High Quality Large Scale Virtual Classroom

Luděk Matyska[1], Eva Hladká[2], Petr Holub[3], Miloš Liška[4]

[1]CESNET, z.s.p.o., Zikova 4, 160 00 Praha, and Masaryk University, Faculty of Informatics and Institute of Computer Science, Botanická 68a, 602 00 Brno, Czech Republic, ludek@ics.muni.cz, [2]Masaryk University, Faculty of Informatics, Botanická 68a, 602 00 Brno, Czech Republic, eva@fi.muni.cz, [3]CESNET, z.s.p.o., Zikova 4, 160 00 Praha and Masaryk University, Institute of Computer Science, Botanická 68a, 602 00 Brno, Czech Republic, hopet@ics.muni.cz, [4]Masaryk University, Faculty of Informatics, Botanická 68a, 602 00 Brno, Czech Republic, xliska@fi.muni.cz

Keywords
Virtual classroom, HDTV video, HD video transmissions.

1. EXECUTIVE SUMMARY

Universities are faced with a difficult problem of simultaneously keep pace with increased number of students and increased extent of knowledge that is to be delivered to students. And students expect to be taught by excellent teachers who have deep understanding of the subject. A *Virtual classroom* concept is a possible solution to the problem of increased number of students, when a set of smaller classrooms is interconnected and a lecture is given from one site to all connected sites.

1.1. Background

Use of videoconferencing technologies brings the virtual classroom concept a step further, giving access to teachers outside one university. Also, even individual students can directly participate, extending thus capacity of the virtual classrooms. However, the major drawback of using classical videoconferencing technologies lies in their limited video resolution. This shortcoming is removed by the High Definition (HD) videoconferencing technologies that are giving an opportunity to create almost natural virtual class environment, with a very realistic illusion of teacher's presence in the remote lecture hall. An HD videoconferencing technology was used to create a virtual classroom for the "Introduction to High Performance Computing" class taught at Louisiana State University.

1.2. Results

The uncompressed HD video processing and multi-point high speed transmission system, developed at MU and in deep technical detail described in our previous papers, deliver high quality teaching environment that has been for the first time used in spring 2007 to deliver the HPC class (total of 26 lectures) to 5 sites in the USA and one in the Czech Republic. The experiment was repeated in 2008, where a combination of HD Polycom system (a lower quality but easier to use and less bandwidth demanding commercial solution) and the uncompressed HD video environment were used to deliver the lecture to even higher number of participating sites.

The participation of MU was further complicated by differences in the start of the spring semester and in 2007 also by not too convenient time of actual lecture. The former issue was solved in 2008 through denser lecture delivery speed (up to 4 lectures per week) in the first half of the semester, the later issue by moving the LSU lecture to more appropriate morning time.

1.3. Conclusions

Students at MU (and their peers at several US universities) were able to attend a high quality lecture given by outstanding teacher with a high quality illusion of direct involvement. The major drawback lies in still high man power needed to setup and run such a lecture and less than expected interactivity among students and with the remote teacher.

2. INTRODUCTION

A novel challenge for contemporary universities is to continue to live with dramatic increase of enrolled students, not all of which are well prepared for the selected field of study on one side, and with a constant pressure of keeping and even increasing the quality of the study programs. The problem is further complicated by the ever increasing extend of knowledge and expectations from the industries about the graduates' skills. Even the richest and best universities has a serious problem of engaging best teachers in all fields of study, while most universities cannot afford the broad coverage and their teachers usually excel in limited number of fields. Keeping into mind that the number of excellent teachers will always limited and their numbers will never keep pace with the number of students, one solution is to *virtualize* teachers, making them available to larger number of students than any time before. A *virtual classroom* concept (Hilts 1986, 1994) was initially introduced to mitigate problem of higher number of student on campus. The idea was to connect several classrooms via standard audio and video equipment (TV) and thus increase number of seats available to students for one lecture without the necessity to build large (and very expensive) lecture halls. The use of analogue technology limited the usability, as only classrooms not too far away (within one campus) could be thus connected. At Masaryk University, Faculty of Informatics, we initially adopted the same approach, connecting three lecture halls to create a virtual classroom with capacity around 500 seats.

