

Morphing distribution trees—On the evolution of multicast states under mobility and an adaptive routing scheme for mobile SSM sources

Thomas C. Schmidt · Matthias Wählisch

Published online: 19 October 2006
© Springer Science + Business Media, LLC 2006

Abstract Source Specific Multicast (SSM) promises a wider dissemination of group distribution services than Any Source Multicast, as it relies on simpler routing strategies with reduced demands on the infrastructure. However, SSM is designed for a priori known and changeless addresses of multicast sources and thus withstands any easy extension to mobility. Up until now only few approaches arose from the Internet research community, leaving SSM source mobility as a major open problem. The purpose of this paper is twofold. At first we analyze characteristic properties of multicast shortest path trees evolving under source mobility. Analytically and by stochastic simulations we derive measures on the complexity of SSM routing under source mobility. At second we introduce a straightforward extension to multicast routing for transforming (morphing) source specific delivery trees into optimal trees rooted at a relocated source. All packet forwarding is done free of tunneling. Multicast service disruption and signaling overhead for the algorithms remain close to minimal. Further on we evaluate the proposed scheme using both, analytical estimates and stochastic simulations based on a variety of real-world Internet topology data. Detailed comparisons are drawn to bi-directional tunneling, as well as to proposals on concurrent distribution trees.

Keywords Multicast routing protocols · Mobile source specific multicast · Shortest path trees · IPv6 mobility management · Internet measurement

1 Introduction

Mobile network technologies today must be seen as one of the major driving forces for multimedia data transmission. Cellular phones and portable paddles are expected to carry individual Internet

T. C. Schmidt (✉)
HAW Hamburg Department Informatik, Berliner Tor 7, D-20099 Hamburg, Germany
e-mail: t.schmidt@ieee.org

T. C. Schmidt · M. Wählisch
FHTW Berlin, Hochschulrechenzentrum, Treskowallee 8, D-10318 Berlin, Germany
e-mail: waehlich@ieee.org

addresses soon. These will be available from IPv6 address space, as is seamless mobility support from the recently released MIPv6 by Johnson et al. [14]. It is the vision that “IPv6 will be pervasive and prevalent across all digital device communications and augurs well for mobility and wireless access on the Internet” Kurup and Sekercioglu [17]. IP multicasting will be of particular importance to mobile environments, where users commonly share frequency bands of limited capacities.

Intricate multicast routing procedures, though, are not easily extensible to comply with mobility requirements. Any client subscribed to a group while in motion, requires delivery branches to pursue its new location; any mobile source requests the entire delivery tree to adapt to its changing positions. 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 Romdhani et al. [26] and Schmidt and Wählisch [29].

Source Specific Multicast (SSM) Bhattacharyya [3] and Holbrook and Cain [11], still in its design process, is considered a promising improvement of group distribution techniques. In contrast to Any Source Multicast (ASM) [7], 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 discovered by out of band channels such as the session directory (SDR) or a Web page. As a consequence, routing simplifies significantly, but invalidates with source addresses changing under mobility. Up until now SSM source mobility remains as an unsolved problem.

Source mobility presents a severe problem for multicast packet distribution. Even though multicast routing itself supports dynamic reconfiguration, as members may join and leave ongoing group communication over time, multicast group membership management and routing procedures are intricate and too slow to function smoothly for mobile users. In addition multicast imposes a special focus on source addresses. Applications commonly identify contributing streams through source addresses, which must not change during sessions, and routing paths in most protocols are chosen from destination to source.

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. Mobile source specific group membership is identified via the logical ID of the sender, typically given by its Home Address (HoA); shortest path delivery trees are erected according to topological information as encoded in the current Care-of Address (CoA) of the Mobile Source. As a mobile source is identified by a (HoA, CoA) pair, any SSM solution that masks the mobility of a multicast source consequently requires intermediate routers to construct a multicast delivery tree with reference to this pair.

In the present paper we start from this observation and analyze in detail the problem complexity of mobile source specific multicast. Special focus is donated to topological correlations of shortest path distribution trees initiated from source movement. Real-world Internet topology data and edge scan measurements are used to quantify characteristic cost measures for the evolution of multicast states under routing schemes adaptive to mobility requirements.

We further on present and thoroughly evaluate an approach to SSM routing adaptive to source mobility, which was primarily introduced in Schmidt and Wählisch [28]. Operating on extended router states, our tree morphing algorithm first extends a given multicast distribution tree to include any new source location. It then transforms the extended tree to a new shortest path tree, thereby re-using all possible previously established branches. This scheme operates fast, without any tunneling and does not cause additional packet loss.

This paper is organized as follows: In Section 2 we review the basic problems of SSM source mobility and related work. The general characteristics of multicast shortest path trees evolving under root mobility are examined in Section 3. Section 4 introduces our new approach and defines the underlying routing algorithms. In Section 5 we present the evaluation, analytical estimates and simulations of the proposed scheme. Finally Section 6 is dedicated to conclusions and an outlook.

2 The mobile multicast source problem and related work

2.1 Problem statement

Any next generation Internet support for multicast source mobility management is required to operate transparently w.r.t. the socket layer. Specific protocol operations or extensions are thus bound to a multicast aware MIPv6 stack and the Internet routing layer. 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 in MIPv6 [14] 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. Bearing in mind characteristic applications, i.e. multimedia distribution, handover delays are to be considered critical. The distance of subsequent points of attachment of the mobile, i.e., the weight of the path from previous to next designated router or 'step size', may serve as an appropriate measure of complexity, cf. Schmidt and Wählisch [30].

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.

2.2 Related work

Two 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 Xylomenos and Polyzos [39]. 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. Though robust and simple, it is well known that bi-directional tunneling may lead to overheads and delays from triangular routing unsuitable for real-time applications.

Following a shared tree approach, Romdhani et al. [25] propose to employ Rendezvous Points of PIM-SM [8] 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 interdomain mobile multicast the authors design a tunnel- or SSM-based backbone distribution of packets between MRPs. These MRPs operate on extended multicast routing tables, which simultaneously hold HoA and CoA. This solution accounts for the ASM interdomain source activation problem, cf. Romdhani et al. [26].

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 [35] 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. In this proposed protocol, data reception subsequent to handovers will be interrupted for the period of address announcement and tree reconstruction. It remains far too slow to be considered seamless. Overheads caused by the construction and maintenance of several trees are significant.

Jelger and Noel [13] suggest handover improvements by employing anchor points within the source network, supporting a continuous data reception during client-initiated handovers. Similar proxy schemes are known from improved unicast mobility by Soliman et al. [33] or ASM by Schmidt and Wählisch [31]. Receiver oriented tree construction in SSM 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.

Somewhat similar concepts have been presented in the MNet approach of Silva et al. [32]. Multicast data is received by and re-distributed through stationary multicast servers. Listeners may be triggered to inter-server handovers by ASM service announcement via the Session Announcement Protocol (SAP).

An additional group of sparse work is of relevance to this paper.

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, Chang and Yen [5] 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 [23] 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.

Recently Lee et al. [18] introduced a state update mechanism for re-using major parts of established multicast trees. 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

[28]. 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.

In the following section we will analyze general, mobility related changes required for multicast topologies. Based on these insights we will specify our approach to the SSM mobile source problem, which falls in this last category of tree modifications.

