Pushing Boundaries of User-Empowered Network-Centric Collaborative Environments

Petr Holub

Habilitation Thesis
The Faculty of Informatics, Masaryk University

Brno, 2013

Except where otherwise indicated, this thesis is my own original work.

Petr Holub
Brno, September 2013

Acknowledgments

First and foremost, my thanks go to my family—my wife Aleška as well as our children Marek and Barbora, who are all very supportive in their own unique ways. I would also like to thank my parents and grandparents, who guided me through many steps of my education and were always good friends of mine.

Second, I have to thank my colleagues and friends: Luděk Matyska, Eva Hladká, Jirka Matela, Miloš Liška, Lukáš Hejtmánek, Martin Pulec, Martin Jirman, Martin Šrom, Andrei Hutanu, and a number of others. Collaborative work as well as stimulating discussions helped me in formation of the basic ideas behind the concepts presented in this thesis. Last but not least, my thanks go to my employers, both Masaryk University and CESNET, where I have a chance to work on the borderline between research and real-world high-performance infrastructures. I am also grateful for the university environment, where teaching students and working with them on research projects helps to keep one's mind fresh and open.

P. H.

Contents

Abstract

Network-centric collaborative environments have become common in many fields of human activity. Emergence of high-speed networks in the last decade enabled implementation of high-quality multimedia communication, which in turn facilitated adoption of the collaborative environments in the fields where their adoption seemed very unlikely or outright unimaginable, such as remote endoscopic operations in medicine or remote cinematographic production in entertainment industry. The overall focus of my research work is on pioneering techniques that enable high-performance network-centric collaborative applications without administrative rights on the networking or computing infrastructures. This is critical for deployment of collaborative systems in real world and we call it "user-empowered" approach.

The goal of this thesis is to provide an overview of my select results that are contributing novelty to the state of the art. The major results have been achieved in the following fields: (a) GPU-accelerated compression of image data (JPEG, JPEG2000), including reformulation of some of the critical algorithms to achieve high level of parallelism, which allows substantially higher performance compared to common CPU-based solutions, while maintaining advantages of commodity hardware; (b) the concept of high-performance programmable user-empowered overlay networks, including optimization of data distribution with flows comparable to capacities of network links; it maintains high performance of the data distribution (despite being not theoretically optimal compared to network-native multicast) and allows advanced features such as per-user processing, high level of security, and operation in restricted networks; (c) self-organization, monitoring, and adaptation of high-performance media applications, including realistic aspects such as limited knowledge of the infrastructure. These approaches are also applicable beyond the collaborative systems, for example in grid environments.

Being one of the architects of national collaborative infrastructure as a part of the National Large Infrastructure for Science and Research operated by CESNET consortium, I have also focused on implementing the results into the infrastructures and practical applications mostly under open-source licenses. These practical experiences further stimulate my research by forcing me to think about various realistic aspects that can be neglected in theoretical research but that are critical in real-world applications.

This habilitation thesis has been written as a collection of peer-reviewed journal and conference papers, a book chapter, and a patent. The commentary part of the thesis provides an integrated view of my work, sets it into a broader context, and states my contribution to each of the papers. My average contribution to these papers is 39%, obviously within the limits of very approximative nature of such a metric. My effort in implementing the results of the research into the practical applications has been recently acknowledged by The Best Open-Source Software Competition Award by ACM Multimedia Special Interest Group during prestigious ACM Multimedia'12 Conference.

Abstrakt (in Czech)

Síťová prostředí pro spolupráci se staly běžnou součástí mnoha lidských činností. Příchod vysokorychlostních sítí v poslední dekádě pak umožnil využití vysoce kvalitní multimediální komunikace, což následně dovolilo nasazení prostředí pro spolupráci i v oblastech, kde se jejich využití zdálo nepravděpodobné nebo přímo nemyslitelné, jako jsou vzdálené endoskopické operace v medicíně či interaktivní vzdálená spolupráce v kinematografické produkci. Celkovým zaměřením mé práce je výzkum technik, které umožňují realizovat výkonná síťová prostředí pro spolupráci, a to bez potřeby mít k infrastruktuře administrativní práva. Tento koncept, který je kritický pro nasazení prostředí pro spolupráci v reálných aplikacích, nazýváme uživatelem řízené prostředí.

Cílem předkládané práce je poskytnout přehled mých vybraných výsledků, které přinášejí posun v tom kterém poli působnosti. Hlavní výsledky byly dosaženy v následujících oblastech: (a) GPU-akcelerovaná komprese obrazových dat (JPEG, JPEG2000), která díky reformulaci algoritmů a výsledné vysoké paralelizaci dosahuje podstatně vyšších výkonů v porovnání s tradičními řešeními na CPU, přičemž zachovává výhody komoditních hardwarových řešení; (b) koncept výkonných programovatelných uživatelem řízených překryvných sítí, a to včetně optimalizace distribuce dat srovnatelných s kapacitami linek; při zachování vysokého výkonu (ačkoli ne zcela optimálního v porovnání s nativním síťovým multicastem) tyto techniky umožňují také specifické vlastnosti jako individuální zpracování dat pro jednotlivé uživatele, vysoké zabezpečení či průchod přes různě omezená síťová prostředí; (c) samoorganizace, monitoring a adaptace výkonných multimediálních aplikací, a to včetně podchycení realistických aspektů jako je například omezená znalost síťové topologie. Navržené postupy najdou své uplatnění i mimo prostředí pro spolupráci, například v oblasti gridových prostředí.

Jako jeden z architektů národní infrastruktury prostředí pro spolupráci v rámci Národní infrastruktury pro vědu a výzkum provozované sdružením CESNET jsem se také zaměřil na implementaci výsledků výzkumu do reálných praktických aplikací, povětšinou uvolněných pod licencemi s otevřeným zdrojovým kódem. Získané praktické zkušenosti pak zpětně stimulují můj výzkum tím, že mne nutí přemýšlet nad realistickými aspekty, které bývají přehlíženy či zanedbávány v teoretickém výzkumu a které jsou nicméně kritické pro uplatnění výsledků v praxi.

Předkládaná práce byla sepsána jako soubor recenzovaných článků z impaktovaných časopisů, konferencí, kapitoly z knihy a patentu. Úvodní část s komentářem poskytuje ucelený pohled na mou práci, zasazuje ji do širšího kontextu a formuluje mé příspěvky k jednotlivým publikacím. Můj průměrný příspěvek k zahrnutým publikacím je 39 %, s přihlédnutím k velmi aproximativní charakteristice takové metriky. Mé úsilí při implementaci výzkumných výsledků do praxe bylo nedávno oceněno cenou The Best Open-Source Software Competition Award od ACM Multimedia Special Interest Group, během prestižní konference ACM Multimedia'12.

Glossary

4K One of the resolutions 3840×2160, 4096×2160 or similar, as common in digital cinema or latest generation of broadcasting standards

8K The resolutions of 7680×4320, which is a proposed future broadcasting standard with estimate first deployments between years 2016–2020.

AC (coefficient) See DC coefficient.

ACM Association of Computing Machinery

CM Context Modeler of JPEG2000

CUDA GPU programming model by NVIDIA company

DC (coefficient) The coefficient expressing zero frequency in both directions in the output of the DCT. The other coefficient, describing alternating frequencies, are called AC.

DCI Digital Cinema Initiative

DCT Discrete Cosine Transform

DWT Discrete Wavelet Transform

EBCOT Embedded Block Coding with Optimal Truncation, part of JPEG2000

JPEG Joint Picture Expert Group and also a name of one of the most common image compression standards

JPEG2000 Image compression standard, a successor to JPEG

GLSL OpenGL Shader Language

GPU Graphics Processing Units

H.323 A videoconferencing standard for IP networks defined by ITU-T

HCI Human-Computer Interaction, a research field bridging computer science with psychology to study interactions of humans with technology.

HD High-Definition (video, audio, etc.)

FEC Forward Error Correction

GPRS General Packet Radio Service, a packet-oriented service in cellular communication networks.

IETF Internet Engineering Task Force, an Internet standardization body

ITU International Telecommunication Union

ITU-T ITU Telecommunication Standardization Sector, formerly known as Comité Consultatif International Téléphonique et Télégraphique (CCITT)

MCU Multi-point Control Unit, a component of H.323/SIP videoconferencing systems that provides multi-point interconnection capability to the end-points, which can create only point-to-point connections.

NAT Network Address Translation

NCCE Network-Centric Collaborative Environments

RLE Run-Length Encoding

SIGMM ACM Multimedia Special Interest Group

SIMD Single Instruction Multiple Data

SIMT Single Instruction Multiple Threads, a slightly more flexible concept compared to SIMD, introduced by NVIDIA

SIP Session Initiation Protocol, a videoconferencing protocol standardized by IETF

USPTO United States Patent and Trademark Office

WDM wavelength-division multiplexing

WFS Wave Field Synthesis

Part I

Commentary

Chapter 1

Introduction

Network-centric collaborative environments (NCCE) have found a foothold in many fields of human activity, ranging from highly respected medical and rescue applications, though many professional, scientific and educational applications, to entertainment industry. NCCEs span a very broad range of purposes, from low quality communication optimized to work even in very adverse conditions, e.g., medicine in battlefield, to high-performance systems that are the major focus of my work.

1.1 History Overview

Early experiments with practical video transmission systems date back to late 1920s by Bell Laboratories in the USA [45]. First long-distance video transmissions as a service became available during years 1936–1940, when the first closed-circuit television with 180 video lines resolution at 25 frames per second was implemented by dr. Georg Schubert at Reich Postzentralamt [1], eventually spanning about 1000 km. The development at Bell Laboratories proceeded further after the World War II, eventually resulting in a commercial product called AT&T Picturephone in 1960s [67], but it did not achieve commercial success.

Many of the concepts that are currently perceived as commonplace in computer mediated communication were pioneered in 1960s by oN-Line System (NLS) by dr. Douglas Carl Engelbart [48]. These included concept of computer control using a mouse, 2D editing or linking of documents and hypermedia. From the NCCE perspective, his work included videoconferencing with screen sharing, computer-aided meetings with shared document editing, document version control, as well as other features that we well ahead of their time.

Development of computer networks. In the next decade, the foundations of Internet were laid by the ARPA, followed by inception of IP multicast [8, 28, 36] in 1980s for network native scalable group communication. In 1992, the multicast was used to develop MBONE [26, 49, 157], multicast backbone network implemented as an overlay on top of unicast IP network, and MBONE Tools [25, 69, 70, 125], a set of multimedia tools for interactive communication based on the multicast service. Because of relatively slow adoption of multicast and problems with global network-native multicast availability across various administrative domains [39, 43, 44], it gradually became clear that implementation of the multi-point data distribution functionality is necessary on a higher layer for reliable operation of collaborative environments; implementation of this functionality on application layer became incorporated into the industrial H.323 and SIP videoconferencing standards, first versions of which became standardized by 1996.

Since the inception of peer-to-peer systems [159] in late 1990s, this architecture became studied also for real-time communication and video distribution. One of the first systems was Virtual Room Videoconferencing System developed at CALTECH since 1995, which uses a network of packet reflectors to process the data [3, 132]. The system itself, however, was not fully peer-to-peer: beyond the network of peer reflectors, the system had one main coordinating and user interface server and end-point clients also stayed outside of the peer-to-peer network. Also, the users did not have any control of the environment and were unable to setup their own reflector nodes. The field of video distribution using peer-to-peer networks has become hot research topic between years 2000 and 2005 [14, 37, 65, 72, 139, 140, 171, 183], [Appendix E]. An important milestone was development of Skype in 2003, which provided the first implementation of efficient application-level data distribution concept using overlay peer-to-peer network practically available to general users beyond the research community. The system has become a de-facto standard in personal audioconferecing and later also videoconferencing solutions. Because of its large-scale target audience, the system has been lacking the user-empowered characteristics in terms of the control of the environment, in order to achieve stability and minimize problems with misbehavior of the users.

After the research field of peer-to-peer became somewhat exhausted, focus of the community shifted to content delivery networks and their optimizations [11, 82, 113, 134, 135, 181, 182, 184, 186], and streaming using commodity protocols such as HTTP [5, 6, 127, 128, 148] and interaction of streaming with congestion control of the TCP protocol [12, 31, 53, 61, 85, 146, 175], which are however mostly relevant only for unidirectional distribution of video and not for interactive applications. Other investigated area is optimization of streaming simultaneously to many users over wireless network links [58, 83, 107, 116, 160, 167], which introduce new level of problems due to the shared nature of the transport medium and their complex temporal characteristics. A smaller community of the researchers, to which my team belongs, further pursues use of high-performance networks to support high-fidelity NCCEs [162, 170], [Appendix A].

Video coding and compression. Practical deployment of the video transmission has been determined by efficiency of its coding and compression; this fact became clear in 1970s when researchers began to see the potential of long-distance video distribution and video-mediated communication. The research in applying concepts of more general data coding to video resulted in the one of the very first standards, ITU-T H.120 (also known as COST 211) based on differential pulse-code modulation in 1984. This was not particularly efficient because of its per-pixel nature and it became obvious that for efficient image and video distribution, one needs coding resulting in less than 1 b per pixel. This stimulated substantial research and development of block-based compression schemes such as intra-frame ITU-T H.261 compression in 1988 and inter-frame MPEG-1 compression in 1992, which standardized concept of I-frames (intra-coded frames), P-frames (forward predicted frames), B-frames (bidirectionally predicted frames), and also not much used D-frames (intra-coded preview frames using DC coefficients). Relevant standard is also JPEG completed in 1992, which has been applied to the video coding as well. Efficiency of the inter-frame compression was later improved by MPEG-2 standard in 1996 (1995 for audio), developed on the similar fundamental concepts as MPEG-1, while bringing in mostly technical improvements (support for higher resolutions, interlaced video, etc.), optimization for higher compression ratios, and splitting of various features into profiles to simplify implementation. Further research into more-efficient coding resulted in MPEG-4 standard in 1998 and a its follow-up H.264 (MPEG-4 Part 10, AVC) in 2003, which included concepts like variable block-size motion compensation, sub-pixel motion compensation, slice-based structure, object-based composite files, or more efficient entropy coding using Context-based

Adaptive Binary Arithmetic Coder (CABAC). Primary purpose of the MPEG standards is one-way delivery of the video and thus overall image quality was considered much more important than latency. Therefore, another set of standards has emerged for the interactive applications: H.263 in 1996, H.263v2/H.263+ in 1998, and H.263v3/H.263++/H.263 2000 in 2000, roughly following the development in the MPEG standards, while adopting only those features that have low latency impact. For high-end imaging and video applications that suffered from various artifacts of compression schemes relying on discrete cosine transform (DCT), JPEG2000, incorporating discrete wavelet transform and very efficient entropy encoding using adaptive arithmetic coder, was created in 2000. However, its substantial computational requirements hindered its deployment as a JPEG substitute in the following years, limiting its use mostly to professional applications such as medical imaging [4, 54].

Intra-frame compression algorithms, some of which were based on still image compression algorithms such as Motion JPEG and Motion JPEG2000, and others designed specifically for video applications such as HDCAM, DV (inter-field, intra-frame), ProRes422, DCI JPEG2000 [35], CinemaDNG or REDCODE RAW, have been successfully applied for low-latency high-performance applications; their frame independence also simplifies transmission in networked systems prone to data loss, facilitates access to stored media and editing in non-linear editing systems, and provides increased robustness in case of media archive corruption.

The H.264 compression scheme, a successor to both MPEG-4 and H.263, became standardized in 2003 and it was supposed to be usable for both latency-insensitive unidirectional application as well as low-latency interactive applications. However, practical approaches of low-latency H.264 are often limited to certain subsets of full feature set, such as intra-frame coding[1] or single-frame VBV with periodic intra-refresh[2]. In 2013, H.265 standard was approved to further build on top of H.264.

Interesting experiments have also been done with the inter-frame compression based on wavelet transforms, such as Dirac format known also as SMPTE VC-2. However, because of problems with global nature of wavelets for motion estimation in inter-frame coding, others are re-exploring alternative approaches such as lapped transforms [118, 173] in the proposed Daala codec[3].

The video coding, compression and general processing has been greatly supported by advent of general-purpose computing on Graphics Processing Units (GPUs), which have become a kind of widely available commodity vector co-processors. Early developments of this field can be seen in OpenGL, whose Architecture Review Board has been established in 1992. The OpenGL kept adding various type of extension and gradually gained programmable processing pipeline using various types of user-definable shaders based on OpenGL Shader Language (GLSL). These concepts have been generalized by Cg (NVIDIA), HLSL (Microsoft in their Direct 3D), Lib Sh, and BrookGPU [21] languages and later resulted in the currently most common GPU programming concepts: NVIDIA CUDA and more general OpenCL originally conceived by Apple. Because of their vector nature and specific architecture with hierarchy of processing and memory resources, the design and implementation of algorithms for GPUs requires specific considerations in order to perform well. Certain types of algorithms will perhaps never perform particularly well though and some of these can also be found in image coding and compression.

[1] `http://www.cast-inc.com/news/post.php?s=2013-05-16-ultra-low-latency-h264-video-encoding-now-available-from-cast`

[2] `http://x264dev.multimedia.cx/archives/249`

[3] `https://wiki.xiph.org/Daala`

Algorithms such as DCT and DWT, which are often in the core of image compression algorithms, can be for instance easily implemented on GPUs, preferably using coefficient lifting schemes [50, 122, 137, 158, 169]. Also motion estimation, intra-frame prediction, and motion compensation algorithms, which are computationally heavy part of the inter-frame video compression algorithm, have been shown to be efficiently implemented on the GPUs [29, 30, 75, 91, 100, 101, 111, 147], improving performance of the compression algorithms by almost and order of magnitude. On the other hand sequential entropy coding and decoding processes, such as adaptive arithmetic (de)coding, and stream parsing are very unlikely to provide substantially higher performance on GPUs compared to CPUs. Overall, the GPU computing found a strong foothold in this field, although such architectures are often not taken into consideration when designing new algorithms. Also, the most widespread H.264 algorithm has been implemented using specialized chips on the commodity GPUs, most often to accelerate video playback. This is likely to happen again if some other algorithm becomes spread widely again, which may be the case for H.265 within several years from now.

High-performance network media transmissions and NCCEs. The research results and progress of technologies behind the high-performance media processing and distribution and NCCEs has been demonstrated in a number of events since 1990s. MBONE backbone together with the MBONE Tools was used to broadcast concert of Rolling Stones from Cotton Bowl in Dallas over the Internet [157, Chapter 6.4]. In 1999, there was a first demonstration of compressed high-definition (HD) video over Internet in HDCAM format at 270 Mbps by University of Washington team lead by Michael C. Wellings. In 2001, a team from McGill University lead by Jeremy Cooperstock created an audio/video environment with latency low enough (sub-20 ms) to facilitate a distributed violin[4].

An important milestone of high-performance interactive video transmission over the network was the famous Lindbergh Operation, where a patient in Strasbourg was operated by a surgeon in New Yorks, backed up by emergency on-site surgeons [120]. The operation involved remote robotic surgery of gall bladder, performed over trans-Atlantic private ATM network with 10 Mbps bandwidth, 70% of which was dedicated to the video in PAL resolution [23]. This also demonstrated adaptability of the humans to the latency.

In 2005, we were one of first two teams who independently developed first multi-point interactive videoconference using Internet, based on low-latency uncompressed HD video—one team was from the University of Washington and ResearchChannel led by Wellings, and the other was our team from CESNET and Masaryk University led by myself. While both demonstrations included long-distance transmissions (from an off-shore submarine research vessel in case of University of Washington demo and trans-Atlantic transmissions in the CESNET/MU demo), the key difference between these two demonstrations was in the data distribution part: our approach demonstrated advantages of the user-empowered approach based on overlay network of UDP packet reflectors, while the other demo had to struggle with multicast-related issues. In 2007, we have pioneered and successfully demonstrated use of optical multicast for long-distance multipoint transmission of uncompressed HD video, as well as self-organization of the environment for high-bandwidth video distribution.

