Multicast Routing for High-Quality Multimedia Environments: Deployment and New Problems

Pavel Troubil, Hana Rudová
Faculty of Informatics
Masaryk University
Brno, Czech Republic
Email: `pavel@ics.muni.cz, hanka@fi.muni.cz`

Petr Holub
Institute of Computer Science
Masaryk University
Brno, Czech Republic
Email: `hopet@ics.muni.cz`

Abstract—**Motivated by remote interactive collaboration over high-quality multimedia, we solve the multigroup multicast routing problem under critical runtime limits and we focus on realistic aspects crucial for deployment of advanced collaborative environments. We developed solvers based on mixed integer programming and ant colony optimization, which are integrated in currently developed environments to support education of deaf students at Czech universities. They are also deployed for distributed media production for live broadcasting of sporting and cultural events. The paper presents our ongoing research and development aims in deployment of the environments and new arising problems in multicast routing.**

Keywords–multicast routing; planning; multimedia; interactive collaboration

I. INTRODUCTION AND MOTIVATION

The recent years have witnessed rapid performance growth of commodity hardware and its multimedia processing capabilities. The progress stimulated development of interactive applications, which can use high-quality media for communication and remote collaboration. Cheaper deployment of such applications brought remote work and cooperation into new areas, like distributed musical performances [1], or distributed media production in broadcasting. New opportunities arose in other areas, e. g., telemedicine [2] and remote scientific visualization. In the pursuit to create self-organized collaborative tools which will support these novel forms of remote interactive communication, we develop methods for automated planning of concurrent network transmissions. They are essential for intelligent behaviour of the collaborative tools, as they eliminate the need of manual setup. The most recent solvers enable effective response to dynamics of computer networks while scaling to medium-sized collaborating groups, supporting 20 – 30 participants in a videoconference.

We target applications with extensive demands on quality of service provided by computer networks, particularly capacity and end-to-end delay. The high-quality interactive multimedia applications usually transmit video streams with high image resolutions, either compressed by codecs with low compression rates or uncompressed at all. Such videos in UHD (Ultra High Definition) or 4K resolution need a network capacity between 2–24 Gb/s. In order to maintain image without visible quality degradation for applications like telemedicine, the transmitted bandwidths always remain relatively high in comparison to available capacities, although networks may enforce higher compression rates. For example, Digital Cinema standard [3] specifies JPEG2000 compressed formats with up to 250 Mb/s bandwidth for 4K video resolution. Low compression rates do not only increase image quality and required bandwidth, but are also encoded faster than codecs with high compression rates. In addition, compression latencies contribute to the total transmission latency from source to destination, which should be minimized for interactive applications. The International Telecommunication Union in their recommendation G.114 [4] declares that end-to-end latencies above 150 ms introduce negative impact on users' perception of interactivity in remote communication.

On top of the UltraGrid [5], a software for low-latency transmissions of high-quality video, the CoUniverse middleware [6] is used to orchestrate the collaborative environments, i. e., configure multiple point-to-point transmissions to build complex interactive applications. This paper aims to present to the community our ongoing work on transmission planning problems in the CoUniverse and methods which we proposed and implemented to solve them. We will also describe the current development of novel collaborative environments with specialized applications, and new planning problems which they introduce.

The CoUniverse is driven by a user-empowered approach, i. e., any user can run the collaborative environment on her own and has full control over all applications within. Running the environment shall be available at any time, as simple as a Skype call, and allow top-tier quality of multimedia. The transmission planning is then similar to peer-to-peer videoconferencing tools, e. g., Celerity [7]. The common denominator is a need to operate without any administrative privileges to computers nor network. As a result, knowledge of physical topology of an underlying network is incomplete or missing at all. Network administrators do not publish the topologies, since they need to modify them anytime, usually in reaction to load or failures. Only some experienced users might have partial knowledge of the underlying topology, usually their local network and its neighborhood. Planning of transmissions relies often on knowledge of overlay topology, i. e., mutual visibility among the collaborative tools similar to reachability by the ICMP (Internet Control Message Protocol) ping. And since native IP multicast is often unavailable across independently administered networks [8], the user-empowered environments also require user-empowered substitutes for multipoint communication called Application level multicasts (ALM) [9].

