



Center for Computation & Technology

CCT Technical Report Series
CCT-TR-2007-2

Uncompressed HD video for collaborative teaching—an experiment

Andrei Hutanu, Ravi Paruchuri, and Daniel Eiland
Center for Computation & Technology
Louisiana State University, Baton Rouge, LA 70803 USA

Milos Liska and Petr Holub
Masaryk University, Brno, Czech Republic

Steven R. Thorpe and Yufeng Xin
MCNC, Research Triangle, NC 27709 USA



Posted July 2007.

cct.lsu.edu/CCT-TR/CCT-TR-2007-2

The author(s) retain all copyright privileges for this article and accompanying materials. Nothing may be copied or republished without the written consent of the author(s).

Uncompressed HD video for collaborative teaching - an experiment

Andrei Hutanu,
Ravi Paruchuri
and Daniel Eiland
Center for Computation & Technology
Louisiana State University
70809 Baton Rouge, Louisiana
Email: ahutanu, ravi9, deiland@cct.lsu.edu

Milos Liska
and Petr Holub
CESNET z.s.p.o, Zikova 4
162 00 Prague, Czech Republic
Masaryk University, Botanická 68a
602 00 Brno, Czech Republic
Email: xliska@fi.muni.cz,
hopet@ics.muni.cz

Steven R. Thorpe
and Yufeng Xin
MCNC, PO Box 12889,
3021 Cornwallis Road
Research Triangle Park
27709 North Carolina
Email: thorpe, yxin@mcnc.org

Abstract—This article describes a distributed classroom experiment carried out by five universities in the US and Europe at the beginning of 2007. This experiment was motivated by the emergence of new digital media technology supporting uncompressed high-definition video capture, transport and display as well as the networking services required for its deployment across wide distances. The participating institutes have designed a distributed collaborative environment centered around the new technology and applied it to join the five sites into a single virtual classroom where a real course has been offered to the registered students.

Here we are presenting the technologies utilized in the experiment, the results of a technology evaluation done with the help of the participating students and we identify areas of future improvements of the system. While there are a few hurdles in the path of successfully deploying this technology on a large scale our experiment shows that the new technology is sustainable and the significant quality improvements brought by it can help build an effective distributed and collaborative classroom environment.

I. INTRODUCTION

The recent development of technologies for video conferencing supporting uncompressed high-definition video has motivated a group of universities and research institutes to conduct an experiment that would assess the value and applicability of the new technology for distributed and collaborative teaching.

The networking, software and video processing technology, supporting multi-party video conferencing using uncompressed video with a resolution of 1920×1080 pixels, was deployed and successfully demonstrated for three participating sites at iGrid 2005 and Supercomputing 2006. This encouraged us to extend the scope of our experiments and test the applicability of our approach with a larger number of participants, over a longer period of time and having increased sustainability and quality requirements.

The opportunity was given by the introduction of a new course, “Introduction to High Performance Computing”, taught at Louisiana State University by Prof. Thomas Sterling.

With the participation of Masaryk University in Brno (Czech Republic) - MU, University of Arkansas in Fayetteville - UARK, Louisiana Tech University in Ruston - LATECH and later joined by North Carolina State University in Raleigh

(through MCNC, North Carolina) LSU has initiated the “HD Classroom” experiment attempting to create a highly interactive environment to allow students and teachers from all these universities to actively participate in the HPC course. Our goal was to both provide the tools needed for remote teaching of the HPC course and to analyze the requirements for utilizing the research technology used by the HD videoconferencing system in a stable production environment

The class took place in the spring semester (January - May) of 2007, and this article presents the design, technology and evaluation of the technologies utilized for the HD classroom experiment and identifies future development strategies that would help provide a better service for the students.

II. DESIGN OVERVIEW

At the core of our experiment was the decision to use uncompressed high-definition video. HD video offers high detail, allowing students to see detailed facial expressions of the lecturers and providing a realistic remote presence experience. However this comes at the cost of generating large amounts of data that even on dedicated hardware takes a long time to compress. The lag created by compressing and uncompressing video without dedicated hardware (which was unavailable at the time of the experiment) is unacceptable for a real-time collaborative environment and can be eliminated using uncompressed video.

