

Scalable Mobile Multimedia Group Conferencing based on SIP Initiated SSM*

Thomas C. Schmidt
HAW Hamburg, Dep. Informatik
Berliner Tor 7
D-20099 Hamburg, Germany
t.schmidt@ieee.org

Hans L. Cycon
FHTW Berlin, FB I
Allee der Kosmonauten 20-22,
D-10315 Berlin, Germany
hcycon@fhtw-berlin.de

Matthias Wählisch
HAW Hamburg, Dep. Informatik
Berliner Tor 7
D-20099 Hamburg, Germany
waehlich@ieee.org

Mark Palkow
daViKo GmbH
Am Borsigturm 40,
D-13507 Berlin, Germany
palkow@daviko.com

Abstract

Multimedia group communication emerges to focal interest at mobile devices, enriching voice or video conferencing and complex collaborative environments. The Internet uniquely provides the benefit of scalable, dynamic group communication services, vitally aiding commonly limited mobile terminals. The traditional Internet approach of Any Source Multicast (ASM) routing, though, remains hesitant to spread beyond limited, controlled environments. It is widely believed that simpler and more selective mechanisms for group distribution in Source Specific Multicast (SSM) will globally disseminate to many users of multicast infrastructure and services. Mobility management for the recent SSM standard is under debate — SSM group session initiation up until now remains unsupported by the Session Initiation Protocol (SIP).

In this paper we present straightforward extensions to SIP for negotiating SSM sessions. SIP protocol specifications and semantic are compatibly extended without adding new SIP methods. We will introduce a multimedia communication software with distributed architecture as implementation reference. Furthermore mobility management of the underlying routing layer is discussed and evaluated on grounds of real-world Internet topologies.

*This ongoing work is supported by the German Bundesministerium für Bildung und Forschung within the Project *Moviecast*.

Keywords: Multimedia Group Session, Fully Distributed Multiparty Conferences, SIP-based Peer-to-Peer, Source Specific Multicast, SSM Source Mobility, VCoIP

1. Introduction

Mobile multimedia communication becomes more and more a day-to-day companion. Audio-visual devices performing synchronous communication such as voice or video conferencing over IP (VoIP/VCoIP) are now almost ubiquitous. This trend is also reflected in teaching and learning scenarios, as there is a rising number of Internet/ Intranet-backed courses at universities or private companies. Real-time communication and infotainment raise new challenges to the Internet infrastructure, and so do mobile conference users. The availability of a new, truly mobile IP enabled network layer [11] not only offers connectivity to nomadic users at roaming devices, while preserving communication sessions beyond IP subnet changes, but re-raises questions concerning the quality of IP services: The constant bit rate scenarios of voice and video conferencing will appear significantly disturbed by packet loss intervals, delays or jitter exceeding 50-100 ms. IP multicasting will be of particular importance to mobile environments, where users commonly share frequency bands of limited capacities [12]. Thus, when heading towards VoIP/VCoIP as a standard Internet service, important steps for global usability have to be taken with a focus on ease and quality.

In the present paper we address the issue of mobile multimedia group conferencing, exemplarily taking perspective on a VCoIP (Video Conferencing over IP) software with a distributed, peer-to-peer architecture and its applications [17]. Such lightweight solution should receive support from network layer multicast, restricting service to Source Specific Multicast communication for the sake of deployment simplicity. Source Specific Multicast (SSM) [2, 9], just released as an initial standard, is considered a promising improvement of group distribution techniques. In contrast to Any Source Multicast (ASM) [5], optimal (S,G) multicast source trees are constructed immediately from (S,G) subscriptions at the client side, without utilizing network flooding or RendezVous Points. Source addresses are to be acquired by out of band channels, which a SIP [22] session initiation in conferencing scenarios should facilitate.

However, up until now SIP is not prepared to negotiate SSM group sessions. We therefore present a simple and straight forward extension of session initiation handshakes, suitable to establish SSM conferencing within an uncoordinated peer-to-peer model. Compliant with standard unicast and ASM transactions, we propose to solely add on protocol semantics without introducing new SIP methods. Our video conferencing system serves as a platform for reference implementation.

We further on discuss session mobility with the special focus on real-time multicast group communication. Conferencing parties request seamless real-time performance of a mobility aware group communication service, thereby attaining the simultaneous roles of mobile multicast listener and source. Intricate multicast routing procedures, though, are not easily extensible to comply with mobility requirements. Significant effort has been already invested in protocol designs for mobile multicast receivers. Only limited work has been dedicated to multicast source mobility, which poses the more delicate problem [19, 26].

Addresses in Internet mobility carry the dual meaning of logical *and* topological identifiers. While MIPv6 operates dual addresses transparently at end points, SSM routing needs to account for logical subscription *and* topological forwarding. We present and analyze approaches to SSM routing, which adapt to source mobility.

In this paper we first discuss the mobile multimedia group conferencing problem and related work. In section 3 we present our SIP signaling scheme for SSM group initiation. Multicast mobility is discussed and analyzed in section 4. Finally, section 5 is dedicated to a conclusion and an outlook.

