Available online at www.sciencedirect.com
SCIENCE DIRECT
ELSEVIER
Future Generation Computer Systems 22 (2006) 856–861
FGCS
FUTURE GENERATION COMPUTER SYSTEMS
www.elsevier.com/locate/fgcs

High-definition multimedia for multiparty low-latency interactive communication

Petr Holub [a,b], Luděk Matyska [a,b,*], Miloš Liška [a,b], Lukáš Hejtmánek [a,b], Jiří Denemark [a,b], Tomáš Rebok [a,b], Andrei Hutanu [c], Ravi Paruchuri [c], Jan Radil [a], Eva Hladká [a,b]

[a] CESNET z.s.p.o., Zikova 4, 16200 Praha, Czech Republic
[b] Masaryk University, Botanická 68a, 62100 Brno, Czech Republic
[c] Center for Computation and Technology, 302 Johnston Hall, Louisiana State University, Baton Rouge, LA 70803, United States

Available online 30 May 2006

Abstract

We describe a high-quality collaborative environment that uses High-Definition (HD) video to achieve near realistic perception of a remote site. The capture part, consisting of a HD camera, Centaurus HD-SDI capture card, and UltraGrid software, produces a 1.5 Gbps UDP data stream of uncompressed HD video that is transferred over a 10GE network interface to the high-speed IP network. The HD video stream displaying uses either a software-based solution with color depth down-sampling and field de-interlacing, or another Centaurus card. Data distribution to individual participants of the videoconference is achieved using a user-controlled UDP packet reflector based on the Active Element idea. The viability of this system has been demonstrated at the iGrid 2005 conference for a three-way high quality videoconference among sites in the Czech Republic, Louisiana, and California.
© 2006 Elsevier B.V. All rights reserved.

Keywords: High-Definition Video; Video Conference; Uncompressed HD Video; High-Speed Multipoint Data Distribution

1. Introduction

Evolution of collaborative environments, following the development of high-band-width low-latency networks brings new possibilities to the quality and extent of collaboration [1,2]. The truly interactive communication requires high-resolution video and audio transmitted fast over a network, with end to end latency below 100 ms to avoid hearing artifacts. Even if captured and transmitted independently, the video and audio must be kept synchronized and thus both transmitted with the lowest latency. All components of an end to end path contribute to the latency, so using uncompressed media is essential, together with almost no buffering. The network needs to provide support by close to zero packet reordering and loss, low latency, and minimal jitter. Such a network can be efficiently constructed based on dedicated circuits over a fiber optic network.

In this paper we describe a multiparty high-definition (HD) videoconferencing system and experiences gained from its use during the iGrid 2005 conference.

2. HD video transport and distribution

The highest resolution for the HD video defined by the common HDTV standard [3] is the 1080i mode with 1920×1080 points and interlaced line scanning. We use the uncompressed HD digital video defined in SMPTE 292M [4], which is transmitted through the Serial Digital Interface (HD-SDI). The bandwidth for such a video-stream with 60 interlaced fields per second, 10 bits per color plane, and 4:2:2 color space sampling is 1.485 Gbps. This payload is equivalent to around 1.5 Gbps over the IP network with 44 byte header per packet.

* Corresponding address: Institute of Computer Science, Masaryk University, Botanická 68a, 60200 Brno, Czech Republic. Tel.: +420 549492105; fax: +420 541212747.

E-mail addresses: hopet@ics.muni.cz (P. Holub), ludek@ics.muni.cz (L. Matyska), xliska@fi.muni.cz (M. Liška), xhejtman@fi.muni.cz (L. Hejtmánek), jirka@ics.muni.cz (J. Denemark), xrebok@fi.muni.cz (T. Rebok), ahutanu@cct.lsu.edu (A. Hutanu), ravi9@cct.lsu.edu (R. Paruchuri), radil@cesnet.cz (J. Radil), eva@fi.muni.cz (E. Hladká).