The main design decision was choosing the video distribution scheme and choosing the technology for each video stream sent/received the participants. In order to provide an interactive environment it was required that each participant be able to see all other participants at all times.

One of our goals was to maximize the use of uncompressed HD video, however hardware requirements and cost (dedicated fiber-connected machine for each video stream - more details below) prevented us from deploying this technology for all the video streams in our system. Our solution was to have the host site (LSU) send one HD video stream of the lecturer to all other participants and each participant to send an HD video stream to the host site.

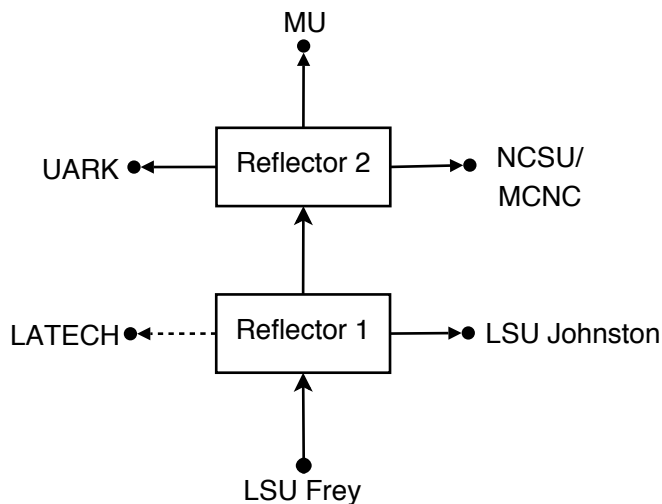


Fig. 1. HD video distribution model

The HD video distribution scheme from the host to the other participants is presented in Figure 1. Since the number of registered students exceeded the capacity of the room equipped with the video and networking equipment necessary for the experiment, LSU participated with rooms in two different buildings of the campus. These rooms are identified by the name of the building in which they are located: Frey was the room/building where the lecturer was present and Johnston was the overflow room that hosted additional students. The arrow to LATECH is shown as dotted because this university was unable to connect to the optical network required for HD video delivery in time for the class.

Lacking HD camera (LSU Johnston), corresponding capture card (NCSU) or network access (LATECH) resulted in only two streams (UARK and MU) being received in high-definition at the lecturer location.

All video streams that were not supported in high definition were transmitted in standard definition using Access Grid (AG). Access Grid software was also chosen for audio transmission and as a back-up for the HD video in case of unexpected failures. A secondary backup was a one-way broadcasting QuickTime server installed at the host site. Lecturers at the remote sites would provide the final backup in case of complete failure of the technology.

The HD videoconferencing system, our network set-up and other technologies used in our experiment are described in the following sections.

III. UNCOMPRESSED HD OVER IP VIDEOCONFERENCE

A. HDTV sending and receiving

The highest effective resolution of HDTV video provides the 1080i mode with 1920×1080 pixels and interlaced line scanning. We used the uncompressed HD video as defined by SMPTE 294 transmitted over HD Serial Digital Interface (HD-SDI). For the purposes of the HPC class we used a video with a total resolution of 2200×1125 , 60 fields per second, 10 b

per pixel and color plane and 4:2:2 subsampling. Thus the necessary bandwidth is 1.485 Gbps for one video stream. The data stream rate is equivalent to around 1.5 Gbps including the overhead generated by packetization into RTP packets using headers of 44 byte per packet (using 8500 bytes Jumbo Ethernet frames).

HDTV video is captured using an affordable HD camera (such as the Sony HVR-Z1E or Sony HVR-A10), transmitted as a component video and converted to HD-SDI using a converter box (AJA HD10A). The HD-SDI is then captured using a DVS Centaurus capture card, encapsulated into RTP/UDP/IP stream and sent to the network via 10G Ethernet NIC card (i.e. Chelsio, Myrinet) using a modified version of the UltraGrid software¹.

The data stream is unpacked on the receiving site and displayed directly. The UltraGrid software was adapted to provide 1080i rendering using Simple DirectMedia Layer (SDL) library and graphics card overlays, 10 b to 8 b per pixel an color plane conversion and linear blend deinterlacing.

The pool of computers deployed for the HD part of the class was rather heterogeneous with some common attributes. Generally we used dual AMD64 (Opteron) computers, motherboards with two PCI-X 133 MHz slots to accommodate both Centaurus and 10GE NIC card. In addition to this the receiving/displaying computers needed fast NVidia graphics card (NVidia 6600 or newer) for video rendering and displaying. All machines were running Linux (kernel 2.6.*).