3 On the evolution of multicast shortest path trees under mobility

3.1 Multicast tree characteristics

Recently multicast distribution trees have been studied well under the focus of network efficiency. Grounded on empirical observations Chuang and Sirbu [6] proposed a scaling power-law for the total number $L_M(m)$ of links in a multicast shortest path tree with m receivers of the form

$$L_M(m) \approx \langle L_U \rangle m^k,$$

where $\langle L_U \rangle$ represents the average number of unicast hops taken by a message between uniformly chosen nodes in the corresponding network. The authors consistently identified the scale factor to attain the independent constant $k = 0.8$. The validity of such universal, heavy-tailed distribution suggests that multicast shortest path trees are of self-similar nature with many nodes of small, but few of higher degrees. Trees consequently would be shaped rather tall than wide.

Subsequent empirical and analytical work of Phillips et al. [24], Van Mieghem et al. [37], Chalmers and Almeroth [4], Adjih et al. [1] and Janic and Van Mieghem [12] debated the applicability of the Chuang and Sirbu scaling law. Van Mieghem et al. [37] proved that the proposed power law cannot hold for an increasing Internet or very large multicast groups, but is indeed applicable for moderate receiver numbers m and a current Internet size of about $N = 10^5$ routers. The authors derive an effective approximation of the scale factor $k = k(N) = \frac{\text{var}[L_U(N)]}{\langle L_U(N) \rangle}$, which slowly increases to 1 with growing Internet size N . Investigating on self-similarity Janic and Van Mieghem [12] semi-empirically substantiated that multicast shortest path trees in the Internet can be modeled with reasonable accuracy by uniform recursive trees (URT) [36], provided m remains small compared to N .

The mobility perspective on shortest path trees focuses on their alteration, i.e., the degree of topological changes induced by source movement, as this may serve as an outer measure for routing complexity. Source specific multicast trees subsequently generated from mobility hand-over steps are not independent, but highly correlated. They most likely branch to the identical receivers and are rooted a 'step-size' apart. It is intuitively tempting to expect a high coincidence in routers and their forwarding states as previous and next distribution tree grow towards receivers. The scaling power-law encourages this vision and gives rise to the following heuristic argument.

A previous and next distribution tree serving the same set of receivers merge at one or several intersection points. By the self-similar nature, the subtrees (of previous *and* next distribution tree), rooted at any such intersection point, exhibit again the scaling law behavior, are tall-shaped with nodes of mainly low degree and thus likely to coincide. If the hop distance between the

roots of previous and next tree remains small, the intersection points may as well be expected to occur in the vicinity of the roots, whence the probability for a dominant coincidence of nodes between subsequent trees should be high.

3.2 Analytical estimates on topological tree changes

To place the heuristic argument outlined above on more rigorous grounds, we comply with the model of Van Mieghem [36]: A multicast shortest path distribution tree (SPT) is generated by uniformly selecting m receivers out of the $m < N$ nodes of the underlying topology, a connected random graph with independent identically and exponentially distributed link weights with mean 1. Asymptotically, the multicast tree then is a uniform recursive tree. From the perspective of source mobility we want to derive measures of those parts of the SPT changing, when the root of the tree is moved, but the leaves remain adjusted.

Consider a pair \mathcal{P} and \mathcal{Q} of such trees with roots p and q and intersection points $\{X_i\}_{i=1}^l$. Furthermore consider the residue graph

$$\mathcal{G} = \mathcal{Q} \setminus \mathcal{P} \cup \{X_i\}_i.$$

\mathcal{G} is the tree rooted at q with leaf nodes $\{X_i\}_i$ and by the recursive property a URT itself (cf. Theorem 16.2.1 of Van Mieghem [36]). To inquire on the dimension of the URT \mathcal{G} , consider the corresponding Markov discovery process within the underlying random graph topology. Starting from root q this process will discover nodes in random direction according to their proximity (here measured in time for link traversal) to previously discovered nodes. This process will terminate, when tree \mathcal{P} is met and all $\{X_i\}_i$ are discovered.

Let $d = \text{dist}(p, q)$ be the link distance of the two tree roots. Then d will estimate the radius of the discovery process and—since link weights are uniformly of mean 1—the number of discovered nodes will roughly be approximated by $\tilde{N} = \frac{\pi}{2} d^2$. Taking the number of nodes N from \mathcal{G} to $N \approx \tilde{N}$ and using the calculation (16.15) of Van Mieghem [36] for the average hop count in a URT, we arrive at the following

Observation 1. Let \mathcal{G} be the residue graph and d defined as above, the hop count $H_{\mathcal{G}}$ between two uniformly chosen nodes in \mathcal{G} , then for the average hop count

$$\langle H_{\mathcal{G}} \rangle \approx \frac{\frac{\pi}{2} d^2}{\frac{\pi}{2} d^2 - 1} \Psi\left(\frac{\pi}{2} d^2\right) + \gamma - 1, \quad (1)$$

where Ψ is the digamma function and the Euler constant is $\gamma = 0.57721\dots$

Returning to the mobile multicast source problem, this observation yields an estimate on the average number of hops, where routing needs to adjust to guide a packet until it reaches a previously established distribution tree. The complexity parameter here is given by the distance of the previous designated router (pDR) to the next designated router (nDR).

Figure 1 compares our theoretical estimate with simulations from real-world Internet topologies, as are described in detail in the next section. Closest and furthest intersection points have been calculated for distribution tree pairs rooted predefined distances apart, representing a shortest and longest path when crossing the residue tree. Theoretical mean values nicely arrange between extreme values, indicating a rather slow increase of hop counts with growing designated router distance.




Fig. 1 Average hop counts of the first and last intersection nodes within the residue graph for different Internet topologies as a function of tree root distances compared to the theoretic mean values

Equation (1) can be generalized to calculate the average number of links in the full residue tree \mathcal{G} (cf. Section 17.2.1 of [36]), leading to

Corollary 1. For the residue tree \mathcal{G} , d and $\{X_i\}_{i=1}^l$ its l leaf nodes defined as above, the hop count $H_{\mathcal{G}}(l)$ to reach the l leaf nodes attains the average

$$\langle H_{\mathcal{G}}(l) \rangle \approx \frac{\pi}{2} l d^2 \frac{\Psi\left(\frac{\pi}{2} d^2\right) - \Psi(l)}{\frac{\pi}{2} d^2 - l}. \tag{2}$$

3.3 Simulation studies on mobile multicast tree evolution

In this section we proceed our analysis by simulating multicast source mobility within real-world network topologies. The effect of source movement on the stability of shortest path trees is directly addressed by constructing and comparing multicast distribution trees. We performed a stochastic discrete event simulation based on the network simulator platform OMNeT++ 3.1 Varga et al. [38] and several realistic topologies of different dimensions. The selection of network data in our simulation must be considered critical, as key characteristics of multicast shortest path trees 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 provided by Heckmann et al. [10] as a large (154 nodes), densely meshed single provider example. For inter-provider data we extracted sub-samples of varying sizes from the SCAN project [27], cf. Govindan and Tangmunarunkit [9], the result of two extensive Internet mapping projects containing 284.805 network nodes connected by 449.246 links. Sub-sampling has been performed with the help of the network manipulator *nem*, employing the generation method “Map Sampling” Magoni [19] and Magoni and Pansiot [20]. Sample sizes, 154, 1.540 and 15.400 nodes, vary by two orders of magnitude. The Boston Generator BRITE of Medina et al. [22] has further been used for topology generation (ATT network) and format

