Virtual Classroom with a Time Shift

Luděk Matyska [1,2] and **Eva Hladká** [1,2] and **Petr Holub** [1,2]

[1]CESNET z.s.p.o, Zikova 4, 162 00 Prague, Czech Republic
[2]Masaryk University, Botanická 68a, 602 00 Brno, Czech Republic

ABSTRACT

In spring 2007, Prof. Sterling's *Introduction to High Performance Computing* has been delivered in a virtual classroom consisting of 5 sites spanning Louisiana, Arkansas, and North Carolina in the USA and the Czech Republic in Europe. By utilizing a multi-point uncompressed audio and High Definition (HD) video capturing, transmission, and presentation system, a high fidelity collaborative environment was available to students, giving them access to the lecturer regardless of the place they attended the lecture. However, the semester start differs at individual participating educational institutions by several weeks and the institutions are also in different time zones. A store and replay technology was used in these cases, together with a larger number of lectures presented per week, to accommodate for this difference. The whole system worked well, the high quality of the HD video allowed students even in the most distant place to get a very good feeling of participating in a live lecture.

INTRODUCTION

Providing high-quality university lectures on specialized advanced topics is increasingly difficult for smaller and less rich universities, as the top field specialist are very scarce resource. To mitigate this situation and to help their less developed peers, some top class universities are providing their lecture material, including captured video of individual lectures, free of charge to everybody interested. The MIT's OpenCourseWare activity is among the most widely known [1]. While helping to improve local teaching quality, this approach does not provide direct access to the teacher and keeps students at individual universities out of contact. Educationally more advanced (and more challenging) is thus the concept of a *virtual classroom*, where students at different places are synchronously attending the lecture. While the concept itself is rather old [1,2] and initially covered one campus only, the recent advances in network and audio/video technology open up access to a really high fidelity environment. The distance between individual halls of such a virtual classroom does not play serious role, students from differing institutions can share a particular lecturer, perceiving the illusion of being in direct contact with him or her.

We have implemented a state of the art multi-point uncompressed HD video over IP transmission system together with other supporting technologies that can be used to provide such a high-fidelity environment. We have expected that our system would provide enough immersion for the students in order to feel being part of a single class, despite being divided by both geographical and time distances. As a pilot experiment, prof. Sterling at the Louisiana State University developed an *Introduction to High-Performance Computing* class that was offered to be shared in the virtual classroom setup at several universities and higher education institutions. The institutions participating in the class were Masaryk University (MU) in the Czech Republic, University of Arkansas (UARK), Louisiana Technical University (LATECH), MCNC and North Carolina State University (NCSU). In this paper, we present the necessary technical details of the setup together with a discussion of the experience on psychological and sociological levels.

From both technical and psychological points of view, the ultimate goal of a system to support true globe-spanning distributed class is to provide minimum disturbance for all the students of a class, while mitigating the separation of students on the remote locations and providing the same opportunities to both local and remote students. In order to achieve this, high-fidelity real-time audio and video transmission system have to be implemented, as further discussed in the following section. However, the real-time transmission and collaboration support is not sufficient, as the distributed classes face time-shift problems. One, simple problem is the time zone shift. While it may be possible to agree on a time that is convenient for two sides, with more sides and full globe coverage, not everybody is able to participate during his usual working hours. More serious problem is caused by differences in start of the class. Order of weeks skew may be found among individual institutions. To overcome both problems, the system must support non-synchronized, non-real time mode of work and the lack of real-time interaction should be taken care of at the educational process level, too.

INTERACTIVE COLLABORATION TECHNOLOGIES

We decided to build our virtual classroom environment of top of the real-time collaboration environment using uncompressed audio and HD video and

[1] See http://ocw.mit.edu/index.html.

over IP [3]. The HD video offers very high detail, allowing, e.g., to capture classical projection screen and to transmit the picture without any distortion (and loss of readability) to the remote location. It is also possible to capture both the projection screen and the lecturer. Using the HD video, students can see detailed facial expressions of the lecturer (and she can see the facial gestures of the remote students), providing a realistic remote presence. However, this high quality resolution does not come without a price—a lot of data is generated, which means the compression takes a long time even on high performance hardware. The several second long lag created by the compression and decompression of the video stream is unacceptable for the real-time collaborative environment. This lag can be easily removed using the uncompressed HD video.