The problem of planning network transmissions is known

ISBN: 978-1-61208-437-4

as multicast routing in literature on combinatorial optimization [10]. The multicast routing is either single-group, i. e., with single data source which delivers data to a set of destinations, or multigroup with multiple data sources and an independent set of destinations for each of them. Mixed integer programming methods were proposed for both single and multiple groups [11]. But for scalability reasons, current research is focused on heuristics, particularly approaches from computational intelligence. For the single group, genetic algorithms are repeatedly revisited, e. g., [12], and ant colony optimization (ACO) techniques are also proposed often [13] [14]. The multigroup problem is harder to solve and is also less studied. Again, the genetic algorithms are the most successful [15].

The rest of the paper is structured as follows. In Section II, we formulate the problems and mention methods applied to solve them. Section III sketches two practical environments, where we implement and deploy solvers proposed in Section II. In Section IV, we analyze the modified routing problems which arise from the deployment.

II. PROBLEMS AND METHODS

We model the computer network by a directed weighted multigraph $G = (V_o \cup V_u, E_o \cup E_u)$. The graph covers two layers of network representation. Overlay nodes V_o are computers running the collaborative tools and overlay links E_o represent mutual reachability of the tools. Knowledge of this layer is always available, including latencies and (potentially inaccurate) capacities of the overlay links. Yet, it is not always sufficient for reliable planning decisions. Overlay links may share underlying physical network infrastructure, but the sharing is not captured in the representation of the overlay links, and may result in overloading of the physical infrastructure when several overlay links are incorrectly considered as independent. The underlay nodes V_u and underlay links E_u represent the underlying network where available — even when it is only part of the topology. In such case, knowledge of association between overlay and underlay links is required. Capacities of some underlay links (but not necessarily all) can be known.

For each data source $s \in S$ $(S \subset V_o)$, there is a set of destinations $D(s)$, which request transmission from the source s. The overlay nodes, which are neither sources nor destinations, are called ALM agents. The agents can receive the data, forward them further, replicate them, and also process contents of the data. They provide infrastructure for multipoint communication by mimicking of native IP multicast in the overlay network. Solution of the multigroup multicast routing problem is a set of multicast trees, one tree per source. Each multicast tree is a subgraph of the overlay network topology graph (V_o, E_o). The tree is directed, rooted in the source s, and set of its leaves equals to the set of destinations $D(s)$.

For each link in all trees, we also have to decide the multimedia format for transmission along the link. Multimedia can be encoded in various formats. Video formats differ, among others, in image resolution, framerate, compression algorithm, and its parameters. For easier representation and simpler solution, we abstract from these attributes and work with a discrete set of formats, which are common in practice. The attributes are expressed in bandwidth required for transmission of the formats, subjective quality perceived by receiving users, latency introduced by on-the-fly encoding or decoding, and amount of hardware resources for the encoding or decoding. The ALM agents are able to transcode multimedia formats on-the-fly, i. e., decode them and encode in a different format. Exploiting this capability, the planning can assign different formats to each incident link of an ALM agent in a multicast tree. Each destination in a multicast tree can receive data in a different format, which shall provide as good quality as possible and be adapted to destination's network capacity.

Solutions have to respect capacities of network links, limits of hardware resources (performance of processor and graphics card or available memory) at overlay nodes, and sometimes also (un)availability of features required for processing of some formats at the nodes. We aim to optimize two components of objective function: minimize total transmission latency, which includes latency of data transmission as well as delay induced by multimedia processing at overlay nodes, and maximize subjective quality of multimedia received at destinations.