2. The Mobile Multimedia Group Conferencing Problem and Related Work

2.1. Group Conference Sessions

Multimedia session-based communication requires a number of initial negotiations to be accomplished. At first, a caller requesting contact to one or several partners will expect to address a personal identifier, but establish the corresponding conference session with the devices currently in use by the callees. Unlike in mobile telephony, the Internet architecture is required to locate users and mask the user-device mapping, following the paradigm of location transparency like in email services. At second, media and service data need to be identified as to meet capabilities of all session members. Once established, sessions need to persist, while mobile devices roam. Intermediate handovers thereby should unnoticeably comply with quality of service measures for real-time communication.

A flexible, fairly general Internet signaling solution has been presented with the Session Initialization Protocol (SIP) [22]. SIP session participants commonly register at some SIP proxy server, which thereby gains awareness of the user's current device. Locating a nomadic user therefore can be considered equivalent to identifying its SIP server or session directory. To enable seamless virtual session naming services, Rosenberg & Schulzrinne [21] proposed to rely on DNS extensions, the SRV and NAPTR records, which are capable of encoding a service-to-server name mapping. Beside user localization, SIP covers negotiations about user capabilities, user availability, the call set-up by Session Description Protocol (SDP) data and the handling of the calls itself. SIP introduces its own infrastructure of proxy servers, which actively perform a proxy-to-proxy routing by using SIP-URLs. SIP is based on an extensible method framework and open to store persistent data.

SIP forms a multi-layered application protocol that interacts between components in a transactional way. Each (asynchronous) request initiates an open transaction state and requires completion by at least one response. SIP operations in unicast management are exemplarily shown in figure 1 for session invitation. Media session parameters are included as SDP parameters within the *INVITE* transaction, which subsequently leads to a media session establishment in separate streams.

Group communication complicates this process significantly. A newly joining member needs to contact an entire group, which requires appropriate addressing *and* transactional state management. Negotiations on




Figure 1. SIP Simple Unicast Flow Setup

media parameters grow complex as common parameter intersections may have to be evaluated for many members. While Any Source Multicast accounts for easy common addressing, Source Specific Multicast presupposes source specific subscriptions. Hence, the use of SSM requires a dedicated distribution of newly joining session member’s addresses, which otherwise remain unnoticeable in previously established SSM groups.

Group conferencing must be considered a generic term for a wide variety of meanings in different contexts [20]. There are essentially three relevant types of multiparty conferences: In *loosely coupled* conferences, no signalling relationship is maintained by conference parties, while meta information are pre-shared out of band or gradually learned by RTCP control streams. *Tightly coupled* conferences rely on a central managing instance, granting a signalling relationship to all participants. Finally, the *fully distributed* multiparty model is built upon peer-to-peer signalling relationships between conference members and allows for infrastructureless, instantaneous establishment and operation. Its scalability can be largely enhanced by the use of multicast packet distribution. Our further discussion will concentrate on this fully distributed approach and its enhancing inclusion of Source Specific Multicast communication at the signalling and the media flow level.

The SIP protocol architecture comprises proxy servers as optional components and thus draws near the general peer-to-peer model. Only recently the IETF has started work on extending the SIP framework to include fully distributed infrastructureless approaches, focussing on potentials to evolve in highly scalable overlay networks, which incorporate Distributed Hash Ta-

bles for look-up purposes [1].

On the occasion of mobile users, SIP provides some mobility support for sessions [30], which requires implementation at the application layer. Employing the regional SIP server as an application specific home agent, handoff notifications are traded via regular SIP messages to the home server (register) and the correspondent node (re-invite). As SIP mobility operates above the transport layer, it remains self-consistent and transparent to the Internet infrastructure, but inherits all underlying delays in addition to its own signaling efforts. SIP application layer handovers have been found in [13] to significantly fall behind Mobile IPv4 performance, which itself is largely outperformed by Mobile IPv6. Only MIPv6 – enhanced by correspondent accelerating protocols – does arrive at real-time compliant handover performance [24]. Consequently we will concentrate on IP-layer mobility and we will discuss extensions to SSM in the following section.

2.2. Mobile Source Specific Multicast

Multicast mobility management has to accomplish two distinct tasks, handover operations for mobile listeners and senders. While many solutions exist for roaming receivers [19], very few schemes have been detailed out for mobile multicast sources. Following a handover, multicast data reception can be fairly easily regained by a remote subscription approach in MIPv6 [11], possibly expedited by agent-based proxy schemes. In contrast, a multicast sender must not change source address while reassociating in a different network, since addresses are associated with media streams via the SIP initiation dialog. Source Specific Multicast on the IP-layer, though, requires active subscription to contributing sources, thereby relying on topologically correct addresses. Recalling the address duality problem, modified multicast routing protocols must be foreseen, as routing at the occurrence of source movement is required to transform any (S, G) state into (S', G) , while listening applications continue to receive multicast data streams admitting a persistent source address. Hence any simple mobility solution such as the remote subscription approach loses its receivers and will no longer function in our context.