B. Video distribution

The video communication using UltraGrid SW is point-to-point only. However the five partners participating in the virtual classroom needed a multi-point distribution of the video stream. The high data transmission rate disqualified the use of native multicast (usually routers do not support multicast reliably at these rates). In order to create a multipoint conference and deliver the class contents to all participants we built and overlay network where specific nodes took care of data distribution. For the distribution nodes (reflectors) we used the generalized Active network Elements (AE) based on UDP packet reflector design [1].

The AE was heavily optimized for the class content distribution so that it provided sustainable UDP packet replication of a incoming video stream into four outgoing streams on a single AMD64 computer. Four AE were deployed in our system. This provided a reasonable level of redundancy in case of failures of individual AE and allowed us to create an optimal distribution tree for all participants on our network (details in the Network section).

C. HDTV videoconferencing quality

The most important measure for an interactive class using any videoconferencing technology is the audio and video end-to-end latency. The end-to-end latency consists of latency given by the network plus latency given by processing of video

¹<http://ultragrid.east.isi.edu/>

in the UltraGrid tool and the AE. Latencies caused by the 10GE network deployed for our class are given in table I. End-to-end latency of one way communication caused by the network is equivalent to one-half of the values.

TABLE I
ROUND TRIP TIMES OF THE 10GE NETWORK END-TO-END LINKS.

End-to-end link	Latency (RTT) in ms
LSU – StarLight	30.631
StarLight – Masaryk university	115.481
LSU – Masaryk university	145.720
StarLight – UARK	19.322
LSU – UARK	49.953
StarLight – MCNC	23.527
LSU – MCNC	53.782

Our previous measurements of end-to-end latency caused by video processing alone showed that capturing, sending, receiving and displaying of the 1080i video takes 175 ± 5 ms. A notable part of which is delay caused by the Centaurus card which buffers 4 fields of the video, displaying of the video using the SDL library on the other end and the delay of LCD. More details on this can be found in [2].

Additional latency is added by multipoint distribution of the video stream using AE. Total increase of the latency induced by the AE is 13 ± 2 ms.

Although the optimal end-to-end latency threshold is 100 ms, it appears that latencies around 200–300 ms are still barely noticeable for human perception and generally not disturbing interactive communication.

D. Lecture recording

One of our challenges of our experiment was the time shift (both time zone and semester starting times) between the various participants. Both were major issues especially for MU which has time zone difference of seven hours to LSU and a five weeks delay in semester start compared to LSU. While this has presented a wide range of practical problems, technically the solution that helped mitigate some of the issues was to record and present some of the lectures to students offline. The recorded videos were available in two formats: post-processed versions of the lectures provided by LSU and a limited number of recordings of the live feed done at MU. The live feed recording was challenging both in terms of capacity (a recorded lecture requires about 1 TB of space) and speed (data rate of 190 Mbytes/s for the video). The solution was a RAID 0 array of 12 disks which provided a write performance of 385 MBps and a read speed of 414 MBps.

IV. NETWORK DEPLOYMENT

The HD video transmission has substantial network requirements. In addition to the raw bandwidth needed (1.5 Gbps for each stream, see above) a real-time application HD video transmission requires minimal transmission latency. The transmission latency is influenced by the wire latency (fixed)

but also by the presence of jitter or data rate fluctuations. Jitter or variable data rate transmission can be compensated by using buffering but this comes at the cost of additional latency and this is undesirable. Additionally, large data packets are necessary in order for the software on the end hosts to be able to sustain the transmission rates, so “jumbo” packets need to be enabled in the entire network.

The Internet cannot meet these requirements and even regular shared services such as those provided by Internet2 are not suitable for this application. In our experience since any unnecessary packet processing comes at a cost (jitter) we feel routing needs to be avoided whenever possible.

Our network was built from a combination of dedicated point-to-point (layer 1) and switched (layer 2) network links offered by various providers (details below) combined into a switched network using minimal bridging over routed (layer 3) networks where necessary.

These networks are able to provide the necessary service but the cost of having all these resources dedicated for a single application is prohibitive. In IV-B we present solutions that help reduce utilization costs using time-sharing by allocating resources on demand based on the application requirements.