First demonstration of low-latency JPEG2000-compressed 4K video at 60 frames per second over the Internet was in 2009, demonstrated using a USA–Japan videoconference across the Pacific by NTT Laboratories. About three months after first long-distance transmissions of 8K video during Summer Olympic Games by BBC and NHK in 2012, we were able to

[4] http://ultravideo.mcgill.ca/project_risq.html

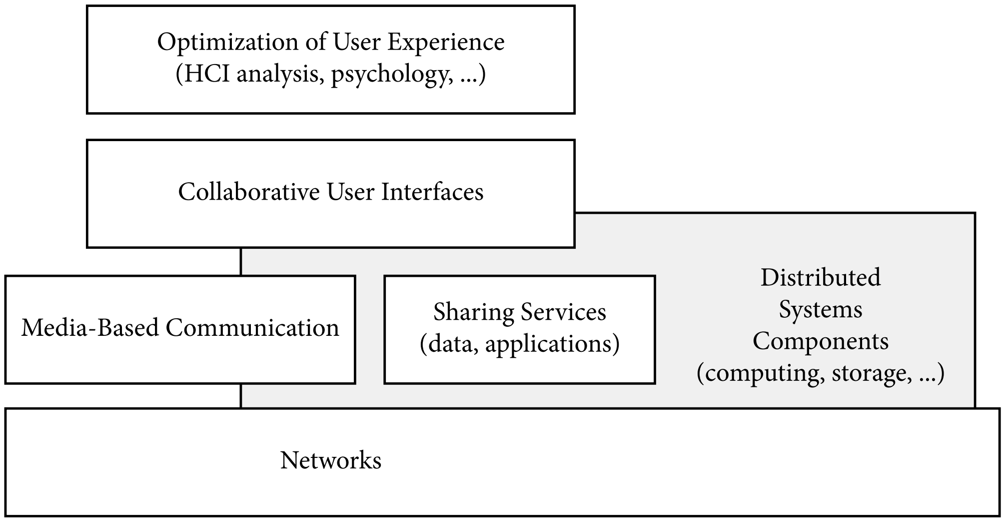

FIGURE 1.1: Architecture of network-centric collaborative systems.

demonstrate interactive multi-point trans-Atlantic 8K video transmissions using commodity hardware and user-empowered data distribution, thanks to our research in GPU-accelerated compression algorithms.

Optimization of NCCEs for human interaction. The way how people interact with user interfaces and the transmitted media is a very important part of the NCCEs and also other types of media applications. This field has been studied for a long time within the human-computer interaction (HCI) research community and it produced a large body of published work. However, a subset of this work is practically relevant to our research into the NCCEs, with focus on high-performance group communication. Since mid 1990s, the HCI community began to explore possibilities of video-mediated communication over the network to increase efficiency of the remote collaboration by preserving various non-verbal communication cues, such as eye contact and gaze awareness, and spatial truthfulness. While relatively simple for communication of individuals, preserving these properties in group-to-group collaboration environments is significantly more complex. Experiments with such environments gave rise to concepts like MAJIC [138] in 1994 for communication of individuals, and then GAZE-2 [177] and even more elaborate MultiView [133] and our GColl [164], [Appendix N] designs. Another research direction is support of communication modes that are impossible in the real-world, but feasible within the virtual environments.

1.2 Structure of NCCEs

In the following text, we define the general structure of NCCEs comprising network data distribution, media applications and data sharing services, and user interfaces. These form the architectural layers as shown in Figure 1.1.

Network Layer. The ideal network layer has to take care of number of tasks related to transmission of the data over the network: (a) it provides data distribution, which is generally multi-point to multi-point by the nature of NCCEs, (b) it improves resistance against packet loss and corruption by applying some type of data-format independent coding techniques to ensure high probability that the data can be reconstructed on the receiver[5], (c) ensures security of the network transmissions, (d) provides operation in various protected or adverse network environments with firewalls and network address translation (NAT), (e) provides fair capacity sharing of the network resources (or otherwise well-defined behavior for specific environments), (f) as well as continuously monitors changes in the network conditions, adapts to them and reports to the upper layers for their own adaptation.

The data distribution can be handled on several different layers, as is further discussed in this thesis: from physical layer (Layer 1 of ISO OSI or TCP/IP models, e.g., using pure optical or opto-electro-optical splitters), through multicast on link and network layer (Layers 2 and 3, implemented, e.g., using IP native multicast or Ethernet broadcast), to user-empowered application-level multicast (Layer 7 of ISO/OSI model and application layer of TCP/IP model, using overlay networks and distributor applications). All these approaches differ in terms of their potential performance, flexibility and their ability to be fully operated by the network end-users (i.e., user-empowered approach).

Future networks are also expected to be better optimized for flow routing in addition to common packet routing, to improve consistency of multimedia data delivery, where each data type has different requirements in terms of quality of service, such as jitter and loss sensitivity. In order to avoid threatening of the core unicast network traffic routing, this may be implemented using specific paradigms such as Software Defined Networking (SDN) [126] and network virtualization [10, 32].

Support for operations in protected network environments became critical in recent years as many networks are hidden behind a single or a few public IP addresses using NAT, which disable direct communication of multimedia applications over the network. Since initiation of the communication from the internal network together with the state tracking of TCP protocol is often used to setup a translation path through the NAT, various tunneling techniques can be applied based on wrapping the original traffic into a TCP connection.[6] Similar techniques can be also used for avoiding need to reconfigure firewalls, which can be hard for administrative reasons and makes NCCEs far from being user-empowered.

To mitigate adverse effects of packet loss in connection-less networks without need for data retransmissions, which is impractical namely for interactive communication over long-distance networks, redundancy-based forward error correction (FEC) techniques can be applied relatively independent on formats of the data. Shifted multiplication [96, 110], 1D/2D XOR-based parity [154, 192], Low-Density Parity Check [59, 115], Low-Density Generator Matrix [152, 170] are examples of schemes that have been successfully applied to various types of traffic [106, 153], while Reed-Solomon code [86, 149] suffers from problems with small block size and relatively high computational complexity. Applicability of advanced FEC schemes depends on availability of sufficient computing power on the encoding and decoding devices and on the acceptable level of introduced latency.

[5]We speak about probability because we assume unidirectional data transfer without retransmission. This is common in low-latency systems as the retransmissions may substantially increase the latency, especially on long distance network links.

[6]Ideally, such tunneling should avoid the TCP congestion control mechanisms which can create unpredictable throughput behavior. While impossible on using standard TCP sockets for stability reasons, for certain applications this can be achieved using emulation of TCP via raw IP sockets; these, however, require super-user rights in normal operating systems. Nagle's algorithm of TCP can be avoided easily using `TCP_NODELAY` option.

Another important functionality of the network layer is monitoring of participating nodes and adaptation to changing performance of the underlying network. In more advanced environments, this can result in well known self-* behavior: self-management, self-optimization, self-adaptation, self-healing, etc. [94, 155] This includes automatic deployment and configuration of the applications, continuous optimization of the data distribution, adaptation of data coding schemes to changing packet loss, or adaptation of employed compression schemes to cope with the changing bandwidth capacity. Self-* trend has been witnessed in the deployment of peer-to-peer systems as a middleware supporting operation of collaborative systems [73, 95], e.g., Skype [15,18], our CoUniverse [Appendix P], or to a very limited extent (transcoding-only) also H.323/SIP systems with multi-point inter-connector (MCU) capable of media format adaptation.

Media Processing Layer The media layer ensures processing of the media based on the needs of each specific NCCE: it can be as simple as acquisition and display of the media that is sent uncompressed over the network to minimize latency and retain the maximum quality, through application of data compression and data-dependent coding techniques to cope with various characteristics of the network layer such as limited capacity and potential corruption of data, to transcoding of the data to fit capabilities of individual devices used by the users to capture and play back the media. Another type of media processing is image anonymization in medical applications, i.e., removal of all the information that could be used to identify the patient (e.g., masking text overlays with the name and day of birth of the patient, providing mosaic overlay over the face of the patient).

Many of the collaborative applications also require video-mediated communication, which is the most bandwidth demanding media of the commonly used media types. Based on the purpose of the collaboration, the requirements on the media and network distribution can vary largely. While a lot of applications is satisfied with what is currently available as commodity video transmissions and videoconferencing systems, i.e., up to full HD resolution (1920×1080 at 30 or 60 frames per second) compressed using an efficient inter-frame compression scheme such as H.264 into 2 Mbps to 10 Mbps network flows, there are many applications that demand more. Fine details in laparoscopic and angiology operations, where the image is used to navigate the instruments inside the human body, require at least full HD resolution with significantly better preservation of image quality, resulting in compressed data flows in the 50 Mbps to 200 Mbps range. These applications are good candidates for upgrading to 4K/SuperHD resolution (3840×2160), when such resolutions become more affordable, to substantially improve detail and accuracy of information available to the surgeon. The high resolution can be also used to aggregate various auxiliary data from instruments in real-time; this is common, e.g., in angiology, where X-ray, OCT, pressure curves, and 3D reconstruction is often used at the same time. Another example is textures and fine details in cinematography and in art restoration, where 4K/SuperHD resolution is already a commonplace. For example, DCI JPEG2000 standard [35] specifies 250 Mbps maximum bitrate for distribution of 4K video to the cinemas, while the intermediate steps of media production work even using higher image quality, up to lossless compression or uncompressed data. Yet another examples are consulting and education of art restoration or interactive navigation across large static imagery coming from the geo-survey applications. In the mid-term future, the video and cinematography applications are likely to increase their resolution to 8K (7680×4320) and up to 120 frames per second. Similarly in audio production, while the resulting streams may be stereo (2 channel) audio with 44.1 kHz sampling rate and 16 b quantization, the production of such media often involves elaborate mixing of tens or hundreds of separate channels of 88.2/96/176.4/192 kHz audio with 24/32 b quantization in a geographically distributed process [20]. The advanced audio environments use 22.2 or even more channels to create immersing perception of surround sound even as a part of the final distribution to the customers [193]. There are even more advanced techniques, such as Wave Field Synthesis (WFS) [16, 17] based on Huygens' principle, which allows spatial audio independent on position of the listener at cost of large number of independent audio channels. For example, the first live WFS transmission to Technical University of Berlin utilized 832 independent audio channels [131]. Technologies based on WFS have become even commercially available, such as IOSONO developed by Fraunhofer Institute.

There are even some applications that do not seem to have a known upper limit on the image quality, such as large images in pathology, astronomy, geological survey and surveillance and intelligence purposes. Some video applications can also utilize different trade-off between spatial and temporal resolution that those commonly available in television and cinematography standards. For example, rescue application, such as images from cameras in visual in infra-red range from the intervention site for fire analysis by the intervention coordinator, strongly prefer spatial resolution of the image at cost of temporal (i.e., reducing the temporal resolution to 1 frame per second or even less). But the data need to be transmitted from the field over very constrained network links such as GPRS; thus the challenge is in a way comparable to transmitting the video with both high spatial and temporal resolution over the high-speed links and problem of streams with bitrates comparable to the capacity of the network links needs to be solved. Yet another example is fast motion analysis, which goes in the opposite direction: sacrificing spatial for temporal resolution, often beyond 1000 frames per second.

Data Sharing Layer Another important aspect is sharing of generic content, which is a very heterogeneous component, depending on the nature of the collaboration. In the HCI context, it is related to deixis [66, 97], how people refer to and interact with the shared subjects. It can range from document, application or computer screen, to remote instrumentation of physical entities on the collaborating sites. In general, the sharing services can be a part of a larger distributed systems, such as large-scale scientific simulations and visualizations. Compared to the media processing and transmissions, these are typically composed of either remote applications residing on dedicated infrastructures, or applications running on the participating nodes, where applications show the content locally and only synchronize over the network via some specific protocol. This allows more complex content interaction schemes compared to "plain image sharing". The research focuses on optimization of complex interactions with the content and it is tightly related to the HCI field.

Collaborative User Interfaces The user interfaces for collaborative environments comprise not only the virtual interfaces, such as layout of windows and types of tools on users computer desktop, but also physical setup such as cameras, sound recording and playback devices, as well as various interactive tools commonly known as tangibles [87]. It is also usually reasonably easy to integrate deixis manipulation into such point-to-point communication schema.

For simple communication of two individuals, standard desktop environment is usually sufficient, and only positioning of camera may need some special attention to implement eye contact; it has been discovered that positioning the camera on top of the video window within 5° results in a well-perceived visual contact [27, 63].

The situation becomes more complex in group-to-group communication scenarios or when a people interact with large display devices. Research in group-to-group communication resulted in number of designs such as GAZE [176], GAZE-2 [177], MultiView [133], and GColl [164], [Appendix N], all of which seek for balance between preserving all the natural communication cues such as gaze awareness and spatial distribution of audio sources, while keeping the system transportable and flexible with respect to number of participants. Interaction of individuals or groups of people with deixis often involves large-scale displays has been explored in proposals such as Continuum [105], JuxtaView [98], SAGE [150] and its interactive extension DIM [88], or CGLX [40], resulting in usable systems that have found its deployment in practical applications world-wide. The work of Jagodic [88] also provided a more generalized view of user interaction with the large-scale display devices, as the interaction modes are much dependent on the distance from the display. Practical interactive systems for groups of people also often include tactile tables [93], such as a number of research projects like ThinSight [76], MightyTrace [77], LambdaTable [99], and even commercial products like Microsoft Surface Table[7] (recently renamed to PixelSense) or MultiTaction[8]. The research pursues technologies for implementing the tables, as well as modes of interaction and optimization of gestures for horizontal surfaces.

Study of Users Behavior On top these layers, there is a substantial component of human-computer interaction, which studies interactions of the human beings with computer systems to provide more efficient and ubiquitous systems. An example of results in this field is the eye contact evaluation with respect to the camera positioning [27, 63] discussed briefly above; such research is inter-disciplinary between psychology and computer technology and goes largely beyond what is the main focus of this thesis. We have only touched it with our studies of efficiency of the GColl user interface [Appendix N].

1.3 Quality and Availability Aspects of NCCEs

An important aspect of the NCCEs is quality of experience (QoE). For media transmissions and processing, translates into the visual or audible quality and interactivity (latency) aspects. The media quality has already been discussed in the previous text and we will focus here more on the interactivity aspect. For voice communication, the general consensus is about 150 ms of unidirectional latency in order to achieve interactive experience as specified in ITU-T G.114 recommendation [191]. For a specific application, this threshold may differ substantially; e.g., for chamber orchestra, the latency needs to be in 10 ms to 40 ms latency range [20, 22] in order to avoid tampering synchronization of musicians, while the medical applications may adapt to higher latencies. This has been demonstrated by the Lindbergh Operation in practice. The Lindbergh Operation team stated that their round-trip latency limit was set to 330 ms for fully interactive work, while latencies above 700 ms would threaten safety of the operation [23], which is in a good agreement. Their practical setup produced about 150 ms of unidirectional latency, where 80 ms was attributed to the network latency and 70 ms to video encoding/decoding [23, 119] with total bitrate of 10 Mbps in each direction. This is in the good agreement with my own unpublished findings based on personal discussions with American surgeons, who have hands-on experiences with remote robotic surgery, namely during the Immersive Medical Telepresence Workshop[9] [80]. It turns out that they can accommodate about 300 ms round-trip latency without changing their habits and up to 500 ms when they adjust their behavior (e.g., by serializing steps that would be performed in parallel otherwise, such as handing over instruments from one robotic hand to the other).

[7] `https://www.microsoft.com/en-us/pixelsense/default.aspx`
[8] `http://www.multitaction.com/products/`
[9] `http://events.internet2.edu/2006/IMT/presentations/20060906-IMT-Wellings.ppt`

The differences between the chamber orchestra requirements and surgeons' requirements clearly demonstrate the scale of interactivity range for audiovisual communication. For haptic feedback, the limit may be much lower, depending on what part of the haptic interaction is mediated over the network and what is computed and interpolated locally. The haptic interaction requires at least 5 kHz input (translates to 200 µs) in order to achieve smooth perception of interaction and hence some level of local interpolation is a must in all but the fastest local networks for haptics.

The above discussed requirements on media quality and interactivity impose non-trivial challenges for media and network processing to to ensure scalability and practical applicability of the environment. An important aspect of practical applicability is what can be implemented using commodity hardware, in contrast to custom hardware solutions. The smaller the niche of such a collaborative system is, the more important the commodity aspect is, because of the cost of the hardware development. Recently rapid increase of commodity processing capacity has been brought about by general-purpose computing on graphics processing units (GPUs); however, they need high degree of fine-grained data parallelism in order to utilize their capabilities. Thus the challenge turns into (re-)formulating the image processing algorithm, most notably image compression and network data coding, as fine-grained parallel algorithms. If this effort succeeds, it can enable a whole new class of applications that were previously not affordable; hence the communication systems may become pervasive and ubiquitous.

Another aspect related to NCCEs is user-empowered nature of the environment, which means that for any specific functionality and services needed for the environment to work, the users have sufficient rights and tools to operate the services on their own. An example of this is the low-latency multi-point data distribution functionality on top of the general-purpose computer networks. History of many advanced network services like network-native multicast has shown that when implementing such services inside the network as a part of the network stack, certain problem patterns tend to recur. Complexity of advanced services makes the core network technology scale worse or have significant financial penalty (nice illustration of this problem is price of ATM vs. Ethernet at the same nominal speed and the resulting prevalence of Ethernet). When a problem occurs, the user of the advanced services is often left behind and the service is shut down in order to avoid threatening the simple service that serves most of the users and services[10]. Both issues result in severe problems with effective availability of such a system to the end users. Therefore it turns out to be more advantageous to provide only the very basic services (such as unicast network service, i.e., sending to from one source to one destination) and leave the advances services up to the applications; thus when a problem occurs, the application can try to solve it or work around it on behalf of the user, even when the user does not have any administrative rights on the network level. Such an approach empowers the user and makes him/her less dependent on the network and system administrators, and increases availability of the service.

1.4 Impact of the Work

In my work, I am trying to push the limits of the user-empowered approach to high-performance data processing and distribution on commodity computing platforms in order to push the limits of video-mediated communication and collaboration. This thesis is based on a collection

[10] This follows the well-known Pareto principle that 80% of the users use only 20% of network capabilities and capacity.

of 18 previously published peer-reviewed research papers that address the described problems and one patent (Czech patent awarded, U.S. patent application pending). The introductory commentary in the Part I, which links together the individual components of my work, comprises the following parts. Chapter 2 describes our contribution to the field of user-empowered data distribution, self-organization and adaptation of the high-performance network systems. Chapter 3 summarizes our results in real-time data processing on commodity hardware, mostly on video compression accelerated using GPUs. Being an infrastructure-oriented person, the last Chapter 4 describes the software tools that were created as practical outcomes of my work, some of which became used world-wide to support advanced multimedia applications. In each of the chapters, there are links to the papers presented in the Collection in the Part II, together with explicit formulation of my contribution to each of the papers.

My research contributes to the following existing directions in the computer science: *(a)* scheduling of network streams with bitrates comparable to the capacity of network links, including realistic aspects such as partially known network topologies and availability of the data processing resources in the networks, *(b)* efficient multi-point data distribution with optional real-time processing of the data, and *(c)* efficient processing of multimedia on commodity computing architectures, with major focus on parallel GPU architectures.

We have contributed novel models for scheduling of real-time multi-point data distribution and possible processing for data streams with bitrates comparable to the capacity of the network links. Results of the scheduling can be implemented using application-layer overlay network based on our programmable network elements. The overall goal of the scheduling is to minimize the data distribution latency in order to achieve best possible interactivity, while there might be some secondary or alternative criteria such as maximization of the received media quality. Our results in this field enable user-empowered self-* behavior for high-performance network applications that utilize streams with bitrates comparable to the capacity of network links [Appendix A], [Appendix B]. Despite the fact that this direction of research is fairly unique, mostly due to the perceived costly equipment and limited availability of high-speed networks, it still has important applications, e.g., in medicine [2, 112, 156], scientific and industrial visualizations [68, 98, 104, 143], and arts [92, 121], where conventional videoconferencing and video-mediated communication systems are not an option due to their limited quality and/or interactivity.