We solved the above mentioned problem by methods of mixed integer programming (MIP). At first, we analyzed the problem without transcoding and knowledge of underlay network topology [16], and evaluated suitability of several cycle-avoidance methods adopted from the travelling salesman problem. Later, we extended our MIP models with the underlay knowledge and support of format selection [17]. The planning of multimedia transcoding is the main distinction of our methods from related work on multicast routing.

Next problem extension introduces uncertain network capacities. The user-empowered approach and inherent inaccuracy or absence of network knowledge can be partially overcome with automated tools for network topology inference [18] and bandwidth estimation [19]. In addition to uncertainty in network parameters caused by continuously changing state of the network, these automated tools provide results with only limited reliability. In some of our tests, the tools misestimated actual bandwidth by more than 50 % on 1 Gb/s links. Instead of assuming that capacities of overlay network links are known exactly, we model them by cumulative distributive functions as is usually the case in network problems with uncertainty [20]. We solved the extended problem by a metaheuristic algorithm based on ant colony optimization, which fits well with the uncertainty and allows sufficient scalability [21].

The iterations of the algorithm adhere to the classical structure of iterations in the ACO. First, the solutions are generated for all data sources in parallel. We adopt the tree-building approach similar to single group algorithm from paper [13], where ants construct the trees by moving in network topology graph. In a next phase, the individual multicast trees are gathered and their load on the network summed. The results are used to set parameters for tree generation in subsequent iterations. Finally, the pheromones are deposited on links used by the trees in an amount proportional to the quality of the solution.

The design of the algorithm also allowed its simple application on the multicast routing problem with dynamic reconfiguration [22]. The dynamic reconfiguration is needed when a solution is deployed, a communication is underway, and the input for planning changes as a result of change in network state or source and destination sets. In such situations, the problem needs to be solved again with ongoing aim to optimize latency and quality, but data transmission paths shall be modified as little as possible. Every change of path from source to a destination interrupts the communication for several seconds and disturbs users. The algorithm in paper [22] is designed to avoid the interruptions, but does not guarantee path preservation.

Since our algorithms target area of interactive applications, they also have to provide interactive response when the user is waiting for establishment of the communication or a dynamic reconfiguration. The ACO algorithm can solve problems with hundreds of overlay nodes within a second [21] [22].

III. CURRENT PROGRESS IN DEPLOYMENT

We are currently working on future deployment of the planning algorithms within the CoUniverse, focusing on two application areas: distributed media production for live broadcasting and education of deaf students at Czech universities.

A. Distributed Media Production for Live Broadcasting

Live broadcasting from sporting and cultural events requires a lot of personnel and specialized technical equipment, which is transported to venues by many production trucks (vehicles equipped with production control rooms). Technology usually used for live broadcasting requires live editing directly at the venue. The director switches in real-time between available cameras, always selecting one source for transmission to the broadcasting center via a satellite connection. High prices of broadcasting transmission routes enforced the presence of large teams and equipment at the venue, since their expensive transport between the venues was still cheaper. These days, computer networks with sufficient capacity cover most of the venues and will allow to substitute original broadcasting routes via satellites with video transmissions over computer networks. Still, many venues will be connected through low-capacity network links in comparison to network core, but on-the-fly multimedia transcoding can overcome the limitations, and will allow to reduce expenses on transport of personnel and the equipment.

Media production in live broadcasts could follow the scheme depicted in Figure 1. Two concurrent sporting events are held at two venues. The director works from the broadcasting control center. Content from all cameras at both venues is delivered in preview quality to the broadcasting control (red arrows). The preview quality requires much lower bandwidth, but is sufficient for video cutting decisions. Based on the decisions (green arrows), video content from one source is delivered to the distribution center in full quality. The editor can switch in real-time among all the sources, and the distribution center feeds the selected content to satellite distribution or terrestrial broadcasting networks (blue dashed arrows). We demonstrated a prototype of such live broadcasting system at The 14th Annual LambdaGrid Workshop [23], where we transmitted an ice hockey match captured by 5 cameras from Czech Republic to New Zealand. Building upon experience from the demonstration, another one will be prepared for the 15th Annual LambdaGrid Workshop in autumn 2015. The demonstration shall present multiple improvements of the system, particularly features important for the broadcasting practice. We expand on them in Section IV.