With SSM an additional address problem needs consideration: A multicast listener, willing to subscribe to an (S, G) state, needs to report for the current location of the mobile source. Concurrently a general intricacy derives from the principle decoupling of multicast source and receivers: A multicast source submits data to a group of unknown receivers and thus operates without feedback channel. Address updates on

handovers of an SSM source have to proceed without means of the mobile source to inquire on properties of the delivery tree or the receivers. As the nature of multicast routing is receiver initiated, whereas source movement is only detectable at the sender side, this leads to a somewhat obstructive interplay. According to common routing procedures, a mobile multicast source on handover should trigger its receivers to re-subscribe to its new address, but remains sightless of an actual fulfillment. All of the above severely add complexity to a robust multicast mobility solution, which should converge to optimal routes and, for the sake of efficiency, should avoid data encapsulation.

Finally, Source Specific Multicast has been designed as a light-weight approach to group communication. Routing complexity remains a major deployment issue. In adding mobility management, it is desirable to preserve the principal leanness of SSM by minimizing additional signaling overheads.

Three principal approaches to SSM source mobility are presently around.

Statically Rooted Distribution Trees: The MIPv6 standard proposes bi-directional tunneling through the home agent as a minimal multicast support for mobile senders and listeners as introduced by [31]. In this approach, the mobile multicast source (MS) always uses its Home Address (HoA) for multicast operations. Since home agents remain fixed, mobility is completely hidden from multicast routing at the price of triangular paths and extensive encapsulation.

Following a shared tree approach, [18] propose to employ Rendezvous Points of PIM-SM [6] as mobility anchors. Mobile senders tunnel their data to these "Mobility-aware Rendezvous Points" (MRPs), whence in restriction to a single domain this scheme is equivalent to the bi-directional tunneling. Focusing on inter-domain mobile multicast the authors design a tunnel- or SSM-based backbone distribution of packets between MRPs.

Reconstruction of Distribution Trees: Several authors propose to construct a completely new distribution tree after the movement of a mobile source. These schemes have to rely on client notification for initiating new router state establishment. At the same time they need to preserve address transparency to the client.

To account for the latter, Thaler [28] proposes to employ binding caches and to obtain source address transparency analogous to MIPv6 unicast communication. Initial session announcements and changes of source addresses are to be distributed periodically to clients via an additional multicast control tree based at the home agent. Source-tree handovers are then

activated on listener requests.

[10] suggest handover improvements by employing anchor points within the source network, supporting a continuous data reception during client-initiated handovers. Receiver oriented tree construction in SSM thereby remains unsynchronized with source handovers and thus will lead to an unforeseeable temporal progress. The authors henceforth are leaving the source in case of its rapid movement with an unlimited number of 'historic' delivery trees to be fed simultaneously.

Tree Modification Schemes: Very little attention has been given to procedures, which modify existing distribution trees to continuously serve for data transmission of mobile sources. In the case of DVMRP routing, [3] propose an algorithm to extend the root of a given delivery tree to incorporate a new source location in ASM. To fix DVMRP forwarding states and heal reverse path forwarding (RPF) check failures, the authors rely on a complex additional signaling protocol.

O'Neill [15] suggests a scheme to overcome RPF-check failures originating from multicast source address changes, by introducing an extended routing information, which accompanies data in a Hop-by-Hop option header.

An extended routing protocol adaptive to SSM source mobility, visualized in figure 2, has been introduced by the authors in [23]. A mobile multicast source (MS) away from home will transmit *unencapsulated* data to a group, using its HoA on the application layer and its current CoA on the Internet layer, just as unicast packets are transmitted by MIPv6. In extension to unicast routing, though, the entire Internet layer, i.e. routers included, will be aware of the permanent HoA. Maintaining address pairs in router states like in binding caches will enable all nodes to simultaneously identify (HoA, G) -based group membership and (CoA, G) -based tree topology. When moving to a new point of attachment, the MS will alter its address from previous CoA (pCoA) to new CoA (nCoA) and eventually change from its previous Designated multicast Router (pDR) to a next Designated Router (nDR). Subsequent to handover it will immediately continue to deliver data along an extension of its previous source tree. Delivery is done by elongating the root of the previous tree from pDR to nDR (s. fig. 2(b)). All routers along the path, located at root elongation or previous delivery tree, thereby will learn MS's new CoA and implement appropriate forwarding states.