However, this implementation of virtual classroom concept does not deal with the more deep problems, namely access to excellent teacher. If the university does not have such a teacher in a particular field, the connection of several on-campus classrooms does not help. The deployment of videoconferencing technologies helped to bring the implementation of the original concept a step further. With them, lecture halls at different campuses of distinct universities can be easily connected, making for the first them the concept of fully remote teacher reality even in a real time. Furthermore, the same technology can be also used to transmit information to individual students, too, making a virtual classroom of theoretically unlimited seating capacity. Scaling the size of the classroom above some threshold creates new problems, namely the issue of direct interaction with the teacher. While in a setup of few large lecture halls it is possible for the teacher to see all the classrooms together on one screen, this is not possible with tens or even thousand (in case of individual students connecting directly sites. While there are solutions for this problem, in most cases such lecture is transformed into streaming, losing the live interactive format.

Another negative aspect of classical videoconferencing systems is the limited quality of transmitted video. With PAL format at most, the systems are suitable to be used for (even large) TV, but not so well for large screen projection. Also, cameras are using lower resolution than the projections system, making direct transmission of the projection screen behind the teacher unsuitable and distorted.

Availability of high quality High Definition (HD) videoconferencing technologies provides new opportunity to create almost naturally impressing virtual classroom environment. It is possible to create a very realistic illusion of teacher's presence in remote lecture hall(s), without any distortion of transmitted material. Even the resolution (1012x768 or even 1280x1024) of standard projection systems is well within the HD area. This way, the remote lecture halls do not need two projection systems—one for the presented materials and the second to make access to teacher, his mimics and body language.

However, capture and transmission of HD video streams is still a challenging task. To keep the data flow at a reasonable speed—but still at the level of tens of Megabits—powerful encoding and decoding systems are needed. These are not only expensive, but also create substantial delay (up to several second in total), again endangering the liveliness and interactivity. This delay can be drastically reduced if we remove the encoding/decoding process completely, using only uncompressed HD video. This results in a very high bitrate of the generated data stream, up to the levels above one Gigabit per second. Even as such throughput is not yet easily available, the recent development and proliferation of global optical networks (see GLIF activities for more reference, http://www.glif.is/) as parts of the academic networks makes it possible to use the uncompressed HD video streams between well equipped universities.

A multipoint HD videoconferencing system, working with uncompressed HD video over optical fibre lines has been developed at Masaryk University (Holub 2006). We decided to use this system as a technical base for an advanced high quality virtual classroom implementation. Together with Louisiana State University (LSU) we started to deliver high quality lecture "Introduction to High Performance Computing" given Prof. Thomas Sterling at LSU, to several US universities and also in a transatlantic experiment to Masaryk University in Brno, Czech Republic. The first full semester course has been given in spring 2007, the second round went during the spring 2008 semester. Each class consisted of 26 bi-weekly lectures of 75—85 minutes each, complemented by extensive home work and a mid-term and final examination. This paper deals with the experience we gained from both runs of this course.