0167-739X/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2006.03.014

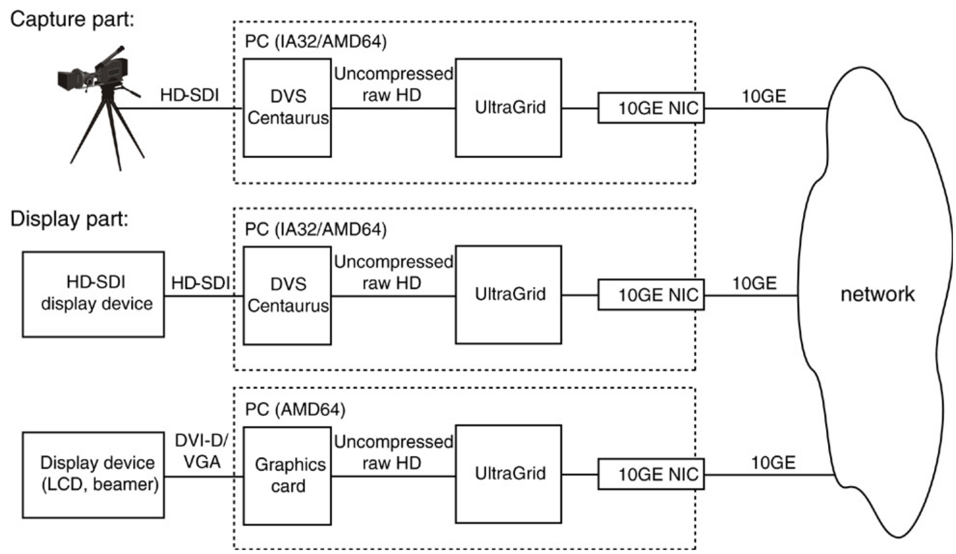

Fig. 1. Site connection scheme for HD transmission.

The capture and display parts of the system able to generate and deal with such data rate are depicted in Fig. 1. The capture part uses the DVS Centaurus (http://www.dvs.de/english/products/oem/centaurus.html) HD capture card, a solution dictated by the selection of Linux operating system environment. We have rewritten the UltraGrid software package [5] to support the 1080i mode. While the same path can be also used for video display (in reverse order), we have extended the UltraGrid with support for software only display, including a field de-interlace algorithm and color space down-sampling from 10 to 8 bits per color plane to avoid the use of expensive Centaurus cards on the receiving end of each HD video path. The computation and data manipulation intensive parts of the UltraGrid software have been optimized for deployment on AMD64 (Opteron) based computers.

The whole capture part starts with the Sony HVR-Z1E camera whose analog output is converted to HD-SDI using the AJA HD10A converter (http://www.aja.com/hd10a.htm). The HD-SDI stream is captured by the DVS Centaurus card, encapsulated into the UDP/IP stream by the Ultra Grid and sent via the Chelsio T110 LR 10GE card (http://www.chelsio.com/T110.htm) to the network. On the receiving end, the IP stream is captured by the same 10GE card, stripped of the IP header, color down-sampled and de-interlaced by the UltraGrid and sent to the graphic card. We use dual AMD64 Opteron 250 computers running at 2.4 GHz with 4 GB RAM and Centaurus and Chelsio cards placed in the PCI-X 133 MHz slots. Linux kernel 2.6.6 is used with drivers and manufacturer provided patches for both cards.

The total end to end latency measured in a laboratory set-up has been 175 ± 5 ms with both computers on the same 10GE Cisco Catalyst 6506 switch. While being above the optimum 100 ms threshold, 175 ms is still acceptable and hardly noticeable on the level of human perception. We have also analyzed part-wise latencies: four fields buffered on the capture card[1] (66.7 ms, calculated), 10 b to 8 b down-sampling ($7 \pm .5$ ms, measured in software), de-interlacing ($7 \pm .5$ ms measured in software), software display ($41 \pm .5$ ms, measured in software), LCD display delay (25 ms, according to manufacturer's specifications). This counts up to 147 ms and we are attributing the remaining delay (28 ms) to in-camera processing, HD-SDI conversion, buffering, and the processing on the graphics card. It is worth noting that the duration of down-sampling and de-interlacing is limited by the memory copying speed (approx. 1 GBps on test-bed machines). The sender and receiver CPU load has been 25% and 72% on average, respectively.