Routers on this extended tree will use RPF checks to discover potential short cuts. Registering nCoA as source address, those routers, which receive the state update via the topologically incorrect interface, will

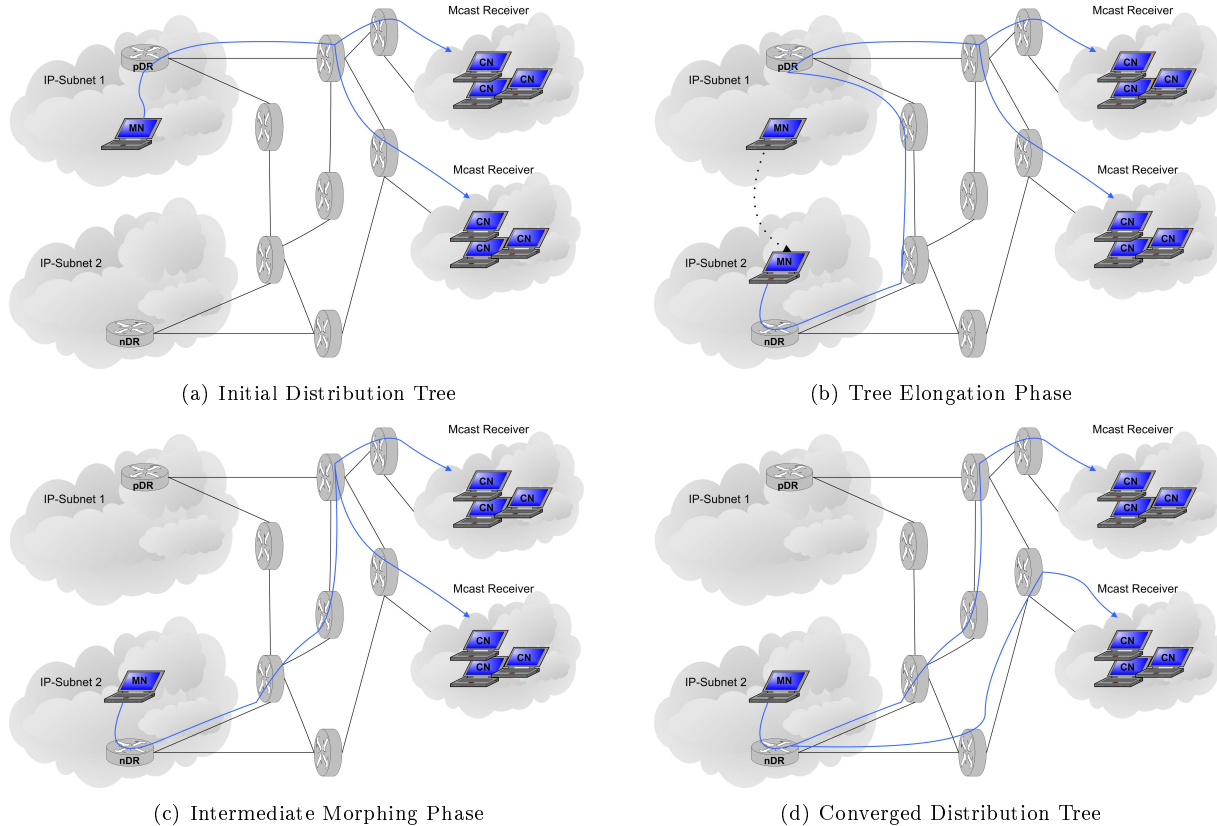


Figure 2. Tree Morphing States

submit a join in the direction of a new shortest path tree and prune the old tree membership, as soon as data arrives at the correct interface. All other routers will re-use those parts of the previous delivery tree, which coincide with the new shortest path tree. Only branches of the new shortest path tree, which have not previously been established, need to be constructed. In this way, the previous shortest path tree will be morphed into a next shortest path tree as shown in figure 2(c). This algorithm does not require data encapsulation at any stage.

Recently [14] introduced a state update mechanism for re-using major parts of established multicast trees, as well. The authors start from initially established distribution states centered at the MS’s home agent. A mobile leaving its home network will signal a multicast forwarding state update on the path to its home agent, similar to the first elongation phase of Tree Morphing [23]. Subsequently, distribution states according to the MS’s new CoA are implemented along the previous distribution tree using multicast forwarding. Multicast data then is intended to natively flow in triangular routes via the elongation and updated tree centered

at the home agent. Consequently this mechanism refrains from using shortest path trees. Unfortunately the authors do not address the problem of RPF-check failures in their paper nor do they present delay estimates based on measurements or simulations.

2.3. The Video Conferencing System

In this section we briefly introduce our reference implementation, a digital audio-visual conferencing system, realised as a server-less multipoint video conferencing software without MCU developed by the authors [17]. It has been designed in a peer-to-peer model as a lightweight Internet conferencing tool aimed at email-like friendliness of use. The system is built upon a fast H.264/MPEG-4 AVC standard conformal video codec implementation [16]. It is a Baseline profile implementation, optimized for real-time decoding and encoding by several accelerating measures like diamond shape motion search, MMX enhanced SAD motion estimation, fast mode selection and a fast subpel search strategy. There is also application-tailored fast wavelet-based video codec [4] used for higher available data rate. By controlling the coding parameters appro-

privately, the software permits scaling in bit rate from 24 to 1440 kbit/s on the fly. All streams can be transmitted by unicast as well as by multicast protocol. Audio streams are prioritized above video since audio communication is more sensitive to distortions in erroneous networks.

