A Scalable Adaptive Multimedia Streaming Middleware

PhD Student Thematic Workshop

Biskupski, Bartosz (Trinity College Dublin, Ireland)
Eide, Viktor S. Wold (Simula Research Laboratory, Norway)
Sacha, Jan (Trinity College Dublin, Ireland)
Schiely, Marc (University of Neuchatel, Switzerland)
Schöffmann, Klaus (Klagenfurt University, Austria)
Serbu, Sabina (University of Neuchatel, Switzerland)
Spielvogel, Christian (Klagenfurt University, Austria)

13-15 June 2007, Klagenfurt, Austria

1 Introduction

Video-on-Demand services are becoming more and more popular. Instead of large dedicated TV sets, people tend to use their workstations or mobile devices, such as PDAs (Personal Digital Assistants) or mobile cellular phones, to consume video and audio content. However, streaming (live) content to these devices introduces a number of difficulties. A small mobile device usually has very limited resource capabilities compared to a standard workstation or a TV set, and therefore, the content has to be adapted to the capabilities of the device, for example by using the MPEG-21 standard.

The main goal of this student workshop is to examine the use of different existing wired and wireless network infrastructures to provide both mobile and stationary devices with the best multimedia streaming quality. Quality improvements (e.g. packet loss reduction) can be achieved by balancing the load among various participating devices, for example by sending the data over alternative paths (multiple source streaming) and merging it at the receiver using the Real-time Transport Protocol (RTP). This could be achieved by building a dynamic adaptive overlay network between the sender and the receiver. A key challenge in creating such an overlay network is *peer selection*, i.e., the selection of peers suitable for media streaming, which requires the introduction of a peer *utility metric* in order to distinguish a "fragile" peer from a "stable" peer. A number of utility metrics can be investigated. One potential approach would be to use the average peer uptime as a utility metric. Other potential approaches may be based on a peer's available bandwidth or average latency. Once the dynamic overlay network has been established, the content can be disseminated to client devices. Different overlays will be considered, including application-level multicast and content-based publish-subscribe services.

2 Participants

The workshop has been organised by six PhD students and a postdoc from four European research institutions. The discussion among the workshop participants was moderated by Prof. Laszlo Böszörmenyi, head of the Information Technologies department at the University of Klagenfurt. The following section contains brief biographies of the workshop speakers.

Bartosz Biskupski is a PhD research student in the department of Computer Science, Trinity College Dublin, Ireland. He received Master of Science degrees in Computer Science from Vrije

University Amsterdam and Warsaw University for his work on transparent fault tolerance for grid and cluster applications. His current PhD research focus is on self-organising peer-to-peer systems for live video streaming. His research is funded by Enterprise Ireland as part of the Mesh TV project, which goal is to provide an efficient scalable platform for live video streaming. He investigates, designs and evaluates approaches for adapting peer-to-peer mesh-based overlays to heterogeneous Internet environments that enable improved network throughput, buffering delay and peer proximity.

Viktor S. Wold Eide is a Post Doc. in the Networks and Distributed Systems department at Simula Research Laboratory, Norway. His interests include distributed systems, communication systems, publish-subscribe communication, multimedia streaming and processing, real-time systems, quality of service, and parallel processing. Eide holds a Dr.Scient. (PhD) from the University of Oslo, Norway, in 2005. The title of his thesis was "Exploiting Event-Based Communication for Real-Time Distributed and Parallel Video Content Analysis". The work was done in the context of the Distributed Media Journaling project, http://www.ifi.uio.no/~dmj/. Eide also holds a Cand.Scient. (MSc) from the University of Oslo, Norway in 1998. The topic of the thesis was automatic parallelisation of recursive algorithms.