3. VIRTUAL CLASSROOM IMPLEMENTATION

A state-of-art low-latency multi-point videoconferencing system based on HD video over IP transmissions together with supporting technologies was implemented (Holub et. al., 2006) and deployed (Matyska et. al., 2007, Hutanu et. al., 2007) to provide a high quality environment for both students and teachers. The system is deploys full HD effective resolution with 60 interlaced fields per second and the resolution of 1920x1080 pixels. We used 4:2:2 color sampling and 10 bit color depth per plane (however, only 8 bit per packet plane, i.e. 24 bits color depth, was used in the projection, as we do not posses full 10 bit per color plane projection equipment). The uncompressed video was send as HD-SDI IP packetized stream, with the total bandwidth requirement of 1.5 Gbps per each stream. This throughput is impossible to achieve with standard Ethernet packet size (MTU) of 1500 bytes, therefore we used 10 GigaEthernet (GE) Jumbo frames with 8500 Bytes each. Video stream is packetized with the modified UltraGrid software (Holub 2006). As it is not possible to use multicast for such high data streams, we used out own implementation of high speed data reflectors to guarantee the multi-point distribution. A dual CPU dual core AMD64 machines with at least 4GB RAM were used as reflectors. As at this transmission speed each reflector is capable just to duplicate the data stream (i.e., sending one input stream to two destinations simultaneously), we had to use a cascade of reflectors to provide fully multi-point distribution.

With the capability of deployed reflectors to only duplicate the uncompressed HD video streams we were not able to create full N:N distribution mesh among all sites. Therefore, we decided to use a simpler 1:N and N:1 models, where the HD videostream from the lecture hall with teacher was transmitted to all other classrooms (the 1:N model), but data from each classroom were sent only to the teacher (N:1). Therefore, only the teacher did see all the students, students themselves were always limited to see just their peers in the same physical classroom.

For audio part, we used a separate logical distribution network. The uncompressed audio (to keep it synchronized with the uncompressed HD video) is much less demanding, with data rates slightly above 1 Mbps. Therefore, we were able to use full N:N distribution mesh, all students were able to hear question posed by any student at any location, The same data reflectors were used in this case, as the smaller bandwidth requirement did not stress them in any way.

Data rates above 1 Gbps are usually not easily transmitted over an Internet, even in the academic environment. Although large network backbones like Internet2 and National Lambda Rail in the USA and GEANT in Europe are able to transmit tens of gigabits per second, usually the last mile within the campus or even the campus connection are not ready to sustain multi gigabits data rates. The physical infrastructure used for this particular virtual classroom implementation used experimental 10 GigaEthernet (10 GE) network which used dedicated links from Brno through Prague to StarLight and then National Lambda Rail line to Arkansas and Louisiana, where it has been further distributed over the Louisiana Optical Network Initiative (LONI) infrastructure. The logical network, implemented with use of packet reflectors, connected in the year 2007 LSU with Masaryk University (MU) in the Czech Republic, University of Arkansas (UARK), Louisiana Technical University (LATECH), Microelectronics Center of North Carolina (MCNC) and North Carolina State University (NCSU). In fact, there has been two physical classrooms at the LSU campus, which added one more stream into the final logical network setup.

The videoconferencing technology was partly changed during Spring 2008 High Performance Computing class as LSU and other US universities deployed commercial Polycom Ultimate HD technology. Polycom HD videoconferencing systems use a 1280x720 at 30 frames per second video compressed using H.264 video compression standard. The bandwidth used for one video stream is in order of megabits per second (compare to 1.5 Gbps of uncompressed 1920x1080 at 30 full frames per second full HD video). Although image quality of the Polycom HD video transmissions is inferior to those of uncompressed HD video transmission, the latency is acceptable and its deployment allowed more universities to participate in the class as they don't need to possess and maintain still rather expensive and not yet widely available 10 Gpbs network connectivity. The truly full HD uncompressed HD video transmissions was maintained only over the longest link, i.e., between Masaryk University and LSU where a 10Gbps link is available. Introduction of compressed HD videoconferencing brought new participants into the HPC class virtual classroom, namely Polytechnic University of New York and University of Arkansas at Little Rock in addition to those who already participated in the previous year.

The decision to support different technologies for the logically homogenous virtual classroom was based also on the last year experience. Using the leading edge technology, for the capture and projection of HD stream combined with the use of experimental network infrastructure lead to problems with stability and especially availability of the service. None of the high speed lines had a backup at the same speed, so any network service interruption meant unavailability of the real time HD stream. While this situation was not too often, it still happened several times during the spring 2007, including once on the transatlantic line. This is unacceptable for real time lecture, so backup technologies were used from the beginning to help to mitigate eventual service disruptions.