An application-sharing facility is included for collaboration and teleteaching. It enables participants to share or broadcast not only static documents, but also any selected dynamic PC actions like animations including mouse pointer movements. All audio/video (A/V) - streams including dynamic application sharing actions can be recorded on any site. This system is equally well suited to intranet and wireless video conferencing on a best effort basis, since the audio/video quality can be controlled to adapt the data stream to the available bandwidth.

The joined use of high bandwidth UDP traffic with TCP updates bound to real-time demands is known to suffer from distortions due to TCP traffic suppression. Application sharing in conferencing applications thus is endangered to encounter disruptions in the event of network congestion. For a service enhanced synchronous use of UDP media sessions and application sharing with reliable data transport requirements, we implemented end-to-end load balancing employing proprietary extensions to UDP, reliable (RUDP), we work on its packet identifiers to control application sharing data flows. On the occasion of a significant number (e.g., 5) of unacknowledged packets, we slow down video packet transmission to reserve required resources for real-time application updates. Audio communication remains undisturbed of load-balancing actions.

As in practice, the allocation of DNS SRV records is very rarely seen, the SIP approach to user location suffers from a limited pervasiveness. To overcome these obstacles, our system restricts call-names to email addresses and takes advantage of the globally established MX server record infrastructure by applying a name convention to session servers as proposed in [27]. This allows for a rollout of session server infrastructure in concordance with email services. In proceeding along this line, session oriented service support for nomadic users can be easily established, while Internet infrastructure remains unchanged.

3. SIP Initiated SSM Group Conferences

3.1. Multicast Capabilities of SIP

The original development of SIP has been inspired by connection oriented telephone services, whence its nature derives from a point-to-point model. Exten-

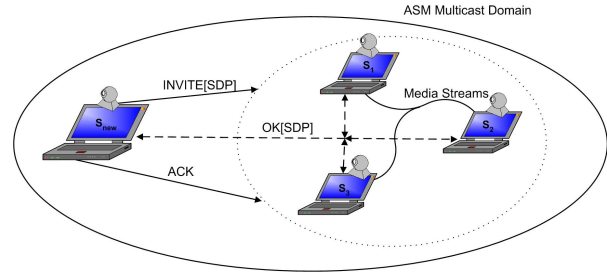


Figure 3. SIP initiating ASM – a callee negotiates its call with a previously defined multicast group

sions to include scalable group communication are not easy to achieve, as was discussed in section 2. Consequently, the basic SIP RFC only defines a minimal message exchange using IP layer multicast: A client wishing to initiate or join into a multiparty conference sends its INVITE request to a multicast group by employing the maddr attribute in the SIP VIA header. Group members subsequently indicate their presence by responding to the same group (cf. figure 3). The transactional nature of SIP dialogs is preserved in the sense that the inviting party interpretes the first arriving OK as the regular completion, while interpreting further messages as irrelevant iterates. Suitable for large, loosely coupled and mutually unknown parties, this simple scheme only operates through Any Source Multicast (ASM) and restrains management from supporting instantaneous peer-to-peer group organisation, as will be subject of the following section.

Additional work is needed to develop peer-to-peer group support within a SIP control plane. Keeping in mind the routing complexity inherent to ASM, it is desirable to rigorously restrict all signalling to Unicast or Source Specific Multicast communication.

3.2. SIP Extensions for SSM

Instantaneous establishment of a fully distributed peer-to-peer conference commonly follows an incremental setup: Some party will initiate a conference by contacting one or several peers via unicast addresses as acquired by a user location scheme like referred to in section 2. Following an initial contact, signalling will then be turned to scalable multicast group communication. Further on new parties will join the conference by either calling or being called by an existing member.

Such group conference initiation scheme is neither covered by the current status of SIP, nor is the employment of Source Specific Multicast for group signalling.

In order to enable SSM, all dialogs must carefully provision addresses of newly arriving senders to all current group members, which need to adapt source specific subscriptions appropriately.

In detail protocol operations of the suggested extensions proceed as follows. A caller, wishing to establish an SSM signalling session with a single peer, will initiate a regular INVITE request to callee’s previously inquired unicast address. Eventually, after the call setup has completed, either party will decide to transfer the established session to group communication. Heading for SSM, it will submit a re-INVITE, i.e., an INVITE carrying the previously established session identifier, announcing its desired multicast group address in the CONTACT field of the SIP header. Concomitantly, the SIP protocol stack will submit a multicast source specific JOIN to its underlying IGMP/MLD stack, thereby subscribing to the group and peer’s source address, which it had both learned from the previous SIP message exchange. Any peer will identify the multicast address in the CONTACT field, designed to specify group contacts for subsequent requests, and proceed along the new protocol semantic for SSM. As ASM multicast address announcement will distinctively appear in a separate VIA header, it will identify the presence of SSM, answer with a regular unicast reply, but will submit a multicast JOIN to the announced group and caller’s source address.

This two-step process purposefully decouples application layer session establishment and underlying multicast routing operations. Temporal progress in IP-layer multicast routing and SIP transactional timers thereby remain independent for the sake of robustly layered protocol operations. Appropriate media session descriptions for source specific multicast distribution of media streams may or may not be submitted along the re-INVITE request.