A. Network topology

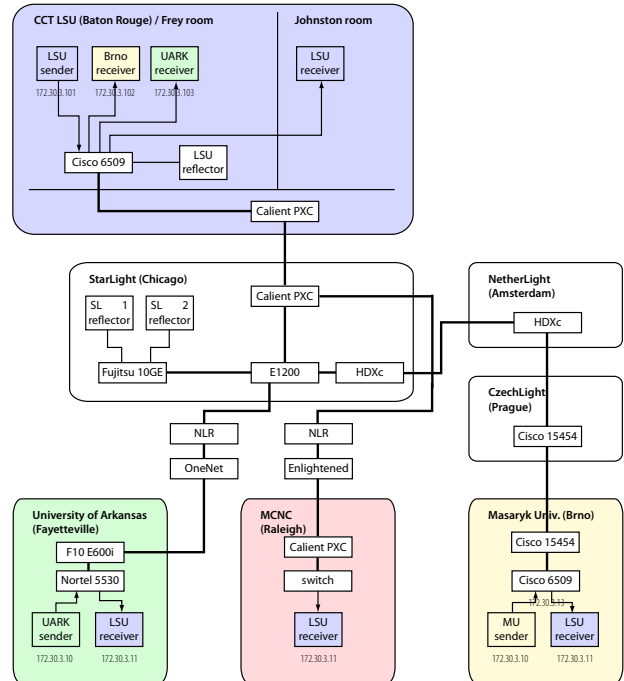


Fig. 2. Network diagram for the experiment

Our network (Figure 2) has been implemented using experimental optical infrastructure. The network topology is star-like with the center in Starlight. LSU was using the Louisiana Optical Network Initiative (LONI) which has an uplink to Chicago over the dedicated Layer 1 link provided by the National Lambda Rail (NLR) and is a part of the Enlightened

testbed. MCNC was connected to Starlight with a similar NLR circuit and UARK was connected via OneNet and NLR. MU was connected using a dedicated circuit Prague-Starlight and the CzechLight infrastructure from Brno to Prague.

The video distribution (shown in Fig. 1) instantiation has changed during the time of our experiment. Initially the two reflectors were both installed in Starlight, but this introduced the unnecessary RTT of the LSU-Starlight link for the second LSU location. Therefore for the last part of the class, one reflector was moved to a new machine at LSU².

B. Automated allocation of resources

Keeping resources dedicated for distributed applications that run occasionally such as the HD classroom experiment is cost ineffective.

The Enlightened project, who provided part of its testbed for our experiment has developed HARC³ [3], a software package used for co-allocating multiple types of resources at the time the application needs them. Using HARC, resource utilization can be maximized as the resources are freed when the application execution time is over and can be allocated by different applications.

HARC can be utilized to allocate both network and compute resources, however in our experiment the compute end nodes were dedicated for our application so HARC was only utilized to allocate network links. The network links available under the control of HARC were the Enlightened testbed links (LSU-Starlight and Starlight-MCNC). The Starlight-Brno and Starlight-Arkansas links were maintained as dedicated for the experiment during its entire duration.

While initially the network links were still manually allocated, in the second half of the experiment HARC was scheduled to allocate and activate the two network links automatically each Tuesday and Thursday well in advance of the scheduled start of class and to de-allocate them after the class was over.

This showed that although not widely deployed, all the network-related components needed to execute such a demanding application in a production environment are available and we proved that they can be effectively utilized together to support a real application today.

V. OTHER TECHNOLOGIES

In order to maintain both the stability of the class as a whole as well as offer students as many interactive opportunities as possible, a number of technologies were employed in tandem with the high definition video as both stopgap measures and experiments in designing a successfully interactive course.

A. Audio

While not necessarily an additional technology to the high definition video, the audio provided enough quirks to warrant

its own section. The audio was both transmitted and received using the robust audio tool (RAT). The hardware used to acquire the audio was varied, but generally consisted of ClearOne's XAP800 as well as a number of both wired and wireless microphones. In order to maintain equal reception between sites, two forms of the RAT were used - one over the private 10Gbit network for those receiving the high definition video, and one made use of by the Access Grid client over regular networks. The audio used 16b quantization and up to 48 kHz sampling rate in stereo. All participants could hear and talk to one another.