While the audio latency is the main driver for the use of uncompressed video streams, audio does not generate data flows of comparable bandwidth. In our application, the audio stream is generated by the `rat` tool [6] and synchronized on the receiving machine via the capability of UltraGrid to utilize time information from the RTP/RTCP packets.

The HD video may also be transmitted in the compressed HDV format at the cost of latency increase. HDV is a proprietary MPEG-2 based compression scheme developed by Sony and it has been designed to be transmitted in the MPEG-TS envelope over the IEEE-1394 interface, similar to the DV format transmission. It uses only 1440×1080 resolution, 50 or 60 interlaced fields per second, 8-bit color space, 4:2:0 color space sampling, and interframe 60:1 compression resulting in approximately 25 Mbps video stream. We have implemented a tool for the FreeBSD operating system [7] to read out the HDV data from the IEEE-1394 encapsulation. The data is sent over the network either using VideoLAN Client (VLC) (http://www.videolan.org/) or some other tool like `netcat` (http://netcat.sourceforge.net/) and it may be rendered using a VLC tool.

The measured end to end latency for the 60i HDV stream in the same laboratory set-up as mentioned above has been 1907 ± 13 ms. As the delay of camera compression is below 1 s, the most of the total latency is accumulated in the VLC buffering, decompression and rendering process. The almost

[1] The 2 fields reported by DVS did not work.

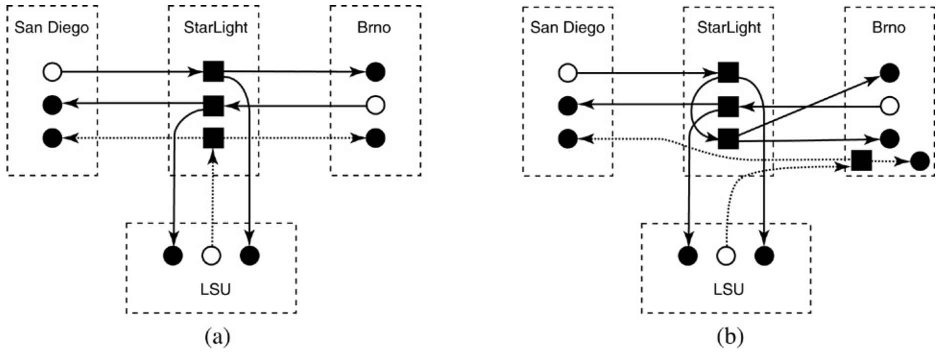

Fig. 2. Connection scheme for (a) the first and (b) the second iGrid experiment. The full lines show uncompressed HD streams; the dotted line shows compressed HDV streams. The empty circles denote sending computers, the full circles stand for display computers, and the full boxes are AEs.

Table 1
Reflector performance on the testbed

(a)		(b)		
Packet size (B)	Max. bandwidth (Mbps)	Bandwidth (Mbps)	Packet loss (%)	CPU load (%)
100	100	1800	0.0	52
500	300	1900	0.0	55
1500	400	2000	0.0099	60
3000	800	2100	0.037	76
6000	1700	2200	1.74	80
9000	2000	2300	7.07	84

(a) Gives maximum bandwidth of a single stream with respect to packet size in use given packet loss <0.01%; (b) gives packet loss with respect to bandwidth of a single stream given 8500 B packet size.

2 s latency practically hinders any real time communication. If it needs to be used, the audio stream must be independent and should not be synchronized with the video. This way, acceptable interactivity level is guaranteed via audio while video visibly lags behind.

To distribute data among more locations in a multipoint videoconference, we used the generalized Active network Element (AE) [8] based on the UDP packet reflector design [9]. This gives us more control over the distribution than a network native scheme (multicast or broadcast), which also may be unavailable for very high-speed networks.

We have optimized the AE to provide a sustainable UDP packet replication rate of up to 2.0 Gbps on high end IA-32 or AMD64 computer. The actual performance on the dual Opteron 250 computer (the same that has been used for the video capture and display) is summarized in Table 1. The results confirm the necessity to use Jumbo frames (long MTU) to achieve the highest performance. Also, the network card to memory bandwidth is important, as the use of an only 100 MHz PCI-X slot reduced the usable reflecting capability to 1.7 Gbps (this is the measured limit above which packets started to be lost significantly). Total latency increase induced by the replication measured on the test bed described above was 13 ± 2 ms.