Being close to real-world applications of the research, this stimulates incorporation of realistic aspects into previously theoretically studied systems, bringing a new level of complexity. Despite being neglected in theoretical works, such realistic aspects often determine practical applicability of the approach. An examples of these realistic aspects is end-to-end view of the network topology for scheduling of data flows, which can be later augmented by partial topology knowledge derived from analysis of the network from the network end-points, or directly from the network administrators [Appendix R]. The reason behind such a model is the fact that in most cases, it is not realistic to assume knowledge of full network topology as an input of the scheduling; the topology is rarely disclosed to the users for both technical and administrative reasons, as further discussed in Chapter 2.

Another example of such realism is the limited availability of network-native multipoint data distribution, or even complete lack of it, unless the functionality on the network layer is separated by some type of network virtualization [32]. Network virtualization is, however, also rather rare and with limited functionality in practice. This stimulated our proposal of programmable active elements that support both application-level data distribution and processing, provide scalability with respect to the number of participants [Appendix E] as well as with respect to the capacity of individual streams [Appendix I].

Similarly in the field of data compression, the theoretical research may easily conclude that some algorithm can not be made scalable on parallel GPU architectures [103, 141, 142], and proposes some other algorithm which would require modification of the original compression standard, as witnessed, e.g., for EBCOT part of JPEG2000 [103]. In practice, this is not very realistic especially for well-established standards, that have many implementations, some of which are built into specialized hardware solutions. Other groups have simply implemented the level of parallelism specified directly in the standard (causal mode in case of JPEG2000) [178] and concluded that entropy coding cannot be efficiently implemented on the GPU architecture, which was consistent with our early conservative findings. Instead, we have put substantial effort into investigating possibilities allowed by the standard. We were able to propose a completely novel reformulation of the computationally most critical part of the JPEG2000 standard that enabled bit-level parallelism and we have mathematically proved its equivalence to the formulation of the algorithm in standard [Appendix G]. We have proposed backward-compatible extensions to the JPEG format that allow substantially more efficient parallel decoding of the image. We have investigated the Huffman and adaptive arithmetic codes to develop techniques that maximize their performance of the GPU architectures. Based on these findings, we have generalized our ideas into what types of image compression algorithms can be made effectively data parallel and what algorithms cannot [Appendix H], at least according to our best current understanding of the problems. The results of our work brings about the highest-performance JPEG2000 and JPEG implementations on commodity GPU and CPU architectures, only comparable to dedicated hardware solutions [Appendix G], [Appendix H], [Appendix C]. Such performance allowed us to implement, e.g., the first multi-point trans-Atlantic 8K video transmissions on commodity PCs, as well as 4K video transmission of JPEG2000 on off-the-shelf PCs in 2011. The impact of our work resulted, e.g., into invitation to SPIE conference ("Applications of Digital Image Processing" track), where people from ISO JPEG committee (`http://www.jpeg.org`) unofficially meet, to discuss new algorithm designs. Our work has also become cornerstone of Comprimato Systems spin-off company, which focuses on commercial applications of GPU acceleration of image compression.

Results of our work have been adopted by many research, development, and application communities around the world (USA, Spain, Portugal, Latin America, China). The both concepts and their implementations have been integrated into other systems, such as SAGE[11] [150] adopting the UltraGrid [Appendix A], [Appendix C] for video transmissions. Only minor portion of these applications resulted into classical scientific citations, since far from everybody publishes their work in this field. My team has recently also received the The Best Open-Source Software Competition Award by ACM Multimedia Special Interest Group for implementation of the research results in UltraGrid as practically usable leading-edge open-source software.

1.5 Future of the Research Field

State of the research related to NCCEs. The research related to NCCEs pursued a number of directions as demonstrated by papers on a number of relevant conferences such as *ACM Multimedia*, *Data Compression Conference*, *ACM Multimedia Systems*, *NOSSDAV Workshop*, *SPIE Applications of Digital Image Processing*, *IEEE INFOCOM*, *ACM SIGCOMM*, or even more practically oriented *CineGrid Workshop*, and journals such as *IEEE Transactions on Communications*, *IEEE Transactions on Communications*, *ACM Transactions on Multimedia Computing, Communications and Applications*, *IEEE Signal Processing Letters*, *IEEE Transactions on Parallel*

[11] `www.sagecommons.org`

and Distributed Systems, and also somewhat surprisingly *Future Generation Computer Systems* which hosts also the research community focusing on high-performance networked media applications. The HCI community publishes in *ACM Transactions on Human-Computer Interactions* and at well-established conferences such as *ACM Conference on Human Factors in Computing Systems (CHI)*, *ACM Conference on Computer Supported Cooperative Work (CSCW)*, *ACM Symposium on User Interface Software and Technology (UIST)* and number of others.

The topics compiled in the following list are based on my observations of research directions relevant for my own research:

- *Improvements in video coding and compression,* which includes incremental improvements of current video compression algorithms with focus on temporal coding, efficient multi-channel (audio channel, video view) coding, improvements of computation efficiency of coding and compression algorithms.

- *Optimization of network protocols for media distribution to large-scale audiences,* which focuses mostly on low-bandwidth media distribution and distribution to large audiences, delivery to mobile devices, TCP-friendly congestion control mechanisms for video distribution, video streaming over HTTP and interaction of various layers in video streaming (HTTP/TCP protocols).

- *Processing and analysis of multimedia data*, which includes integration of video in various social networks and relates search in the high-volumes of data, similarity analyses and user tracking, high-performance low-latency media processing, use of GPU, programmable hardware and custom hardware design. Some research also investigates generalization of efficient programming using hardware-software co-design and automatic optimizations.

Our perspective of future research. Our focus is on applications of high-speed networks and high-performance computing to processing and distribution of high-quality video and audio, mostly for interactive applications and advanced collaborative environments. This field lies a bit outside of the mainstream multimedia research, which focuses mainly on distribution of lower-quality multimedia to large-scale audiences. Our goals are given by the opportunity to support the most advanced remote applications in fields like medicine, scientific visualization and digital cinema and broadcasting, and by our capabilities to integrate multimedia applications with high-performance networks.

On the networking level, we will continue with modeling the networked systems from the perspective of users and network end-nodes, which is important to capture the realistic aspect of limited access range of each user (and recursively also administrator in most cases). We will try to incorporate various uncertainties that necessarily arise, such as imprecise estimates/measurements of network capacities or possible errors in discovery of network topology, in order to optimize scheduling of networks streams to maximize the quality while minimizing probability that the deployed scenarios will not work because of fuzzy or outright erroneous inputs. The system has to react effectively with the changes, while minimizing negative impact on the users. This is another important step to make even the high-performance systems self-adaptive and self-managing to maximum extent.

Another field in which we plan to work in the near future is Software Defined Networks (SDN) based on OpenFlow [64, 126]. Compared to the active networks [168, 179, 185], upon which our work on user-empowered data distribution was largely based, the SDN approach is more constrained in terms of its functionality: the data may not be arbitrarily processed on the network element, but rather the routing and switching tables can be modified and if more complex processing is needed, the data can be redirected to another node or the OpenFlow controller. Also, the SDN does not assume in-the-band programmability of the network elements directly by the packets passing through the network devices. This constrained nature makes it significantly more adoptable by the vendors of production network equipment. The data duplication functionality in the SDN-enabled network elements can be easily implemented using the capabilities of the OpenFlow mechanism itself. On the contrary the data processing functionality will be hard to implement efficiently; the data packets can be redirected to the OpenFlow controller or some other computer for further processing, but this is far from being as efficient as the original idea of active networks to embed the processing into an active router or switch. However, based on our work, it can be easily combined with the application-level data distributors to build continuous data distribution trees in a similar user-empowered way like multicast was combined with unicast to bridge the subnetworks without the multicast functionality and also to provide data transformation as a part of the data distribution tree.

Our work in the field of GPU accelerations proceeds in the direction of domain-specific automatic code optimizations focused on data coding and compression algorithms. While the generic automatic optimizations are very hard to design and implement in practice, we believe that domain-specific optimizations are much more realistic. We want to utilize our experiences with manual continuous improvements of the GPU-specific compression and coding implementations to cope with the developments of the GPU architectures themselves. There is a number of recurring optimization patterns that could be explored automatically: from low-level modifications of how the data is stored (registers versus shared memory and their different ratios), through examining performance of several possible implementations on the instruction level, to changes in overall layout of data and the processing thread-blocks. As the GPU architectures evolve, the impact of individual optimizations can be estimated, but there are often significant deviations from those estimates and a lot of manual experimentation with the code is actually required to maximize throughput performance of the code. Therefore an approach, that implements automatic exploration of the space of possible optimizations as a part of the compilation of the code for the specific hardware platform could significantly improve the speed of the development of best-performing code. Machine-learning approaches can be further employed to make use of past experiences with the results of optimizations for different hardware systems to improve the search strategies. Compared to the previous approaches to automatic optimizations that focused on searching for optimum composition and merging of computational kernels, e.g., to reduce overhead of data transfers [55], this proposal comprises more levels: optimizations inside each of the kernel, composition of the kernels, as well as overall partitioning of the data that influences thread-block structure of the computation and preprocessing and segmentation of input data even on the CPU.

There are other open spaces to explore that bridge GPU computing and network data processing. Several authors have proposed network packet routing, analysis and filtering using GPUs [24, 84, 129, 136, 165]. Our experiments with applying GPUs to forward error correction coding and decoding based on Low-Density Matrix Generator have good potential for these applications, despite very preliminary nature of our experiments. The interesting approach in this field is development of parallel algorithms that provide as high protection strength as possible and that are parallel enough to utilize massively parallel architectures. Since GPU vendors already work on minimizing overhead of data transfers between the GPUs and I/O cards including network interfaces, as witnessed by NVIDIA GPUDirect for Mellanox network interface cards to enable high-performance GPU-to-GPU communication [161].

1.6 Context and Impact into Education

The field of NCCEs integrates computer networks and design of network protocols, data coding and compression and more generally also signal processing, computer graphics, efficient processing of data on parallel and distributed systems, design of user interfaces with possible extensions to human-computer interaction (HCI) field. Thus it is well-aligned with standard curricula of computer science / informatics with specialization in any of the fields mentioned above. It can be also studied from the inter-disciplinary HCI perspective on the borderline between computer science and psychology.

Such curricula are commonly found on many universities, while some of them are very close to the areas which we are focusing on in this field; some specific examples follow. EVL, Cal-it2 and SNE are probably closest to the focus of my team and our Laboratory of Advanced Networking Technologies (SITOLA).

- *Electronic Visualization Laboratory (EVL)* at *The University of Illinois at Chicago* led by prof. Jason Leigh offers both Masters of Science (MS) and Doctor of Philosophy (PhD) in the Computer Science and Engineering with specialization in Computer Graphics[12]. This laboratory has been our great inspiration when founding the SITOLA at Faculty of Informatics Masaryk University.
- *California Institute for Telecommunications and Information Technology (Cal-it2)* led by prof. Larry Smarr and prof. Thomas DeFanti offers educational opportunities for undergraduate and graduate courses at *University of California San Diego*[13] and *University of California, Irvine*[14]. The institute focuses on is core research fields (computer-centric, such as visualization and computer graphics, as well as inter-disciplinary, such as computer-artistic, computing systems with nano-technologies) and the students get hands-on research experience by being parts of larger teams and graduate students are also individually lead by advisers from the Cal-it2. It has its spin-off laboratory also at King Abdullah University of Science and Technology (KAUST) with similar research and education focus[15].
- Systems & Network Engineering[16] (SNE) research group led by prof. Cees de Laat at University of Amsterdam focuses on hybrid systems and network architectures, research of related protocols, resource description frameworks, path computation and network scheduling, network security, etc. for high-speed networks. It offers SNE Master's curriculum[17], which has been recently voted the best Dutch Master's in informatics. It also offers positions to graduate students.
- Digital Media and Content[18] (DMC) at *Keio University* led by prof. Hideo Saito focuses on applied research in high-performance multimedia processing and network distribution for entertainment (cinema, arts, broadcasting), medicine and other fields. It has close collaboration with Japanese national telecommunication company NTT to apply their research results. The laboratory offers collaboration to limited number of graduate students.

[12] `http://www.evl.uic.edu/core.php?mod=9&type=6&cat=17`

[13] `http://www.calit2.net/education/ucsd/index.php`

[14] `http://www.calit2.net/education/uci/index.php`

[15] `http://www.vis.kaust.edu.sa/`

[16] `https://www.os3.nl/`

[17] `http://www.uva.nl/en/education/master-s/master-s-programmes/masters-programmes/masters-programmes/content/folder/system-and-network-engineering/study-programme/study-programme.html`

[18] `http://www.dmc.keio.ac.jp/en/`

- *Shared Reality Lab*[19] at *McGill University* led by assoc. prof. Jeremy Cooperstock is one of the laboratories that offers specialization into the HCI with focus on high-performance systems for high-fidelity of distributed interaction. Limited information about the curricula is available online[20].
- *Multimedia Communication* group[21] at *Klagenfurt University* led by prof. Hermann Hellwagner offers classical specialization on processing of multimedia, its distribution (adaptive streaming, peer-to-peer systems), as well as quality of service and quality of experience issues. It offers lecture-based curricula for undergraduate and graduate students[22] as well as specialized lab classes[23].

[19] http://www.cim.mcgill.ca/sre/

[20] http://www.cim.mcgill.ca/Academic/Academic/ProspectiveStudents, http://www.cim.mcgill.ca/Academic/Courses

[21] http://www.uni-klu.ac.at/tewi/inf/itec/mmc/

[22] http://www.uni-klu.ac.at/tewi/inf/itec/mmc/2424.htm

[23] http://www.uni-klu.ac.at/tewi/inf/itec/mmc/2425.htm

Chapter 2

User-Empowered Data Distribution in Computer Networks

2.1 User-Empowered Data Distribution Based on Overlay Networks

An example of a service that provides theoretical optimum behavior is network-native multicast in IP (IPv4/IPv6) networks [28, 36]. However, the multicast is also a service that lacks the user-empowered design. While the multicast is known to work reliably in individual administrative domains, as witnessed by many commercial video delivery networks, it is rarely a dependable service when crossing multiple administrative domains as common in the Internet. When any problem with multicast occurs, the users are depending on their network administrator or even a number of administrators implementing a coordinated problem analysis. Because in most networks, the multicast routing has much lower priority compared to the unicast, the users frequently face problems that are never solved properly. Therefore, the availability of IP multicast, or more specifically lack of its availability, has been studied in 2000–2003 [39, 43, 44], and the situation basically remains the same up to now.

Multi-point data distribution has been proposed as application-level multicast, explicitly or implicitly using overlay network concepts [46, 47, 52, 60, 74, 144, 162]. The overlay networks also allow implementation of other specific features, such as faster convergence and finding of alternative paths better suited to the specific purpose compared to the common IP routing (e.g., Resilient Overlay Networks [9]), or per-user processing [74].

We have contributed to the state of the art by proposing a system based on programmable active network concept called Active Elements, which supported also more advanced features such as per-user processing, operation in adverse networking environments (e.g., NAT and firewall traversal).

In order to ensure scalability of the distribution with respect to the number of connected participants, the Active Elements can also establish scalable interconnecting networks. This capability has been designed using loadable plug-ins to enable incorporating various distribution models. A number of suitable models has been proposed previously by many independent groups in the past, most of which fall into one of the two categories: (1) mesh first distribution models like Narada [33], Delaunay triangulation [109], Bayeux [189], and (2) tree first models like YOID [56], TBCP [123], HMTP [188], SHDC [124], NICE [14], Overcast [90], ZIGZAG [172]. Some other models may also be found in [47, 108].

We have also developed an architecture for parallel real-time processing of the data with possible limited data reordering, in order to ensure scalability with respect to the bitrate of individual streams, i.e., to support streams with bitrates beyond capacity of any single Active Element. This utilizes the fact that the real-time media applications need to be able to correct limited packet reordering using buffering on the receiver side; our proposed Fast Circulating Token protocol [Appendix I] allowed us to prove the resulting maximum packet reordering so that the applications can set up their receiving buffers appropriately.

[Appendix J] Eva Hladká, Petr Holub and Jiří Denemark. *User Empowered Programmable Network Support for Collaborative Environment.* In Universal Multiservice Networks: Third European Conference, ECUMN 2004, Porto, Portugal, October 25-27, 2004. Proceedings. Lecture Notes in Computer Science 3262, Heidelberg: Springer-Verlag Berlin, 2004. 10 p. ISSN 0302-9743.

Contribution: I have co-authored the idea of a user-empowered programmable data processor and distributor, based on active router concept. (40%)

[Appendix E] Petr Holub, Eva Hladká and Luděk Matyska. *Scalability and Robustness of Virtual Multicast for Synchronous Multimedia Distribution.* In 4th International Conference on Networking, ICN 2005, Reunion Island, France, April 2005. Proceedings. Lecture Notes in Computer Science 3421, Heidelberg: Springer Berlin, 2005. 8 p. ISSN 0302-9743.

Contribution: I have developed the concept of peer-to-peer network of data processors/distributors with separated data plane and control plane. (60%)

[Appendix I] Petr Holub and Eva Hladká. *Distributed Active Element for High-Performance Data Distribution.* IFIP International Conference on Network and Parallel Computing, NPC 2006, Tokio, Japan, October 2006. Proceedings. pp. 27–36, 10 p.

Contribution: I have created the concept of parallel processing with limited synchronization, developed the Fast Circulating Token protocol, investigated its theoretical characteristics, implemented its prototype and performed the experimental evaluation. (90%)

[Appendix K] Petr Holub and Eva Hladká. *Virtual Multicast.* In Trends in Telecommunications Technologies. Edited by Christos J. Bouras, ISBN 978-953-307-072-8, pages 63–86. Hard cover, 768 pages, Publisher: InTech, Published: March 01, 2010. doi: 10.5772/180

Contribution: Being a summary paper of our effort on user-empowered data distribution and processing, I have authored and co-authored concepts and was the principal editor of the paper. (60%)

2.2 Data Distribution and Application Orchestration for Data Streams Comparable to Network Link Capacities

Setup of high-performance NCCEs is very tedious work when done manually. It requires setup of many components, their correct configuration. In practice it also requires handling of dynamic behavior of the system, such as addition of new participants and reaction to failures of various components. This requires continuous monitoring and reactions to events occurring in the system. Since our first demonstration of multi-point videoconferencing system based on uncompressed HD video in 2005, we have spent several years supporting advanced collaborative environments based on HD video and high-fidelity audio, sometimes comprising tens of participating sites with a number of components on each site. This motivated us to develop a system that could orchestrate such a system automatically. We have realized that one of the critical differences compared to most of the research done in this field by other groups, is the need to support multimedia flows with bitrates comparable to the capacity of the network links. For example, uncompressed HD video requires 1.5 Gbps and thus only up to 6 streams can practically fit one 10 Gbps network link, the maximum bandwidth per link available at that time.

We have proposed a user-empowered orchestration framework called CoUniverse, which relies on an overlay peer-to-peer network for its control plane (i.e., discovery of nodes, setup of communication channels for exchange of control messages, security, etc.), while the data plane (i.e., distribution of the data of the multimedia applications) runs natively on the network. Because the problem of scheduling streams with bitrates comparable to the capacity of network links is $\mathcal{NP}$-complete, as shown in [114, Section 4.5], the design of the system uses one scheduler application per collaborative environment, which either directly schedules or coordinates scheduling of the streams using constraint programming [Appendix B], mixed integer programming [Appendix R], or local search heuristics [Appendix M]. Because of the *raison d´etre* of CoUniverse—management of highly interactive collaborative applications—the primary objective function of the scheduling was to minimize latency of the stream distribution. However, alternative objective functions are also available, such as compound latency minimization and quality maximization, or minimum perturbation if the system reacts to some changes in the environment and the previous schedule has already been in place. Yet another objective function can minimize differences among the data distribution latencies from a data source (called "producer") to the receivers (called "consumers") in order to equalize opportunities of user to react. The scheduler can take into the account also capabilities and resources of individual participating nodes, namely those nodes that handle the data multiplication/distribution task (called "distributors"), and setup transcoding for individual receiving nodes as a part of the distribution tree. Even though originally designed for high-bandwidth applications running in high-performance networks, this model is also applicable to low-bandwidth networks where even low-bitrate streams can easily become comparable to the capacities of the network links.