In multiparty environments, the straight forward generalization for switching a previously established unicast conference into SSM group communication is shown in figure 4. A newly arriving party will contact some member of the established session via regular unicast INVITE. The callee might decide to accept this request and forward it to its partner, thereby initiating unicast sessions among the three. Eventually the callee will decide to select multicast for the conference signalling and will submit the corresponding re-INVITE procedure.

If a new source S_{new} contacts an established SSM group conference, it will do so by inviting some member S . If S decides to accept the caller, it will redistribute its INVITE to the SSM group and acknowledge the initial call by placing the group address in the CON-

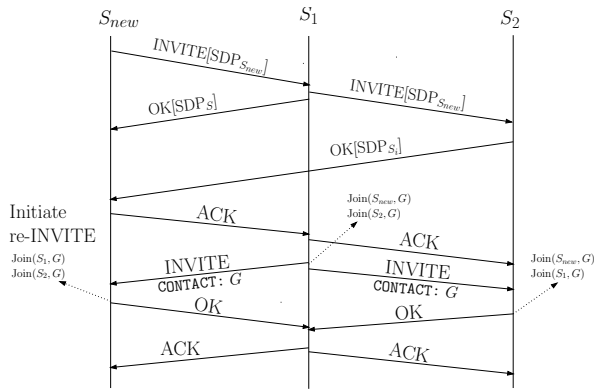


Figure 4. Switching Unicast to SSM: re-INVITE




Figure 5. SIP initiating SSM: Extending established group sessions to a newly arrived party

TACT header field. As displayed in figure 5, all group members will immediately add S_{new} to their source specific multicast filters. S_{new} subsequently will learn about all group members from (unicast) OK-messages as needed for its own multicast subscriptions. Note that call redistribution will remain a point-to-point request of S_{new} at the application layer, but a transmission of source S at the network layer and therefore compliant with previous SSM group establishment.

Proceeding along this incremental way, a callee will never be required to redistribute messages to more than one party or group. This scheme thus remains fully scalable and fairly transparent to group sizes. Multicast initiation of media sessions may be led correspondingly.

4. Evaluation of Adaptive Routing for Mobile SSM Sources

Mobility initiated handovers may in general lead to packet loss and delay. Disturbances will derive from the subnetwork and the network layer. The tree morphing multicast routing scheme described in section 2.2 will not produce any packet loss in addition to mobile IP handovers, as can be easily concluded from primary packet forwarding relying on unicast source routes. For a brief evaluation measure we will therefore concentrate on protocol overhead and convergence. Based on real-world Internet topologies we simulate the packet distribution and compare our results to an idealized handover scheme with HA-based control tree as introduced by Thaler (cf. section 2.2).

We performed a stochastic discrete event simulation based on the network simulator platform OMNeT++ 3.1 [29] and several real-world topologies of different dimensions. The selection of network data in our simulation must be considered critical, as key characteristics of multicast routing only make an impact in large networks, and as topological setup fixes a dominant part of the degrees of freedom in routing simulations. We chose the ATT core network [8] as a large (154 core nodes), densely meshed single provider example. For inter-provider data we extracted sub-samples of varying sizes from the "SCAN + Lucent" map [7]. Sample sizes, 154, 1.540 and 15.400 core nodes, vary by two orders of magnitude.

Mean values for the convergence of the routing protocol to optimal packet forwarding have been calculated for these sample networks. Convergence time is evaluated per tree in units of router hops under the assumption of homogeneous link delays. Comparison is drawn to an expedited, idealized HA-Handover scheme derived from [28]: As the work of Thaler has not been detailed out, we disregard any messaging overhead to be defined therein and assume that the mobile node subsequent to handover will immediately signal its new address to its receivers down its permanent HA-based tree. On the reception of updates, multicast listeners are then expected to immediately join the tree towards the new source location, such that optimal packet flow is reached with shortest path tree's routing completion. Following this idealized setting we calculate a lower bound for the time to optimal packet forwarding of the handover procedure proposed in [28].

Figure 6 compares the convergence times of the tree morphing protocol at a designated router distance of 5 with the corresponding results of the idealized HA scheme derived from [28] as functions of receiver multiplicities. In general, both schemes nicely reproduce


their multicast nature by showing very little dependence on receiver numbers. For TM this approves the previously assumed parallel processing of short cuts. While results for the single provider ATT network remain close, TM significantly outperforms the HA scheme for multiprovider Internet topologies of all sizes. This basically reflects its ability to re-use increasing parts of the wider branched trees, whereas the inter-tree handover approach always requires the re-creation of all routing states.

For the same reasons overheads therefore can be expected to reliably remain below costs of the reconstruction of a distribution tree. In general signaling overhead and suboptimal packet distribution remain low in case of local movement. However, they may increase with the number of temporary short cuts in use, even though the rapid algorithmic convergence as shown above will keep intermediate distribution periods short.