The basic backup technology used was a low-bandwidth AccessGrid videoconferencing system. All sites run also AccessGrid software on dedicated computers connected via standard Internet. The LSU lecture hall was among other things equipped with a camera connected all the time to the local AccessGrid system and the data were transmitted along with the HD streams. This helped to switch to the backup practically immediately after service disruption, this switch was fully manual and required one person at each remote site to follow the HD transmission and take care in case of service disruption. Due to rather high demands and novelty of technology used it was not possible to involve one of the students who took the lecture to also care of eventual problem and to perform all the necessary steps. Therefore, a dedicated person had to be employed.

The same AccessGrid technology was used also in 2008, providing the necessary backup. Although the experience from 2007 lead to the development of CoUniverse, a new framework for at least semi-automatic service switch in case of problems, the man power demand for the backup deployment remained also in 2008.

As additional backup solution, all classes were recorded in both raw uncompressed HD (one lecture requires almost 1 TB disk space in this case) and compressed Quicktime formats. The classes were also webcasted for those who were not able to participate in the classroom.

4. VIRTUAL CLASSROOM ISSUES

4.1. Technology issues

The HD video stream capture and projection system based on UltraGrid software proved to be mature enough not to create any serious problems during the whole class. The network, despite several disruption and unavailability at the precise time of lecture, was stable. The setup clearly demonstrated that current equipment and network infrastructure are capable to sustain regular production level stress at very high data rates. The major drawback was the necessity to involve experts in the network, high speed multi-point data distribution and the HD video manipulation, which is not a scalable solution for higher number of sites or more regular deployment (e.g., many lectures weekly). This extreme man power requirements for the use of purely experimental equipment and infrastructure lead to the second year decision to use the Polycom HD video system, as their deployment does not require high expertise and they can use Internet as their network requirements are moderate to low. While the video quality is lower, this is compensated by easy to use and “from box ready” equipment.

However, the audio caused much more problems than originally expected. As we focused on the HD video processing and transmission, including the multi-point distribution of very high data streams, we did not consider the audio a challenge. But human hearing is much less forgiving than the vision. The uncompressed video very well compensated for eventual small data loss, and human vision is able to deal with short loss of the visual input (in fact, the vision system itself trains our brain to compensate and re-create parts of picture). The audio system, on the other hand, does not tolerate even small loss at the input. Also, any background noise (generated e.g. by the simple computer audio system) makes the hearing much more difficult. Noise or audio glitches can lead to loss of focus for students, making the virtual classroom experience very unpleasant. Most of the audio related problems were removed using the professional Polycom HD systems that also integrate high quality microphone and echo canceling system. However, further improvements of the audio subsystem are necessary for fully immersive virtualized classroom experience.

4.2. Other Lessons Learned

Although the virtual classroom technology provides a low-latency videoconferencing capabilities and as such creating an interactive environment, the overall level of interactivity between the students and teachers or teaching assistants was usually very low. There might be many reasons for this lack of interaction and we tried to identify at least some of them. The occasional audio problems between remote sites mentioned in the previous section definitely preclude a full remote interactivity. Instead of feeling as part of the large virtual classroom, students were disturbed and perceived the distinction between “them” and “us”.

Another problem was lack of training of the use of videoconferencing systems. The use of uncompressed HD video reduced the latency, but it was still at the level of up half a second. Without training, it was not easy for students to interrupt teacher and ask questions remotely—they felt lagged, not being able to react fast enough to interrupt at the right moment.

We also felt that the distribution of tasks among individual sites—all homework was collected locally, hands-on were done using local equipment—prohibits extended collaboration among students at different universities. Probably also other effects of the type and the field of the class, as discussed in (Gudzial et. al, 2001), were responsible for the less then expected interaction.