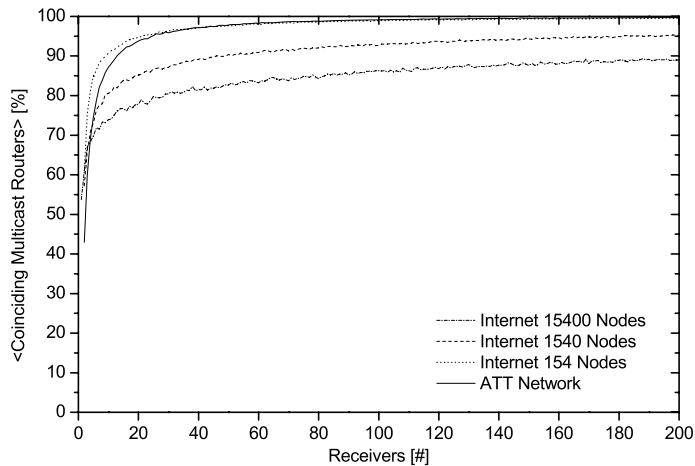


Fig. 2 Relative router coincidence between subsequent multicast distribution trees for mobile senders at pDR -to- nDR distance 5

transformation. These network samples and topological setups will be used for all further simulations throughout this paper.

In detail we uniformly sample receivers as attached to, and designated multicast routers from the edge nodes of the given topology data sets. Edge nodes are identified as routers of degree one and represent transition points to ‘customer’ networks within the Internet core systems. Previous designated router (pDR) and next designated router (nDR) are chosen at varying predefined distances, giving rise to a pair of shortest path trees for the identical set of receivers. Trees are then analyzed regarding their intersections and coinciding routers.

Results for the relative change of distribution trees as a function of receiver multiplicity for a medium step size of 5 are shown in Fig. 2. It is noteworthy that even in large networks and for moderate receiver numbers more than 80% of multicast tree routers remain fixed under a mobility step. Figure 3 visualizes the same relative router coincidence as a function of $pDR - nDR$ distance for a fixed number of 40 receivers. For small mobility step sizes the changes imposed on distribution trees remain negligible and almost independent of the underlying network topology. Only for larger network samples the tree correlation eventually degrades with increasing distances.

These empirical comparisons of multicast shortest path trees generated subsequently from source handovers strongly reproduce the above stated expectations, derived from arguments on self-similar multicast trees. All tree autocorrelation functions in Fig. 3 attain surprisingly low decay, indicating high probabilities for wide ranging agreements of the trees. Results thereby seem to indicate a weak dependence on solely the size of the underlying network. Network sizes here are a measure of the topological dimension multicast group members are spread across. In reconsidering Fig. 2 it can be concluded that tree correlations are low only in the event of very few widespread receivers. Whenever multicast listeners populate the topological region of distribution, multicast forwarding states will largely coincide and the inherent complexity of mobile multicast routing will remain limited.

3.4 First estimates on geographic proximity distributions at Internet edges

Throughout this analysis it has become evident that the proximity of subsequent points of attachment serves as a distinct measure on the complexity of mobility management. This obser-

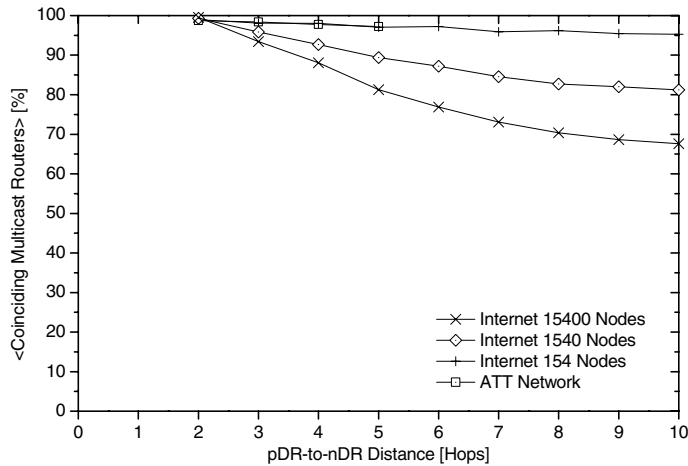


Fig. 3 Relative router coincidence between subsequent mobile multicast distribution trees serving 40 receivers as a function of the ‘mobility step size’

vation likewise holds for $pDR - nDR$ distances in the event of multicast source mobility, as for $pAR - nAR$ access router distances, when mobility handover performance is considered, cf. Schmidt and Wählisch [30]. It is therefore of particular concern to inquire on distance distributions a user could face when moving within the real-world.

Real-world mobility, at first, occurs within geographic vicinity. With the exception of satellite links, a mobile performs handovers to geographically neighboring access points. While intra-network/provider access router distances obviously range between zero and very few hops, edge distance distributions at inter-provider traversals are rather unpredictable. In an ongoing project we thus started to explore the correlation of geographic and network proximity by evaluating regional delay distributions from scanned inter-provider Internet data. We proceeded as follows. Initially geographic clusters of IP-ranges were identified, covering a limited region such as the German cities of Berlin or Hamburg. Geographic coordinates were taken from MaxMind LLC [21], which delivered the most reliable data in our spot tests. After operational hosts from each IP-range had been identified, routes were traced using the *traceroute* utility [34]. Then for each pair of paths, the coinciding hop closest to destinations was identified as transit point, interconnecting the edge nodes. Thus, under the assumption of symmetric routing at Internet border areas, an upper bound for edge node distances can be derived from a single source scan of the accessible parts of a geographically clustered Internet. Optimal distances can be approximated by employing appropriately varying sources of the scans.

First preliminary results for the European cities Berlin and Hamburg are shown in Fig. 4, scanned from single sources located within the towns. For each region 200 out of approximately 5.000 networks have been selected at random, giving rise to 19.900 distance pairs, out of which about one third admitted hosts accessible by UDP or ICMP. Empirical distributions displayed for both cities gain significant weight to small distances. Events at distance zero result from hosts out of identical networks, which occasionally are split by MaxMind into artificially separated IP ranges. Hamburg, attaining $mean = 8.5$, $\sigma = 3.6$ and Berlin, $mean = 9.6$, $\sigma = 3.2$ thereby indicate a regional Internet topology, where points of mobile attachment are likely to be less than

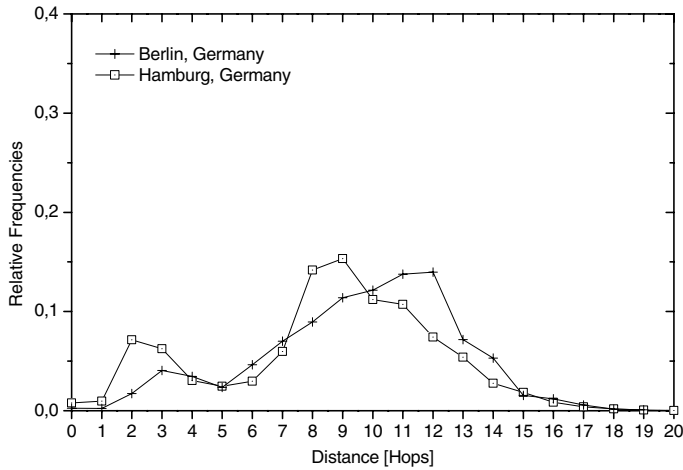


Fig. 4 Internet measurement of upper bounds for inter-provider edge distances in Berlin and Hamburg

10 router hops apart. For further, more detailed discussions and results we refer to a forthcoming publication.