The original model based on constraint programming suffered from rather limited scalability [Appendix B]. Thus the model was later converted to mixed integer-programming paradigm [Appendix R] to improve its performance while still looking through the full space of possible solutions. An order of magnitude better scalability has been achieved using local search techniques [Appendix M], at the cost of possibly not finding the optimum solution(s). The local search approach also supports dynamics of the system in a better way: it enables fine control of the strategy how the previous schedule is taken into consideration when reacting to some change in the environment. Minimizing amount of perturbations is an important property in order to minimize the impact of the changes on the users participating in the collaborative session.

When applying the user-empowered view of the infrastructure onto the process of building a collaborative environments, there are inherent limitations, which need to be taken into the account. On the network level, the user may not have access to the topology knowledge for various administrative reasons. The network administrators usually do not share this information with the users as the network topology is dynamic and may change abruptly, which would imply notifying all the users the information had been shared with. Also, the internal topology is sometimes considered a competitive advantage by some network operators. Some information about the topology may be inferred during the operation of the infrastructure as demonstrated, e.g., by [19, 38, 117], but the default initial view must not require any prior knowledge. All the knowledge must be obtained entirely from the end-nodes of the network, as non-administrators may not install their own internal boxes. Also, in the global private Layer 2 networks (ISO OSI) such as GLIF[1], the information beyond point-to-point latency and bandwidth may be particularly hard to infer reliably with low overhead.

Therefore, CoUniverse works primarily with logical view of the network, which represents only the end-to-end view of the infrastructure [Appendix P],[Appendix B]. The network nodes are the computers connected to the network and each node has one or more network interfaces. The nodes are hosting individual applications, each of which falls into one of the three following categories: data producers, consumers, and distributors. The logical links between pairs if node interfaces represent their possible mutual reachability. If the link is not present, the interfaces are connected to different networks that have no inter-connection. The link has a state (up or down), has some end-to-end latency and capacity. While link state and end-to-end latency can be easily determined directly from the rather unobtrusive measurements, the bandwidth is more complex to determine since the logical links may share the same physical link. As a result, the primary estimate comes from the capacity of network interfaces, which can be set lower than the nominal link rate by the node user if the interface is known to be unable to process data at link rate. During operation of the CoUniverse, passive packet loss measurements from the orchestrated applications can be subsequently used to determine the real bandwidth available. This approach can be augmented by active bandwidth measurements or estimation techniques, but they all have certain disadvantages: active measurements can create substantial load on the network, possibly deteriorating performance of already running transmissions (esp. if a loss-based congestion control algorithm with relatively slow recovery is used as is the case in TCP), and the estimative techniques such as inter-packet dispersion measurements [13, 41, 81, 130, 151, 166, 187] are known to have precision limitations [7, 62, 89, 163]. Our recent work therefore focuses on embedding partial topology knowledge into the problem [Appendix R]: this reflects the fact that only topology of certain regions in the network may be known to the users—but in many cases, it is the site-topology that aggregates links and introduces bandwidth sharing among the logical links and thus even partial topology can make the scheduling process significantly more precise in practice.

[Appendix P] Miloš Liška, Petr Holub. *CoUniverse: Framework for Building Self-organizing Collaborative Environments Using Extreme-Bandwidth Media Applications.* In Lecture Notes in Computer Science vol. 5415 Euro-Par 2008 Workshops - Parallel Processing. Las Palmas de Gran Canaria, Spain : Springer Berlin / Heidelberg, 2008. ISBN 978-3-642-00954-9, pp. 339-351. 2008.

Contribution: I have co-authored the fundamental ideas behind the scheduling framework for high-bandwidth applications and implemented substantial part of the CoUniverse framework in order to validate the system in practice. (50%)

[Appendix B] Petr Holub, Hana Rudová, Miloš Liška. *Data transfer planning with tree placement for collaborative environments.* Journal of Constraints, Springer, 16, 3, pp. 283–316, 34 p. ISSN 1383-7133 (print) 1572-9354 (online). 2011.

Contribution: I have created the idea of the constraint-based model and developed its fundamentals. I have also implemented the model into the CoUniverse application and made its evaluations. (50%)

[Appendix M] Jiří Marek, Petr Holub, Hana Rudová. *Local Search Heuristics for Media Streams Planning Problem.* Proceedings of The 27th IEEE International Conference on Advanced

[1] `http://www.glif.is`

Information Networking and Applications (AINA-2013) Barcelona, Spain, March 25-28, 2013.

Contribution: I have co-authored the scheduling model with focus on its realistic aspects, and was co-editor of the paper. (30%)

[**Appendix R**] Pavel Troubil, Hana Rudová, Petr Holub. *Media Streams Planning with Transcoding.* In Proceedings of IEEE 12th International Symposium on Network Computing and Applications 2013 (NCA 2013). ISBN 978-0-7695-5043-5, s. 41–48. Cambridge, Massachusetts, USA.

Contribution: I worked mainly on the concepts of partial topology knowledge and adaptation of the model to the related uncertainties. (25%)

This approach has also been shown to allow end-user to orchestrate the network based on dynamically allocated network circuits, by adding some partial topology information into the system—basically the information of what end-to-end logical links are supported by a dynamic circuits that need to be allocated (initialized) prior to being used, and torn down once the transmissions are over. Practical implementation of these concepts has been demonstrated by integrating CoUniverse with Internet2 Dynamic Circuit Network (DCN) service[2] [Appendix Q], being one of the very first applications to utilize such a service programmatically [78, 79].

[**Appendix Q**] Miloš Liška, Petr Holub, Andrew Lake, and John Vollbrecht. *CoUniverse Orchestrated Collaborative Environments with Dynamic Circuit Networks.* In Proceedings of 2010 Ninth International Conference on Networks, 2010. ISBN 978-0-7695-3979-9, pp. 300-305. France.

Contribution: I am author of the general idea of controlling the circuits from the frameworks orchestrating the applications and guided the research in this field. (40%)

We have designed the same user-empowered approach originally for monitoring of the Grid infrastructures using so called "worms", that mimic the user's behavior and that work under user's identity. This allows to discover more problems compared to traditional monitoring based on administrative processes introduced into the system. This system has been designed, prototyped, and experimentally evaluated within the EU GridLab project[3].

[**Appendix F**] Petr Holub, Martin Kuba, Luděk Matyska, Miroslav Ruda. *Grid Infrastructure Monitoring as Reliable Information Service.* Second European AcrossGrids Conference, AxGrids 2004, Nicosia, Cyprus, January 28-30, 2004. Revised Papers, p. 220–229, Lecture Notes in Computer Science 3165, ISBN 978-3-540-22888-2 (print) 978-3-540-28642-4 (online), Series ISSN 0302-9743. 2004

Contribution: I have co-developed the idea of the worm-based approach and was the principal editor of the paper. (35%)

2.3 Optical Multicast

For rather special applications, such as extreme bandwidth applications or applications using multiple WDM channels, we have invented a dynamic optical-level multicast device based on

[2] `http://www.internet2.edu/network/dc/`

[3] `http://www.gridlab.org/`

specific broad-bandwidth (in terms of optical wavelength) splitters, that enables user-controlled multicasting of incoming optical signals on pure optical level (Layer 1 of ISO OSI model) without opto-electro-optical conversion. The advantage of this system is that it works with any form of the signal and nominal rate of the Layer 2 protocol (e.g., 1/10/40/100 Gbps), since the only more significant modification of the signal is its attenuation. It can also work with multiple wavelengths at the same time, as is common for wavelength-division multiplexed (WDM) signals. Being a physical device, we have patented this system.

[**Appendix S**] Josef Vojtěch, Miloslav Hůla, Jan Radil, Miroslav Karásek, Stanislav Šíma, Jan Nejman, Petr Holub. *Device for Multicast of Optical Signals in the Internet and Other Networks.* U.S. Patent Application has been approved as of late September 2013, but not yet published by USPTO; the application has been filed as US 2010/0310254 A1. Czech Patent granted, filed as no. 300811.

Contribution: I have invented the basic idea of using controllable broad-bandwidth splitters to implement the multicast, which has been then further developed by specialists in construction of optical distribution systems. Later, I worked on experimental validation of the prototype devices for real applications and first public demonstration of the technology during the GLIF 2007 workshop. (20%)

Chapter 3

Processing of Multimedia

Advent of general-purpose computing on graphics processor units (GPUs) has substantially changed the landscape of many data processing and high-performance computing applications. Despite the fact that the GPUs require very high degree of data parallelism because of their single instruction multiple data/thread (SIMD/SIMT) model and feature a number of caveats in terms of efficient utilization of their memory and bandwidth resources, they are able to bring 10–100× acceleration of many applications compared to their CPU counterparts. They have also a number of other advantages beyond raw computing power. If used effectively, their computational performance to energy consumption balances extremely well, as witnessed by their prevalence on the top position in Green500 supercomputing list[1]. For different types of applications, the GPUs are usually available both as consumer-grade graphics cards for workstations, desktop computers, or even laptops, while server applications can be accelerated using GPU-based accelerators without the graphics output.

As the GPUs supporting general-purpose computing became commodity, we have started to look into their applications in video-mediated communication. We have focused on acceleration of two standard compression algorithms JPEG and JPEG2000, because of their applicability to both static images and video. Both of them are used in the movie compression, known as Motion JPEG (M-JPEG) and Motion JPEG2000 (M-JPEG2000) respectively. Because of their good balance between image quality and compression ratio, as well as many advanced features (e.g., regions of interest, direct addressing and decompression of selected parts of the image, error resiliency), the JPEG2000 and M-JPEG2000 has become a standard in many professional applications, such as medical applications (DICOM standard in Supplement 106 [190], [42]), digital cinematography, or long-term preservation of large imagery in archiving and mapping applications. However, high computational complexity of JPEG2000 still hinders its wider deployment and the less-demanding JPEG standard still keeps its foothold in many commodity applications from web to movie recording in consumer devices.

JPEG is a well-known image compression standard defined by Joint Picture Expert Group as ITU T.81 standard, based on Discrete Cosine Transform (DCT) and Huffman and Run-Length Encoding (RLE) entropy coding in its baseline profile[2]. The block-level parallelism, which is built into the JPEG standard, does not achieve sufficient level of parallelism to saturate the GPUs. Parallelization of DCT transformations has been studied extensively [137, 145, 169] and the results were insufficient for efficient implementation on the GPU architectures. We have proposed a novel parallelization of the Huffman and RLE coding, utilizing the fact that the

[1] `http://www.green500.org/lists/green201211`, as of November 2012.

[2] The ITU T.81 allows also other modes, such as utilization of arithmetic coding instead of Huffman and RLE coding, but these parts of the standard are not royalty-free.

coding uses static symbol tables and thus it can be parallelized using an approach similar to well-known parallel prefix sum [71]. Our proposal thus achieves pixel-level parallelism. We have also proposed a backward-compatible extension of the JPEG stream with auxiliary indexes designed to work together with the restart intervals defined in the standard; this extension enables parallel processing of the decoding of the compressed JPEG stream. Parallelization of DCT is a well-known task and together with our proposed algorithms, it allows encoding and decoding of up to 8K video (i.e., 16× 1080i/p HD video) in real-time on a single state-of-the-art GPU.

JPEG2000 is a significantly more complex standard based on Discrete Wavelet Transformation (DWT) and adaptive arithmetic coding. Similar to the JPEG, the JPEG2000 also includes block-level parallelism as a part of the standard; but this granularity is not sufficient to achieve good performance on the GPUs, even when using pipelining, as each thread would consume more resources and thus limiting the amount of parallelism on one card anyway. Pipelining would also deteriorate the low-latency nature of the codec. The DWT has been studied in several research papers [57, 102, 122, 174, 180] and shown to be implementable on a GPU very efficiently using so called lifting schema [34]. We have employed a sliding window algorithm in combination with the lifting schema, to achieve a good performance. Our major break-through result stems from reformulation of so called Context Modeling (CM), which precedes the adaptive arithmetic coding and creates the contexts for each bit in the image that are used by the adaptive coding to develop its state. The CM used to be the computationally most demanding part of the compression standard and the major obstacle for efficient implementation on massively parallel architectures of GPUs. The CM had never been made effectively parallel because of its serial definition, and alternative scheme was proposed to work around this problem [103]. We have reformulated the original CM algorithm in a bit-parallel way and mathematically proved its equivalence with the original definition in the standard. This is done using a fixed point algorithm and at the cost of some over-computation, but the level of achieved parallelism of the GPUs compensates for this overhead. We have also proposed technical optimization of the arithmetic encoding as well, but compared to the CM, this part of the processing does not consume as much computation time. Our implementation of the JPEG2000 has been shown to exceed any CPU-based solutions by far and achieves performance comparable only to dedicated hardware solutions.

[Appendix G] Jiří Matela, Vít Rusňák, Petr Holub. *Efficient JPEG2000 EBCOT Context Modeling for Massively Parallel Architectures.* In Data Compression Conference (DCC), 2011. Washington, DC, USA : IEEE Computer Society, 2011. ISBN 978-0-7695-4352-9, p. 423–432. 2011, Snowbird, Utah, USA. 2011.

Contribution: I have co-authored the formal formulation of the algorithm and the proof of equivalence. I have co-edited the paper. (25%)

[Appendix H] Jiří Matela, Petr Holub, Martin Jirman, Martin Šrom. *GPU-Specific Reformulations of Image Compression Algorithms.* Proc. SPIE 8499, Applications of Digital Image Processing XXXV, 849901 (October 15, 2012). 2012

Contribution: I have co-authored several parallelization techniques described in the paper. (30%)

[Appendix C] Petr Holub, Martin Šrom, Martin Pulec, Jiří Matela, Martin Jirman. *GPU-Accelerated DXT and JPEG Compression Schemes for Low-Latency Network Transmissions of HD, 2K, and 4K Video.* Future Generation Computer Systems, Amsterdam, The Netherlands: Elsevier Science, 29, 8, pp. 1991–2006, 16 p. ISSN 0167-739X. October 2013.

Contribution: I have co-authored the principles of acceleration of parallel JPEG entropy coding and parallel decoding, I have designed the overall architecture of the system. I was also the principal editor of the paper. (40%)

These research results also stimulate our future research into the automatic optimizations of the code. There are certainly the research component that require human invention, such as reformulation of the CM or formulation of lifting schema for the DCT and DWT. On the other hand, there was still a lot of benchmarking and manual exploring of the space of possible technical optimizations in order to maximize the performance, such as distribution of data among various types of on-chip and off-chip memory resources, optimization of data pre-fetching, various types of parallel partitioning of the problems, or use of new specialized instructions and their idioms. Such technical optimizations have also very significant impact on the performance. For example, almost whole implementation of the adaptive arithmetic encoder of the JPEG2000 is based on those technical optimizations that could be automatically explored. Another motivation for the automation is the fact that these technical optimizations commonly change with new generations of GPU hardware and thus optimization explorations have to recur to find the best configuration. Hence we see great potential in domain-specific language extensions or event complete languages that enable optimization automation.

By proposing and implementing the parallelization of JPEG and JPEG2000, we have enabled implementation of many advanced applications on commodity hardware, that could only be supported by custom hardware previously. Our results have found their applications in medicine (e.g., processing of Gigapixel imagery in pathology [51], real-time transmissions of endoscopic and laparoscopic operations for remote consulting and education), digital cinematography (e.g., critical screening in remote production and geographically distributed post-production, packaging of the video for projection in digital cinemas) [Appendix C], and distributed scientific and industrial visualizations [Appendix D]. The results have also become the cornerstone of the Comprimato Systems spin-off company.

Chapter 4

Software Tools

4.1 User-Empowered Reflectors

The concept of user-empowered programmable data distribution has been implemented as family of "rum" software products. They were used with a number of different applications, e.g., MBONE Tools, VideoLAN Client, UltraGrid. The reflectors have been implemented in three flavors:

- rum2 – A full-featured data distributor based[1] on the published active router concept, with modular architecture and modules loadable in run-time. A number of modules has been implemented, ranging from image processing applications to specialized control of data distribution. rum2 is currently also part of the GColl system.
- rum-ultra – A high-performance data distributor designed for maximum throughput performance, without modular architecture, currently part of the UltraGrid software.
- rum-ultra-transcode – A high-performance data distributor and processor that uses UltraGrid as a backend for the processing. It is also distributed as a part of UltraGrid software.

The reflectors are implemented using C/C++. rum2 is available under GPL license, while rum-ultra and rum-ultra-transcode are available under BSD license.

4.2 UltraGrid

An early version of UltraGrid software was created in 2002 at ISI EAST within an NSF-funded project led by Colin Perkins and Ladan Gharai, supporting uncompressed 720p video transmission over IP networks to experiment with high-bandwidth multimedia flows. As of late 2004, the development of the software ceased and was only used for experimentation with congestion control protocols by Perkins [146].

In early 2005, we have adopted this software for the world's first demonstration of international low-latency multi-point uncompressed high-definition 1080i video during the GLIF'05 workshop [Appendix A]. Because of private preparation and competitive nature of the demos for the workshop, there was another similar demonstration of multi-point uncompressed high-definition video based on iHDTV software by ResearchChannel from University of Washington,

[1] `https://gitorious.org/rum2/rum2`

lead by Michael Wellings. Therefore, the GLIF'05 meeting saw two world-first demonstrations on this topic. One of the principal differences was the fact, that while ResearchChannel team used multicast (with all its advantages and drawbacks), our team was able to implement this functionality in entirely user-empowered way using reflectors.

We have continuously developed UltraGrid since 2005, focusing namely on pioneering of what is possible to implement in terms of interactivity and video quality using affordable commodity hardware. Advent of general-purpose computing allowed efficient processing and transmissions not only for 1080/720 HD video, but also for 4K (2011) and 8K (2012) video.

UltraGrid can be used both with network native IP multicast and user-empowered reflectors, which allows us to directly compare the two approaches. UltraGrid was also successfully used with pure optical multicast, but this setup exhibits non-trivial complications stemming from having exactly same duplicate packets sent to several different destinations. Therefore, it is suitable only for very specific applications, which are not feasible otherwise.

The application is written mostly in C/C++, available as open-source under BSD-license[2]. It works on Linux, Apple MacOS X, and recently also on Microsoft Windows. The compressions are implemented using NVIDIA CUDA (for maximum performance) and OpenGL Shader Language (for maximum portability).

The UltraGrid software has been successfully deployed worldwide to support various applications from sciences (namely remote visualizations) and tele-medicine (remote consulting as well as training and education of physicians) to arts (interactive theater and music performances) and education. UltraGrid has also become one of the key components within the Czech national infrastructure for research and education, developed by CESNET association, to support advanced applications that cannot be sufficiently supported by commodity videoconferencing and video recording and streaming applications.

[Appendix L] Petr Holub, Jiří Matela, Martin Pulec, Martin Šrom. *UltraGrid: Low-Latency High-Quality Video Transmissions on Commodity Hardware.* Proceedings of the 20th ACM international conference on Multimedia, p. 1457–1460. ACM New York, NY, USA. ISBN: 978-1-4503-1089-5. 2012.

Contribution: I am architect and coordinator of the team that develops UltraGrid software since 2005. I was the principal editor of this paper. Ultragrid has received the Best Open-Source Software Competition Award by ACM Multimedia Special Interest Group at the prestigious ACM Multimedia conference in 2012. (40%)

[Appendix A] Petr Holub, Luděk Matyska, Miloš Liška, Lukáš Hejtmánek, Jiří Denemark, Tomáš Rebok, Andrei Hutanu, Ravi Paruchuri, Jan Radil and Eva Hladká. *High definition multimedia for multiparty low-latency interactive communication.* Future Generation Computer Systems, Amsterdam, The Netherlands: Elsevier Science, 22, 8, pp. 856–861, 6 p. ISSN 0167-739X. 2006.

Contribution: I was the principal author of the ideas behind the paper (the concept of user-empowered low-latency transmissions of the uncompressed video), was the leader of the development team and principal editor of the paper. (50%)

[Appendix D] Andrei Hutanu, Gabrielle Allen, Stephen D. Beck, Petr Holub, Hartmut Kaiser, Archit Kulshrestha, Miloš Liška, Jon MacLaren, Luděk Matyska, Ravi Paruchuri, Steffen Prohaska, Ed Seidel, Brygg Ullmer, Shalini Venkataraman. *Distributed and collaborative visualization of large data sets using high-speed networks.* Future Generation Computer

[2] `http://ultragrid.sitola.cz/`

Systems, Amsterdam, The Netherlands: Elsevier Science, 22, 8, pp. 1004–1010, 7 p. ISSN 0167-739X. 2006.