The large geographic distance between LSU and MU together with different rules governing the beginning and end of a semester was a reason for the largest problems encountered. The spring semester starts in mid January at LSU, but only at the second half of February at MU. This is 6 week difference, with almost half of the lectures presented at LSU before the first lecture could be delivered to students at MU. In the first year of the experiment we somehow overlooked this issue and in fact were not able to catch up with the exception of several last lecture. All the lectures were recorded, with a remote recording. All the data streams were sent in real time to Brno, where one of our PhD students had a vigil. The raw uncompressed HD was directly stored in the high speed RAID system and the recorded lectures were replayed when the semester at MU started. With This “store and replay” system we actually checked the infrastructure for the interactive virtual classroom setup, but MU students did not directly benefit from it. On the other hand, it helped us to overcome the time zone problem. Lectures at LSU started at 3:30pm, which translated to 9:30pm (or even 10:30pm with non-synchronized summer time) at MU. These late night lecture are out of scope of the allowed lecture time, we needed an explicit consent from all enrolled students to participate in the interactive real-time lectures (it did not pose a real problem as this was the case only near the end of the semester).

As the start of the semester is based on a university wide policy, it was not possible to resolve this problem for year 2008. However, the remote delivery of the lectures helped to find more suitable time for lectures at LSU, as in 2008 they started at 10:30am (4:40pm at MU). Also, we proposed students to speed up the first half of the semester, with up to 4 classes per week (from the 2 classes per week delivered by Prof. Sterling). This imposed some additional logistic problems on the home work, as we had to modify the original schedule and desynchronize it with the lectures. Prof. Sterling, being aware of this problem, refrained from discussing the home work assignments during regular lectures and we were not forced to “censor” any lecture (removing discussion on the home work results). The high quality of HD video helped students to sustain the higher information load during the first half of the semester, giving them an opportunity to follow the lectures in the interactive real-time mode later on.

5. USER EXPERIENCE

Students attending the spring 2007 HPC class in the virtual classroom were asked to evaluate their experience with the virtual classroom capabilities and the HPC class in the virtual classroom. Students generally expressed high satisfaction with the quality of the HD video, as it allowed practically full immersion into the classroom environment. Unfortunately, students also expressed less positive valuation of the audio quality. Only few students rated the audio quality as high quality and not-disturbing. Other felt rather large discrepancy between the video and audio quality, with a negative impact of the lower audio quality on the overall experience with the HPC class.

Students were in general satisfied with the amount of supplementary material and the quality of prepared hands on. None of the students at MU had problems with the language—no translation of materials has been provided, even the local communication between the supervising local teacher (first author of this paper), local TAs and students went in English. They appreciated not only the immersion during the lecture—result of the deployed technology—but also the effort to create a realistic copy of a US university class environment.

To create such a realistic environment consumed a tremendous amount of man power of teaching assistants and other support personnel. In 2007, a local supervisor plus 5 other people (two PhD graduates, two PhD students and one Master student) were involved to keep the lecture “flying”. In 2008, the numbers did not change, while the focus and work distribution changed. In the first year, most effort went into the technology, making sure that all the components are working properly all the time. In the second year, the technology has been more mature (also as use of the HD Polycom systems in the US decreased the responsibility of MU people to maintain the whole environment—in 2007 one PhD student spent the whole spring semester at LSU to support all the technology also from the US side). As the result, only two persons—one PhD and one Master student—took care of the technology setup (with initial help of one assistant professor and another PhD student), the remaining members of the team were involved in direct support of students (the home work assignment preparation, checks and evaluations). The local teacher was responsible of the mid term and final exams and the evaluation of theoretical parts of home works as well as the small student’s projects.

We evaluated the man power needed to fully run the course as much higher than for any locally given lecture. However, we were able to reduce the man power responsible for the technology, with most of the remaining man power associated with a different way of running a lecture at LSU and at MU (the Masaryk University runs higher number of less demanding lectures).