4 Tree morphing: An algorithm to source mobility

4.1 General idea

In the present section we will give a first overview of the new concept of multicast routing, adaptive to source mobility. 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. 5). 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 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 Fig. 6. This algorithm does not require data encapsulation at any stage.

Fig. 5 Elongation of the tree root

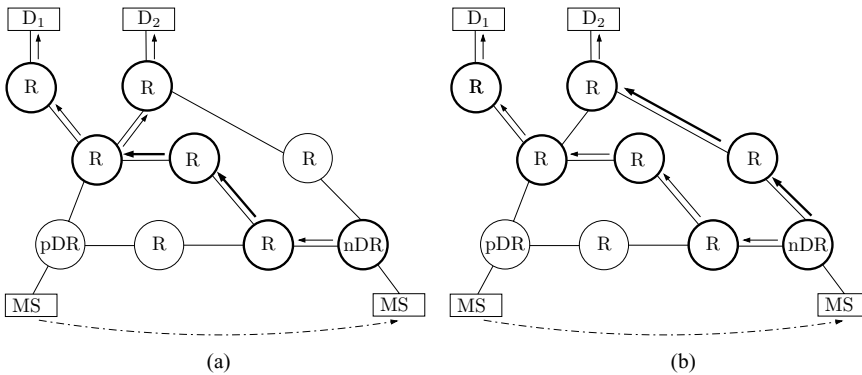



Fig. 6 Morphing states: routers on tree intersections will subsequently join towards nDR and prune towards pDR. D_i indicate receiver domains not affected by the tree morphing protocol

4.2 Routing requirements and protocol extensions

The tree morphing algorithm is not built upon a specific multicast routing protocol, but will require the following functional mechanisms, compliant with current protocols such as PIM-SM [8]:

- Outgoing router interfaces need to maintain (S, G) states to denote their participation in the distribution tree. These states will be extended to include the Home Address identifier.
- Routers need the ability to explicitly *join* an (S, G) state.
- Routers need the ability to explicitly *prune* an (S, G) state. Alternatively, but with lower efficiency, routing states may time out.
- Finally, the computation of standard *Reverse Path Forwarding (RPF)* check is used.

As a first principal extension, all router states describing delivery trees for mobile sources need extension to include both, the transient CoA and the permanent HoA. They will be further denoted by (CoA, HoA, G) . These augmented states account for the address duality problem discussed in Section 1 and will serve for identification of states during forwarding and updates.

Further mobility extensions to existing SSM routing protocols required by our algorithm are quite limited. Based on the standard functions mentioned above, protocols need to interpret a Hop-by-Hop multicast mobility option header for signaling. Routers then are required to process

two algorithmic extensions described in Section 4.4, the STATE INJECTION ALGORITHM and the EXTENDED FORWARDING ALGORITHM.

The details of signaling, MS and routing operations under mobility will be described in the following sections.

4.3 Mobile source handover initiation and signaling

We consider a multicast sender operating the mobility protocol MIPv6 [14], or accelerating schemes s.a. Fast MIPv6 by Koodli [16] or Hierarchical MIPv6 by Soliman et al. [33], see Schmidt and Wählisch [30] for multicast details. Prior to handover the MS is assumed to submit data to an intact source specific delivery tree, which does not need to be optimal. Data packets carry MS's current CoA as source address in concordance with (CoA, HoA, G) states of the source tree. A mobility destination option header is included with data to signal the HoA to receiver applications. A multicast binding cache at the receiver site preserves the (CoA, HoA) association.

The MS eventually may perform an instantaneous handover and fulfill MIPv6 reconfiguration. As soon as reassociated, the MS may immediately return to transmit unencapsulated data to the multicast group, using its topological correct nCoA as source address. To inform the routing infrastructure about its new location, it adds a state update message in a Hop-by-Hop option header to the first data packet(s) and uses source routing through the previous designated router pDR. This state update message will contain previous $(pCoA, HoA, G)$ and next $(nCoA, HoA, G)$ routing states, a sequence identifier and may include a security credential.

Following this update signaling, the MS continues to deliver multicast data as done prior to handover. It may proceed to a subsequent handover, as well.

4.4 Router operations

A multicast routing infrastructure for SSM provides the capability to construct and prune source specific delivery trees for stationary nodes. Our tree morphing algorithm will extend underlying routing protocols to include tree adaptability to mobile sources. We assume the abstract functions described in 4.2 to be present in the routing protocol and describe in detail the additional operations needed under source mobility now.

The tree morphing algorithm will proceed in three phases.

Phase 1—Tree elongation. As a Mobile Source moves, its designated multicast router may change and the root of the previous distribution tree may invalidate. To reconnect the established tree to a newly located root, a path from the nDR to the pDR is added according to unicast routing (s. Fig. 5). This tree elongation is initiated by the MS's state update message (s. Section 4.3), which is received by intermediate routers through the Hop-by-Hop option header on a unicast source route from nDR to pDR. On the reception of the update, any router will implement the new $(nCoA, HoA, G)$ state on its (unicast) forwarding interface. On arrival of the state update packet, pDR will—according to the source route—forward it to multicast group G, the group of the delivery tree rooted at pDR. At that instance, a multicast branch of $(nCoA, HoA, G)$ states has been established from nDR to pDR.

*Phase 2—Multicast state injection and forwarding.*¹ Once the state update packet has arrived at pDR, the previous root of the delivery tree, its Hop-by-Hop option header is processed along the

¹ Further on we will denote “some state with group address G and home address HoA ” by (\cdot, HoA, G) , whereas $(*, HoA, G)$ stands for all such states.

previous tree. On each hop the new $(nCoA, HoA, G)$ state is implemented on the forwarding interfaces of the previous $(pCoA, HoA, G)$ —state tree. Previous states are kept only if the update packet was received on a topological incorrect interface. In detail this algorithm runs as follows:

STATE INJECTION ALGORITHM

```

  ▷ Upon receiving an  $(nCoA, HoA, G)$ 
  ▷ state update for multicast group  $G$ 
1  for all  $(\cdot, HoA, G)$  Forwarding Interfaces
2    do if  $(RPF-CHECK(nCoA) = TRUE)$ 
3      then REPLACE all  $(\cdot, HoA, G) - states$ 
          by  $(nCoA, HoA, G)$ 
4      else ADD  $(nCoA, HoA, G) - state$ 
5          INIT TREE_OPTIMIZATION

```

After the update has been processed, the packet is passed along the newly implemented forwarding states. At this stage, the delivery tree does not need to be optimal and packets may fail at standard RPF check. To prevent discarding, incoming packets need to be accepted from any interface, which is a topological member of the current or a previous distribution tree of (\cdot, HoA, G) state. Therefore an extended forwarding, which accounts for all source address states (\cdot, HoA, G) , has to be applied until local tree optimization has completed. Packets thereby will be forwarded along an (CoA, HoA, G) tree, provided they arrived at the topologically correct interface for this CoA . A tree will be locally optimal, as soon as packets arrive at the topological correct interface. The details of this extended forwarding algorithm read:

EXTENDED FORWARDING ALGORITHM

```

  ▷ Upon arrival of packet with source address  $nCoA$  and
  ▷ in the presence of multiple  $(*CoA, HoA, G) - states$ 
1  for each  $(\cdot CoA, HoA, G)$  Forwarding Interfaces
2    do if  $(RPF-CHECK(nCoA) = TRUE)$ 
3      then FORWARD_PACKET_ON_INTERFACE
4          REMOVE  $(*, HoA, G) - states$ 
          except  $(nCoA, HoA, G)$ 
5      else if  $(RPF-CHECK(\cdot CoA) = TRUE)$ 
6          then FORWARD_PACKET_ON_INTERFACE
7          else DISCARD_PACKET

```

In applying this forwarding algorithm, the delivery tree thus will not only transport intermediate detouring packets, but will continuously degenerate branches dispensable due to optimization incidences. As soon as (\cdot, HoA, G) forwarding states have reduced to singular entries, the router operations continue as defined by its standard multicast routing protocol without mobility extension.

Finally, state update packets will arrive at the receivers of the (\cdot, HoA, G) SSM group. The mobile IPv6 stack capable of multicast tree morphing will interpret the Hop-by-Hop option header as a binding update and alter its multicast binding cache entry. Thereafter the standard destination option header is processed and data is transparently passed as (HoA, G) to the transport layer.

Phase 3—Tree optimization. As a result of source movement with successive tree elongation, but also from any intermediate morphing state, the delivery tree may cease to be optimal. Any router will observe suboptimal routes from packets arriving at a topological incorrect interface (w.r.t.

packet’s source address). As part of the algorithm it will then dynamically attempt to join to an optimal shortest path tree. When receiving a multicast packet for group (\cdot, HoA, G) with source address $nCoA$ at the wrong interface, a router will immediately submit a join to $(nCoA, G)$. The underlying SSM routing protocol will initiate the construction of a shortest path source specific branch. The router will learn about its completion by $(nCoA, HoA, G)$ traffic arriving at the correct interface and will then prune (\cdot, HoA, G) on all incoming interfaces corresponding to previous CoA addresses.

Once an intermediate router learned about suboptimal routes, this algorithm will perform optimization as rapid as possible. The scheme is self-healing and robust, but will construct any possible short cut, even though not part of the final shortest path tree.

4.5 Sample operations

To illustrate the mode of tree morphing operations we now discuss two examples in detail. The first scenario consists of a typical ‘local movement’ and is visualized in Fig. 7. The MS moves from pDR to nDR , which are directly connected and share the forwarding Router R_1 towards the receiver domains. For simplicity reason it is assumed that time can be measured in (homogeneous) router hops and that the MS submits packets at the constant rate of 1 per hop. The detailed evolution of the routing states are laid out in Table 1, where we use the abbreviated notation $CoA \mapsto R$ to denote state (CoA, HoA, G) present at the interface directed towards R . At each time step router states are displayed after processing of the received packet. Individual packet processing is omitted from the table for the sake of comprehensiveness.

In detail the tree elongation is completed at time step 1, at time 2 the crossover router R_1 initiates a short cut. Starting at time instance 3 packets are already forwarded down the optimal delivery tree. Steps 4 and 5 are used to eliminate the dispensable path via the previous Designated Router, so that the tree morphing according to source mobility has converged at a time of 6 router hops. Forwarding states in (possibly large) delivery domains D_1, D_2 are simply replaced according to packet diffusion.

An example of enhanced complexity is presented in Fig. 8. Previous and new delivery trees are basically disjoint and the tree transformation has to proceed via several transition steps. Excerpts

Fig. 7 An example of local source movement

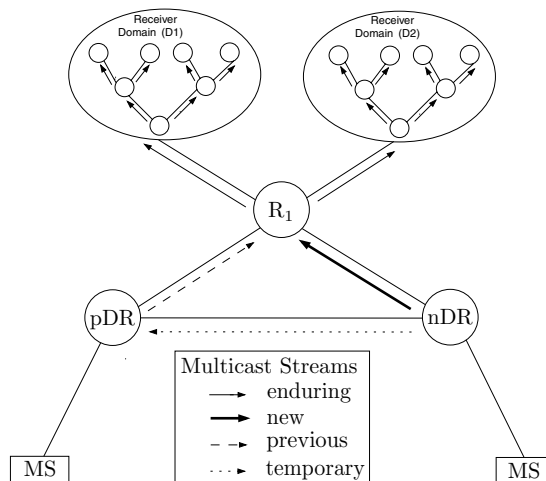
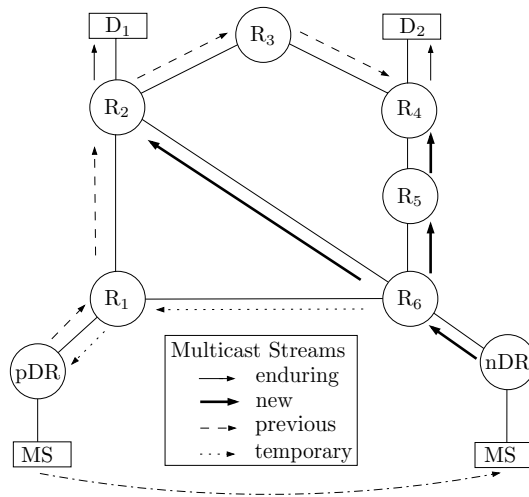


Table 1 Detailed evolution of router states and actions for the handover example displayed in Fig. 7

Time [hops]	Forwarding states at			Subsequent action
	<i>nDR</i>	<i>pDR</i>	R_1	
0	$nCoA \mapsto pDR$	$pCoA \mapsto R_1$	$pCoA \mapsto D_1, D_2$	<i>nDR</i> : submit state update packet
1	$nCoA \mapsto pDR$	$nCoA \mapsto R_1$	$pCoA \mapsto D_1, D_2$	<i>nDR</i> : submit pkt 2 $\mapsto pDR$ <i>pDR</i> : forward update
2	$nCoA \mapsto pDR$	$nCoA \mapsto R_1$	$pCoA \mapsto D_1, D_2$ $nCoA \mapsto D_1, D_2$	R_1 : forward update <i>nDR</i> : submit pkt 3 $\mapsto pDR$ R_1 : join ($nCoA, HoA, G$) $\mapsto nDR$
3	$nCoA \mapsto pDR, R_1$	$nCoA \mapsto R_1$	$pCoA \mapsto D_1, D_2$ $nCoA \mapsto D_1, D_2$	<i>nDR</i> : submit pkt 4 $\mapsto pDR, R_1$ D_1, D_2 : state replacement depth 1
4	$nCoA \mapsto pDR, R_1$	$nCoA \mapsto R_1$	$nCoA \mapsto D_1, D_2$	R_1 : prune ($*, HoA, G$) $\mapsto pDR$ <i>nDR</i> : submit pkt 5 $\mapsto pDR, R_1$ D_1, D_2 : state replacement depth 2
5	$nCoA \mapsto pDR, R_1$	–	$nCoA \mapsto D_1, D_2$	<i>pDR</i> : prune ($*, HoA, G$) $\mapsto nDR$ <i>pDR</i> : discard dupl. pkt 5 from <i>pDR</i> R_1 : discard pkt 4 from <i>pDR</i> <i>nDR</i> : submit pkt 6 $\mapsto pDR, R_1$ D_1, D_2 : state replacement depth 3
6	$nCoA \mapsto R_1$	–	$nCoA \mapsto D_1, D_2$	<i>pDR</i> : discard pkt 6 from <i>nDR</i> <i>nDR</i> : submit pkt 7 $\mapsto R_1$ D_1, D_2 : state replacement depth 4
7+	$nCoA \mapsto R_1$	–	$nCoA \mapsto D_1, D_2$	R_1, R_1 : regular pkt transport D_1, D_2 : state replacement depth 5+

