.
- The audio codec typically used in TANGO (GSM) provide roughly telephone-quality sound -- not broadcast quality. You should be careful to speak clearly and perhaps a little more slowly than usual for face-to-face teaching to insure that you are understood by the remote students.
- Before setting up a TANGO-based class, give careful consideration to the content of your class, and be sure that TANGO will support your needs, and provide a way to share all of the content you plan on using. If you're not sure, talk to the TANGO developers.
- Our experience is that the attention of remote students wanders a little more easily than in face-to-face classes. Probably because there is less stimulus, and less of a sense of connection with the instructor, so we suggest...
  - Don't be afraid to make course materials colorful and visually interesting.
  - Try to make the class a little more interactive than you would a regular face-to-face class to help hold student interest/attention.
- If there is a support person or technician available at the delivery site, don't be afraid to make use of them if it will make you more comfortable. They might be able to help monitor student questions on chat or assist in other ways. Talk to them before the class to plan roles, signals, etc.

## For Delivery and Recipient Sites

#### Hardware Requirements/Recommendations

- We recommend the use of two separate machines to deliver classes where possible, one handling TANGO and the course materials (shared browser or WebWisdom), and the other handling audio and video. This is because the audio and video codecs can require significant amounts of CPU time, and the two system configuration prevents this from adversely impacting the rest of TANGO.

- In classroom environments beware of the microphones (and speakers) often built into computer systems. Mics are typically omnidirectional -- they pick up sound from all directions. In the distance education setting, this makes them prone to producing feedback, as they pickup the output of a speaker and feed it back into the audio system. Cheap, lower sensitivity mics with more selective pickup patterns can help reduce feedback problems. Headsets with boom microphones are another alternative.
- A wireless lapel microphone can give instructors a little more mobility without sacrificing audio quality. A wireless handheld mic can be helpful in allowing students to ask questions.
- Here is a list of audio and video hardware that we have used successfully with TANGO/Buena Vista. This is not an exclusive list!
    - Video
        - Winnov Videum Conference Pro card and camera (Windows 95/98/NT)
        - Intel Smart Video Recorder III with any compatible camera (Windows 95/98/NT)
        - Panasonic EggCam
        - SGI O2 standard camera
    - Audio
        - SoundBlaster64 audio card with full-duplex driver
        - Parrot ST Gold headphones
        - VXI boom mic systems
- Please see the Buena Vista Manual for detailed guidelines for setting up audio and video.
- A good, *well-tested* audio setup is crucial to a successful class. Feedback, in particular, can be not only annoying but can also damage ears and equipment.
- Because it has to send so many a/v streams, the instructor's workstation generally has the greatest demand placed on it of any computer involved in the training. Give some thought to the available bandwidth (probably 10 MB/s ethernet), and CPU power (for the a/v codecs) in planning for the class.

**Network Considerations**

- In addition to the HTTP server from which it downloads Java applets (usually on port 80), TANGO uses ports 11000, 11015, 12567, and 34999 (at present). If any of these ports are blocked by a firewall, some or all TANGO functionality will be lost.
- Some thought and testing should be given to the **quality of the network** between all of the sites involved (primarily the path between the delivery site and each of the recipient sites). We have found the DREN to generally provide good quality connections, but the commercial Internet can be highly variable. If network quality is questionable or highly variable, it might be advisable to try to schedule class to avoid the highest network traffic periods (typically late afternoon on the east coast).
- In cases where the network quality is questionable or unreliable, using speaker phones (on a conference call, if there are multiple recipient sites) can be a useful alternative to using TANGO's Buena Vista to transmit audio, especially when coupled with proxy web servers at recipient sites to make the best use of the available network capability. Audio transmission is the most susceptible to network quality.
- Minimal bandwidth requirements (per stream) for TANGO are as follows: Audio (GSM) 13 kb/s, video (H.263) 15 kb/s, courseware 120 kb/s. The number of streams required during a training class varies greatly depending on details of the setup, which are often at least partly under the control of local training room staff.
    - Audio: The instructor's workstation must send one stream to each *Buena Vista audio client*.

If students are listening individually on headphones, there would be one client per *student*. If sites are broadcasting audio to the training room using speakers, there would be one client per *site*. Each open student microphone which might be used for asking questions will also send a stream to *all* of the other Buena Vista audio clients. Since open mics tend to pick up and transmit environmental noise, we strongly recommend that student mics be used in a "click to talk" mode. *Note that a multicast version of Buena Vista is under development, which will eliminate the stream multiplication caused be the current point-to-point architecture, but will require a working multicast backbone between sites.*

- Video: This is currently transmitted "point-to-point", just like the audio, so the number of streams is determined the same way. In TANGO trainings, video usually presents a "talking head", not actually conveying the course materials. Therefore, it is possible to get away with many fewer streams. Instead of separate streams of individual faces, the instructor would probably rather see an overview of the entire classroom. While courseware may not work so well on a centralized display at recipient sites, a video of the instructor probably will (as opposed to each student running their own video client).
- Courseware: This is a web server delivering slides/pages to student browsers. In general, one stream is required to each receiving workstation. This can be reduced to one per site through the use of proxy web servers at recipient sites (see below). It can also be reduced by using a centralized large display at recipient sites, however these may have a less-than-acceptable image quality for some types of course material.