The uncompressed HD video was HD-SDI over IP transmission based on UltraGrid software [3], featuring full 1080i SMPTE 274M/292M video with effective resolution of 1920×1080, 60 interlaced fields per second, 4:2:2 color space sampling, and 10 b per color plane.[2] The stream was packetized into the 8,500 B Jumbo Ethernet frames with 44 B IP/UDP/RTP header overhead and the resulting data rate was about 1.5 Gbps. With this data rate, we decided to distribute the video asymmetrically, in 1:N way from site where the class is given to all the other sites, and in N:1 way back, i.e. all sites are sending back to the lecturer only (this way the lecturer can see all his students but students can not see other lecture halls apart from the one with a lecturer). This asymmetry was implemented due to bandwidth constraints at participating sites.

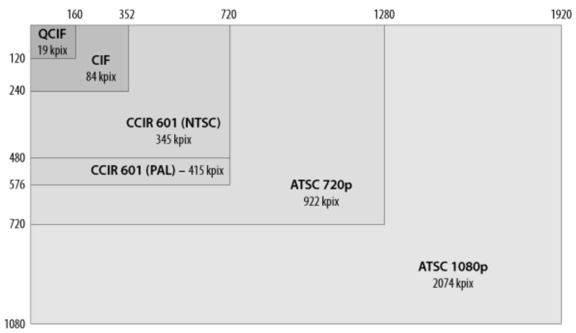

Figure 1: HD video formats

The audio part has been implemented by RAT [4] at 16 b quantization and up to 48 kHz sampling rate in stereo. Thus data rates up to 1.6 Mbps were used. The audio was distributed in full N:N way, i.e., all participants can hear and talk to one another. While the compression and decompression of the audio is not so computationally demanding as the processing of HD video, we used uncompressed audio to provide highest quality possible.

The high data rate needed for the uncompressed HD video streams transmission is not widely available on the shared production networks. We used a dedicated high speed network described in the following section, but we have also implemented a backup solution based on AccessGrid and webcasting in QuickTime format for those unable to participate at full network data rate.

Network and Data Distribution Setup

The five partners participating in the virtual classroom needed a multi-point distribution of the video and audio streams. The very high data transmission rate disqualified use of native multicast—the routers usually do not support such high multicast distribution rate. We built an overlay network where specific nodes—the reflectors—took care of multi-point data distribution (see Figure 2). While a shared production networks available in academic environment already provide data rates above 1 Gbps at some locations, we used a dedicated 10 Gbps network as the underlying transmission medium.

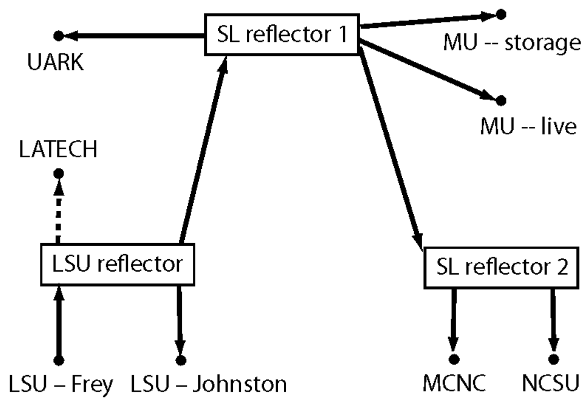

Figure 2: Overlay data distribution network for the 1.5 Gbps uncompressed HD video streams