Jan Sacha is a PhD candidate at Trinity College Dublin, Ireland. He is a member of Distributed Systems Group and is working in the area of peer-to-peer systems and self-organising systems. His main research contribution is a novel peer-to-peer topology, called gradient topology, that exploits resource heterogeneity in a peer-to-peer system in order to improve the overall system stability and performance. He is currently jointly funded by the Irish Research Council for Science, Engineering and Technology, and by Microsoft. Previously, he was involved in a pan-European project Digital Business Ecosystem and he completed a 3-month internship at Microsoft Research Cambridge, UK. He obtained an MSc degree at both Vrije University Amsterdam and Warsaw University as he participated in a student exchange program between the two universities. He wrote a MSc thesis on the Globe distributed system under Andy Tanenbaum's supervision at Vrije University Amsterdam, and he wrote a bachelor thesis at Warsaw University.

Marc Schiely is writing his PhD on P2P media streaming at the University of Neuchatel in Switzerland. He finished his studies in computer science at ETH Zurich in 2004 with a diploma thesis on routing algorithms in mobile ad-hoc networks. His research is part of the CrossFlux project which is funded by Swiss National Fonds. His research focuses on the construction of efficient overlay networks for live media streaming in heterogeneous environments with high node failure rates. Further, existing approaches to P2P streaming are analysed and conclusions are used to enhance the CrossFlux system. Therefore an analysis of tree- and mesh-based systems is part of future work. A working prototype has been implemented in Java and is currently evaluated in larger environments such as ModelNet and PlanetLab.

Sabina Serbu is a PhD student in the Computer Science department at the University of Neuchatel, Switzerland. Her research interests include load balancing, fault tolerance, routing strategies, and gossip-based protocols in P2P systems. On load balancing, Sabina worked on a solution for load balancing that alleviates the load on the over-loaded peers in the system, by sharing it among under-loaded nodes. The idea of this solution could be applied for multimedia streaming, with the purpose not to overload some peers with routing too many streams, or by redefining load as for example taking into consideration the available bandwidth of the peers. The stream traffic should be evenly shared among the peers participating in the system. Sabina also did a study on the existing solutions of gossipping protocols that are used when multicasting/multimedia streaming. The main usage of gossipping is for membership management and reliable delivery of the multimedia messages. The conclusions of her study are that gossipping is very efficient as a background task for multicasting, and that these algorithms need further adaptation to the time constraint demanded by multimedia streaming.

Christian Spielvogel is a PhD candidate at Klagenfurt University, Austria. As teaching and research assistant he is a member of the Distributed Multimedia Systems Group. He works in the area of Peer-to-Peer systems and Content Delivery Networks. In his PhD thesis he strives for Quality of Service aware ways for delivering audio-visual content. He proposes an innovative architecture he calls Proxy-to-Proxy. The goal of the Proxy-to-Proxy architecture is to enable Quality-of-Service based multimedia delivery over best-effort networks.

Klaus Schöffmann is a PhD candidate at Klagenfurt University, Austria. He is a member of the Distributed Multimedia Systems group and is working in the area of video coding and video summarisation. Klaus Schöffmann finished his studies in computer science at Klagenfurt University in 2005 with a diploma thesis on the Design and Implementation of a Video Session Migration System. His current research focuses on the creation and visualisation of video abstracts based on video segmentation in the compressed domain of H.264/AVC. In his recent research he developed a shot-detection algorithm for H.264/AVC compressed bit streams. Currently, he is working on shot-classification and fine-grained video segmentation beside of developing innovative user interfaces for video-browsing.

3 Presentations

The workshop participants gave presentations on their current research topics. Apart from the students organising the workshop (i.e. the authors of this report), there were five additional guest: Markus Fauster, Michael Ransburg and Udo Schilcher from Klagenfurt University, and Pengpeng Ni and Håkon Stensland from Simula Research Laboratory.

This section contains the titles and short descriptions of the talks given at the workshop. Full version of slides presented during the talks are available at *http://www.minema.di.fc.ul.pt/ /reports/phdws/ScalableAdaptiveMultimediaDeliveryMiddleware_Appendix.pdf*

3.1 Sabina Serbu: Load Balancing and Gossipping for Multimedia Streaming

The neighbours of a node in a multicast overlay are selected based on a peer metric, which may be their uptime or load in terms of used/available bandwidth. Under a dynamic neighbourhood (nodes may join/leave the system or they may change their load in time) a mechanism would be required to keep the traffic load balanced in the system. A solution that was proposed for structured P2P systems is presented to analyse its possible applicability.