Contribution: I have designed the tools for remote transmissions of the visualizations. (15%)

FIGURE 4.1: ACM Best Open-Source Software Competition Award for UltraGrid. Photo from the ceremony (on the right) by Pavel Korshunov.

4.3 CoUniverse

CoUniverse framework [Appendix P] was implemented for the two reasons: (1) to provide a tool that simplifies setup and maintenance of larger infrastructures based on high-performance interactive media applications and multi-point data distribution, and (2) in order to experimentally validate concepts of data flow scheduling for streams with bitrates comparable to the capacities of the network links.

CoUniverse has been implemented in Java in order to ease portability, and it utilizes JXTA[3] middleware for basic peer-to-peer functionality. It is available as open-source software under BSD-style license.

[Appendix P] Miloš Liška, Petr Holub. *CoUniverse: Framework for Building Self-organizing Collaborative Environments Using Extreme-Bandwidth Media Applications.* In Lecture Notes in Computer Science vol. 5415 Euro-Par 2008 Workshops - Parallel Processing. Las Palmas de Gran Canaria, Spain : Springer Berlin / Heidelberg, 2008. ISBN 978-3-642-00954-9, pp. 339-351. 2008.

Contribution: I have co-authored the fundamental ideas behind the scheduling framework for high-bandwidth applications and implemented substantial part of the CoUniverse framework in order to validate the system in practice. (50%)

4.4 GColl

Collaborative environments for groups of participants tend to mask away certain communication cues, such as gaze awareness, i.e., the fact that a person usually looks at the other

[3] http://jxse.kenai.com/

person who he/she is communicating with. In traditional videoconferencing systems, the person either looks at everybody (when looking into the camera) or at nobody (when looking sideways). This problem gets even more pronounced in group-to-group communication: in most videoconferencing setups, there is only a single camera for the group and thus it may even difficult for remote participants to find out who is talking to whom.

In order to overcome these limitations, we have proposed GColl, a specific group-to-group communication system with preservation of partial gaze awareness. The partial gaze awareness means that each person is able to tell who is looking at him/her, while full gaze awareness means that each person is able to tell at whom is each other person looking. GColl software has been extensively used for psychological research in how the trust develops between local and remote groups within different types of collaborative environments—from face-to-face environments to traditional videoconferencing [Appendix N], [Appendix O]. The experiments show that GColl is actually closer to the face-to-face meetings in terms of the results than traditional videoconferencing systems.

Originally, the GColl had been based on MBONE Tools for the client tools and the user-empowered reflector RUM. In order to ensure long-term maintainability and to overcome limits of now defunct MBONE Tools, the functionality of GColl has been merged into the UltraGrid in 2012.

[Appendix N] Petr Slovák, Pavel Troubil, Petr Holub. *GColl: Enhancing Trust in Flexible Group-to-Group Videoconferencing.* In Proceeding CHI EA '10 CHI '10 Extended Abstracts on Human Factors in Computing Systems. Pages 3607–3612 ACM New York, NY, USA. table of contents ISBN: 978-1-60558-930-5. 2010

Contribution: I have co-authored the ideas behind GColl principles and architecture. (20%)

[Appendix O] Petr Slovák, Peter Novák, Pavel Troubil, Vít Rusňák, Petr Holub, and Eric C. Hofer. *Exploring collaboration in group-to-group videoconferencing.* In From Research to Practice in the Design of Cooperative Systems: Results and Open Challenges; Proceedings of the 10th International Conference on the Design of Cooperative Systems. London : Springer, 2012. ISBN 978-1-4471-4092-4, s. 229-244. 30.5.2012, Marseille, France.

Contribution: I have co-authored the ideas behind GColl principles and architecture. (20%)

4.5 Summary of World Firsts

2005, iGrid'05 User-empowered multi-point videoconference using uncompressed 1080i video over trans-Atlantic networks.

Utilized applications: UltraGrid, high-performance reflectors

2005, iGrid'05 Distributed interactive remote visualization using uncompressed 1080i video over trans-Atlantic networks.

Utilized applications: UltraGrid, high-performance reflectors

2007, GLIF'07 Optical multicast for multi-point distribution of uncompressed 1080i video.

Utilized applications: UltraGrid

2007, GLIF'07 Self-organizing uncompressed/compressed interactive 1080i video transmissions.

Utilized applications: UltraGrid (with CPU-based DXT compression), high-performance reflectors, CoUniverse

2011, CineGrid Trans-Atlantic interactive 4K video transmissions for real-time movie post-production review/approval process.

Utilized applications: UltraGrid (with GPU-accelerated JPEG)

2012, GLIF Multi-point interactive trans-Atlantic video transmissions of 8K video.

Utilized applications: UltraGrid (with GPU-accelerated JPEG and DXT compression), high-performance reflectors

2012, CineGrid Trans-Atlantic interactive 4K video transmissions for real-time movie post-production review/approval process, using GPU-accelerated JPEG2000.

Utilized applications: GPU-accelerated JPEG2000 (in FlashNET software developed jointly by CESNET and Felix International for remote critical screening in production and post-production in cinematography)

Bibliography

[1] "Public television in Germany," *Nature*, 137(3462):391, 1936.

[2] M. Ackerman and C. Locatis: "Advanced networks and computing in healthcare," *Journal of the American Medical Informatics Association*, 18(4):523–528, 2011.

[3] D. Adamczyk, D. Collados, G. Denis, J. Fernandes, P. Galvez, I. Legrand, H. Newman, and K. Wei: "Global platform for rich media conferencing and collaboration." In *2003 Conference for Computing in High-Energy and Nuclear Physics (CHEP 03)*, La Jolla, California, 2003.

[4] A. Agarwal, A. H. Rowberg, and Y. Kim: "Fast JPEG 2000 decoder and its use in medical imaging," *IEEE Transactions on Information Technology in Biomedicine*, 7(3):184–190, 2003.

[5] S. Akhshabi, L. Anantakrishnan, A. C. Begen, and C. Dovrolis: "What happens when http adaptive streaming players compete for bandwidth?." In *Proceedings of the 22nd international workshop on Network and Operating System Support for Digital Audio and Video*, NOSSDAV '12, pages 9–14, New York, NY, USA, 2012. ACM, ISBN 978-1-4503-1430-5. `http://doi.acm.org/10.1145/2229087.2229092`.

[6] S. Akhshabi, A. C. Begen, and C. Dovrolis: "An experimental evaluation of rate-adaptation algorithms in adaptive streaming over HTTP." In *Proceedings of the second annual ACM conference on Multimedia systems*, MMSys '11, pages 157–168, New York, NY, USA, 2011. ACM, ISBN 978-1-4503-0518-1. `http://doi.acm.org/10.1145/1943552.1943574`.

[7] A. A. Ali, F. Michaut, and F. Lepage: "End-to-end available bandwidth measurement tools : A comparative evaluation of performances," *CoRR*, abs/0706.4004, 2007.

[8] K. C. Almeroth: "The evolution of multicast: From the MBone to inter-domain multicast to Internet2 deployment," *IEEE Network*, 14(1):10–20, January/February 2000.

[9] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris: "Resilient overlay networks." In *18th ACM Symp. on Operating Systems Princpiles (SOSP)*, Banff, Canada, October 2001.

[10] T. Anderson, L. Peterson, S. Shenker, and J. Turner: "Overcoming the internet impasse through virtualization," *Computer*, 38(4):34–41, 2005.

[11] J. Apostolopoulos, T. Wong, W. tian Tan, and S. Wee: "On multiple description streaming with content delivery networks." In *IEEE INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings.*, volume 3, pages 1736–1745, 2002.

[12] M. Azad, R. Mahmood, and T. Mehmood: "A comparative analysis of DCCP variants (CCID2, CCID3), TCP and UDP for MPEG4 video applications." In *Information and Communication Technologies, 2009. ICICT '09. International Conference on*, pages 40–45, 2009.

[13] S. Banerjee and A. K. Agrawala: "Estimating available capacity of a network connection." In *Proceedings of IEEE International Conference on Networks*, 2001.

[14] S. Banerjee, B. Bhattacharjee, and C. Kommareddy: "Scalable application layer multicast.." In *Proceedings of the ACM SIGCOMM 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, August 19-23, 2002, Pittsburgh, PA, USA*, pages 205–217. ACM, 2002, ISBN 1-58113-570-X.

[15] S. A. Baset and H. Schulzrinne: "An analysis of the Skype peer-to-peer internet telephony protocol." In *IEEE INFOCOM 2006*, Barcelona, Spain, May 2006. `http://www1.cs.columbia.edu/~salman/publications/skype1_4.pdf`.

[16] A. J. Berkhout: "A holographic approach to acoustic control," *Journal of the Audio Engineering Society*, 36(12):977–995, 1988.

[17] A. J. Berkhout, D. de Vries, and P. Vogel: "Acoustic control by wave field synthesis," *The Journal of the Acoustical Society of America*, 93:2764, 1993.

[18] P. Biondi and F. Desclaux: "Silver needle in the Skype." In *BlackHat Europe*, March 2006. `http://www.blackhat.com/presentations/bh-europe-06/bh-eu-06-biondi/bh-eu-06-biondi-up.pdf`.

[19] L. Bobelin and T. Muntean: "Algorithms for network topology discovery using end-to-end measurements." In *International Symposium on Parallel and Distributed Computing*, pages 267–274, Los Alamitos, CA, USA, 2008. IEEE Computer Society, ISBN 978-0-7695-3472-5.

[20] N. Brock, M. Daniels, S. Morris, and P. Otto: "A collaborative computing model for audio post-production," *Future Generation Computer Systems*, 27(7):935–943, 2011, ISSN 0167-739X. `http://www.sciencedirect.com/science/article/pii/S0167739X11000082`.

[21] I. Buck, T. Foley, D. Horn, J. Sugerman, K. Fatahalian, M. Houston, and P. Hanrahan: "Brook for GPUs: Stream computing on graphics hardware," *ACM Transactions on Graphics*, 23(3):774–783, 2004.

[22] N. Buso and C. Allocchio: *LOLA: LOw LAtency audio visual streaming system*, October 2011. `http://www.internet2.edu/presentations/fall11/20111004-ALLOCCHIO-LOLA.pdf`, Internet2 Fall Member Meeting.

[23] S. E. Butner and M. Ghodoussi: "Transforming a surgical robot for human telesurgery," *IEEE Transactions on Robotics and Automation*, 19(5):818–824, 2003.

[24] N. Cascarano, P. Rolando, F. Risso, and R. Sisto: "iNFAnt: NFA pattern matching on GPGPU devices," *ACM SIGCOMM Computer Communication Review*, 40(5):20–26, 2010.

[25] S. Casner and S. Deering: "First IETF internet audiocast," *ACM Computer Communication Review*, 22(3):92–97, July 1992.

[26] S. Casner: *Frequently asked questions (FAQ) on the multicast backbone (MBONE)*, December 1994. ftp://ftp.isi.edu/mbone/faq.txt.

[27] M. Chen: "Leveraging the asymmetric sensitivity of eye contact for videoconference." In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI'02)*, pages 49–56. ACM, 2002.

[28] D. R. Cheriton and S. E. Deering: *Host groups: a multicast extension for datagram internetworks*. Technical report, Stanford University, Stanford, CA, USA, 1985.

[29] N.-M. Cheung, X. Fan, O. C. Au, and M.-C. Kung: "Video coding on multicore graphics processors," *Signal Processing Magazine, IEEE*, 27(2):79–89, 2010.

[30] N.-M. Cheung, O. C. Au, M.-C. Kung, P. H. Wong, and C. H. Liu: "Highly parallel rate-distortion optimized intra-mode decision on multicore graphics processors," *IEEE Transactions on Circuits and Systems for Video Technology*, 19(11):1692–1703, 2009.

[31] S. H. Choi and M. Handley: "Fairer tcp-friendly congestion control protocol for multimedia streaming applications." In *Proceedings of the 2007 ACM CoNEXT conference*, CoNEXT '07, pages 54:1–54:2, New York, NY, USA, 2007. ACM, ISBN 978-1-59593-770-4. http://doi.acm.org/10.1145/1364654.1364717.

[32] N. M. K. Chowdhury and R. Boutaba: "Network virtualization: state of the art and research challenges," *IEEE Communications Magazine*, 47(7):20–26, 2009.

[33] Y.-H. Chu, S. G. Rao, and H. Zhang: "A case for end system multicast." In *ACM SIGMETRICS: Measurement and Modeling of Computer Systems*, pages 1–12, 2000.

[34] I. Daubechies and W. Sweldens: "Factoring wavelet transforms into lifting steps," *J. Fourier Anal. Appl*, 4:247–269, 1998.

[35] DCI: *Digital Cinema System Specification v. 1.2*. http://www.dcimovies.com/DCIDigitalCinemaSystemSpecv1_2.pdf, 2008.

[36] S. E. Deering and D. R. Cheriton: "Multicast routing in datagram internetworks and extended LANs," *ACM Transactions on Computer Systems*, 8:85–110, 1990.

[37] H. Deshpande, M. Bawa, and H. Garcia-Molina: *Streaming live media over a peer-to-peer network*. Technical report 2001-30, Stanford InfoLab, April 2001. http://ilpubs.stanford.edu:8090/501/.

[38] K. Dichev, F. Reid, and A. L. Lastovetsky: "Efficient and reliable network tomography in heterogeneous networks using bittorrent broadcasts and clustering algorithms," *CoRR*, abs/1205.1457, 2012.

[39] C. Diot, B. N. Levine, B. Lyles, H. Kassem, and D. Balensiefen: "Deployment issues for the IP multicast service and architecture," *IEEE Network*, 14(1):78–88, 2000.

[40] K. Doerr and F. Kuester: "CGLX: a scalable, high-performance visualization framework for networked display environments," *IEEE Transactions on Visualization and Computer Graphics*, 17(3):320–332, 2011.

[41] C. Dovrolis, P. Ramanathan, and D. Moore: "Packet-dispersion techniques and a capacity-estimation methodology," *IEEE/ACM Trans. Netw.*, 12(6):963–977, December 2004, ISSN 1063-6692. `http://dx.doi.org/10.1109/TNET.2004.838606`.

[42] D. Dragan and D. Ivetić: "Architectures of DICOM based PACS for JPEG2000 medical image streaming," *Computer Science and Information Systems*, 6(1):185–203, 2009. `http://www.comsis.org/archive.php?show=ppr122-0810`.

[43] F. Dressler: "Availability analysis in large scale multicast networks." In *15th IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS2003)*, volume I, pages 399–403, November 2003.

[44] F. Dressler: *Monitoring of Multicast Networks for Time-Synchronous Communication*. PhD thesis, University of Erlangen-Nuremberg, May 2003. `http://www7.informatik.uni-erlangen.de/~dressler/publications/dissertation.pdf`.

[45] O. E. Dunlap: *The Outlook For Television*. Harper & Brothers Publishers, New York And London, 1932.

[46] A. El-Sayed: *Application-Level Multicast Transmission Techniques Over The Internet*. PhD thesis, INRIA Rhône-Alpes, France, 2004. `http://www.inrialpes.fr/planete/people/elsayed/phd/phd_final_080304.pdf`.

[47] A. El-Sayed, V. Roca, and L. Mathy: "A survey of proposals for an alternative group communication service," *IEEE Network*, 17(1):46–51, January/February 2003.

[48] D. C. Engelbart and W. K. English: "A research center for augmenting human intellect." In *Proceedings of the December 9-11, 1968, Fall Joint Computer Conference, part I*, AFIPS '68, pages 395–410, New York, NY, USA, 1968. ACM. `http://doi.acm.org/10.1145/1476589.1476645`.

[49] H. Eriksson: "The multicast backbone," *Communications of the ACM*, 8:54–60, 1994.

[50] B. Fang, G. Shen, S. Li, and H. Chen: "Techniques for efficient DCT/IDCT implementation on generic GPU." In *IEEE International Symposium on Circuits and Systems, 2005. ISCAS 2005.*, pages 1126–1129. IEEE, 2005.

[51] J. Feit, L. Matyska, V. Ulman, L. Hejtmánek, H. Jedličková, M. Ježová, M. Moulis, and V. Feitová: "Virtual microscope interface to high resolution histological images," *Diagnostic pathology*, 3 Suppl 1:S10, January 2008, ISSN 1746-1596. `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2500118\&tool=pmcentrez\&rendertype=abstract`.

[52] R. Finlayson: *The UDP multicast tunneling protocol*, November 2003. `http://tools.ietf.org/id/draft-finlayson-umtp-09.txt`, IETF Draft <draft-finlayson-umtp-09.txt>.

[53] S. Floyd, M. Handley, and E. Kohler: *Datagram Congestion Control Protocol (DCCP)*. RFC 4340, Internet Engineering Task Force, March 2006. `http://tools.ietf.org/html/rfc4340`.

[54] D. H. Foes, E. Mukab, R. M. Sloneb, B. J. Erickson, M. J. Flynr, D. A. Clunie, K. K. Lloyd Hildebrand, and S. Younga: "JPEG 2000 compression of medical imagery." In *Proc. SPIE*, volume 3980, San Diego, California, 2002.

[55] J. Fousek, J. Filipovič, and M. Madzin: "Automatic fusions of CUDA-GPU kernels for parallel map," *ACM SIGARCH Computer Architecture News*, 39(4):98–99, 2011.

[56] P. Francis: *Yoid: Extending the internet multicast architecture*. `http://www.aciri.org/yoid/docs/index.html`, Unpublished report for ACIRI, April 2000.

[57] J. Franco, G. Bernabé, J. Fernández, and M. E. Acacio: "A parallel implementation of the 2D wavelet transform using CUDA." In *Euromicro Conference on Parallel, Distributed, and Network-Based Processing*, pages 111–118, Los Alamitos, CA, USA, 2009. IEEE Computer Society.

[58] F. Fu and M. van der Schaar: "A systematic framework for dynamically optimizing multi-user wireless video transmission," *IEEE Journal on Selected Areas in Communications*, 28(3):308–320, 2010, ISSN 0733-8716.

[59] R. Gallager: "Low-density parity-check codes," *IRE Transactions on Information Theory*, 8(1):21–28, 1962.

[60] P. Galvez: "From VRVS to EVO (Enabling Virtual Organizations)." In *TERENA Networking Conference 2006*, Catania, Italy, May 2006.

[61] L. Gharai and C. Perkins: "Implementing congestion control in the real world." In *Multimedia and Expo, 2002. ICME '02. Proceedings. 2002 IEEE International Conference on*, volume 1, pages 397–400, 2002.

[62] E. Goldoni and M. Schivi: *End-to-end available bandwidth estimation tools, an experimental comparison*. In F. Ricciato, M. Mellia, and E. Biersack (editors): *Traffic Monitoring and Analysis*, volume 6003 of *Lecture Notes in Computer Science*, pages 171–182. Springer Berlin Heidelberg, 2010, ISBN 978-3-642-12364-1. `http://dx.doi.org/10.1007/978-3-642-12365-8_13`.

[63] D. M. Grayson and A. F. Monk: "Are you looking at me? Eye contact and desktop video conferencing," *ACM Trans. Comput.-Hum. Interact.*, 10(3):221–243, September 2003, ISSN 1073-0516. `http://doi.acm.org/10.1145/937549.937552`.

[64] A. Greenhalgh, F. Huici, M. Hoerdt, P. Papadimitriou, M. Handley, and L. Mathy: "Flow processing and the rise of commodity network hardware," *ACM SIGCOMM Computer Communication Review*, 39(2):20–26, 2009.

[65] Y. Guo, K. Suh, J. Kurose, and D. Towsley: "P2Cast: peer-to-peer patching scheme for VoD service." In *Proceedings of the 12th international conference on World Wide Web*, WWW '03, pages 301–309, New York, NY, USA, 2003. ACM, ISBN 1-58113-680-3. `http://doi.acm.org/10.1145/775152.775195`.