Fig. 8 An example of source movement in a complex source network topology



of the router states are shown in Table 2: Tree elongation has completed after time 3. At instance 5 *pDR* has been pruned out and R_2 initiates the first short cut, which becomes effective at time 6 and causes optimal packet delivery into domain D_1 and enhanced forwarding into D_2 . At time 7 R_4 triggers the second short cut, which becomes effective at time 9. Hereafter packets flow

Table 2 Excerpt of router state evolution for the handover example displayed in Fig. 8

Time [hops]	Forwarding states at			
	R_1	R_2	R_4	R_6
0	$pCoA \mapsto R_2$	$pCoA \mapsto R_3, D_1$	$pCoA \mapsto D_2$	–
4	$pCoA \mapsto R_2$ $nCoA \mapsto R_2, pDR$	$pCoA \mapsto R_3, D_1$	$pCoA \mapsto D_2$	$nCoA \mapsto R_1$
5	$nCoA \mapsto R_2$	$pCoA \mapsto R_3, D_1$ $nCoA \mapsto R_3, D_1$	$pCoA \mapsto D_2$	$nCoA \mapsto R_1$
8	–	$nCoA \mapsto R_3, D_1$	$pCoA \mapsto D_2$ $nCoA \mapsto D_2$	$nCoA \mapsto R_1, R_2$
11	–	$nCoA \mapsto R_3, D_1$	$nCoA \mapsto D_2$	$nCoA \mapsto R_2, R_5$
13	–	$nCoA \mapsto D_1$	$nCoA \mapsto D_2$	$nCoA \mapsto R_2, R_5$

along the shortest path tree. Finally dispensable router states are eliminated at instances 11 and 13, when the algorithm has fully converged.

5 Evaluation and simulation

To evaluate the performance of the proposed scheme, we investigate the following aspects of significant relevance.

5.1 Handover initiated packet loss and delay

Regular MIPv6 handovers may in general lead to packet loss and delay, which can be minimized by using protocols such as FMIPv6 [16] or HMIPv6 [33]. There will be no additional packet loss caused by the tree morphing multicast handover, since a Mobile Source is enabled to immediately transmit multicast data. The initial tree elongation may result in triangular routing according to the distance between pDR and nDR . Subsequent tree optimization will monotonically reduce suboptimal routes.

It should be noted that state update messages injected subsequent to handover may immediately override router states from the previous distribution tree. This in principle may lead to dropping of delayed and overrun packets sent prior to handover. However, taking into account that update packets are issued only after the delay of layer 2 and layer 3 handover and—after reaching the previous distribution tree through pDR —are forwarded along the identical, possibly congested path and traffic class, only significant malfunctions of a network may possibly lead to packet overrun and the additional loss from that cause.

To judge on performance quality of the tree morphing (TM) scheme, we now analyze its delay effects within realistic Internet topologies. Therefore a stochastic discrete event simulation based on the network simulator platform OMNeT++ 3.1 Varga et al. [38] was performed, using the same real-world topologies as described in Section 3.3. It is worth noting that this selection of realistic Internet data covers the two significant aspects, i.e., intra-provider to inter-provider discrepancies, as well as scaling factors of two orders of magnitude for self-similar data sets.

The delay excess relative to optimal routes has been calculated as characteristic performance measure under the assumption of homogeneous link delays. Extreme values, i.e. maximal delays at initial elongation phase and minimal after convergence, were evaluated for tree morphing

(TM) as functions of the distance from pDR to nDR . In detail, designated routers within a given topology were randomly chosen from edge routers (node degree = 1) according to their predefined distances. For each pair of edge routers at the mobile source a uniformly distributed set of 20 receivers was established and delay values were taken from average reception time. Sampling of source positions was repeated 20 times for each parameter set in order to better explore the large phase space. Comparisons are drawn with bi-directional tunneling (BT), which does not depend on designated router distances, but on HA position. The delay excess in BT as function of HA position does not converge to a characteristic value, but rather admits a broad distribution. The latter has been derived from scattering HA positions uniformly from core routers (node degree > 1) within the sample networks.

It should be noted that these simulations concern delays for all three distribution trees in presence and thus qualitatively cover the solutions discussed in Section 2. Aside from additional signaling overhead, BT reflects the delays of Thaler [35], TM those of Jelger and Noel [13].

The results of our simulations are displayed in Fig. 9. pDR to nDR distances were chosen between 2 and 10, except for the ATT network, which exhibits a maximal edge router separation of 5. Error bars indicate the standard deviation of initial TM delay excess, as calculated from events differing in location of the mobile source. Plotted lines indicate the linear regression curves derived from this result set. Delay excess distributions for scattered HAs in BT are laid underneath TM curves in grey dots.

It can be observed that initially maximal delays of the tree morphing scheme tend to remain below the average of permanent BT packet retardation. Convergence of the TM then will lead to

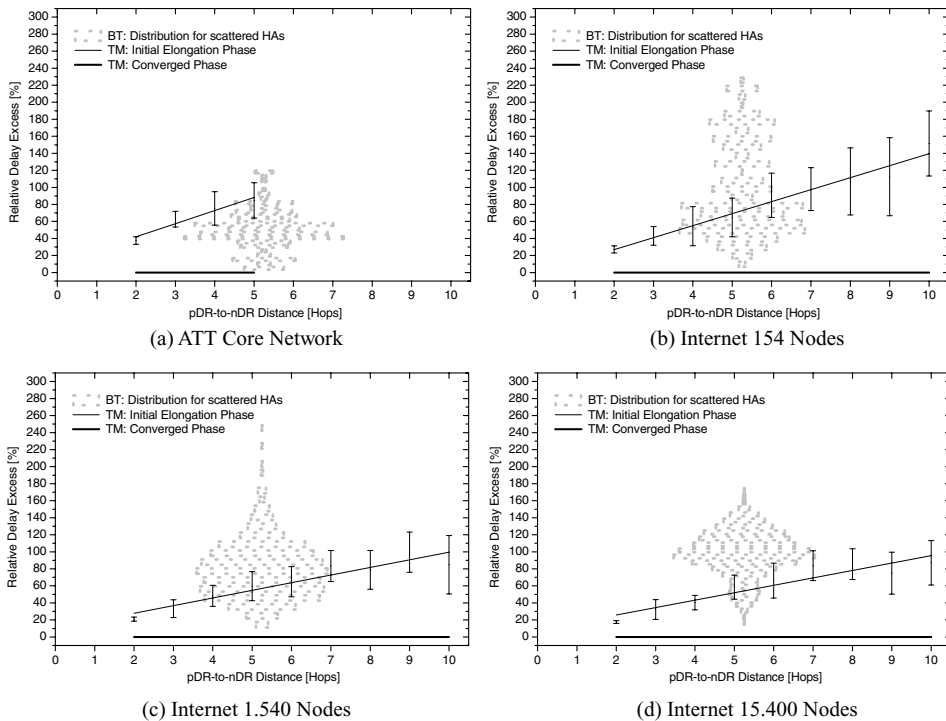


Fig. 9 Excess delay of optimal routes: comparison of bt and tm, initial and converged phase, for different network topologies