One of the mostly used techniques for membership management and reliability in multicast/multimedia streaming is based on gossip protocols. For membership management, information is gossiped about peer neighbours. For reliability, each node gossips information about the chunks or the messages that it has already received to be used to detect lost messages, which are then requested by the peers that have not received them. A summary is presented for some of the related work in the field.

3.2 Marc Schiely: CrossFlux: Tit-for-Tat Revisited: Trading Bandwidth for Robustness

In the last few years many different approaches for distributing content over peer-to-peer (P2P) networks have been discussed and the most important characteristics have been identified: scalability, efficiency, reliability and fairness. With CrossFlux we propose a P2P system for media streaming which incorporates these properties starting from the design. In our approach reliability is coupled with fairness by rewarding peers that contribute more with a higher number of backup links. This coupling can be achieved by using links (1) for content distribution in one direction and (2) as backup in the opposite direction. Further for maximising the throughput and distributing the load among the participating nodes an adaptive join procedure and reorganisation algorithms

are being used. Our evaluation of CrossFlux shows that recovery of node failures is fast due to the backup links, efficiency is increased with the help of self-adaptive techniques and the load is well distributed among the peers.

3.3 Bartosz Biskupski: Peer-to-Peer Mesh-based Live Video Streaming in Heterogeneous Environments

In this talk I present MeshCast, a peer-to-peer (p2p) multicast protocol for applications requiring high bandwidth (such as live video streaming) from a server to a large number of receivers. Traditional tree-based approaches to overlay multicast inefficiently utilise the outgoing bandwidth of participating nodes and poorly adapt to node membership churn. In contrast, MeshCast is based on Chainsaw mesh-based approach to data delivery that better utilises bandwidth and provides excellent adaptation properties. I identify properties that enable mesh-based overlay multicast protocols to better utilise the available bandwidth and consequently support higher data stream rates in heterogeneous environments. MeshCast uses a gossip-based algorithm to adapt the overlay to peer heterogeneity, while still preserving the advantages of a mesh-based overlay. Experiments show that MeshCast can support 68% higher stream rates and provides a 22% improvement in buffering delay over the recently proposed Chainsaw protocol for a heterogeneous node bandwidth distribution.

3.4 Jan Sacha: Gradient Topology

Peer-to-peer environments exhibit a very high diversity in individual peer characteristics ranging by orders of magnitude in terms of uptime, available bandwidth, and storage space. Many systems attempt to exploit this resource heterogeneity by using the highest performing and most reliable peers, called super-peers, for hosting system services. However, due to inherent decentralisation, scale, and complexity of peer-to-peer environments, self-managing super-peer selection is a challenging problem. In this talk, we present an approach to super-peer election and discovery, based on the gradient peer-to-peer topology, that allows a peer-to-peer system to efficiently elect and exploit a set of super-peers with globally highest utility. We discuss a number of peer utility metrics and describe a neighbour selection algorithm that generates a P2P topology with gradient utility structure. The utility information contained in the topology enables an efficient routing algorithm, called gradient search, for high utility peer discovery. In our approach, peers use decentralised aggregation techniques in order to estimate system-wide peer utility properties and to calculate super-peer election thresholds. This approach allows peers to adapt and maintain the set of super-peers to the changes in the running system. The approach can also reduce the number of switches between super-peers and ordinary peers by relaxing super-peer utility requirements.

3.5 Viktor S. Wold Eide: Exploiting Content-Based Networking for Fine Granularity Adaptive Multi-Receiver Video Streaming

Efficient delivery of video data over computer networks has been studied extensively for decades. Still, multi-receiver video delivery is challenging, due to heterogeneity and variability in network availability, end node capabilities, and receiver preferences.