Making use of two audio clients created a number of problems in regards to echo cancellation that was a constant drawback of the course. The audio itself needed daily adjusting to assure that it was both coherent as well as undisrupted throughout the course. Most agree that the audio problems encountered throughout the course were the most difficult to manage. While video services could generally be left to run once started, the audio needed continued maintenance. These adjustments would often change mid-course depending on the state of the audio network or the volume level of the presenter.

B. Access Grid

The Access Grid Toolkit (AGTK 3.02)⁴ was used extensively throughout the course both as a low-bandwidth solution to classes who could not access or send the high definition stream, and also as the first stopgap measure to providing continuity throughout the class. The Access Grid is a widely employed video conferencing tool which makes use of multicasting or unicast bridges to send multiple video and audio streams to participating groups. Unlike the HD video, the Access Grid trades video quality for video quantity. Most Access Grid venues make use of multiple cameras which allows any number of participants to make use of an active camera view without having one camera constantly moving from face to face. Using the Access Grid offered a number of key benefits.

Because of its already extensive development and employment, universities could make use of the Access Grid client often without the purchase of more equipment or without need of further training. The Access Grid became the main access solution for Louisiana Tech University, since they could not employ their HD solution in time for the class to begin. Their students were still able to take part using the Access Grid that, while not offering the same video quality as the HD stream, did offer certain levels of interaction between sites.

C. Webex

While it would have been possible to transfer power-point slides and demos using an HD camera, it was determined that the use of another HD stream would not be practical. As such, a commercial software known as Webex⁵ was used to transfer both the professors laptop display as well as demos from other computers to the various sites. Unlike making use of other

²For the final lecture, the NCSU set-up was separated from MCNC so a third reflector needed to be installed at MCNC to accommodate the additional receiver

³<http://www.cct.lsu.edu/~maclaren/HARC>

⁴<http://www.accessgrid.org/>

⁵<http://www.webex.com/>

display sharing software, Webex allowed for quick switching of the current stream being viewed while also providing a rapid setup interface that allowed sites easy access to the material. While this allowed for a very clear view of the presenter monitor, there was admittedly enough lag between the sending and the reception of the information that quick moving demos often trailed behind the speakers intentions.

D. NCast

The Ncast⁶ telepresenter provided the third level of stopgap in case of network failure. The telepresenter is a webstreaming device which combines both a camera view of the instructor, their audio, and their power-point or demo files into one compressed QuickTime video stream. While the telepresenter is best known as a webstreaming box, most students did not use it to watch live streams. Instead, it was used to view recordings of the classes both after class as study guides for the exams and also as a method for clearing up confusion brought when audio or video issues were present. According to a survey of the students, most made extensive use of these recordings to refresh their memory before exams and in some ways the recordings became a video based textbook to be used in place of the written form.

VI. CLASS SETUP (ROOMS)



Fig. 3. Local set-up showing remote sites in HD at Louisiana State University

A. Louisiana State University - Class Host Site

The class was primarily hosted at Louisiana State University in the Frey Computing Services building with the support of Center for Computation & Technology. There were three Opteron workstations setup in this room, one to send the video and two to display the video from Masaryk University Brno (see left of Fig. 3) and University of Arkansas(see right of Fig. 3) classrooms. Two 30 inch display were setup to display the HD video from the above two sites. Dr. Sterling can be seen lecturing the class as the podium(see Fig. 5).

⁶<http://www.ncast.com/>



Fig. 4. Access Grid Display set-up at Louisiana State University



Fig. 5. Dr. Thomas Sterling Lecturing at Louisiana State University

The class was also transmitted in HD to another classroom on campus at Johnston Hall, Louisiana State University and also to MCNC. There was a 42 inch plasma screen (see Fig. 4) setup at the back of the room to display the video from LSU - Johnston Hall, LATECH and MCNC through Access Grid.

B. Louisiana State University - Johnston Hall Site

The set-up at LSU - Johnston Hall classroom, one of the recipient sites (see Fig. 6) includes HD video of the lecturer and slides (center), AG video from all other sites (right) and presentation material through Webex (left) on three different screens. An Opteron workstation is used to receive and display the HD stream. Access Grid nodes are used to display the presentation through Webex and AG video from other sites.