(relatively) undelayed packet delivery, which is never met in BT. Little dependence on network size becomes visible for TM—relative delays more strongly change with topologic characteristics. In a densely meshed provider network such as the ATT core, packet transitions are rapid and therefore initial delays from tree elongation account more dominantly for our relative measure. In the contrary it is interesting to note that delays from BT admit a systematic dependence on network size: BT average delay excess increases from 45% in the small ATT network to about 120% in the 15.400 node multiprovider Internet. From these observations it can be concluded that bi-directional tunneling attains appropriate performance for small communities within a densely meshed core network, but becomes infeasible in large inter-provider domains. The tree morphing even in its initially weakest phase exhibits fairly uniform performance, no matter how large the underlying network is.

5.2 Robustness and protocol convergence

The tree morphing scheme is robust in the sense that it transforms any intact, not necessarily optimal distribution tree into a new SSM shortest path tree rooted at the new source location. This can be easily observed from $(*, HoA, G)$ -states being only completely removed by the underlying multicast routing protocol, whose correctness is assumed. All intermediate router states conduct loop-free packet forwarding, as packets are sent down a coherent concatenation of shortest path trees. At any stage this resulting distribution tree does not attain an overall loop, but connects all receivers with the current source. Even though tree branches may intersect, any packet forwarding decision is based on only one underlying branch. This branch is identified by the corresponding extended RPF check, which tests for activation of each forwarding state current packet's incoming interface. The algorithm is self-healing through its autonomous elimination of invalidated or incorrect states.

Robustness of signaling in our scheme is equivalent to the degree of reliability in packet distribution. As there is no acknowledgement in multicast and as strongly asymmetric short cuts may lead to packet overrun, the Hop-by-Hop state update message should be piggybacked not only with the first, but rather with a first sequence of data packets. In QoS domains update packets should be classified with lowest drop precedence as not to be discarded by routers.

The tree morphing algorithm is robust under rapid movement. This can be concluded from observing that elongation at the tree root will equally function in multiple steps, while tree optimization will work on any distribution tree. Routing convergence under rapid MS's movement is assured even in the case of ping-pong mobility, as long as the tree elongation step can complete, i.e. the frequency of motion remains above the packet traveling distance between pDR and nDR . The latter assumption must be considered weak.

The convergence of the routing algorithms attains two measures, one for the time to reach optimal packet forwarding and the other for the final convergence to an optimal shortest path source tree.

Observation 2. For any receiver D_i the Tree Morphing Protocol has converged to optimal packet forwarding, iff multicast packets submitted by the mobile source are forwarded to D_i along the shortest (reverse unicast) path from D_i to MS.

Observation 3. For any receiver D_i and group $(nCoA, HoA, G)$ the Tree Morphing Protocol has converged to minimal shortest path tree, iff all router states (\cdot, HoA, G) forwarding towards D_i are member of the shortest path $(nCoA, HoA, G)$ source tree.

Consequently the total time to convergence T_{conv} decomposes into the time to optimal packet forwarding F_{conv} and the additional time M_{conv} needed for reshaping the tree to the minimal, i.e.

$$T_{\text{conv}} = F_{\text{conv}} + M_{\text{conv}}.$$

To evaluate the rate of convergence to optimal packet forwarding after source handover, consider any receiver D_i and the crossover router X_i between the previous and the next shortest path delivery tree towards D_i . If $X_i = pDR$, i.e. tree elongation alone has formed an optimal tree, forwarding is always optimal. Otherwise the time $F_{\text{conv}}^{(i)}$ to forwarding convergence will be bound by the successive signaling from the MS to pDR , following down the previous forwarding tree and returning to nDR via the last initiation point of short cut, which is given by the tree intersection X_i . This will be equivalent to the unicast forwarding time from nDR via pDR via X_i back to nDR . Denoting by dist_u^v the distance from u to v , then

$$F_{\text{conv}}^{(i)} \leq \text{dist}_{nDR}^{pDR} + \text{dist}_{pDR}^{X_i} + \text{dist}_{X_i}^{nDR}. \quad (3)$$

As all tree optimizations are performed in parallel, the total time to optimal forwarding convergence is given by

$$\begin{aligned} F_{\text{conv}} &= \max_i \{ F_{\text{conv}}^{(i)} \} \\ &\leq \text{dist}_{nDR}^{pDR} + \max_i \{ \text{dist}_{pDR}^{X_i} + \text{dist}_{X_i}^{nDR} \}. \end{aligned} \quad (4)$$

As an example consider nDR and pDR to be located in adjacent domains connected by a single peering point P . Aside from local branches the only tree intersection point will be P , which simultaneously lies on the route between pDR and nDR . Thus $F_{\text{conv}} = \text{dist}_{nDR}^{pDR} + \text{dist}_{pDR}^P$ and Router signaling attains a communication overhead close to the possible minimum dist_{nDR}^{pDR} between the two multicast domains.

Subsequent to optimal packet forwarding, distribution trees are reduced to minimal shortest path trees. For any receiver D_i the router X_i of intersection between the previous and the next shortest path delivery tree towards D_i will initiate prunes on the (possibly already degenerate) former distribution branches, as soon as optimally forwarded packets arrive. Signaling thereby follows a path inverse to state injection from X_i to pDR . Hence the convergence time M_{conv}^i to tree minimization after optimal forwarding does never exceed $F_{\text{conv}}^{(i)}$. As in general all possible short cuts are used and the previous delivery tree may have been partially deconstructed, the inequalities

$$M_{\text{conv}} \leq F_{\text{conv}} \quad \text{and} \quad T_{\text{conv}} \leq 2 F_{\text{conv}} \quad (5)$$

rigorously hold.

In our above example $M_{\text{conv}} = \text{dist}_P^{pDR}$, whence the total convergence to final tree geometry $T_{\text{conv}} = \text{dist}_{nDR}^{pDR} + 2 \text{dist}_{pDR}^P$ will be in the order of one roundtrip time between pDR and nDR .