Here we demonstrate that content-based networking is a promising technology for efficient multi-receiver video streaming. The contribution of this work is the bridging of content-based networking with techniques from the fields of video compression and streaming. In the presented approach, each video receiver is provided with fine grained and independent selectivity along the different video quality dimensions region of interest, signal-to-noise ratio for the luminance and the chrominance planes, and temporal resolution. This allows each client to individually adapt the quality of the received video according to its currently available resources and own preferences. Experimental results show that efficient delivery, in terms of network utilisation and end node processing requirements, is maintained.

3.6 Christian Spielvogel: A ProXy-to-ProXy QoS Aware Multimedia Delivery

Client/Server architectures do not scale well. Moreover, if the connection between the client and the server has not the necessary level of QoS then video streaming may be impossible, even if the server is far from being overloaded. P2P communication has good features for data transmission with limited requirements, but it is too much fragile to serve as a basis for video streaming. CDNs are rigid, static, proprietary and expensive. In my PhD thesis I suggest an intermediate solution, based on dynamic overlay networks of proxies. This technology is called a Proxy-to-Proxy middleware for media streaming. The basic idea is that content, typically but not necessarily originating from large, high-quality servers and from live broadcast sources, is dynamically replicated (fully or partially) onto logical networks of proxies supporting a small geographic region and a certain area of interests. For example, a network of proxies of hotels in one area, another network of proxies of basket ball clubs in another area may build such proxy groups. Such groups share besides computing and networking resources also access rights and get thus an inexpensive and extremely efficient streaming service. The streaming is realized using RTP over UDP. Therefore a proper error treatment mechanism is required. Proper error treatment is achieved by (1) applying error avoidance, (2) applying error concealment or (3) combining both approaches. If none of the three treatments is able to achieve error freeness, the request has to be rejected. In order to avoid comparing all possible media stream resolutions, numbers of streaming servers and amounts of redundant network packets, I use the A* algorithm. Using this algorithm it is possible to achieve the highest end-user satisfaction by examining a minimal set of alternatives.

3.7 Klaus Schöffmann: Video Analysis in the Compressed Domain of H.264

In the last decade the consumption of digital video data on many even mobile computer devices has experienced a massive growth. YouTube, video-on-demand download portals or IP-TV broadcast services are only a few popular examples which show the ubiquity of digital video data.

However, the usage, or let's say handling, of digital videos still lacks of interactivity. For instance, searching for a small part of a video in a very large file may become a very exhausting and annoying task when using a standard video player. Furthermore, with high resolutions (e.g. HDTV or Cinema-HD) this task may also include noticeable delays due to the high decoding effort when quickly browsing through a video file.

We need a video tool which allows us to fast and efficiently navigate through a video file and gives us an impression of what is going on in a video and what it is about. Indeed video analysis is the first step for such a tool. In our current research, we are concentrating on how to use the compressed data of H.264/AVC encoded video files for such an analysis in order to save decoding time. Furthermore, we are investigating alternative presentation types of a video in order to quickly impart the content of a video to an end user.

3.8 Pengpeng Ni: User-friendly H.264/AVC for Remote Video Browsing

With the growing popularity of variable network technologies, it is highly desirable to enable effective and quick browsing of remote multimedia content, especially when the amount of digitalised video materials is expanding dramatically. However, increasing the temporal resolution of traditional MPEG compliant video will easily lead to network congestion. To solve the problem, we propose a transcoding scheme for H.264/AVC video that fully utilises the benefits of recently proposed SP- and SI-frames to facilitate user-friendly remote stream browsing. Based on a quadratic rate-distortion model, we apply our scalable rate control algorithm to adaptively determine related encoding parameters respect to different characteristics of the video streams. The proposed scheme is expected to be effectively used to implement a full digital Video Cassette

Recording(VCR) functionality in multimedia streaming applications such as Video-On-Demand, remote video editing and video broadcasting.