- Make sure each site's network connection are sufficient to handle necessary traffic. Likewise, check how the sites are connected to each other and especially to the delivery site. For example, a T-1 line (1.5 Mb/s) can carry at most 100 a/v streams, or about 10 simultaneous pages of courseware *in the absence of any other traffic*.

**Software Requirements/Recommendations**

- To reduce network bandwidth requirements, we recommend that all recipient sites setup a **proxy web server** with caching capability (i.e. Netscape Proxy Server; we would welcome reports of experience with others) and configure their browsers to use it. During a class, each page of web-based course material will be downloaded to every participating browser, resulting in many simultaneous "hits" on the web server providing the class materials, and each request consuming bandwidth (a simple 20 kB GIF image distributed to 50 clients requires that the web server send 1 MB of data in total). Properly configured, all requests at a given site can instead be handled by a local proxy server, reducing the demand on the class web server to one copy per site instead of one per client.
- If possible, all sites should use the same version of Netscape Navigator. This helps to avoid "site-specific" problems caused by TANGO interacting with bugs specific to a particular version of Navigator.
- Pre-class testing should exercise all TANGO tools that will be used during class delivery to insure that they work correctly at all sites. This is all the more important if all sites are not using the same version of Netscape Navigator.
- Use separate chat tools for communications between the instructor and students (i.e. questions), and among the support personnel.
- Use a **robust HTTP server** to provide the course materials. Our experience suggests that some servers (in particular Netscape FastTrack normally provided on SGI boxes and Microsoft Peer Web Server) will hang easily when subjected to the demands of a typical class. We recommend the free HTTP server Apache as a more robust alternative.

**Other Guidelines**

- The delivery site should try to arrange the instructor workstation and camera(s) so that the instructor can easily shift the focus of their attention between the screen, the local class, or the remote classes (camera) without having to completely change position.

## General Comments

- Agree in advance on the web site and TANGO server that will be used to start TANGO for the class (and all tests). Using the same web site will help insure that everyone is using the same version of all Java-based TANGO tools, and of course everyone must connect to the same TANGO server. TANGO itself requires very little bandwidth (as opposed to the audio, video, and courseware streams), and it makes little practical difference which TANGO server is used.
- Testing in advance of the the actual class is very important to its success. The amount of testing needed depends on the novelty of the particular event planned, and on the level of experience of the participating sites. Care should be taken to exercise *all* of the tools (within TANGO or outside of it) that will be used during the class to check for any version conflicts that might cause problems. The audio and video setup at all sites should also be tested. Testing should be done far enough in advance to allow time to fix or work around problems that might be uncovered, but close enough that the configuration will not "decay" between the test and the actual event. Instructor practice to familiarize themselves with TANGO is a separate matter, but it can be useful to have the instructor "drive" the test sessions too -- they should be the most familiar with what tools they intend to use, and how they'll use them.
- Scheduling tests and classes at multiple sites can sometimes prove to be more trouble than actually delivering the class! Unfortunately, TANGO doesn't have a solution for this (yet?). Don't give up too early.
- In all discussions involving multiple sites, be sure to make clear which **time zone** is intended when setting schedules.
- Responses so far are mixed on the importance of video to the distance learning process. Since the courseware is assumed to be web-based, video in this case is typically the instructor ("talking head") and/or an overview of the recipient classroom. We have conducted many successful classes without video in either direction, but some have commented that the video helps make the connection between instructor and students, and may increase student attention levels.
- It is helpful to spend a few minutes at the beginning of the class to describe TANGO and its use in the class. This might be done by site support staff, as part of helping the students login to TANGO, or by the instructor as soon as everyone is signed on. Please see http://www.osc.edu/~troy/tango/ for a student-oriented introduction to TANGO Interactive.
- Prior to the beginning of the class, the instructor and the site support staff should should agree on how the students will log on to TANGO (i.e. who will help them) and how far to go before the instructor takes over (i.e. should students try to launch any applications after class starts, or will instructor launch everything required?)
- Audio/video setup is usually somewhat specialized and should be done by site support staff in cooperation with the instructor in advance of the class. Apart from a/v (Buena Vista) and a support chat tool, most applications can be easily launched to all users by the instructor.