[66] C. Gutwin and S. Greenberg: "Design for individuals, design for groups: tradeoffs between power and workspace awareness." In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*, CSCW '98, pages 207–216, New York, NY, USA, 1998. ACM, ISBN 1-58113-009-0. `http://doi.acm.org/10.1145/289444.289495`.

[67] A. D. Hall: "Developing picturephone service," *Bell Telephone Magazine*, pages 14–21, 1964.

[68] S. Han, N. Kim, K. Choi, and J. W. Kim: "Design of multi-party meeting system for interactive collaboration." In *Communication Systems Software and Middleware, 2007. COMSWARE 2007. 2nd International Conference on*, pages 1–8. IEEE, 2007.

[69] V. Hardman, M. A. Sasse, and I. Kouvelas: "Successful multiparty audio communication over the internet," *Communications of the ACM*, 41(5):74–80, 1998.

[70] V. Hardman, I. Kouvelas, M. Sasse, and A. Watson: *A packet loss robust audio tool for use over the mbone*, 1996. Department of Computer Science Research Note, RN/96/8.

[71] M. Harris, S. Sengupta, and J. D. Owens: *Parallel prefix sum (scan) with CUDA*. In H. Nguyen (editor): *GPU Gems 3*, chapter 39, pages 851–876. Addison Wesley, August 2007. http://www.idav.ucdavis.edu/publications/print_pub?pub_id=916.

[72] M. Hefeeda, A. Habib, B. Botev, D. Xu, and B. Bhargava: "PROMISE: peer-to-peer media streaming using CollectCast." In *Proceedings of the eleventh ACM international conference on Multimedia*, MULTIMEDIA'03, pages 45–54, New York, NY, USA, 2003. ACM, ISBN 1-58113-722-2. http://doi.acm.org/10.1145/957013.957022.

[73] M. V. J. Heikkinen and S. Luukkainen: "Technology evolution of mobile peer-to-peer communications." In *Proceedings of the 4th Annual International Conference on Wireless Internet*, WICON '08, pages 43:1–43:9, ICST, Brussels, Belgium, Belgium, 2008. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ISBN 978-963-9799-36-3. http://dl.acm.org/citation.cfm?id=1554126.1554179.

[74] E. Hladká, P. Holub, and J. Denemark: "An active network architecture: Distributed computer or transport medium." In *3rd International Conference on Networking (ICN'04)*, pages 338–343, Gosier, Guadeloupe, March 2004. ISBN 0-86341-325-0.

[75] C.-W. Ho, O. C. Au, S.-H. Gary Chan, S.-K. Yip, and H.-M. Wong: "Motion estimation for H.264/AVC using programmable graphics hardware." In *IEEE International Conference on Multimedia and Expo, 2006*, pages 2049–2052. IEEE, 2006.

[76] S. Hodges, S. Izadi, A. Butler, A. Rrustemi, and B. Buxton: "ThinSight: versatile multi-touch sensing for thin form-factor displays." In *Proceedings of the 20th annual ACM symposium on User interface software and technology*, pages 259–268. ACM, 2007.

[77] R. Hofer, P. Kaplan, and A. Kunz: "MightyTrace: multiuser tracking technology on lc-displays." In *Proceedings of the SIGCHI Conference on on human factors in computing systems (CHI'08)*, pages 215–218, 2008. http://dl.acm.org/citation.cfm?id=1357091.

[78] P. Holub, R. Arnold, L. Flournoy, P. Fontello, M. Gill, and J. McCarthy: *Interfacing CoUniverse & Internet2 DCN*, April 2009. http://events.internet2.edu/2009/spring-mm/agenda.cfm?go=session&id=10000479&event=909, At Internet2 Joint Tech Winter Meeting 2009, Texas A&M University, USA.

[79] P. Holub, M. Liška, A. Lake, and J. Vollbrecht: *Interfacing CoUniverse & Internet2 DCN*, February 2009. At Internet2 Joint Tech Winter Meeting 2009, Texas A&M University, USA.

[80] P. Holub and L. Miloš: *Uncompressed HD technology*, 2006. Immersive Telemedicine Workshop, 2006. Phoenix, AZ, USA. Invited talk.

[81] N. Hu and P. Steenkiste: "Evaluation and characterization of available bandwidth probing techniques," *IEEE Journal on Selected Areas in Communications*, 21(6), August 2003.

[82] C. Huang, A. Wang, J. Li, and K. W. Ross: "Understanding hybrid CDN-P2P: why limelight needs its own red swoosh." In *Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, NOSSDAV'08, pages 75–80, New York, NY, USA, 2008. ACM, ISBN 978-1-60558-157-6. `http://doi.acm.org/10.1145/1496046.1496064`.

[83] J. Huang, Z. Li, M. Chiang, and A. Katsaggelos: "Joint source adaptation and resource allocation for multi-user wireless video streaming," *Circuits and Systems for Video Technology, IEEE Transactions on*, 18(5):582–595, 2008, ISSN 1051-8215.

[84] N.-F. Huang, H.-W. Hung, S.-H. Lai, Y.-M. Chu, and W.-Y. Tsai: "A GPU-based multiple-pattern matching algorithm for network intrusion detection systems." In *22nd International Conference on Advanced Information Networking and Applications-Workshops, 2008. AINAW 2008.*, pages 62–67. IEEE, 2008.

[85] A. Huszák and S. Imre: "DCCP-based multiple retransmission technique for multimedia streaming." In *Proceedings of the 6th International Conference on Advances in Mobile Computing and Multimedia*, MoMM '08, pages 21–28, New York, NY, USA, 2008. ACM, ISBN 978-1-60558-269-6. `http://doi.acm.org/10.1145/1497185.1497195`.

[86] K. A. S. Immink: *Reed-Solomon codes and the compact disc*. In S. B. Wicker and V. K. Bhargava (editors): *Reed-Solomon codes and their applications*, pages 41–59. New York: IEEE press, 1994.

[87] H. Ishii and B. Ullmer: "Tangible bits: towards seamless interfaces between people, bits and atoms." In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems (CHI'97)*, pages 234–241. ACM, 1997.

[88] R. Jagodic, L. Renambot, A. Johnson, J. Leigh, and S. Deshpande: "Enabling multi-user interaction in large high-resolution distributed environments," *Future Generation Computer Systems*, 27(7):914–923, 2011.

[89] M. Jain and C. Dovrolis: "Ten fallacies and pitfalls on end-to-end available bandwidth estimation." In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, IMC '04, pages 272–277, New York, NY, USA, 2004. ACM, ISBN 1-58113-821-0. `http://doi.acm.org/10.1145/1028788.1028825`.

[90] J. Jannotti, D. K. Gifford, K. L. Johnson, M. F. Kaashoek, and J. J. W. O'Toole: "Overcast: reliable multicasting with on overlay network." In *OSDI'00: Proceedings of the 4th conference on Symposium on Operating System Design & Implementation*, pages 197–212, Berkeley, CA, USA, 2000. USENIX Association.

[91] G. Jin and H.-J. Lee: "A parallel and pipelined execution of H.264/AVC intra prediction." In *The Sixth IEEE International Conference on Computer and Information Technology, 2006. CIT'06.*, pages 246–246. IEEE, 2006.

[92] J. Jo, J. Chae, M. Lee, J. Kim, and O. Byeon: "Uncompressed high definition visual sharing system towards interactive networked collaboration." In *Proceedings of the First International Conference on Immersive Telecommunications*, page 10. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2007.

[93] J. Karat and J. Vanderdonckt: *Tabletops - Horizontal Interactive Displays*. Human-Computer Interaction Series. Springer London, London, 2010, ISBN 978-1-84996-112-7. `http://www.springerlink.com/index/10.1007/978-1-84996-113-4`.

[94] J. O. Kephart and D. M. Chess: "The vision of autonomic computing," *Computer*, 36(1):41–50, 2003.

[95] S. C. Kim, M. G. Kim, and B. H. Rhee: "Seamless connection for mobile p2p and conventional wireless network." In *Advanced Communication Technology, The 9th International Conference on*, volume 3, pages 1602–1605, 2007.

[96] I. Kouvelas: *A Combined Network, System, and User Based Approach to Improving the Quality of Multicast Audio*. PhD thesis, University College London, London, England, May 1998.

[97] A. Kranstedt, P. KĂĽhnlein, and I. Wachsmuth: *Deixis in multimodal human computer interaction: An interdisciplinary approach*. In A. Camurri and G. Volpe (editors): *Gesture-Based Communication in Human-Computer Interaction*, volume 2915 of *Lecture Notes in Computer Science*, pages 112–123. Springer Berlin Heidelberg, 2004, ISBN 978-3-540-21072-6. `http://dx.doi.org/10.1007/978-3-540-24598-8_11`.

[98] N. K. Krishnaprasad, V. Vishwanath, S. Venkataraman, A. G. Rao, L. Renambot, J. Leigh, A. E. Johnson, and B. Davis: "Juxtaview-a tool for interactive visualization of large imagery on scalable tiled displays." In *Cluster Computing, 2004 IEEE International Conference on*, pages 411–420. IEEE, 2004.

[99] C. Krumbholz, J. Leigh, A. Johnson, L. Renambot, and R. Kooima: "Lambda table: high resolution tiled display table for interacting with large visualizations." In *Proceedings of 5th Workshop on Advanced Collaborative Environments, Redmond, Washington*, 2005.

[100] M. C. Kung, O. Au, P. Wong, and C.-H. Liu: "Intra frame encoding using programmable graphics hardware." In *Advances in Multimedia Information Processing–PCM 2007*, pages 609–618. Springer, 2007.

[101] M. Kung, O. C. Au, P. Wong, and C. H. Liu: "Block based parallel motion estimation using programmable graphics hardware." In *International Conference on Audio, Language and Image Processing, 2008. ICALIP 2008.*, pages 599–603. IEEE, 2008.

[102] W. J. van der Laan, A. C. Jalba, and J. B. Roerdink: "Accelerating Wavelet Lifting on Graphics Hardware Using CUDA," *IEEE Transactions on Parallel and Distributed Systems*, 22(1):132–146, January 2011, ISSN 1045-9219. `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5539768`.

[103] R. Le, I. R. Bahar, and J. L. Mundy: "A novel parallel Tier-1 coder for JPEG2000 using GPUs." In *IEEE SASP'11*, pages 129–136, June 2011.

[104] J. Leigh, A. E. Johnson, M. Brown, D. J. Sandin, and T. A. DeFanti: "Visualization in teleimmersive environments," *Computer*, 32(12):66–73, 1999.

[105] J. Leigh, A. Johnson, K. Park, A. Nayak, R. Singh, V. Chowdhry, and T. A. DeFanti: "Amplified collaboration environments." In *VizGrid Symposium*, Tokyo, Japan, 2002. `http://www.evl.uic.edu/cavern/papers/ContinuumPaper.pdf`.

[106] A. Li: *RTP payload format for generic forward error correction*, December 2007. `http://tools.ietf.org/html/rfc5109`.

[107] Y. Li, Z. Li, M. Chiang, and A. Calderbank: "Content-aware distortion-fair video streaming in congested networks," *Multimedia, IEEE Transactions on*, 11(6):1182–1193, 2009, ISSN 1520-9210.

[108] Z. Li and Y. Shin: *Survey of overlay multicast technology*. `http://networks.cs.ucdavis.edu/~lizhi/papers/ECS289.pdf`, Unpublished report for Networks Lab, UC Davis, June 2002.

[109] J. Liebeherr and M. Nahas: "Application-layer multicast with delaunay triangulations." In *Global Telecommunications Conference, 2001. GLOBECOM '01. IEEE*, pages 1651–1655 vol.3, 2001.

[110] S. Lin, D. Costello, and M. Miller: "Automatic-repeat-request error-control schemes," *Communications Magazine, IEEE*, 22(12):5–17, 1984.

[111] Y.-C. Lin, P.-L. Li, C.-H. Chang, C.-L. Wu, Y.-M. Tsao, and S.-Y. Chien: "Multi-pass algorithm of motion estimation in video encoding for generic GPU." In *Proceedings of IEEE International Symposium on Circuits and Systems, 2006. ISCAS 2006.*, pages 4–pp. IEEE, 2006.

[112] W.-L. Liu, K. Zhang, C. Locatis, and M. Ackerman: "Internet-based videoconferencing coder/decoders and tools for telemedicine," *Telemedicine and e-Health*, 17(5):358–362, 2011.

[113] Y. Liu, Y. Guo, and C. Liang: "A survey on peer-to-peer video streaming systems," *Peer-to-Peer Networking and Applications*, 1(1):18–28, 2008, ISSN 1936-6442. `http://dx.doi.org/10.1007/s12083-007-0006-y`.

[114] M. Liška: *Self-organizing Collaborative Environments*. PhD thesis, Faculty of Informatics, Masaryk University Brno, Czech Republic, 2010.

[115] D. J. MacKay and R. M. Neal: *Good codes based on very sparse matrices*. In *Cryptography and Coding*, pages 100–111. Springer, 1995.

[116] A. Majumda, D. Sachs, I. Kozintsev, K. Ramchandran, and M. Yeung: "Multicast and unicast real-time video streaming over wireless lans," *Circuits and Systems for Video Technology, IEEE Transactions on*, 12(6):524–534, 2002, ISSN 1051-8215.

[117] A. Malekzadeh and M. H. MacGregor: "Network topology inference from end-to-end unicast measurements." In *Advanced Information Networking and Applications Workshops (WAINA), 2013 27th International Conference on*, pages 1101–1106. IEEE, 2013.

[118] H. S. Malvar: "Extended lapped transforms: Properties, applications, and fast algorithms," *IEEE Transactions on Signal Processing*, 40(11), November 1992.

[119] J. Marescaux: *Code name: Lindbergh operation*, 2001. `http://www.ircad.fr/event/lindbergh/lindbergh_editorial_en.pdf`, WebSurg Editorial.

[120] J. Marescaux, J. Leroy, M. Gagner, F. Rubino, D. Mutter, M. Vix, S. E. Butner, and M. K. Smith: "Transatlantic robot-assisted telesurgery," *Nature*, 413(6854):379–380, 2001.

[121] T. Margolis, S. Brown, T. Cornish, H. Bracho, M. Stanton, T. C. Melo de Brito Carvalho, F. F. Redigolo, F. C. de Castro, K. Kaneko, J. de Almeida, *et al.*: "Tri-continental premiere of 4K feature movie via network streaming at FILE 2009," *Future Generation Computer Systems*, 27(7):924–934, 2011.

[122] J. Matela: "GPU-based DWT acceleration for JPEG2000." In *Annual Doctoral Workshop on Mathematical and Engineering Methods in Computer Science*, pages 136–143, 2009, ISBN 978-80-87342-04-6.

[123] L. Mathy, R. Canonico, and D. Hutchison: "An overlay tree building control protocol," *Lecture Notes in Computer Science*, 2233:76–87, 2001.

[124] L. Mathy, R. Canonico, S. Simpson, and D. Hutchison: "Scalable adaptive hierarchical clustering." In *NETWORKING '02: Proceedings of the Second International IFIP-TC6 Networking Conference on Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; and Mobile and Wireless Communications*, pages 1172–1177, London, UK, 2002. Springer-Verlag, ISBN 3-540-43709-6.

[125] S. McCanne and V. Jacobson: "vic: A flexible framework for packet video." In *Proceedings of ACM Multimedia'95*, pages 511–512, San Francisco, CA, USA, November 1995.

[126] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner: "OpenFlow: enabling innovation in campus networks," *ACM SIGCOMM Computer Communication Review*, 38(2):69–74, 2008.

[127] R. K. Mok, E. W. Chan, X. Luo, and R. K. Chang: "Inferring the QoE of HTTP video streaming from user-viewing activities." In *Proceedings of the first ACM SIGCOMM workshop on Measurements up the stack*, W-MUST '11, pages 31–36, New York, NY, USA, 2011. ACM, ISBN 978-1-4503-0800-7. `http://doi.acm.org/10.1145/2018602.2018611`.

[128] R. Mok, E. Chan, and R. Chang: "Measuring the quality of experience of HTTP video streaming." In *Integrated Network Management (IM), 2011 IFIP/IEEE International Symposium on*, pages 485–492, 2011.

[129] S. Mu, X. Zhang, N. Zhang, J. Lu, Y. S. Deng, and S. Zhang: "IP routing processing with graphic processors." In *Proceedings of the Conference on Design, Automation and Test in Europe*, pages 93–98. European Design and Automation Association, 2010.

[130] S. Y. Nam, S. J. Kim, S. Lee, and H. S. Kim: "Estimation of the available bandwidth ratio of a remote link or path segments," *Comput. Netw.*, 57(1):61–77, January 2013, ISSN 1389-1286. `http://dx.doi.org/10.1016/j.comnet.2012.08.015`.

[131] J. Nettingsmeier: *Birds on the wire – olivier messiaen's livre du saint sacr é ment in the world's first wave field synthesis live transmission*, 2008. `http://stackingdwarves.net/public_stuff/event_documentation/wfs_live_transmission_2008/WFS-Report-web.pdf`.

[132] H. B. Newman and P. Galvez: *Virtual room videoconferencing system*, April 20 2001. US Patent App. 09/839,847.

[133] D. Nguyen and J. Canny: "Multiview: spatially faithful group video conferencing." In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI'05)*, pages 799–808, New York, NY, USA, 2005. ACM, ISBN 1-58113-998-5.

[134] T. Nguyen and A. Zakhor: "Distributed video streaming over Internet." In *Proc. of Multimedia Computing and Networking (MMCN'02)*, 2002.

[135] J. Ni and D. H. K. Tsang: "Large-scale cooperative caching and application-level multicast in multimedia content delivery networks," *IEEE Communications Magazine*, 43(5):98–105, 2005, ISSN 0163-6804.

[136] A. Nottingham and B. Irwin: "Parallel packet classification using GPU co-processors." In *Proceedings of the 2010 Annual Research Conference of the South African Institute of Computer Scientists and Information Technologists*, pages 231–241. ACM, 2010.

[137] A. Obukhov and A. Kharlamov: *Discrete cosine transform for 8x8 blocks with CUDA*. Technical report, NVidia, October 2008. `http://developer.download.nvidia.com/compute/DevZone/C/html/C/src/dct8x8/doc/dct8x8.pdf`.

[138] K.-I. Okada, F. Maeda, Y. Ichikawaa, and Y. Matsushita: "Multiparty videoconferencing at virtual social distance: MAJIC design." In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'94)*, pages 385–393, New York, NY, USA, 1994. ACM, ISBN 0-89791-689-1.

[139] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai: "Distributing streaming media content using cooperative networking." In *Proceedings of the 12th international workshop on Network and operating systems support for digital audio and video*, NOSSDAV '02, pages 177–186, New York, NY, USA, 2002. ACM, ISBN 1-58113-512-2. `http://doi.acm.org/10.1145/507670.507695`.

[140] V. Padmanabhan, H. Wang, and P. Chou: "Resilient peer-to-peer streaming." In *11th IEEE International Conference on Network Protocols, 2003. Proceedings.*, pages 16–27, 2003.

[141] I. K. Park, N. Singhal, M. H. Lee, and S. Cho: "Efficient design and implementation of visual computing algorithms on the GPU." In *Image Processing (ICIP), 2009 16th IEEE International Conference on*, pages 2321–2324. IEEE, 2009.

[142] I. K. Park, N. Singhal, M. H. Lee, S. Cho, and C. W. Kim: "Design and performance evaluation of image processing algorithms on GPUs," *IEEE Transactions on Parallel and Distributed Systems*, 22(1):91–104, 2011.

[143] K. S. Park, A. Kapoor, and J. Leigh: "Lessons learned from employing multiple perspectives in a collaborative virtual environment for visualizing scientific data." In *Proceedings of the third international conference on Collaborative virtual environments*, pages 73–82. ACM, 2000.

[144] P. Parnes, K. Synnes, and D. Schefstrom: "Lightweight application level multicast tunneling using mtunnel," *Computer Communication*, 21(15):1295–1301, April 1998.

[145] P. Patel, J. Wong, M. Tatikonda, and J. Marczewski: *JPEG compression algorithm using CUDA*. Technical report, Department of Computer Engineering University of Toronto, 2009. `http://www.eecg.toronto.edu/~moshovos/CUDA08/arx/JPEG_report.pdf`.