3.9 Markus Fauster: A New Approach to the Adaptation of Multimedia Scenes

Automatic adaptation and personalisation is very difficult and restricted in usefulness, if content is only available in monolithic representations like video streams. Object-based videos, multimedia scenes and presentations have been research and standardisation topics for many years, but practical application lags behind expectations despite the potential advantages of these representations. Correctly authored multimedia scenes allow for much more sophisticated content processing due to their structured representations and the optimised encoding options for many elementary media types. However, adaptation as well as authoring is difficult in today's low level scene description formats, which is why I am working on a higher level, constraint based language that supports scene description, adaptation constraints and media element dependencies with a single vocabulary. With this approach it will be possible to derive an optimised scene structure satisfying all conditions automatically.

3.10 Michael Ransburg: Treatment of Extensive Multimedia Metadata in Multimedia Content Streaming Scenarios

Today's increasing variety of media data results in a great diversity of XML-based metadata, which describes the media data on semantic or syntactic levels, in order to make it more accessible to the user. This metadata can be of considerable size, which leads to problems in streaming scenarios. Other than media data, XML metadata has no concept of "samples" thus inhibiting streamed (and timed) processing, which is natural for media data. In order to address the challenges and requirements resulting from this situation, the concept of streaming instructions is introduced. In particular, streaming instructions address the problem of fragmenting metadata, associating media segments and metadata fragments, and streaming and processing them in a synchronised manner. This is achieved by enriching the metadata with additional attributes to describe media and XML properties. Alternatively, a style sheet approach provides the opportunity to dynamically set such streaming properties without actually modifying the XML description. We show the usefulness of our work by implementing an adaptation node which uses our mechanisms to extend a static adaptation approach (MPEG-21 Digital Item Adaptation) towards dynamic and distributed usage scenarios.

4 Connection Points

The workshop participants have been divided into the following two groups according to their research interests: the multimedia coding group, and the peer-to-peer streaming group. The former group discussed topics related to multimedia content encoding, streaming, and displaying, while the latter group was more focused on peer-to-peer overlays and their usage for multimedia streaming. In spite of this seemingly strict division, more ad-hoc discussions took place among people having common interests in both of these two main groups.

The following two tables describe the group participants, the topics discussed in each group, and the connections between the participants based on their research interests.

4.1 Media coding group

	Håkon Stensland	Klaus Schöffmann	Marcus Fauster	Michael Ransburg	Pengpeng Ni
Scalable H.264/AVC	×	×	×		×
Parallelisation	×	×	×		
Video Analysis		×			×
Video Streaming	×	×		×	×

4.2 Peer-to-peer streaming group

	Bartosz Biskupski	Christian Spielvogel	Jan Sacha	Marc Schiely	Sabina Serbu	Viktor Eide
Mesh and tree topologies	×			×		×
Topology adaptation		×	×		×	
Gossipping protocols	×	×	×	×	×	×
Utility metrics		×	×		×	
Peer heterogeneity	×		×	×		×
Fairness	×			×	×	
Load balancing		×			×	
Network locality	×	×		×	×	×
Multicast	×			×	×	×
Multiple description coding	×	×				×
Network simulators	×	×	×	×	×	

5 Open Issues

This section contains a list of relevant topics that were mentioned during the workshop but were not directly addressed by any of the participants. These topics need further study in the future.

Security Any peer-to-peer system, in order to be deployed in a real-life environment, needs to include at least basic security mechanisms. In particular, a peer-to-peer system must deal with malicious nodes and malicious users that may attempt to attack and compromise the system from the inside.

Incentives for cooperation It is a well known fact that users of peer-to-peer applications prefer to consume resources of the system rather than contribute their own resources. Mechanisms must be developed in open peer-to-peer systems that encourage resource sharing by users and punish so called "free riders". This can be also described as the fairness problem.

Combine mesh and tree topologies It has been observed that mesh and tree approaches to data streaming have many similarities. In particular, data sent throughout a mesh overlay forms logical multicast trees. These trees are similar to multicast trees maintained by peers in tree-based approaches to data streaming. However, logical trees in mesh overlays do not need to be explicitly maintained by peers as well as self-adapt to varying bandwidth, peer churn, etc. We have identified that a number of algorithms developed for tree-based p2p streaming approaches can be applied to mesh-based streaming and vice versa.