3. iGrid 2005 experiment

The goal of the CZ101 iGrid 2005 experiment was providing a low-latency multiparty collaborative environment with HD video. Three sites participated: iGrid premises at San Diego UCSD campus, Masaryk University/CESNET in Brno, Czech Republic, and Louisiana State University (LSU) in Baton Rouge. Three networking circuits (in fact L3 networks) met at StarLight, Chicago, each coming from one participating site: an iGrid circuit from San Diego (RTT $78.2 \pm .2$ ms, 4 hops), an NLR circuit from LSU (RTT $31.09 \pm .04$ ms, 2 hops), and a circuit from Brno spanning CzechLight and NetherLight (RTT $126.7 \pm .3$ ms, 2 hops).

The video capturing and displaying set-up described above was used at all the sites with the exception of LSU, where the unavailability of the Centaurus card[2] led to the fallback HDV solution. The video was displayed using 1920×1200 resolution at all sites (HD video was wrapped with black borders on top and bottom of the screen, while the HDV had to be up-scaled first).

Three reflectors were set up at StarLight, each replicating a stream from one site – two 1.5 Gbps and one 25 Mbps streams – to the remaining two sites (see Fig. 2(a)). To stress the infrastructure, in the second part of the demo we cascaded two of the reflectors so that Brno was receiving two identical streams from San Diego (see Fig. 2(b)). This way we achieved total network flow of 4.5 Gbps on the Brno–StarLight circuit. This set-up has also proven the feasibility of combining AEs into an AE network to provide scalable data distribution [8] even at very high data rates.

The total unidirectional end-to-end latency for the uncompressed HD distributed using a single reflector was 290 ms in the worst case (San Diego–Brno), while in the best case (San Diego–LSU) it was 242 ms. The total latency was noticeable during the experiment, but it was barely disturbing.

The WAN-PHY interface of the Cisco Catalyst 6506 switch in Brno was used to gather network statistics shown in Fig. 3 during the last demonstration. This was the uplink interface that aggregated all the traffic from and to the Brno site. The outgoing traffic was rather regular, while the incoming traffic displayed some burstiness. We attribute this behavior to the use of the L3 network, where the packets are still buffered at some intermediate network elements and thus the latency is rather uneven. We expect utilization of a pure end to end optical network with no intermediate buffers to remove this problem.

[2] It has been stuck at US customs because of hurricanes.