C. Masaryk University

The local classroom for Masaryk University was hosted by the Laboratory of Advanced Networking Technologies, and the set-up was centered around the large projection system (Design Cineo3+ 1080i) with a camera pointed at the students



Fig. 6. Set-up at Louisiana State University - Johnston Hall



Fig. 7. Local set-up at Masaryk University

(see Fig. 7). MU did not usually participate in the Access Grid sessions, so they were unable to see sites other than the LSU Frey site which hosted the lecturer.

D. University of Arkansas

The set-up at University of Arkansas (Fig. 8) included two Opteron workstations: one to send the local HD video stream to LSU, and the other to display the HD video stream from LSU. University of Arkansas used a SONY 60 inch Rear Projection HDTV to display the LSU HD video stream. Access Grid was a back up solution that the class room was setup and ready to switch any time. The Webex transmission of the course material was projected on a screen next to the HDTV that displayed the lecturer from LSU.

E. MCNC/North Carolina State University

The set-up at MCNC/NCSU site was only equipped with receiving capability for HD video. Access Grid and Webex were also deployed. A team of NC State senior students participated this class in the form of senior design project. During the course, the MCNC/NCSU team experimented and



Fig. 8. Local set-up at the University of Arkansas

evaluated various aspects of the three technologies in terms of video/audio quality, cost, and classroom environment. The system deployment actually consists of two Opteron workstations, one in the MCNC campus and the second one in the NCSU campus, connected via a dedicated 10GE metro optical network.

VII. STUDENT EVALUATION OF THE TECHNOLOGY

Near the completion of the course, students were asked to evaluate their experiences both with the instructors, graduate assistants, and their views of the technologies used throughout the course. We gathered 37 responses from volunteering students, 21 of them that participated in the class locally (at LSU Frey) and 16 from remote students (UARK and LSU Johnston). The students were asked to grade the quality various technologies on a scale from one (poor) to five (good).

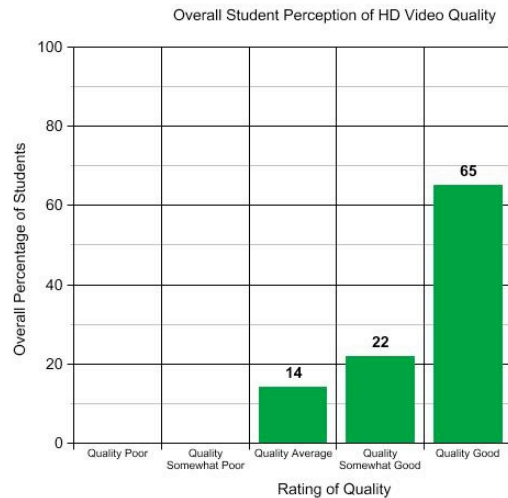


Fig. 9. Video quality

In regards to technological quality, almost all students agreed that the high definition video offered marked improvement over past used avenues of video based courses. This is

illustrated in Fig.9 that shows the vast majority of students rating the video quality as good.

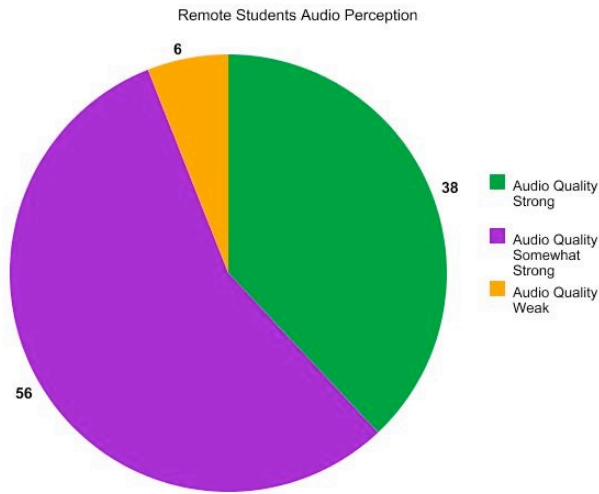


Fig. 10. Audio quality perception (percentages)

However, because of both a lack of consistent audio (see Fig.10 - although most students rate the audio positively, only 38 % rate it as good), poor interactivity training, and a number of other issues, students felt that their interaction with both the instructors as well as students from other sites were severely lacking. The student rating is shown in Fig.11 and shows a complex mix of positive and negative ratings and considering the overall positive response to the technology this needs to be interpreted as a negative evaluation.