Network locality Many system designers characterise peers by their available bandwidth or network latency. However, the latency or bandwidth are properties of connections between two communicating peers. Consequently, a peer-to-peer system should adapt its topology to the underlying network infrastructure in order to maximise the total throughput or to minimise the total latency.

Streaming with the gradient topology The gradient topology has been designed as a general-purpose peer-to-peer architecture that can be tailored for different applications and domains. One of the challenges discussed at the workshop is the application of the gradient topology to the multimedia streaming systems.

Rate-distortion and Scalable Video Coding The flexibility provided by Scalable Video Coding techniques comes often at the price of decreased video compression efficiency. Video encoder needs to balance the trade-off between the bit-rate of the video stream and its quality loss by determining the optimal choice of encoding parameter settings. However, the rate-distortion optimisation problem can be very complex since multiple coding options are involved in the encoding process and applications may also put time constraints on the encoding environment. For codec designers, the challenge is how to evaluate and improve the effectiveness of Scalable Video Coding for real-time video streaming applications.

Automatic code generation for multicore-architectures Nowadays, CPUs are using a multi-core design rather than a single-core design. Thus, in order to fully benefit from the available power of a CPU future applications will have to use parallel code (e.g. multi-threaded). As manually parallelising code is a difficult and error-prone task, there is an obvious need for doing this task automatically. Although preprocessor hints for some compilers have been proposed in the past, we think that a more appropriate solution should be developed. For instance, one solution could be to enhance the statements of programming languages for that purpose.

6 Future Collaboration

The following section describes common research interests discovered during the workshop between the participants and potential topics for future collaborations between the participants.

J. Sacha, C. Spielvogel, and S. Serbu: Topology adaptation and utility metrics Jan, Christian and Sabina identified common interests in the topics of peer-to-peer topology adaptation based on utility metrics. Affinity metrics, based on physical proximity, semantical closeness, and current system load, are used by Christian in his work for the proxy and stream selection. Jan is interested in applying similar metrics (specific for multimedia streaming) for the generation of the gradient topology. Neighbour selection, and hence topology adaptation, is usually based on gossiping. Sabina is working on applications of gossiping protocols in peer-to-peer networks. Christian is planning to use a gossip-based approach in his system in order to decentralise group formation, while Jan is using gossip-based algorithms for the neighbour selection in the gradient topology.

M. Schiely, C. Spielvogel, and B. Biskupski: Multiple Description Coding (MDC) The three researchers have a common interest in Christian's research on Multiple Description Coding (MDC) based on ffmpeg open-source tool for video encoding. The use of the MDC can have significant importance in building future peer-to-peer video streaming protocols that adapt to the available bandwidth in overlay. The MDC video encoding enables clients with different bandwidth capacity to receive different video stream rates. This is in contrast to regular video encoding where clients with low bandwidth receive distorted video. The MDC implementation by Christian has an important advantage over many other implementations that it can operate in real-time, which is essential for live peer-to-peer video streaming.

Further collaboration could focus on (1) exchanging knowledge on Multiple Description Coding, (2) providing feedback to Christian on his implementation and (3) extending the current implementation to enable temporal and SNR scalability to achieve better performance.

B. Biskupski, M. Schiely, and V. Eide: Peer-to-peer overlays for media streaming

The three researchers and their respective research groups focus on peer-to-peer overlay networks for media streaming. Different structures for disseminating media data between nodes are considered, including trees and meshes. While the researchers from Ireland and Switzerland so far have mostly considered multicast distribution, the researcher from Norway have worked on publish-subscribe communication for data dissemination. The publish-subscribe approach is driven by a need for supporting heterogeneity and variability, where different nodes may receive different subsets of the data in the media streams.

A number of similarities between our approaches have been identified including a need for (1) dynamic adaptation to changing bandwidth, (2) exploitation of node heterogeneity, (3) providing robustness to membership churn (nodes joining and failing), and (4) adaptation to node bandwidth constraints by using some kind of scalable video coding, including multiple description coding and layered coding.