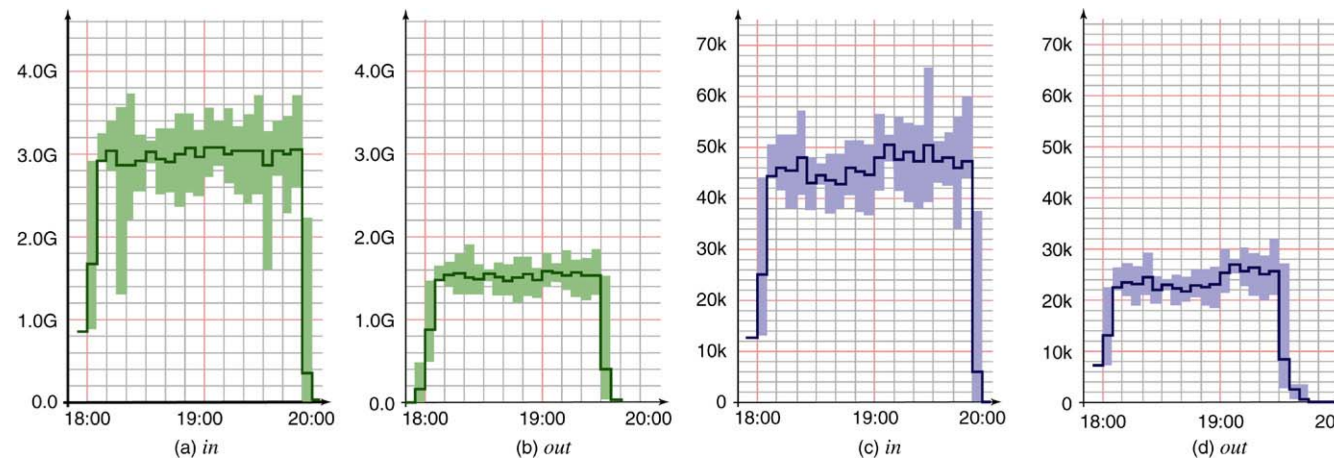

Fig. 3. Network statistics gathered on the Brno uplink. The lines are five minute averages, while the surface is an envelope curve with min/max values within the five minute interval. (a) gives *in* throughput in Gbps, (b) gives *out* throughput in Gbps, (c) gives *in* throughput in packets per second (pps), and (d) gives *out* throughput in pps.

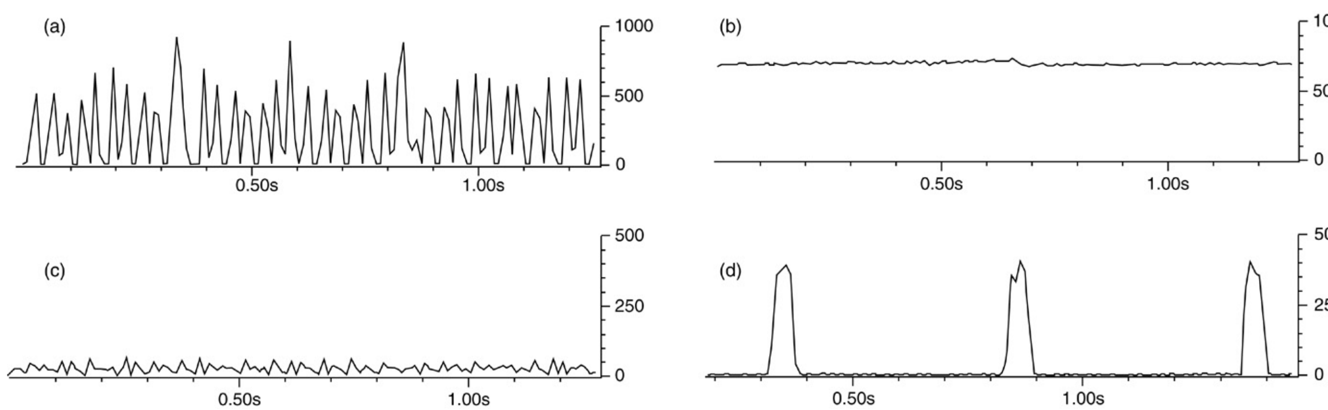

Fig. 4. Bursty traffic produced by the uncompressed HD transport (a) compared to smooth traffic by our measurement probes (b). Bursty traffic produced by the HDV transport with VLC (c) and `netcat` (d). Vertical axis denotes number of packets per 20 ms interval.

4. Experiences, problems, and related work

The UltraGrid produces very bursty traffic (see Fig. 4), that reduced the duplicating capacity of the AE to approximately 200–600 Mbps with large fluctuations; no packet loss is observed inside the AE software, as the data was actually lost during reading from the received packet queues of the underlying operating system. Similar burstiness occurs also with the HDV transport as shown in Fig. 4. We modified the AE to use the non-blocking `read()` function giving explicit precedence to the sending thread when no input data is available. While this increased CPU load to 100%, the duplication runs at the target speed.

Another source of problems we encountered was the overheating and instability of dual Opteron computers with Chelsio and especially with both Chelsio *and* Centaurus cards. The fast assemblage and set-up of dual Opteron computers at San Diego even resulted in one of the machines being unable to receive or send data with the Chelsio card above 800 Mbps (the same card worked perfectly if moved into any of the remaining two machines).

Related Work. Our work follows development of UltraGrid [5], enhancing it with 1080i resolution support and extending its software display including de-interlacing and color space scaling. We also further extend our work on Active Elements for the efficient data distribution under strict user control (see discussion in [8]). During the iGrid 2005, ResearchChannel had a similar demonstration [10] of multipoint HD videoconference. In contrast to their set-up, we used the software display instead of HD-SDI, independent audio stream (externally synchronized), and Active Elements instead of multicast for data distribution.