The network has been implemented using experimental 10 Gigabit Ethernet (10GE) infrastructure and both statically and dynamically allocated λ-services. The network topology resembles star with center in StarLight (SL) in Chicago, IL. LSU and LATECH[3] were using 10GE backbone Louisiana Optical Network Initiative (LONI), which has an uplink to SL using a dedicated Enlightened wave running 10GE on National Lambda Rail (NLR). The Enlightened wave

[2] For more details on the HD formats see e.g. *Society of Motion Picture and Television Engineers:* "The Bit-serial Digital Interface or High Definition Television Systems", SMPTE-292M-1998 and "1920x1080 Scanning and Analog and Parallel Digital Interfaces for Multiple Picture Rates", SMPTE-274M-1998.

[3] LATECH is not part of the actual 10GE infrastructure though, because of the last-mile problems from the LONI point of presence at the LATECH campus to the room used for the teaching and thus they are using AccessGrid as the backup solution.

was being dynamically provisioned for the class and for the testing windows using a preliminary version of HARC [5], the software stack developed by the Enlightened project. UARK was connected via OneNet network and a wave on NLR running 10GE; again, the circuit was automatically setup before the class begins and shut down after it ends. MCNC and NCSU were connected to SL in the same way using NLR waves. MU, the only transatlantic partner, was connected using a dedicated permanent circuit Prague–SL leased by CESNET and the Brno–Prague implemented using CzechLight experimental infrastructure based on leased dark fiber lit by CESNET equipment. While the rest of the infrastructure runs 10GE, Brno–SL circuit runs OC192 protocol. The whole infrastructure is switched L2 network and is transparent on the IP level; only a few important network devices have been assigned IP addresses for monitoring and debugging purposes (see Figure 3).

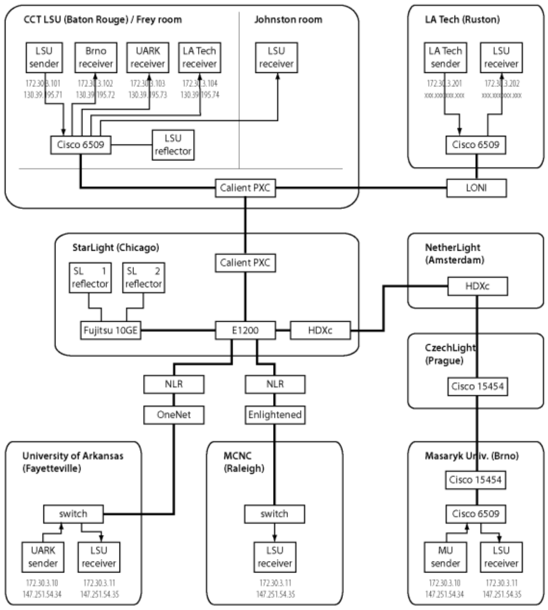

Figure 3: Detailed data distribution overlay network

Latencies are the major problem in real-time collaborative environments—with higher latencies it is not possible have a true dialog. The network latencies shown in Table 1 are sufficiently low not to create interactivity problems even in case of the transatlantic connection.

Table 1: Network round-trip latencies (RTTs) in the 10 Gigabit infrastructure

Ping from	To StarLight	to LSU
LSU	30.6 ms	-
MU	115.4 ms	145.7 ms
MCNC	23.5 ms	53.8 ms
UARK	19.3 ms	49.6 ms

The actual data distribution was implemented using user-empowered software UDP packet reflectors [6,7]. The reflectors ran on dual-Opteron computers equipped with either Chelsio or Myrinet 10GE NICs. For the audio, although distributed in the full N:N pattern, a single reflector distributing is sufficient.
The video was distributed in 1:N and N:1 way and the schematics of the 1:N video distribution is shown in Figure 2. Note that there are actually two sites participating at LSU as not everybody can fit in a single room—thus the same technology was used for distributing the data locally on the campus between two buildings. There were also two streams being sent to MU—one of them was used for live video feed, while the other was used for the full-quality recording.