A number of open issues have been identified that require further investigation. Issues include a comparison of mesh-based, tree-based, and hybrid approaches, proximity issues where costs can be reduced by adapting the overlays to node proximity, incentives issues, and security issues (to prevent malicious node behaviour).

The participants also discussed a number of evaluation approaches used for validating their research work, including ModelNet, EmuLab, PlanetLab, and ns2.

The research area of peer-to-peer overlay networks for media streaming is quite large. However, we have identified several opportunities for further collaboration. As starting points we propose to investigate and compare tree-based, mesh-based, and hybrid approaches. As an example, it may be possible to apply tree-based adaptation algorithms to mesh overlays and vice versa. Due to the complexity of tools and systems for validating research in this area it seems fruitful to exchange experiences on the use of the above mentioned evaluation tools. Similarly, available software for scalable video coding are complex. To further strengthen the evaluation of the suitability of proposed methods for peer-to-peer overlays for media distribution it seems highly beneficial to exchange knowledge on the use of scalable video streaming systems.

K. Schöffmann, P. Ni, and V. Eide: Video analysis based on scalable video coding

Klaus is currently working on semantic-related video analysis in the compressed domain of H.264/AVC in order to provide a more satisfying interactivity of video data usage. Viktor has previously worked on real-time video analysis (but not in the area of H.264/AVC). Pengpeng and Viktor are working on scalable video coding. As video analysis is obviously a first step of efficient scalable video coding in the sense of rate-distortion performance, there should be some collaborative work within this domain. Both research areas could benefit from each other: scalable video coding could help to perform information-retrieval more efficiently and faster. Moreover, video analysis could help to perform scalable video coding in a more efficient way.

H. Stensland and K. Schöffmann: Video-coding parallelisation

Since Håkon has just started his PhD thesis in the area of video coding parallelisation and Klaus has already some experience on parallelising H.264/AVC decoding, there is a very high potential to work together on this topic.

V. Eide and M. Ransburg: Meta-data driven adaptation

Ransburg and Eide have worked on adaptation of (scalable) media for several years, and both have published a number of papers within this research topic. Their work has lead to implementations of different systems which provide similar functionality. It should be noted that the research activities on media adaptation in both Klagenfurt and at Simula Research Laboratory are not single-person activities. On the contrary, a number of people is associated to this research topic at both locations.

For their work, the researchers from Klagenfurt have relied on open standards from the MPEG/JVT group and also actively contributed to them. On the other hand, the researchers from Norway developed an individual solution for scalable video coding, starting in 2003, due to lack of sufficient functionality and support in the open standards.

Furthermore, researchers from Norway follow a middleware managed approach to multimedia adaptation. In this approach, applications are composed from reusable components and the adaptation process is handled by the middleware itself. Researchers from Klagenfurt have not yet worked on middleware for adaptation.

The two groups have used different media codecs and different meta-data (again, open standards versus individual solutions). The Klagenfurt system provides interoperability, as it relies on open standards. On the other hand, the system developed in Norway provides some interesting new approaches to adaptation. An example is support for adaptive behaviour where the play-out rate is (slightly) decreased in order to reduce bandwidth requirements. The resulting temporal displacement may get reduced by another adaptation where the play-out rate is increased. The Klagenfurt system cannot provide such kind of adaptation so far.

Additionally, the Klagenfurt group have investigated in a codec agnostic approach for multimedia adaptation which enables to deploy generic adaptation engines. They also have done work on session mobility (descriptors for indicating this are standardised in MPEG-21 DIA). Both topics are of interest for the researchers from Norway.

Given the differences and similarities, it would be interesting to combine the strengths of the approaches taken by the two research groups. A further analysis of both approaches would allow a better understanding of potential synergies.

Therefore, it would be beneficial to meet and discuss further in order to explore the potential synergies in more detail. Future cooperations in for example projects, papers, or student exchanges seem likely.

Acknowledgements

The workshop organisers and participants would like to acknowledge the MiNEMA and ESF for the funding and the support which made this workshop possible.