5. Conclusions and future work

We have demonstrated that a high-quality multiparty videoconference based on transmission of uncompressed HD streams is already achievable provided adequate networking resources are available. The same data distribution and HD video transmission set-up has also been used in combined visualization demo during the conference. Although utilizing software reflector, L3 network instead of dedicated optical circuits, and excessive frame buffering on Centaurus capture card led to higher than theoretically minimal latency and jitter, the whole set-up provided a realistic high-quality collaborative environment.

However, the experience has also demonstrated deficiencies and open challenges suggesting future research and development. At the network level, we plan to repeat the experiment over L2 and L1 networks to achieve smaller latency and jitter [11]. We also plan to replace software based Active Elements with optical splitters that would further reduce the latency for multiparty transmission at the L1 level. The rather complicated set-up of the network for the demonstration also proved the necessity for a special control plane over (optical) networks spanning several administrative domains if dedicated circuits are to be provided effectively. On the application level, we will focus on employing hardware with lower buffering requirements for latency reduction and on implementing smoothing algorithms for packet sending to mitigate problems with bursty traffic.

Acknowledgments

We would like to thank Tom Košnar, Martin Míchal and Josef Vojtěch from CESNET for network statistic gathering and European optical line provisioning, and Alan Verlo, Pieter de Boer, Paola Grosso and other people at iGrid 2005 NOC for their support. Also, the help of Cisco and Chelsio (both lending some equipment) is highly appreciated. This project has been supported by a research intent "Optical Network of National Research and Its New Applications" (MŠM 6383917201) and "Parallel and Distributed Systems" (MŠM 0021622419).

References

[1] T.A. DeFanti, M.D. Brown, C. de Laat, iGrid 2002: The international virtual laboratory, Future Generation Computer Systems 19 (6) (2003) 803–804.

[2] R. Singh, J. Leigh, T.A. DeFanti, F. Karayannis, TeraVision: A high resolution graphics streaming device for amplified collaboration environments, Future Generation Computer Systems 19 (6) (2003) 957–971.

[3] Society of Motion Picture and Television Engineers, 1280 × 720 scanning, analog and digital representation and analog interfaces, SMPTE 296M-1998.

[4] Society of Motion Picture and Television Engineers, Bit-serial digital interface for high-definition television systems, SMPTE 292M-1998.

[5] C. Perkins, L. Gharai, T. Lehman, A. Mankin, Experiments with delivery of HDTV over IP networks, in: 12th International Packet Video Workshop, Pittsburgh, PA, USA, 2002.

[6] V. Hardman, A. Sasse, M. Handley, A. Watson, Reliable audio for use over the internet, in: Proceedings of INET'95, Honolulu, Hawaii, 1995.

[7] P. Holub, HDV capture for FreeBSD 5/6 operating system. http://sitola.fi.muni.cz/~hopet/HDV/, http://docs.freebsd.org/cgi/getmsg.cgi?fetch=37672+0+archive/2005/freeb%sd-firewire/20050206.freebsd-firewire.

[8] P. Holub, E. Hladká, L. Matyska, Scalability and robustness of virtual multicast for synchronous multimedia distribution, in: Networking—ICN 2005: 4th International Conference on Networking, 17–21 April, Reunion Island, France, 2005, Proceedings, Part II, in: Lecture Notes in Computer Science, vol. 3421/2005, Springer-Verlag, Heidelberg, La Réunion, France, 2005, pp. 876–883.

[9] E. Hladká, P. Holub, J. Denemark, An active network architecture: Distributed computer or transport medium, in: Proceedigns of 3rd International Conference on Networking, ICN'04, Gosier, Guadeloupe, 2004, pp. 338–343.

[10] M. Wellings, J. DeRoest, A. Philipson, J. Eveleth, J. Brown, C. Latham, R. Johnson, G. McLaughlin, A. Howard, J. O'Callaghan, M. Lack, E. Verharen, D. Devereaux-Weber, A. Kato, Global N-way interactive high-definition video conferencing over long-pathway, high-bandwidth networks, US118 iGrid 2005 demo. http://www.igrid2005.org/program/applications/videoservices_nwayconf.html.