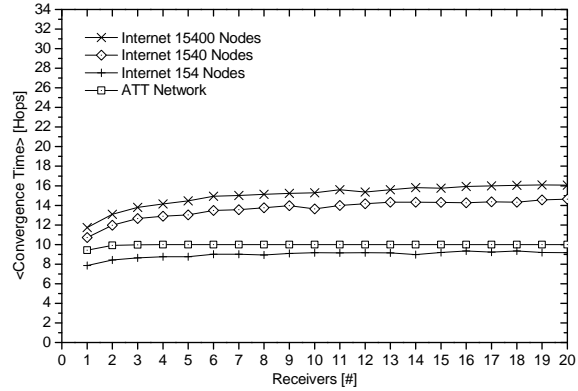
To evaluate the routing cost of the tree morphing scheme, we calculated the mean number of newly established multicast states per handover, including all temporary short cuts needed during convergence. Previously established forwarding states, which are re-used by immediate modifications from TM state updates without involvement of regular multicast state establishment, are disregarded. Results are again compared to the idealized Thaler scheme, which is cost equivalent to any tree reconstruction scheme.

Results from stochastic simulations as described above are displayed in figure 7. The tree morphing protocol significantly outperforms the HA-based scheme, which, according to its full construction of a new distribution tree, admits almost linear cost increase. While the latter results from a minimal tree extension by one new edge router per receiver¹, TM strongly demonstrates its ability to decouple from receiver network topology. Almost all newly established forwarding states remain in the vicinity of the designated routers. Temporary short cuts do not seem to have remarkable impact, as on average only very few (partial) branches are constructed as impermanent forwarding paths. For state maintenance cost as well as for protocol convergence the adaptive tree morphing thus proved its efficient source mobility management by re-using the global and changing only local parts of the distribution tree. For further evaluations and a thorough discussion we refer the reader to [25].

¹From the networking perspective we omit to sample multiple receivers per edge node, as they remain effectless for both routing schemes.




(a) Expedited HA Scheme

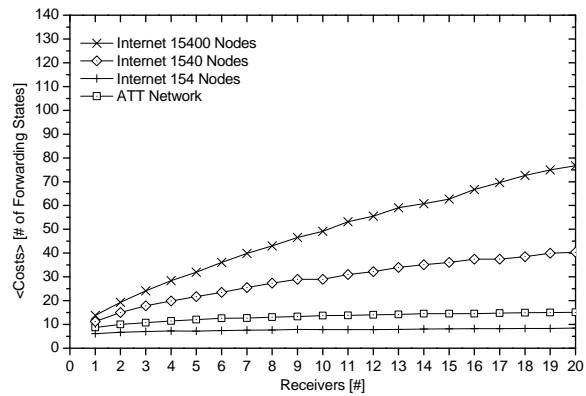


(b) Tree Morphing at pDR-to-nDR distance 5

Figure 6. Mean Convergence Time to Optimal Packet Forwarding: Comparison of an Idealized HA-based Handover Scheme and TM as Functions of Receiver Multiplicity



(a) Expedited HA Scheme



(b) Tree Morphing at pDR-to-nDR distance 5

Figure 7. Routing Cost per Handover: The Number of Newly Established Forwarding States for the Idealized HA-based Handover Scheme and TM as Functions of Receiver Multiplicity

5. Conclusions & Outlook

In this paper we addressed the essential issues of scalable multimedia group conferences, negotiated by the Session Initiation Protocol and relying on Source Specific Multicast distribution among equal peers. We introduced, discussed and analyzed signalling extensions of SIP and routing methods to adapt SSM to source mobility. Protocol operations could be shown to proceed in a fully scalable, efficient way, preserving utmost flexibility at communicating parties.

Within this ongoing project work, we will head further for optimizations and empirical testing of our solution in large-scale, realistically perturbed networks. We intend to achieve this goal by deploying a video conferencing test client onto the planetlab platform.

References

- [1] S. Baset, H. Schulzrinne, E. Shim, and K. Dhara. Requirements for SIP-based Peer-to-Peer Internet Telephony. Internet Draft - work in progress 00, IETF, March 2006.
- [2] S. Bhattacharyya. An Overview of Source-Specific Multicast (SSM). RFC 3569, IETF, July 2003.
- [3] R.-S. Chang and Y.-S. Yen. A Multicast Routing Protocol with Dynamic Tree Adjustment for Mobile IPv6. *Journ. Information Science and Engineering*, 20:1109–1124, 2004.
- [4] H. L. Cycon, M. Palkow, T. C. Schmidt, M. Wählich, and D. Marpe. A fast wavelet-based video codec and its application in an IP version 6-ready serverless videoconferencing system. *International Journal*