[146] C. Perkins and L. Gharai: "RTP and the Datagram Congestion Control Protocol." In *Proceedings of the IEEE International Conference on Multimedia and Expo*, pages 1521–1524, Toronto, Canada, July 2006.

[147] B. Pieters, D. Van Rijsselbergen, W. De Neve, and R. Van de Walle: "Motion compensation and reconstruction of H.264/AVC video bitstreams using the GPU." In *Eighth International Workshop on Image Analysis for Multimedia Interactive Services, 2007. WIAMIS'07.*, pages 69–69. IEEE, 2007.

[148] L. Plissonneau and E. Biersack: "A longitudinal view of HTTP video streaming performance." In *Proceedings of the 3rd Multimedia Systems Conference*, MMSys '12, pages 203–214, New York, NY, USA, 2012. ACM, ISBN 978-1-4503-1131-1. http://doi.acm.org/10.1145/2155555.2155588.

[149] I. S. Reed and G. Solomon: "Polynomial codes over certain finite fields," *Journal of the Society for Industrial & Applied Mathematics*, 8(2):300–304, 1960.

[150] L. Renambot, A. Rao, R. Singh, B. Jeong, N. Krishnaprasad, V. Vishwanath, V. Chandrasekhar, N. Schwarz, A. Spale, C. Zhang, *et al.*: "SAGE: The Scalable Adaptive Graphics Environment." In *Proceedings of WACE*, pages 2004–09, 2004.

[151] V. J. Ribeiro, R. H. Riedi, R. G. Baraniuk, J. Navratil, and L. Cottrell: "pathChirp: Efficient available bandwidth estimation for network path." In *4th Passive and Active Measurement Workshop (PAM 2003, 6 April 2002)*, 2003.

[152] V. Roca, Z. Khallouf, and J. Laboure: "Design and evaluation of a low density generator matrix (LDGM) large block FEC codec." In B. Stiller, G. Carle, M. Karsten, and P. Reichl (editors): *Group Communications and Charges. Technology and Business Models. 5th COST 264 International Workshop on Networked Group Communications, NGC 2003 and 3rd International Workshop on Internet Charging and QoS Technologies, ICQT 2003.*, pages 193–204. Springer, 2003.

[153] V. Roca and C. Neumann: "Design, evaluation and comparison of four large block FEC codecs, LDPC, LDGM staircase and LDGM triangle, plus a Reed Solomon small block FEC codec," *INRIA, Rapport de Recherche*, 2004. http://hal.inria.fr/inria-00070770/.

[154] J. Rosenberg and H. Schulzrinne: *An RTP payload format for generic forward error correction*. RFC 2733, Internet Engineering Task Force, December 1999. http://tools.ietf.org/html/rfc2733.

[155] M. Salehie and L. Tahvildari: "Self-adaptive software: Landscape and research challenges," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 4(2):14, 2009.

[156] E. Samei, A. Badano, D. Chakraborty, K. Compton, C. Cornelius, K. Corrigan, M. J. Flynn, B. Hemminger, N. Hangiandreou, J. Johnson, *et al.*: "Assessment of display performance for medical imaging systems: executive summary of AAPM TG18 report," *Medical Physics*, 32:1205, 2005.

[157] K. Savetz, N. Randall, and Y. Lepage: *MBONE: Multicasting Tomorrow's Internet*. IDG Books Worldwide, Inc., Foster City, CA, USA, 1st edition, 1995, ISBN 1568847238. http://www.savetz.com/mbone/.

[158] A. Saxena, F. C. Fernandes, and Y. A. Reznik: "Fast transforms for intra-prediction-based image and video coding." In *Data Compression Conference (DCC), 2013*, pages 13–22, 2013.

[159] R. Schollmeier: "A definition of peer-to-peer networking for the classification of peer-to-peer architectures and applications." In *First International Conference on Peer-to-Peer Computing, 2001. Proceedings.*, pages 101–102. IEEE, 2001.

[160] E. Setton, T. Yoo, X. Zhu, A. Goldsmith, and B. Girod: "Cross-layer design of ad hoc networks for real-time video streaming," *Wireless Communications, IEEE*, 12(4):59–65, 2005, ISSN 1536-1284.

[161] G. Shainer, A. Ayoub, P. Lui, T. Liu, M. Kagan, C. R. Trott, G. Scantlen, and P. S. Crozier: "The development of Mellanox/NVIDIA GPUDirect over InfiniBand—a new model for GPU to GPU communications," *Computer Science-Research and Development*, 26(3-4):267–273, 2011.

[162] T. Shimizu, D. Shirai, H. Takahashi, T. Murooka, K. Obana, Y. Tonomura, T. Inoue, T. Yamaguchi, T. Fujii, N. Ohta, S. Ono, T. Aoyama, L. Herr, N. van Osdol, X. Wang, M. D. Brown, T. A. DeFanti, R. Feld, J. Balser, S. Morris, T. Henthorn, G. Dawe, P. Otto, and L. Smarr: "International real-time streaming of 4K digital cinema," *Future Generation Computer Systems*, 22(8):929–939, October 2006, ISSN 0167-739X.

[163] A. Shriram and J. Kaur: "Empirical evaluation of techniques for measuring available bandwidth." In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications*, pages 2162–2170. IEEE, May.

[164] P. Slovák, E. Hladká, and P. Troubil: "Videoconferencing design for remote groups." In *Proceedings of the 3rd International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2007)*, pages 437–441, 2007.

[165] R. Smith, N. Goyal, J. Ormont, K. Sankaralingam, and C. Estan: "Evaluating GPUs for network packet signature matching." In *IEEE International Symposium on Performance Analysis of Systems and Software, 2009. ISPASS 2009.*, pages 175–184. IEEE, 2009.

[166] J. Strauss, D. Katabi, and F. Kaashoek: "A measurement study of available bandwidth estimation tools." In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, IMC '03, pages 39–44, New York, NY, USA, 2003. ACM, ISBN 1-58113-773-7. http://doi.acm.org/10.1145/948205.948211.

[167] G.-M. Su, Z. Han, M. Wu, and K. Liu: "Joint uplink and downlink optimization for real-time multiuser video streaming over WLANs," *Selected Topics in Signal Processing, IEEE Journal of*, 1(2):280–294, 2007, ISSN 1932-4553.

[168] D. L. Tennenhouse, J. M. Smith, W. D. Sincoskie, D. J. Wetherall, and G. J. Minden: "A survey of active network research," *Communications Magazine, IEEE*, 35(1):80–86, 1997.

[169] S. Tokdemir: *DCT implementation on GPU*. Master's thesis, Georgia State University, 2006. http://digitalarchive.gsu.edu/cgi/viewcontent.cgi?article=1032&context=cs_theses.

[170] Y. Tonomura, D. Shirai, M. Kitamura, T. Nakachi, T. Fujii, and H. Kiya: "Low-Density Generator Matrix codes for IP packet video streaming with backward compatibility." In *Communications (ICC), 2011 IEEE International Conference on*, pages 1–5. IEEE, 2011.

[171] D. Tran, K. Hua, and T. Do: "A peer-to-peer architecture for media streaming," *IEEE Journal on Selected Areas in Communications*, 22(1):121–133, 2004, ISSN 0733-8716.

[172] D. A. Tran, K. A. Hua, and T. T. Do: "ZIGZAG: An efficient peer-to-peer scheme for media streaming." In *IEEE INFOCOM '03*, 2003.

[173] T. D. Tran, J. Liang, and C. Tu: "Lapped transform via time-domain pre- and post-filtering," *IEEE Transactions on Signal Processing*, 51(6), June 2003.

[174] M. Ujaldon and U. V. Catalyurek: "High-performance signal processing on emerging many-core architectures using CUDA." In *2009 IEEE International Conference on Multimedia and Expo*, pages 1825–1828. IEEE, June 2009, ISBN 978-1-4244-4290-4. `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5202878`.

[175] J. Van Velthoven, K. Spaey, and C. Blondia: "Performance of constant quality video applications using the DCCP transport protocol." In *Local Computer Networks, Proceedings 2006 31st IEEE Conference on*, pages 511–512, 2006.

[176] R. Vertegaal: "The GAZE groupware system: mediating joint attention in multiparty communication and collaboration." In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 294–301. ACM, 1999.

[177] R. Vertegaal, I. Weevers, C. Sohn, and C. Cheung: "GAZE-2: conveying eye contact in group video conferencing using eye-controlled camera direction." In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI'03)*, pages 521–528, New York, NY, USA, 2003. ACM, ISBN 1-58113-630-7.

[178] A. Weiß, M. Heide, S. Papandreou, N. Fürst, and A. Balevic: *CUJ2K: a JPEG2000 encoder in CUDA*. Technical report, IPVS, Universitaet Stuttgart, 2009.

[179] D. Wetherall: "Active network vision and reality: lessons from a capsule-based system." In *DARPA Active Networks Conference and Exposition, 2002. Proceedings*, pages 25–40. IEEE, 2002.

[180] T.-T. Wong, C.-S. Leung, P.-A. Heng, and J. Wang: "Discrete wavelet transform on consumer-level graphics hardware," *IEEE Trans. Multimedia*, 9(3):668–673, April 2007, ISSN 1520-9210.

[181] C. Wu, B. Li, and S. Zhao: "On dynamic server provisioning in multichannel P2P live streaming," *IEEE/ACM Trans. Netw.*, 19(5):1317–1330, October 2011, ISSN 1063-6692. `http://dx.doi.org/10.1109/TNET.2011.2107563`.

[182] C. Wu, B. Li, and S. Zhao: "Diagnosing network-wide P2P live streaming inefficiencies," *ACM Trans. Multimedia Comput. Commun. Appl.*, 8(1S):13:1–13:19, February 2012, ISSN 1551-6857. `http://doi.acm.org/10.1145/2089085.2089090`.

[183] D. Xu, M. Hefeeda, S. Hambrusch, and B. Bhargava: "On peer-to-peer media streaming." In *22nd International Conference on Distributed Computing Systems, 2002. Proceedings.*, pages 363–371, 2002.

[184] D. Xu, S. Kulkarni, C. Rosenberg, and H.-K. Chai: "Analysis of a CDN–P2P hybrid architecture for cost-effective streaming media distribution," *Multimedia Systems*, 11(4):383–399, 2006, ISSN 0942-4962. `http://dx.doi.org/10.1007/s00530-006-0015-3`.

[185] Y. Yemini and S. Da Silva: "Towards programmable networks." In *IFIP/IEEE International Workshop on Distributed Systems: Operations and Management*, pages 1–11. Citeseer, 1996.

[186] H. Yin, X. Liu, T. Zhan, V. Sekar, F. Qiu, C. Lin, H. Zhang, and B. Li: "Design and deployment of a hybrid CDN-P2P system for live video streaming: experiences with LiveSky." In *Proceedings of the 17th ACM international conference on Multimedia*, MM '09, pages 25–34, New York, NY, USA, 2009. ACM, ISBN 978-1-60558-608-3. `http://doi.acm.org/10.1145/1631272.1631279`.

[187] Y. Yin and W. Wu: "A real-time measurement algorithm for available bandwidth," *Int. J. Communications, Network and System Sciences*, 2(8):746–753, 2009.

[188] B. Zhang, S. Jamin, and L. Zhang: "Host multicast: A framework for delivering multicast to end users.." In *IEEE INFOCOM '02*, New York, NY, 2002.

[189] S. Q. Zhuang, B. Y. Zhao, A. D. Joseph, R. H. Katz, and J. D. Kubiatowicz: "Bayeux: an architecture for scalable and fault-tolerant wide-area data dissemination." In *NOSSDAV '01: Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video*, pages 11–20, New York, NY, USA, 2001. ACM Press, ISBN 1-58113-370-7.

[190] *Digital Imaging and Communications in Medicine (DICOM). supplement 106: JPEG 2000 interactive protocol*, January 2006. `ftp://medical.nema.org/medical/dicom/final/sup106_ft.pdf`, DICOM Standards Committee, Working Group 4.

[191] *ITU-T recommendation G.114 – One-way transmission time, general recommendations on the transmission quality for an entire international telephone connection*, May 2003. published by International Telecommunication Union, Telegraph and Telephone Consultative Committee.

[192] *SMPTE 2022-1-2007: Forward error correction for real-time video/audio transport over IP networks*, May 2007. published by Society of Motion Picture and Television Engineers.

[193] *SMPTE 2036-2-2008: Ultra High Definition Television – Audio characteristics and audio channel mapping for program production*, July 2008. published by Society of Motion Picture and Television Engineers.

Part II

Collection of Papers

Appendix A

High definition multimedia for multiparty low-latency interactive communication

by Petr Holub, Luděk Matyska, Miloš Liška, Lukáš Hejtmánek, Jiří Denemark, Tomáš Rebok, Andrei Hutanu, Ravi Paruchuri, Jan Radil and Eva Hladká

Future Generation Computer Systems, Amsterdam, The Netherlands: Elsevier Science, 22, 8, pp. 856–861, 6 p. ISSN 0167-739X. 2006.

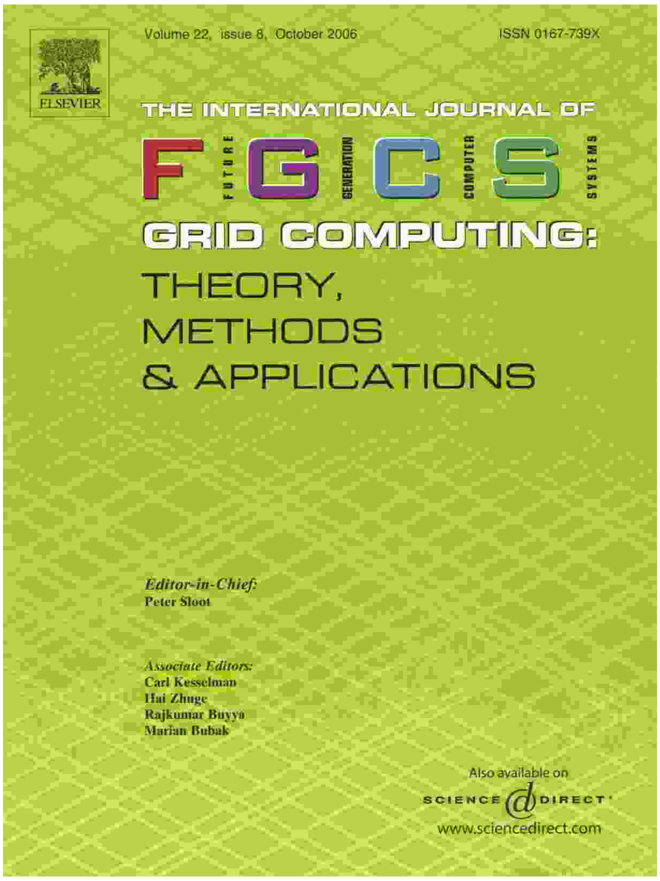

A

papers/FGCS2006-final.pdf

A

papers/FGCS2006-final.pdf

A

papers/FGCS2006-final.pdf

A

papers/FGCS2006-final.pdf

A

papers/FGCS2006-final.pdf

A

papers/FGCS2006-final.pdf

Appendix B

Data transfer planning with tree placement for collaborative environments

by Petr Holub, Hana Rudová, Miloš Liška

Journal of Constraints, Springer, 16, 3, pp. 283–316, 34 p. ISSN 1383-7133 (print) 1572-9354 (online). 2011.

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

papers/CONS2011-final.pdf

papers/CONS2011-final.pdf

B

B

Appendix C

GPU-Accelerated DXT and JPEG Compression Schemes for Low-Latency Network Transmissions of HD, 2K, and 4K Video

by Petr Holub, Martin Šrom, Martin Pulec, Jiří Matela, Martin Jirman

Future Generation Computer Systems, Amsterdam, The Netherlands: Elsevier Science, 29, 8, pp. 1991–2006, 16 p. ISSN 0167-739X. October 2013.

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

C

papers/FGCS2013-final.pdf

papers/FGCS2013-final.pdf

C

Appendix D

Distributed and collaborative visualization of large data sets using high-speed networks

by Andrei Hutanu, Gabrielle Allen, Stephen D. Beck, Petr Holub, Hartmut Kaiser, Archit Kulshrestha, Miloš Liška, Jon MacLaren, Luděk Matyska, Ravi Paruchuri, Steffen Prohaska, Ed Seidel, Brygg Ullmer, Shalini Venkataraman

Future Generation Computer Systems, Amsterdam, The Netherlands: Elsevier Science, 22, 8, pp. 1004–1010, 7 p. ISSN 0167-739X. 2006.

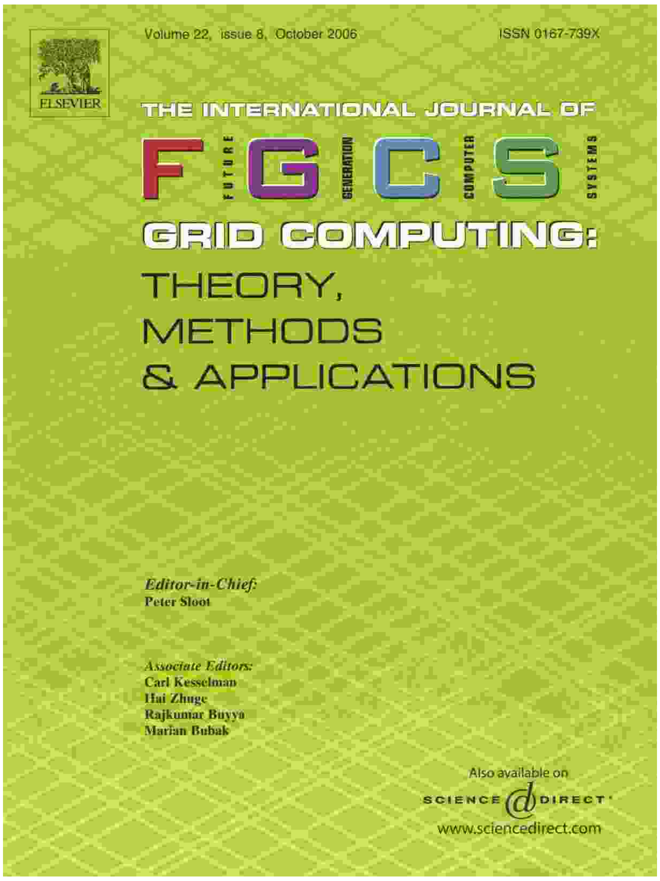

D

papers/FGCS2006-2-final.pdf

D

papers/FGCS2006-2-final.pdf

D

papers/FGCS2006-2-final.pdf

papers/FGCS2006-2-final.pdf

D

papers/FGCS2006-2-final.pdf

D

papers/FGCS2006-2-final.pdf

D

papers/FGCS2006-2-final.pdf

Appendix E

Scalability and Robustness of Virtual Multicast for Synchronous Multimedia Distribution

by Petr Holub, Eva Hladká and Luděk Matyska

In 4th International Conference on Networking, ICN 2005, Reunion Island, France, April 2005. Proceedings. Lecture Notes in Computer Science 3421, Heidelberg: Springer Berlin, 2005. 8 p. ISSN 0302-9743.