Results of the 2008 student’s evaluation will be available for the presentation at the conference.

6. CONCLUSIONS

We implemented a high quality virtual classroom using the uncompressed HD processing and transmission system developed at MU. The virtual classroom has been used to deliver HPC class taught by Prof. Sterling to students at several US universities and at MU. We dealt with the technological as well as “political” and physical challenges, extending the virtual classroom concept to cover geographically very distant sites.

As a result, much higher number of students was able to attend an excellent lecture given by outstanding teacher with a high quality illusion of direct involvement. Students at MU generally positively evaluated the quality of the environment, confirming the usability of the HD video to create realistic immersion for teaching purposes. The experimental wide area high speed 10 GE network used also confirmed its near maturity, while initiating a new stream of research to better manage backup solutions in case of some failure. Audio was the most problematic part, but this problem was largely removed in the second year through more experience and the use of commercial products (however, the audio system still needs further development to match with the HD video quality). On the other hand, event the perceived high quality did not initiate enough interaction between students and the remote teacher, leaving space for further improvements.

The two year experiment revealed a major drawback of too high man power needed to setup and a high quality virtual lecture. To deliver a lecture and keep its quality in an environment that differs from the original one (as is the case for Czech Republic and USA university systems), high number of

TAs and their dedication is needed to guarantee not only the technological, but also the pedagogical quality of the virtual lecture.

We believe that changes in the lecture layout and the accompanied home work and hands on materials are need both to decrease the necessary man power and to increase the interactivity during the lecture. We plan to continue with the HPC lecture and similar experiments to find a suitable framework where the potential of the technology is met with the pedagogical results while keeping the TA's man power involvement at the level similar to (or even lower than) locally delivered lecture.

7. ACKNOWLEDGEMENTS

The financial support of CESNET's Research Intent MSM6383917201 is highly appreciated, as well as support at Masaryk University. Also, the authors would like to thank all the involved students and namely Tomáš Rebok from the Laboratory of advanced networking technologies at MU (a joint laboratory of Faculty of Informatics and Institute of Computer Science at MU and CESNET) for their help with running the lecture. And last but not least, the work would be impossible without dedication of Prof. Sterling and his team at LSU.

8. REFERENCES

Guzdial, M., Ludovice, P., Realff, M., Morley, T., Carroll, K., and Ladak, A. (2001). *"The challenge of collaborative learning in engineering and math".* In Frontiers in Education Conference, 2001. 31st Annual, vol. 1.

Hiltz, S.R. (1986). *"The "virtual classroom": Using Computer-mediated Communication for University Teaching"*, Blackwell Synergy.

Hiltz, S.R. (1994). *"The virtual classroom: learning without limits via computer networks"*, Ablex Publishing Corp. Norwood, NJ, USA.

Holub, P., Matyska, L., Liška, M., Hejtmánek, L., Denemark, J., Rebok, T., Hutanu, A., Paruchuri, R., Radil, J., Hladká, E. (2006). *High-definition multimedia for multiparty low-latency interactive communication.* Amsterdam, The Netherlands: Elsevier Science. Future Generation Computer Systems, vol. 22(8), p. 856-861.

Hutanu, A., Xin, Y., Thorpe, S., Paruchuri, R., Eiland, D., Liška, M, Holub, P. (2007). *Uncompressed HD Video for Collaborative Teaching—An Experiment.* White Plains, NY: ICST. Proceedings of the *3rd International Conference on Collaborative Computing: Networking, Applications and Worksharing*, 6p.

Matyska, L., Hladká, E., Holub, P. (2007). *Virtual Classroom with a Time Shift.* Kumamoto, Japan: Kumamoto university. Proceedings of the Eigth Conference on International Technology Based Higher Education and Training, Kumamoto, Japan. p. 131-135.