Site Setup

Each site was equipped with one sender and one or more receivers for uncompressed HD video over IP transport system. The machines were essentially the same as used for the reference UltraGrid 1080i implementation [3], i.e., dual-Opteron computers with Myrinet or Chelsio 10GE NICs. The sender computers were equipped with DVS Centaurus HD-SDI capture card, while the receivers used NVidia graphics card to render the video on attached LCD screen or projector. 24" to 30" LCD screens and plasmas were used and installation at MU also used the Projection Design Cineo3+ 1080i projectors. The video part runs flawlessly with very good user perception.

Though less challenging from the networking and processing perspective, the audio part created many more problems due to various factors namely on audio capture side. The worst problems are caused by bad audio installations (ranging from wiring problems to echo canceling, gain control problems, and over-processing of the sound by various components before it gets to the computer—all this resulting in noises and distorted sound) used within the lecturing rooms primarily at LSU and wireless microphone interferences. Problems have also been encountered when using some on-board integrated sound cards, as their obviously half duplex behavior resulted in clicks and sound distortions. After appropriate corrective actions were undertaken, the sound quality has become very high and pleasing to hear.

Data Storage

As stated above, the lecture had to be presented to students also in the asynchronous way. This was achieved through the recording of the live stream for a later use. The media streams (1 video and N audio streams) from the network were captured and stored as raw data on a disk array. Varying slightly depending on the size, each 1.5 hr lecture was about 1 TB of stored data.

The data storage has been implemented using the very fast local disk array comprising 12 disks (actually 6+6 disks attached to two physical channels of one SATA to SCSI disk array controller) in RAID 0, as it required both read and write throughput of at least 190 MBps. The actual write performance of the array was 385 MBps and read performance was 414 MBps. The tested parallel filesystem PVFS did not provide sufficient robustness with the inexpensive hardware we used. The stored streams were indexed so that seeking capability was also available. The data was archived to a slower RAID 5 disk array immediately after recording and later also to a tape archive so that there was always a backup copy available. When needed, the data was replayed from the same RAID 0 disk array.

TIME-SHIFT MITIGATION

The time warping had to be implemented for two scales of time shifting—the short term, as the students may not be forced for instance to stay up to very late night due to different timezones, and the long term, because of different semester start dates. Both cases actually happened for the *Introduction to HPC* class.

The time difference between LSU and MU is 7 hours, which meant that the lecture starting at 3pm local time in Louisiana was received at 10pm at MU in Brno. The two hour lecture ended slightly before a midnight Czech Republic local time[4]. The time zone differences with other partners were at most one hour, which did not cause a real problem. In theory, the time zone problem could be solved by careful planning, but as we found, the LSU internal rules made it impossible to move the lecture to the more appropriate morning time and we had to account for this.

There is unfortunately no way how to directly overcome the second problem, differences in semester start. The only possible solution is to relax the real-time requirement and use recorded data instead of the live lecture. The semester at LSU started already in mid January, while the MU starts only after mid February. With two lectures per week this meant a delay of 10 lectures (full 15 hours of recorded material). To catch this delay students agreed to take up to twice the number of hours per week (usually there has been three instead two lectures). However, this created another problem with the problem sets and other prescribed homework.

Other supporting material. The students were also provided supplementary material like all the slides, additional reading, problem sets and homework, etc. The supporting material must be uncovered individually for each student's group, to follow the differing speed of presenting the recorded lectures. The major problem has been associated with the problem sets and the homework. Usually 4 to 6 weeks were available to solve a problem set and return the report. However, with the "faster than real time" lecture presentation the time available could be reduced, as some "future" lecture might provide hints for the problem set solution before they are actually submitted by students. We maintained supplementary materials in an archive local to each of the desynchronized institutions. The eventual shortening of time for homework has been agreed with students. In some cases, students were given new problem set before the previous was submitted—we checked that the "future" lectures are harmless, not containing any unwanted hints. All the homework as well as mid-term and final exams were evaluated locally.