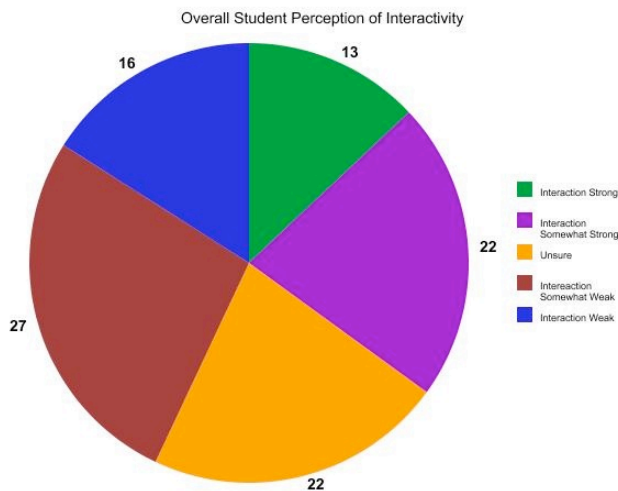


Fig. 11. Interactivity perception (percentages)

Reasons for this lack of interaction were many-fold. Some students felt that because of a lack of a written textbook, the professor did not allow enough time within his lecture for adequate questioning. Oftentimes, he was so engaged in presenting the material that questions - if they were present - were glossed over. Many students at remote sites felt uncomfortable

asking questions during a lecture from fear of interrupting the instructor or because of their distrust of the less than amicable audio system.

Furthermore, this lack of interaction was also caused by a lack of connection between sites. Because of cost limitations, sites were limited to receive and send only one HD stream. The difference in quality between the stream received from the host site and the streams received from the other participants and the fact that some locations were not set-up to display all the remote participants made it difficult for sites to connect to each other during the class. The result was impairment of the communal environment the technology was intended to create.

VIII. EXPERIENCES, ISSUES AND LESSONS LEARNED

A. HD Video

Overall, the HD Video was the best and most reliable component of our system. There were no significant failures and the quality of this component was reflected positively by the students. The hardware supporting it also proved to be very stable. However there is one area where the HD system could be improved in the future and this is the video distribution mechanism and the system configuration.

Currently the distribution network and the software at the end points are statically configured for the fixed set of participants. The data reflector network, video sender and receiver units are manually configured well in advance of the beginning of the class for this static configuration.

This is not a scalable solution and is highly labor intensive to set-up. The existing system works for the number of participants in our experiment but would not work with a significantly higher number of participants.

The solution here is to move towards an automatically configurable collaborative environment that would control the configuration of the data distribution network and also the software on the endpoints. On top of this, a floor control system such as that of Isabel [4] could be implemented. The configuration of the system would be changed dynamically if another speaker (or a student having a question) takes the floor. One possible application of this would be a multi-discipline, or multi-professor class where different professors (possibly at globally distributed locations) have input at different times during a single session, on an interactive basis.

Also, as the number of participants increases the capacity of the lecturing site to receive video from all participants will be exceeded. We already experienced this situation at LSU where hardware constraints limited us from receiving more than three HD video streams. Downscaling needs to be supported if the number of participants is increased.

B. Networking

The network was also a valuable component of our system, managing to transfer the roughly 200 Terabytes of total data during the approximately 30 hours of live lecturing. However physical failures on one of our network links for two weeks resulted in about four lectures (approximately 20%) not being transmitted in HD video to the other participants. This shows

that redundancy in the network paths is important for this to be available as a production system.

Another issue is monitoring. In our system, the testing before the class start and monitoring during the class was done manually by operators who watched the video streams or looked at various performance metrics on their console while the class was running. Many times failures were only detected very late, when there was little time available to fix them.

Some amount of in-advance testing needs to happen automatically and monitoring should be automatically coordinated with the operator only receiving warnings or errors. Monitoring and testing needs to happen on multiple levels: network (can data flow through at the specified rates and within the packet loss limits), end hosts and software.