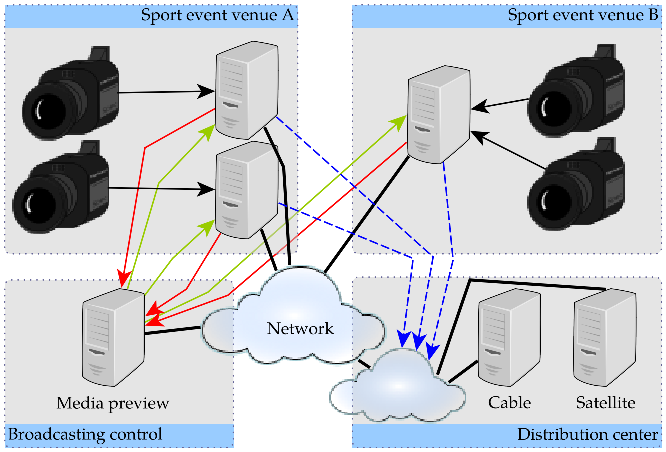

Figure 1. Scheme of live broadcasting infrastructure

B. Remote Interpreting of Sign Language for Education

Integrated education of deaf and hard-of-hearing students poses a challenging problem for any educational facility, but becomes especially difficult at universities. The narrow specialized fields of study use a lot of technical terms and courses are often taught in a foreign language. Should the deaf students have an option to visit lectures and seminars, highly skilled interpreters of sign language are required, who know the specialized vocabulary and/or are able to interpret the foreign language. Since such skilled interpreters are scarce, they are often required to inefficiently travel among universities. Videoconferencing can offer an effective and efficient alternative comfortable for students, interpreters, and teachers. Better education could be provided to more deaf students with lower costs.

The videoconference remote interpreting (VRI) has been studied mainly for legal purposes at the police station, in court, or healthcare. For example, see the Avidicus project [24]. Current research is focused rather on psychological implications of the remote interpreting than technological aspects. VRI usually uses only a single video and audio channel, but deployment in university lectures requires advanced solutions. The Support Center for Students with Special Needs at Masaryk University (SCSSN) designed the CoUnSiL, a specialized environment for remote interpreting in lectures, where each student has his own screen with picture of teacher, her presentation, interpreter, and other students. Since students need to follow and understand long lectures and seminars, the visual quality of the videos has to be maximized. As illustrated by picture-in-picture interpreting on TV, the sign language is comprehensible even if the picture is low quality. Yet, the low quality is not suitable for education, since it wears the deaf out very quickly and diminishes their ability to comprehend and retain information. Also due to a difficult setup of the environment, deployment of the CoUniverse is natural. In cooperation with SCSSN, we are adapting the CoUniverse for the education of deaf students. Testing of the prototype received a positive response from future users, and deployment for regular teaching across several Czech universities is expected during 2016.

IV. NEW PROBLEMS

The development of the environment for live broadcasting brought to light several new extensions of the problem. The camera switching feature requires a change of approach to transition between two solutions during a dynamic reconfiguration. Until now, our algorithms did not support any elaborate transition procedure. Yet, the live broadcasting has special requirements on cuts when switching between two sources. For good visual experience of spectators, the cut has to take place exactly between two video frames. But on the packet level, the frames cannot be so easily separated, and borders between frames are not aligned among individual sources. The problem of seamless switching can be solved by concurrent transmission of data from both sources to a specialized switching ALM agent. While receiving both data sources, such ALM agent can decide the right moment to stop transmission of the original stream and start the new one, when the cut shall be visually clear. But the concurrent reception doubles requirements on incoming network capacity of the processing node. The capacity can be reserved for two sources at a time during the entire broadcast, decreasing quality of received media all the time. Another solution of the problem could involve a multistep transition. First, the original source would transmit lower bandwidth (and quality as well) to the switching ALM agent. Second, transmission of the new source to the agent would start. Third, the original source would stop transmitting to the switching agent, and finally the freed capacity would be filled by the new source. As a result, spectators would see the lower quality picture only during the transition phase, not all the time. From the routing point of view, the multiphase transitions require not only switch between original and new solution, but four solutions deployed in relatively quick succession.