The lectures at MU were replayed at the same days (Tuesday and Thursday) as at LSU, but at early evening hours. In the middle of semester (also due to the different holiday times and other reasons that lead to few than expected lectures at LSU) we were able to synchronize the MU and LSU classes. We still continued to present the recorded lectures at 7pm but invited students to stay longer and watch also the live lessons given by prof. Sterling or his colleagues.

Most of the time the whole system worked well, but at several occasions (3 times during the 13 week long semester) the live transmission has not been possible—either the transatlantic or other part of the path from LSU to Brno was broken and not operational. All the lectures were regularly recorded also at LSU and post-processed there. We used these post-processed lectures if the recording at MU was not available. The post-processing created another problem—usually it had taken one week before the material was ready for presentation. When MU and LSU were not synchronized, we had enough local material to present, but after the mid-term synchronization, there was no safety margin (the LSU has just one lesson ahead). Once we had to use a different lecture (recorded keynote from an international conference).

LESSONS LEARNED

As we are expecting to repeat the experience next spring with substantially higher number of participants and also for everybody interested in similar educational setup, we summarize the major issues here. Two most important lessons are: (1) it is possible to run a lecture in HD quality at the virtual classroom setup even over transatlantic distances, and (2) it is not easy and needs a lot of man power.

The amount of necessary manual work surprised us. Part of it was due to a lot of homework and problem sets students were expected to work on during the semester. All the problem sets were prepared with the LSU resources available and we had to adjust our

[4] We may consider it irony that both time zones use the same abbreviation—CET. It denotes Central European Time in Europe and also Central Time in the USA.

computing environment to provide a similar setup to our students. Also, the teaching assistants at LSU had to learn how to share their experience (including the process of evaluation student's work) with the staff local at other sites. Without this, students would be graded differently.

To operate the whole system and especially the network setup took another portion of man power. As stated above, most lines were dynamically activated and de-activated, some of them through the use of experimental HARC software. All the transmissions were real-time and there was just a little time before the start of a lecture—the setup was usually initiated just half hour earlier, leaving very limited time for reaction to any problem encountered. We have learnt that when using experimental network, it is imperative to have a fall back solution available. An alternative capturing and presentation system must be available at all sites, having less stressing demands on the network and thus enabling use of a shared academic production networks instead of the dedicated experimental one. If more than the 5 sites are to be connected, much more automation is also needed, including a constant monitoring of the quality of network, drop rate (esp. for the audio) and other parameters that influence the quality of perception. Ideally, the system should be able to react immediately to network problems and reconfigure either the network or the collaborative environment itself on the fly.

The importance of good audio can not be overstated. While the full HD video quality provided details as expected, when a lower quality backup solution was used, it did not have tremendous negative impact. However, when the audio quality deteriorated, the lecture became completely useless. Even small noise, clicks and other artefacts negatively influence the student's focus and ability to follow the lecture (and learn from it). New microphones, audio cards and other equipment were tried during the first lectures to provide the best audio experience, but some problems repeated the whole semester. We are considering to study the studio quality audio systems to provide a flawless audio; it is worth even the higher price of studio systems.

When the semester starts at different time, it is very important to plan ahead how to mitigate it. We "learned by doing", adapting to the student's requirements and ability to follow the lectures (e.g., we did not know in advance whether students will be able to take more than the two lectures per week). Next time, we would like to synchronize with LSU as soon as possible, as it lessens the problem of material revealing and, most importantly, allow direct collaboration among students at different institutions. Without attending same lectures students have a little in common (they solve different problem sets, they have different experience from the last lecture seen etc.).

We also learned that the lecture must be prepared with the virtual classroom setup in mind. While Prof. Sterling prepared a lot of homework and other material, none of them directly encouraged remote student's collaboration. For the next semester, we would like to see also team homework, when students from different institutions will work together to solve a particular problem set. More interactivity between the lecturer and students is also very much needed. Students must be explicitly asked to participate, because even with the advanced technology the ability to ask question remotely is something that people do not know automatically and it must be learnt. Encouraging direct interaction between students and asking for feedback during or after lectures will teach students how to live well in a virtual classroom environment.