C. Audio

The sound quality was the biggest issue of our experiment, and this was considered a distraction for many of the participating students. The bridging of two audio loops did not function properly and the sound installation of many participants was of poor quality. The major problems were not network related but rather caused by audio installations (components, wires and sound cards). The solution for improving the quality of this important component of our system should be simple and is to have either a single or two completely separated audio loops (separate echo canceling) and having strict quality requirements for the audio installations at the participating sites.

D. Classroom issues

Having rooms set-up for a collaborative classroom is perhaps the most difficult issue that needs to be solved. The room set-up was the single component which was inadequate at every single participating site. Some participants have used meeting rooms that are suitable for interaction but not very appropriate for lecturing (no room for displays, cameras), other participants have used large classrooms that were not suitable for remote interaction (large noisy projectors, hard to install microphones for large number of people). Building a collaborative environment for both large number of sites and large number of people could be a unsolvable problem; there is probably a trade-off between how interactive an environment can be and the number of participants and sites. However, we believe that careful reorganization of the rooms (large rooms where teacher should be able to see all participants at the same time, sound insulation of projectors, and installing theater-like microphones) will improve the quality and should alleviate some of the issues we experienced.

E. Other tools

We realized early in our experiment that audio and video applications are not enough to build an effective distributed classroom. A desktop sharing application (Webex) was rapidly adopted (second lecture) and utilized throughout the experiment for slide sharing and demonstrations. We also identified other resources that are available in a normal classroom and

are also required for a distributed environments. First, a “hand-raising” application needs to be available for all participants to notify the lecturer of outstanding questions. Also, a shared whiteboard or smart board application is required in order to distribute content created on-the-fly during the lecture.

IX. CONCLUSION

Our experiment showed that despite the high development, deployment and maintenance costs plus a wide range of technical difficulties we were able to provide the necessary service for students around the world to effectively participate in the Introduction to HPC class.

We believe that the necessary components for high-quality collaborative teaching and distance education are available. When the issues we have identified are solved, this system will be a feasible option for future course delivery and it will provide opportunities for new types of classroom environments.

ACKNOWLEDGMENT

This work was supported by the Center for Computation & Technology, Louisiana Optical Network Initiative (LONI), the Enlightened project, the National Lambda Rail (NLR), CESNET, Starlight and the Arkansas Research and Education Optical Network (ARE-ON).

The authors thank all participants to the HD classroom experiment, in particular the following: Stacey Simmons, Jon MacLaren, Gabrielle Allen, Edward Seidel, Steven Beck, Ben Blundell, Lonnie Leger, Jorge Ucan, Boyd Bourque, Matthew Eastman, Chirag Dekate (LSU), Joe Mambretti, Fei Yeh, Alan Verlo, Linda Winkler (Starlight), Tomáš Rebok, Lukáš Hejtmánek, Eva Hladká, Luděk Matyska (MU), David Merrifield, Amy Apon (UARK), Lina Battestilli, John Moore (MCNC), Chris Womack, Box Leangsuksun (LAT-ECH), Mladen Vouk (NCSU) and most importantly the author of the course Thomas Sterling (LSU) and all the participating students from all sites.

REFERENCES

- [1] E. Hladká, P. Holub, and J. Denemark, “An active network architecture: Distributed computer or transport medium,” in *3rd International Conference on Networking (ICN’04)*, Gosier, Guadeloupe, Mar. 2004, pp. 338–343.
- [2] P. Holub, L. Matyska, M. Liska, L. Hejtmánek, J. Denemark, T. Rebok, A. Hutanu, R. Paruchuri, J. Radil, and E. Hladká, “High-definition multimedia for multiparty low-latency interactive communication.” *Future Generation Comp. Syst.*, vol. 22, no. 8, pp. 856–861, 2006.
- [3] J. MacLaren, “Co-allocation of Compute and Network resources using HARC,” in *Proceedings of “Lighting the Blue Touchpaper for UK e-Science: Closing Conference of the ESLEA Project.* PoS(ESLEA)016, 2007, http://pos.sissa.it/archive/conferences/041/016/ESLEA_016.pdf.
- [4] J. Quemada, T. de Miguel, S. Pavón, G. Huecas, T. Robles, J. Salvachúa, D. A. A. Ortiz, V. Sirvent, F. Escibano, and J. Sedano, “Isabel: an application for real time collaboration with a flexible floor control.” in *CollaborateCom.* IEEE, 2005.