Visually clear cuts require concurrent reception of two sources at a switching ALM agent. There could be also additional requirements on such node, e. g., a specialized hardware or software equipment for video processing. For every source and subset of its destinations, there can be a predefined ALM agent (or a set of agents), which have to lie on path between the source and destination subset. Our ACO algorithm is open to such modification, as we only split tree construction in two phases: one to search for the predefined ALM agents, another to build the rest of the paths down to destinations.

The cuts between cameras occur in reaction to user-induced request. The dynamic reconfiguration of multicast trees is not caused by change of network topology, e. g., failure of a network link. Therefore, the planning algorithm should certainly guarantee preservation of transmission paths from a previous solution where a user requires so. The path preservation would not be only favoured by the objective function, but required by problem constraints. Again, the ACO algorithm can be run with the preserved path as an unmodifiable basis for further growth of the multicast tree.

Although the previous problems are inspired by the application in live broadcasting, their solution may find future use in other collaborative environments. The following problem of planning under strong network uncertainty is common to almost all our future applications. As mentioned in Section II, our algorithms support two layers of network representation: always available knowledge of overlay and optional (possibly partial) knowledge of underlying physical topology. The ACO algorithm is also designed for cooperation with tools for bandwidth estimation and inference of underlying network topology with uncertain results. Unfortunately, the tools are not ready for deployment. According to our testing, their runtimes are often high, and results on high-capacity networks unreliable beyond admissibility. In order to partially overcome the issues, we are searching for methods of partial topology inference from information on data loss. Most often, our issues with network congestion are linked to incorrect assumptions that overlay links are independent while they actually share capacity of underlay links. In contrast to the commonly proposed methods, which are unobtrusive and aim to load the network as little as possible, we can use the data about actual congestion when one occurs. And once it occurs, we need to respond quickly and establish or restore transmissions, no matter if we are able to infer anything about the underlying network or not. Therefore, we need to develop strategies for the responses in various situations, like initialization of the collaborative environment, arrival of a new participant, or sudden drop of delivery rates without any user-induced cause. The strategies have to include selection of links in multicast trees where rates should be decreased, the decrease rate, and procedures for detection and utilization of free network capacity. We assume that the strategies will not modify the multicast routing algorithms, but rather find suitable parameters of the input network topologies for their execution.

ACKNOWLEDGMENT

We thank the Czech Science Foundation for support by project GAP202/12/0306. Work of Petr Holub is also supported by CESNET Large Infrastructure project LM2010005.

REFERENCES

[1] C. Drioli, C. Allocchio, and N. Buso, “Networked performances and natural interaction via LOLA: Low latency high quality A/V streaming system,” in Information Technologies for Performing Arts, Media Access, and Entertainment. Springer, 2013, pp. 240–250.

[2] K. Zhang, W.-L. Liu, C. Locatis, and M. Ackerman, “Uncompressed high-definition videoconferencing tools for telemedicine and distance learning,” Telemedicine journal and e-health : the official journal of the American Telemedicine Association, vol. 19, no. 8, 2013, pp. 579–584.

[3] Digital Cinema Initiatives, “Digital Cinema System Specification v. 1.2,” http://dcimovies.com/specification/DCI_DCSS_v12_with_errata_2012-1010.pdf, p. 155, 2012, accessed: 26/08/2015.