CONCLUSIONS

The collaborative system described in this paper had been implemented by the beginning of year 2007 and has been used since then for supporting the HPC class. Due to running on highly experimental networking infrastructure, some sites had to resort to using backup solutions at times, but the infrastructure worked flawlessly most of the time. Another experience gathered from this class is how important is high-quality audio installation at participating sites and how difficult it is to implement it properly. The low sound quality is much less tolerated when other means of communication are of high quality compared to common point-to-point communication tools.

We are currently undertaking a number of studies on subjective user perceptions in order to map the technology improvements to user perception improvements. As more sites want to participate, the whole system needs to get revamped in order to support much higher degree of monitoring and self organization.

RELATED WORK

We are not aware of any e-learning system that uses the uncompressed HD video for real time lecture delivery. The virtual classroom concept is widely used, but with the more classical collaboration technologies like H.323 videoconferencing systems and simple shared work space.

Interest in HD video is growing, also due to the increased availability of adequate equipment designated for the consumer market. A recent presentation at EUNIS conference [8] gives a good introduction to the use of commercially available products that use compression and introduce thus latencies unacceptable for the live collaborative environments of the real-time virtual classrooms. The HD video is currently used for lecture streaming only.

ACKNOWLEDGMENT

This project has been kindly supported by the research intent “Optical Network of National Research and Its New Applications” (MSM 6383917201). We would also like to thank other people working with us on this project, namely Tomáš Rebok, Miloš Liška, Lukáš Hejtmánek from Laboratory of Advanced Networking Technologies at Masaryk University, and our partners at Center for Computational Technology at Louisiana State University and other partnering sites.

REFERENCES

[1] S.R. Hiltz, “The “virtual classroom”: Using Computer-mediated Communication for University Teaching”, Blackwell Synergy, 1986

[2] S.R. Hilts, “The virtual classroom: learning without limits via computer networks”, Ablex Publishing Corp. Norwood, NJ, USA, 1994

[3] P. Holub, L. Matyska, M. Liška, L. Hejtmánek, J. Denemark, T. Rebook, A. Hutanu, R. Paruchuri, J. Radil, E. Hladká, “High Definition Multimedia for Multiparty Low-latency Interactive Communication”, Future Generation Computer Systems, Vol. 22(8), pp. 856–861, 2006

[4] V. Hardman, A. Sasse, M. Handley, A. Watson, “Reliable audio for use over the Internet”, Proceedings of INET’95, Honolulu, Hawai, 1995

[5] L. Battestelli, A. Hutanu. G. Karmous-Edwards, D.S. Katz, J. MacLaren, J. Mambretti, J.J. Moore, S.-J. Park, H.G. Perros, K.S. Sundar, S. Tanwir, S.R.Thorpe, Y. Xin, “EnLIGHTened computing: An architecture for co-scheduling and co-allocating network, compute, and other grid resources for high-end applications”, 2007. Available from http://enlightenedcomputing.org/index.php?n=Main.AboutEnLIGHTenedPage?action=dowload&upname=EnlightenedGrid07.pdf

[6] E. Hladká, P. Holub, J.Denemark, “An Active Network Architecture: Distributed Computer of Transport Medium”, Proceedings of 3rd International Conference on Networking (ICN’04), pp. 338–343, Gosier, Guadeloupe, 2004

[7] P. Holub, E. Hladká, L. Matyska, “Scalability and Robustness of Virtual Multicast for Synchronous Multimedia Distribution”, Proceedings of 4th International Conference on Networking, La Réunion, LNCS vol. 3431, pp. 876–883, Springer-Verlag Heidelberg, 2005

[8] J. Oxenford, “Using HD Videoconferencing for Teaching and Learning”, Proc. EUNIS 2007 Conference, Grenoble, Paris, 6 pages, 2007