[11] C. de Laat, E. Radius, S. Wallace, The rationale of the current optical networking initiatives, Future Generation Computer Systems 19 (6) (2003) 999–1008.

Petr Holub graduated at Faculty of Sciences, Masaryk University in Brno and received the Ph.D. in computer science from Faculty of Informatics MU in informatics, focusing at high-speed networks, multimedia, and parallel and distributed systems. Currently he works at Institute of Computer Science MU in the Laboratory of Advanced Networking Technologies and participates on its scientific leadership. He is also a researcher with CESNET. His professional interests include high-speed networks and suitable protocols, active networks, user-empowered overlay networks, advanced collaborative environments, grid environments, and computational quantum chemistry and its implementation on distributed systems.

Luděk Matyska is an Associate Professor in Informatics at Faculty of Informatics, and he also serves as a vice-director of Institute of Computer Science, both at Masaryk University in Brno, Czech Republic. He got a Ph.D. in Chemical Physics from Technical University Bratislava, Slovakia. His research interests lie in the area of large distributed computing and storage systems, with a specific emphasis on their management and monitoring. He also works in high speed network applications, with a specific emphasis on collaborative work support and use of all these technologies in various e-learning activities. He lead national Grid infrastructure projects and participates in several EU funded international projects including the CoreGRID Network of Excellence.

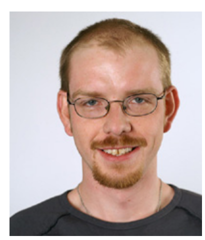

Miloš Liška is a Ph.D. student at Faculty of Informatics at Masaryk University. He is also a CESNET researcher working on projects concerning multimedia, tools for collaborative environments and collaborative workflows. He is interested especially in tools for multimedia processing and video transmission.

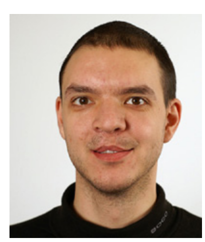

Lukáš Hejtmánek graduated in computer science and currently he is a Ph.D. student at Faculty of Informatics, Masaryk University. His main research interests include: high speeds networks, network data storage and peer to peer storage, distributed storage with client–server semantics, data replicas, replicas management, and data indices.

Jiří Denemark graduated in computer science at Faculty of Informatics, Masaryk University. Currently he is a Ph.D. student at Masaryk University and participates in CESNET research projects focused on multimedia data transmissions and processing and environment for virtual collaboration. His research interests are active overlay networks, virtual multicast, network support for collaborative systems, virtual machines and utilizing their technologies in grid environments.

Tomáš Rebok graduated in computer science at the Faculty of Informatics at the Masaryk University in Brno and currently he is a Ph.D. student at the same faculty. He is also a researcher with CESNET, working on problems of synchronous multimedia distribution and network performance measurement technologies. His research interests include active routers and distributed routers with QoS support on high-speed networks.

Andrei Hutanu graduated in 2002 and received his engineering diploma in Computer Science from Politehnica University of Bucharest. Currently he is a researcher in the Grid and Visualization departments of the Center for Computation and Technology at Louisiana State University. Prior to this he was employed for two years as a researcher in the Visualization department of the Zuse Institute Berlin. His research interests are distributed interactive visualization, data management and optical network applications.

Ravi Paruchuri received his B.Tech. from University of Madras in computer science and engineering and his M.Sc. from Louisiana State University in systems science. Currently he is the manager of IT Operations within the Center for Computation and Technology at Louisiana State University.

Jan Radil received the M.Sc. and Ph.D. degrees in electrical engineering and electrotechnology from the Czech Technical University, Prague, in 1996 and 2004, respectively. He joined the Research and Development Department, CESNET, Prague, in 1999, where he is responsible for optical networking and development of next generation of the experimental research networks.

Eva Hladká is an Assistant Professor in Informatics at Faculty of Informatics, she is a head of Laboratory of Advanced Networks Technology. She is also a senior researcher with CESNET. She got Ph.D. in Computer Science from Masaryk University in Brno. Her research interests include user empowered network technology and active networks, network support for collaborative environments, virtual multicast, and P2P networks of communication elements. She also works on collaborative environments for e-learning and telemedecine.