[4] International Telecommunication Union, “One-way transmission time: ITU-T recommendation G.114,” http://www.itu.int/rec/T-REC-G.114/en, p. 20, 2013, accessed: 26/08/2015.

[5] “UltraGrid,” http://www.ultragrid.cz/en, accessed: 26/08/2015.

[6] “CoUniverse,” http://couniverse.sitola.cz, accessed: 26/08/2015.

[7] X. Chen, M. Chen, B. Li, Y. Zhao, Y. Wu, and J. Li, “Celerity: towards low-delay multi-party conferencing over arbitrary network topologies,” in Network and Operating System Support for Digital Audio and Video. ACM, 2011, pp. 123–128.

[8] C. Diot, B. N. Levine, B. Lyles, H. Kassem, and D. Balensiefen, “Deployment issues for the IP multicast service and architecture,” IEEE Network, vol. 14, no. 1, 2000, pp. 78–88.

[9] C. Yeo, B. Lee, and M. Er, “A survey of application level multicast techniques,” Computer Communications, vol. 27, no. 15, 2004, pp. 1547–1568.

[10] C. A. S. Oliveira and P. M. Pardalos, “A survey of combinatorial optimization problems in multicast routing,” Computers & Operations Research, vol. 32, no. 8, 2005, pp. 1953–1981.

[11] C. A. Noronha Jr. and F. A. Tobagi, “Optimum routing of multicast streams,” in Proceedings of the IEEE INFOCOM, 1994, pp. 865–873.

[12] Y. Yen, H. Chao, R. Chang, and A. Vasilakos, "Flooding-limited and multi-constrained QoS multicast routing based on the genetic algorithm for MANETs," Mathematical and Computer Modelling, vol. 53, no. 11-12, 2011, pp. 2238–2250.

[13] H. Wang, H. Xu, S. Yi, and Z. Shi, "A tree-growth based ant colony algorithm for QoS multicast routing problem," Expert Systems with Applications, vol. 38, no. 9, 2011, pp. 11 787–11 795.

[14] P.-Y. Yin, R.-I. Chang, C.-C. Chao, and Y.-T. Chu, "Niched ant colony optimization with colony guides for QoS multicast routing," Journal of Network and Computer Applications, vol. 40, 2014, pp. 61–72.

[15] L. Sanna Randaccio and L. Atzori, "Group multicast routing problem: A genetic algorithms based approach," Computer Networks, vol. 51, 2007, pp. 3989–4004.

[16] P. Troubil and H. Rudová, "Integer linear programming models for media streams planning," in International Conference on Applied Operational Research, 2011, pp. 509–522.

[17] P. Troubil, H. Rudová, and P. Holub, "Media streams planning with transcoding," in IEEE Network Computing and Applications (NCA), 2013, pp. 41–48.

[18] A. Malekzadeh and M. H. MacGregor, "Network topology inference from end-to-end unicast measurements," in Advanced Information Networking and Applications Workshops (WAINA), 2013, pp. 1101–1106.

[19] F. Thouin, M. Coates, and M. Rabbat, "Large scale probabilistic available bandwidth estimation," Computer Networks, vol. 55, no. 9, 2011, pp. 2065–2078.

[20] D. Lorenz and A. Orda, "QoS routing in networks with uncertain parameters," IEEE/ACM Transactions on Networking, vol. 6, no. 6, 1998, pp. 768–778.

[21] P. Troubil, H. Rudová, and P. Holub, "Media streams planning with uncertain link capacities," in IEEE Network Computing and Applications (NCA), 2014, pp. 197–204.

[22] ——, "Media streams planning with uncertain link capacities," Networks: Special issue on metaheuristics in network optimization, 2014, under review.

[23] "14th annual global lambdagrid workshop," http://www.glif.is/meetings/2014/, accessed: 26/08/2015.

[24] AVIDICUS Project, "Videoconference interpreting," http://www.videoconference-interpreting.net/, accessed: 26/08/2015.