- of Wavelets, Multiresolution and Information Processing, 2(2):165–171, June 2004.
- [5] S. E. Deering. Host Extensions for IP Multicasting. RFC 1112, IETF, Aug. 1989.
- [6] B. Fenner, M. Handley, H. Holbrook, and I. Kouvelas. Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised). RFC 4601, IETF, August 2006.
- [7] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. In *Proceedings IEEE INFOCOM 2000*, volume 3, pages 1371–1380, Tel Aviv, Israel, March 2000. IEEE Computer Society.
- [8] O. Heckmann, M. Piringler, J. Schmitt, and R. Steinmetz. On Realistic Network Topologies for Simulation. In *MoMeTools '03: Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research*, pages 28–32, New York, NY, USA, August 2003. ACM Press.
- [9] H. Holbrook and B. Cain. Source-Specific Multicast for IP. RFC 4607, IETF, August 2006.
- [10] C. Jelger and T. Noel. Supporting Mobile SSM sources for IPv6 (MSSMSv6). Internet Draft – work in progress (expired) 00, individual, January 2002.
- [11] D. B. Johnson, C. Perkins, and J. Arkko. Mobility Support in IPv6. RFC 3775, IETF, June 2004.
- [12] G. Kurup and Y. A. Sekercioglu. Source Specific Multicast (SSM) for MIPv6: A Survey of Current State of Standardisation and Research. In *Proceedings of Australian Telecommunications, Networks and Applications Conference (ATNAC 2003)*, Melbourne, December 2003.
- [13] T. T. Kwon, M. Gerla, S. Das, and S. Das. Mobility Management for VoIP Service: Mobile IP vs. SIP. *IEEE Wireless Communications*, 9(5):66–75, October 2002.
- [14] H. Lee, S. Han, and J. Hong. Efficient Mechanism for Source Mobility in Source Specific Multicast. In K. Kawahara and I. Chong, editors, *Proceedings of ICOIN2006*, volume 3961 of *LNCS*, Berlin, Heidelberg, 2006. Springer-Verlag. (in press).
- [15] A. O'Neill. Mobility Management and IP Multicast. Internet Draft – work in progress (expired) 01, IETF, July 2002.
- [16] J. Ostermann, J. Bormans, P. List, D. Marpe, N. Narroschke, F. Pereira, T. Stockhammer, and T. Wedi. Video Coding with H.264/AVC: Tools, Performance and Complexity. *IEEE Circuits and Systems Magazine*, 4(1):7–28, April 2004.
- [17] M. Palkow. The daViKo homepage, 2006. <http://www.daviko.com>.
- [18] I. Romdhani, H. Bettahar, and A. Bouabdallah. Transparent handover for mobile multicast sources. In P. Lorenz and P. Dini, editors, *Proceedings of the IEEE ICN'06*. IEEE Press, April 2006.
- [19] I. Romdhani, M. Kellil, H.-Y. Lach, A. Bouabdallah, and H. Bettahar. IP Mobile Multicast: Challenges and Solutions. *IEEE Comm. Surveys & Tutorials*, 6(1):18–41, 2004.
- [20] J. Rosenberg. A Framework for Conferencing with the Session Initiation Protocol (SIP). RFC 4353, IETF, February 2006.
- [21] J. Rosenberg and H. Schulzrinne. Session Initiation Protocol (SIP): Locating SIP Servers. RFC 3263, IETF, June 2002.
- [22] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. SIP: Session Initiation Protocol. RFC 3261, IETF, June 2002.
- [23] T. C. Schmidt and M. Wählisch. Extending SSM to MIPv6 — Problems, Solutions and Improvements. *Computational Methods in Science and Technology*, 11(2):147–152, November 2005. Selected Papers from TERENA Networking Conference, Poznań, May 2005.
- [24] T. C. Schmidt and M. Wählisch. Predictive versus Reactive – Analysis of Handover Performance and Its Implications on IPv6 and Multicast Mobility. *Telecommunication Systems*, 30(1–3):123–142, November 2005.
- [25] T. C. Schmidt and M. Wählisch. Morphing Distribution Trees – On the Evolution of Multicast States under Mobility and an Adaptive Routing Scheme for Mobile SSM Sources. *Telecommunication Systems*, October 2006. in print, appeared online.
- [26] T. C. Schmidt and M. Wählisch. Multicast Mobility in MIPv6: Problem Statement. IRTF Internet Draft – work in progress 01, MobOpts, October 2006.
- [27] T. C. Schmidt, M. Wählisch, H. L. Cycon, and M. Palkow. Global serverless videoconferencing over IP. *Future Generation Computer Systems*, 19(2):219–227, February 2003.
- [28] D. Thaler. Supporting Mobile SSM Sources for IPv6. Proceedings of ietf meeting, individual, December 2001.
- [29] A. Varga et al. The OMNeT++ discrete event simulation system. <http://www.omnetpp.org>, 2005.
- [30] E. Wendlund and H. Schulzrinne. Mobility Support using SIP. In *Proc. of the 2nd ACM Intern. Workshop on Wireless Mobile Multimedia*, pages 76–82, New York, NY, USA, 1999. ACM Press.
- [31] G. Xylomenos and G. C. Polyzos. IP Multicast for Mobile Hosts. *IEEE Comm. Mag.*, 35(1):54–58, 1997.