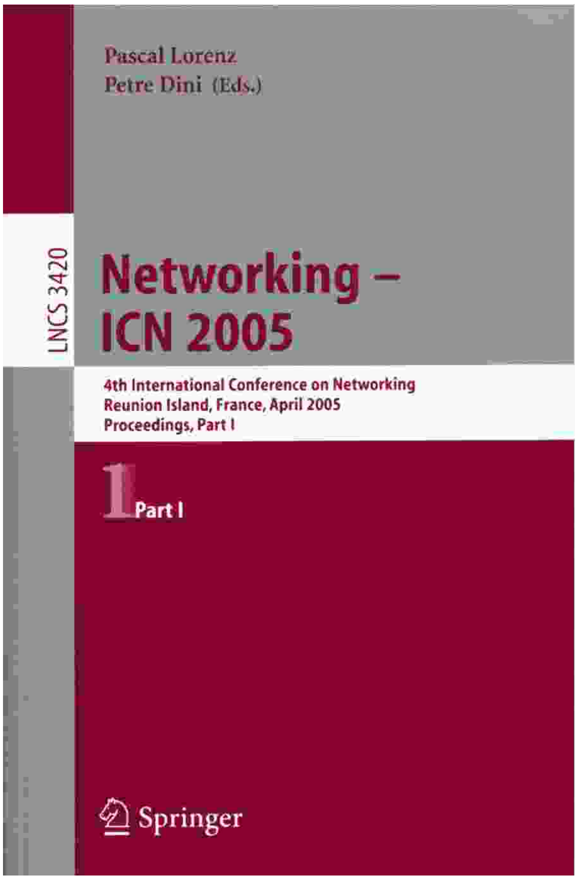

E

papers/ICN2005-final.pdf

papers/ICN2005-final.pdf

E

E

papers/ICN2005-final.pdf

E

papers/ICN2005-final.pdf

E

papers/ICN2005-final.pdf

E

papers/ICN2005-final.pdf

E

papers/ICN2005-final.pdf

E

papers/ICN2005-final.pdf

Appendix F

Grid Infrastructure Monitoring as Reliable Information Service

by Petr Holub, Martin Kuba, Luděk Matyska, Miroslav Ruda

Second European AcrossGrids Conference, AxGrids 2004, Nicosia, Cyprus, January 28-30, 2004. Revised Papers, p. 220–229, Lecture Notes in Computer Science 3165, ISBN 978-3-540-22888-2 (print) 978-3-540-28642-4 (online), Series ISSN 0302-9743. 2004

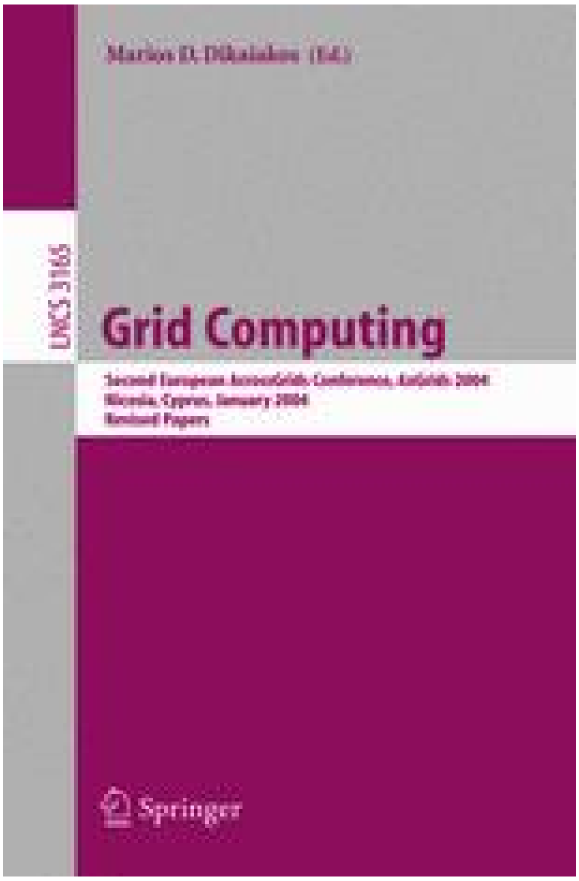

F

papers/AxGrids2004-final.pdf

F

papers/AxGrids2004-final.pdf

F

papers/AxGrids2004-final.pdf

papers/AxGrids2004-final.pdf

F

F

papers/AxGrids2004-final.pdf

F

papers/AxGrids2004-final.pdf

F

papers/AxGrids2004-final.pdf

papers/AxGrids2004-final.pdf

F

F

papers/AxGrids2004-final.pdf

papers/AxGrids2004-final.pdf

F

F

Appendix G

Efficient JPEG2000 EBCOT Context Modeling for Massively Parallel Architectures

by Jiří Matela, Vít Rusňák, Petr Holub

In Data Compression Conference (DCC), 2011. Washington, DC, USA : IEEE Computer Society, 2011. ISBN 978-0-7695-4352-9, p. 423–432. 2011, Snowbird, Utah, USA. 2011.

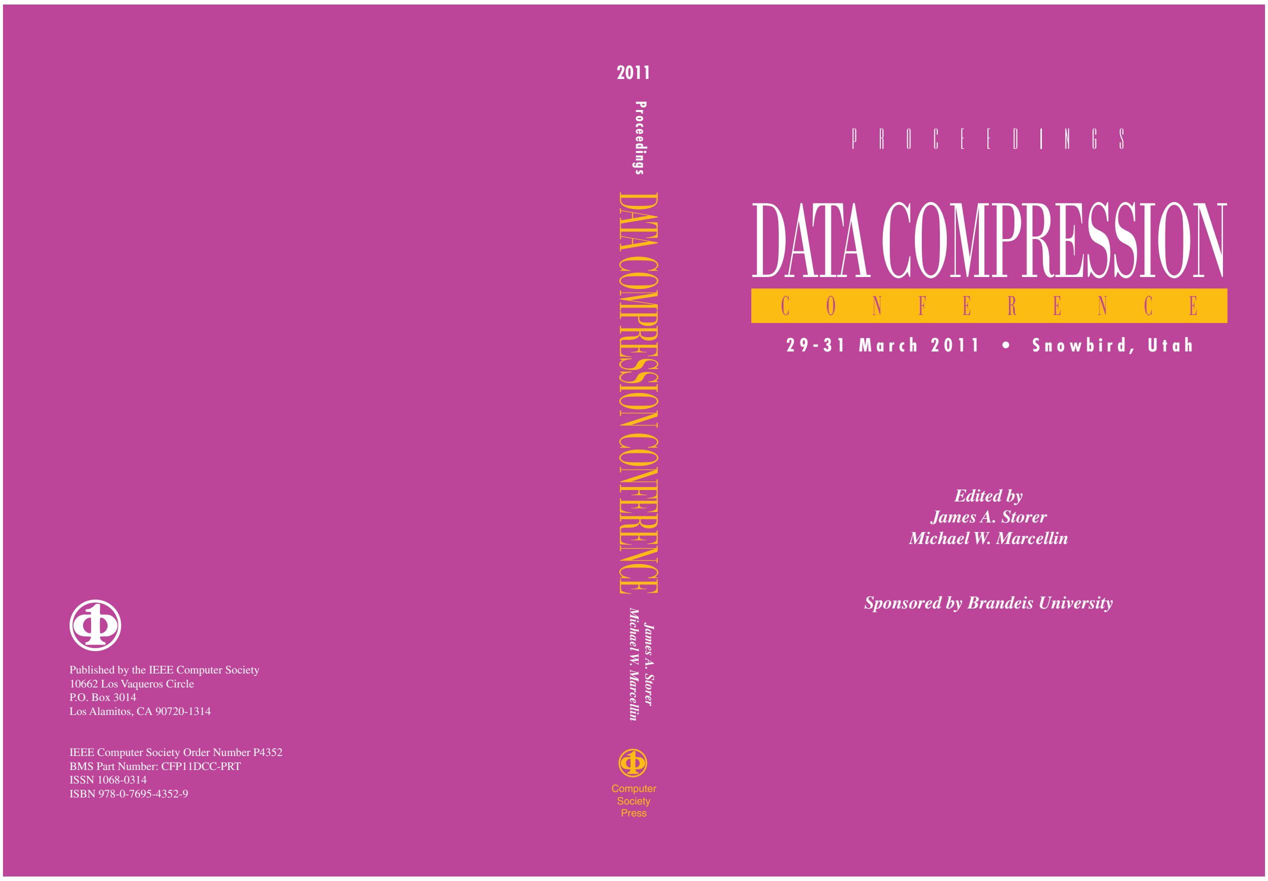

G

papers/DCC2011-final.pdf

papers/DCC2011-final.pdf

G

G

papers/DCC2011-final.pdf

papers/DCC2011-final.pdf

G

G

papers/DCC2011-final.pdf

papers/DCC2011-final.pdf

G

G

papers/DCC2011-final.pdf

papers/DCC2011-final.pdf

G

G

papers/DCC2011-final.pdf

papers/DCC2011-final.pdf

G

Appendix H

GPU-Specific Reformulations of Image Compression Algorithms

by Jiří Matela, Petr Holub, Martin Jirman, Martin Šrom

Proc. SPIE 8499, Applications of Digital Image Processing XXXV, 849901 (October 15, 2012). 2012

PROCEEDINGS OF SPIE

Applications of Digital Image Processing XXXV

Andrew G. Tescher
Editor

13–16 August 2012
San Diego, California, United States

Sponsored and Published by
SPIE

Volume 8499

Proceedings of SPIE 0277-786-786X, V.8499

SPIE is an international society advancing an interdisciplinary approach to the science and application of light.

Applications of Digital Image Processing XXXV, edited by Andrew G. Tescher, Proc. of SPIE
Vol. 8499, 849901 · © 2012 SPIE · CCC code: 0277-786/12/$18 · doi: 10.1117/12.2009271

H

papers/SPIE2012-final.pdf

papers/SPIE2012-final.pdf

H

H

papers/SPIE2012-final.pdf

papers/SPIE2012-final.pdf

H

H

papers/SPIE2012-final.pdf

papers/SPIE2012-final.pdf

H

H

papers/SPIE2012-final.pdf

papers/SPIE2012-final.pdf

H

Appendix I

Distributed Active Element for High-Performance Data Distribution

by Petr Holub and Eva Hladká

IFIP International Conference on Network and Parallel Computing, NPC 2006, Tokio, Japan, October 2006. Proceedings. pp. 27–36, 10 p.

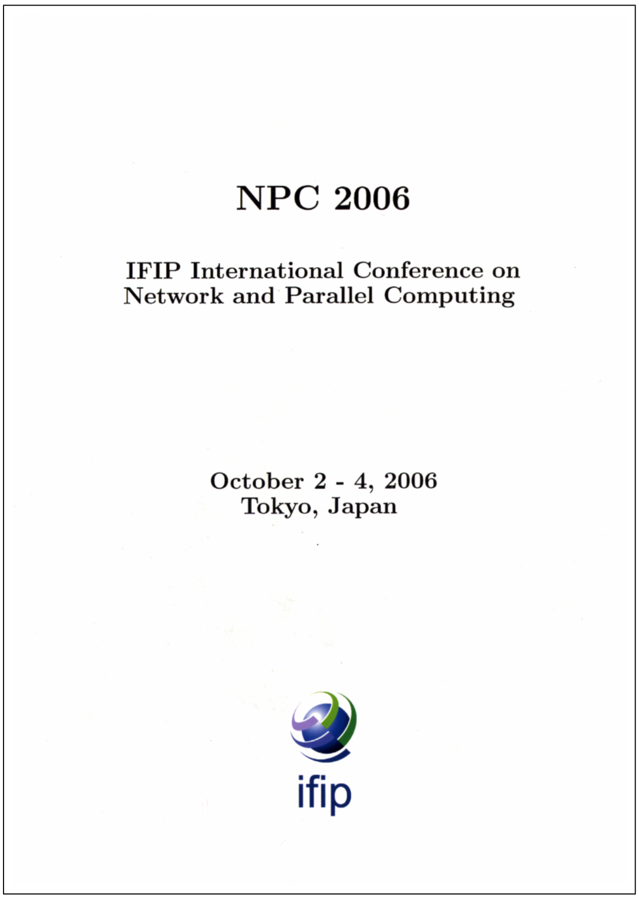

I

papers/NPC2006-final.pdf

papers/NPC2006-final.pdf

I

I

papers/NPC2006-final.pdf

papers/NPC2006-final.pdf

I

I

papers/NPC2006-final.pdf

papers/NPC2006-final.pdf

I

I

papers/NPC2006-final.pdf

papers/NPC2006-final.pdf

I

I

papers/NPC2006-final.pdf

papers/NPC2006-final.pdf

I

I

Appendix J

User Empowered Programmable Network Support for Collaborative Environment

by Eva Hladká, Petr Holub and Jiří Denemark

In Universal Multiservice Networks: Third European Conference, ECUMN 2004, Porto, Portugal, October 25-27, 2004. Proceedings. Lecture Notes in Computer Science 3262, Heidelberg: Springer-Verlag Berlin, 2004. 10 p. ISSN 0302-9743.

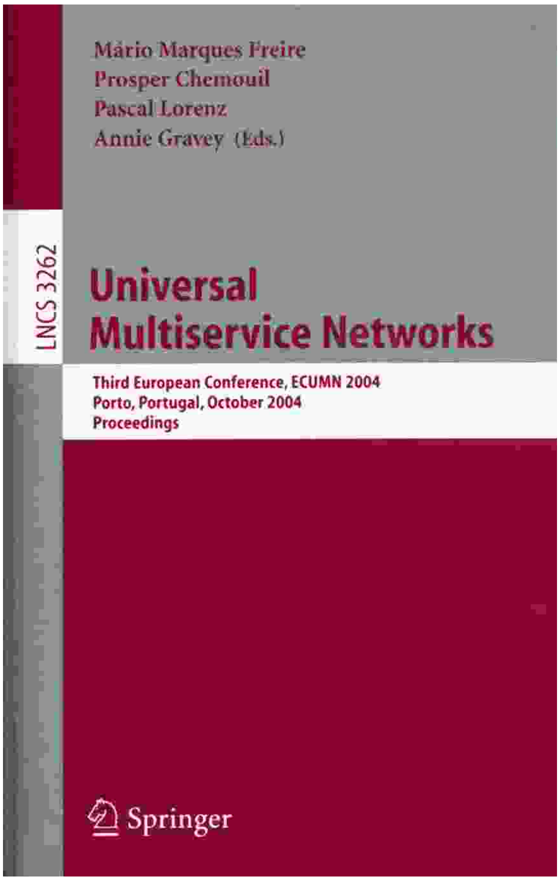

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

J

papers/ECUMN2004-final.pdf

Appendix K

Virtual Multicast

by Petr Holub and Eva Hladká

In Trends in Telecommunications Technologies. Edited by Christos J. Bouras, ISBN 978-953-307-072-8, pages 63–86. Hard cover, 768 pages, Publisher: InTech, Published: March 01, 2010.
doi: 10.5772/180

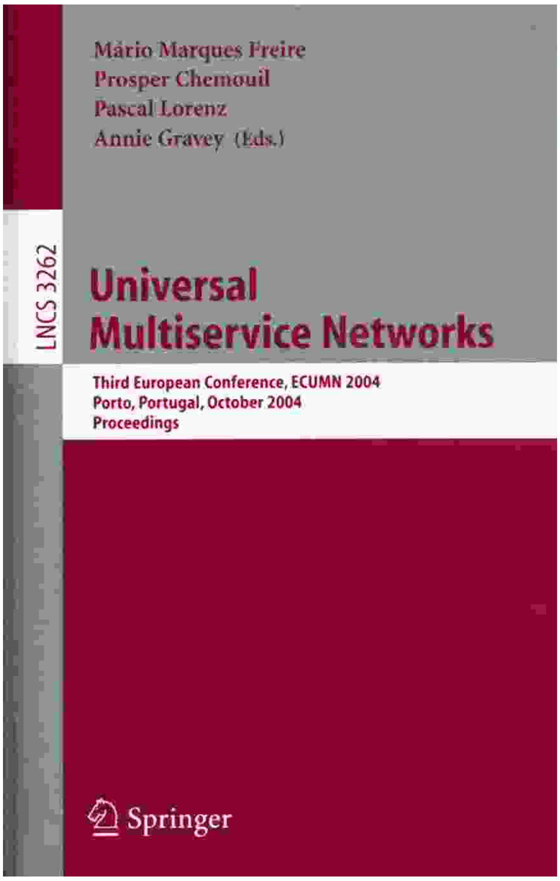

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

papers/ECUMN2004-final.pdf

K

Appendix L

UltraGrid: Low-Latency High-Quality Video Transmissions on Commodity Hardware

by Petr Holub, Jiří Matela, Martin Pulec, Martin Šrom

Proceedings of the 20th ACM international conference on Multimedia, p. 1457–1460. ACM New York, NY, USA. ISBN: 978-1-4503-1089-5. 2012.

papers/ACMMM2012-final.pdf

L

LMMM2012-final.pdf

papers/ACMMM2012-final.pdf

L

Appendix M

Local Search Heuristics for Media Streams Planning Problem

by Jiří Marek, Petr Holub, Hana Rudová

Proceedings of The 27th IEEE International Conference on Advanced Information Networking and Applications (AINA-2013) Barcelona, Spain, March 25-28, 2013.

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

papers/AINA2013-accepted.pdf

M

Appendix N

GColl: Enhancing Trust in Flexible Group-to-Group Videoconferencing

by Petr Slovák, Pavel Troubil, Petr Holub

In Proceeding CHI EA '10 CHI '10 Extended Abstracts on Human Factors in Computing Systems. Pages 3607–3612 ACM New York, NY, USA. table of contents ISBN: 978-1-60558-930-5. 2010

N

papers/CHI2010-final.pdf

papers/CHI2010-final.pdf

N

N

papers/CHI2010-final.pdf

papers/CHI2010-final.pdf

N

papers/CHI2010-final.pdf

papers/CHI2010-final.pdf

N

Appendix O

Exploring collaboration in group-to-group videoconferencing

by Petr Slovák, Peter Novák, Pavel Troubil, Vít Rusňák, Petr Holub, and Eric C. Hofer

In From Research to Practice in the Design of Cooperative Systems: Results and Open Challenges; Proceedings of the 10th International Conference on the Design of Cooperative Systems. London : Springer, 2012. ISBN 978-1-4471-4092-4, s. 229-244. 30.5.2012, Marseille, France.

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

papers/COOP2012-final.pdf

O

Appendix P

CoUniverse: Framework for Building Self-organizing Collaborative Environments Using Extreme-Bandwidth Media Applications

by Miloš Liška, Petr Holub

In Lecture Notes in Computer Science vol. 5415 Euro-Par 2008 Workshops - Parallel Processing. Las Palmas de Gran Canaria, Spain : Springer Berlin / Heidelberg, 2008. ISBN 978-3-642-00954-9, pp. 339-351. 2008.

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

papers/EuroPar-ROIA2009-final.pdf

P

Appendix Q

CoUniverse Orchestrated Collaborative Environments with Dynamic Circuit Networks

by Miloš Liška, Petr Holub, Andrew Lake, and John Vollbrecht

In Proceedings of 2010 Ninth International Conference on Networks, 2010. ISBN 978-0-7695-3979-9, pp. 300-305. France.

papers/ICN2010-final.pdf

Q

papers/ICN2010-final.pdf

Q

papers/ICN2010-final.pdf

Q

papers/ICN2010-final.pdf

Q

papers/ICN2010-final.pdf

Q

papers/ICN2010-final.pdf

Q

Appendix R

Media Streams Planning with Transcoding

by Pavel Troubil, Hana Rudová, Petr Holub

In Proceedings of IEEE 12th International Symposium on Network Computing and Applications 2013 (NCA 2013). ISBN 978-0-7695-5043-5, s. 41–48. Cambridge, Massachusetts, USA.

papers/NCA2013-final.pdf

R

papers/NCA2013-final.pdf

R

papers/NCA2013-final.pdf

R

papers/NCA2013-final.pdf

R

papers/NCA2013-final.pdf

R

papers/NCA2013-final.pdf

R

papers/NCA2013-final.pdf

R

papers/NCA2013-final.pdf

R

Appendix S

Device for Multicast of Optical Signals in the Internet and Other Networks

by Josef Vojtěch, Miloslav Hůla, Jan Radil, Miroslav Karásek, Stanislav Šíma, Jan Nejman, Petr Holub

U.S. Patent Application has been approved as of late September 2013, but not yet published by USPTO; the application has been filed as US 2010/0310254 A1. Czech Patent granted, filed as no. 300811.

papers/OM-US.pdf

S

papers/OM-US.pdf

S

papers/OM-US.pdf

S

papers/OM-US.pdf

S

papers/OM-US.pdf

S

papers/OM-US.pdf

S

papers/OM-US.pdf

S

papers/OM-CZ.pdf

S

papers/OM-CZ.pdf

papers/OM-CZ.pdf

S

papers/OM-CZ.pdf

S

papers/OM-CZ.pdf

S

papers/OM-CZ.pdf

S

papers/OM-CZ.pdf

S