To treat more complex scenarios we performed stochastic simulations of the tree morphing protocol signaling within the OMNeT++ platform and Internet topology data as described in Sections 3.3 and 5.1. Mean values of F_{conv} , the convergence of the routing protocol to optimal packet forwarding, have been calculated for samples as obtained in the previous section. 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

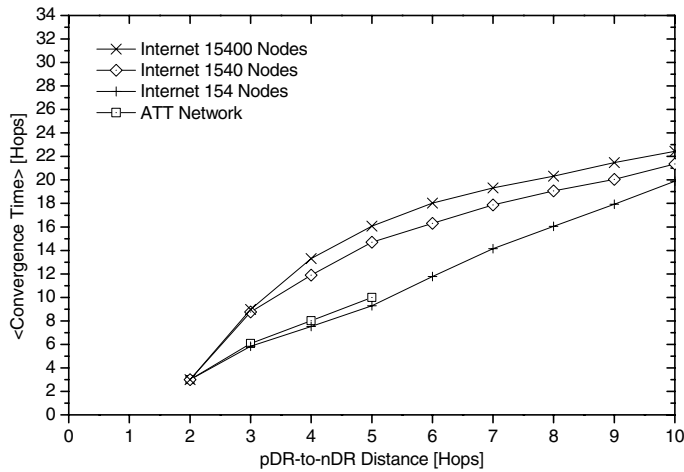


Fig. 10 Mean convergence time to optimal packet forwarding for the tree morphing as function of DR distance

from Thaler [35]: 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 CoA 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 completion. Following this idealized setting we calculate a lower bound for the time to optimal packet forwarding of the handover procedure proposed in [35].

The mean convergence time for the tree morphing as a function of *pDR* to *nDR* distance is displayed in Fig. 10. A sharp minimum can be observed for small designated router distances, where the tree morphing protocol admits its best performance. Curves for all topologies coincide in their minimum at distance two, which is due to the geometric property that only the tree intersecting crossover router connects the two DRs, independent of network topology.

Convergence for close distances of designated routers is attained after only a few router hops and mobility effects then remain almost invisible. Asymptotically its mean dependence on designated router distance as derived from the slopes of curves nicely approaches the roundtrip time between those DRs. Facing Eq. (4), it can be concluded that even in large networks and DRs far apart, crossover routers between the previous and next multicast tree remain within the region of source attachments. This result corresponds to the *a priori* observation discussed in Section 3 that multicast distribution trees obtained from subsequent mobility steps are not uncorrelated, but are likely to significantly overlap.

Figure 11 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 Thaler [35] 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.

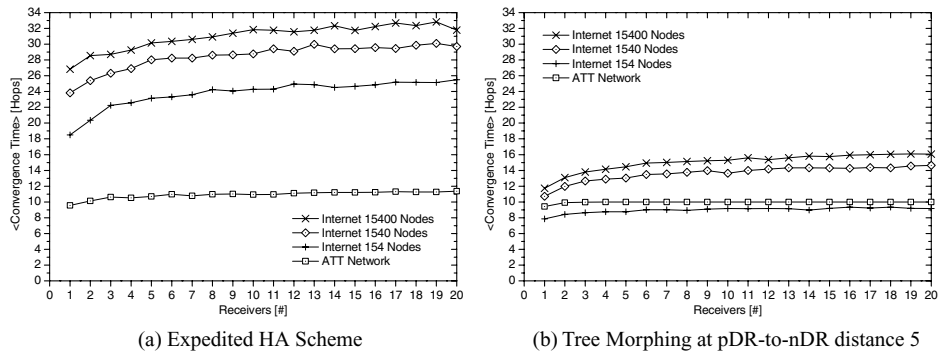


Fig. 11 Mean convergence time to optimal packet forwarding: comparison of an idealized HA-based handover scheme and TM as functions of receiver multiplicity

5.3 Infrastructural cost

The proposed mobile multicast routing scheme re-uses all possible parts of a previously established distribution tree with a single update message. 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 derived in Section 5.2 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 compared to the idealized Thaler scheme as described in Section 5.2, which is cost equivalent to any tree reconstruction scheme.

Results from stochastic simulations as described above are displayed in Fig. 12. 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

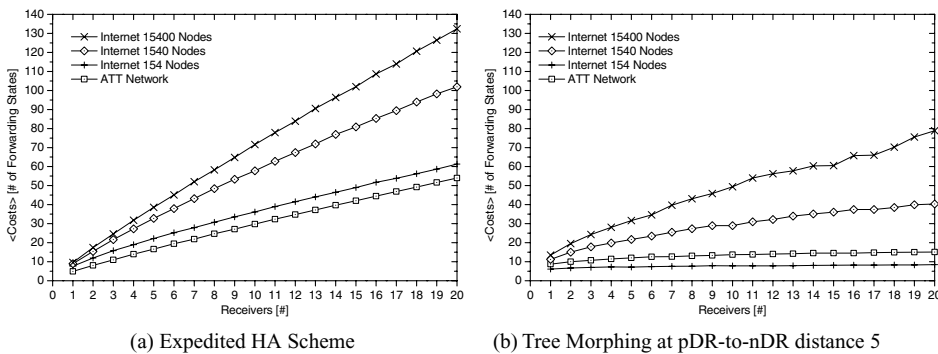


Fig. 12 Routing cost per handover: the number of newly established forwarding states for the idealized HA-based handover scheme and TM as functions of receiver multiplicity

from a minimal tree extension by one new edge router per receiver,² 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 effective source mobility management by re-using the global and changing only local parts of the distribution tree.

5.4 Security

The dynamic routing update procedure bears the risk of malicious mischief. A bogus node may inject state update packets to capture or damage a distribution tree, as well as to perform a distributed denial of service attack to on-tree routers. As known from unicast mobility, the location update procedure can be protected by including a proof of HoA address ownership along with the update message. Such proof can be given by an appropriately distributed security credential or the use of Cryptographically Generated Addresses (Aura [2], cf. Kellil et al. [15] for a comprehensive survey on multicast access control). A detailed definition of a securing layer for the proposed scheme, though, is beyond the scope of this article.

6 Conclusions and outlook

In this paper we presented a comprehensive analysis the source mobility problem in SSM routing and an approach to solve it. This novel scheme of morphing a previous distribution tree into a new shortest path tree deserves its motivation from initially derived properties of shortest path trees under root mobility transform and operates based on common multicast routing protocols with simple algorithmic extensions. After a handover it allows for immediate data transmission and strictly avoids data encapsulation.

Characteristic aspects of this tree morphing algorithm subsequently were analyzed, donating special focus on delay performance and protocol convergence. A rigorous upper bound for convergence was derived from a coincidence measure of the previous and the next distribution tree. All procedures could be shown to be robust and self-healing. Forwarding delays, convergence and state maintenance cost subsequent to handover have been calculated by stochastic simulations using real-world Internet topologies. It was found that maximal delays of the tree morphing scheme on average remain below packet retardation in bi-directional tunneling. Furthermore delays in our algorithm could be shown to perform independently of network size, while bi-directional tunneling performance degrades with network scaling. Protocol convergence times and costs remain systematically below corresponding values of competitive approaches and were found negligible for small mobility ‘step sizes’.

In future work we will concentrate on optimization of the scheme and on evaluation of further characteristic measures in the presence of network disturbances. A formulation of a corresponding security layer will be on schedule, as well.

Acknowledgments Major work on Internet edge scanning (cf. Section 3.4) has been compiled by Yhing Zhang, which is gratefully acknowledged. We would like to thank a referee.

² From the networking perspective we omit to sample multiple receivers per edge node, as they remain effectless for both routing schemes.

References

1. C. Adjih, L. Georgiadis, P. Jacquet and W. Szpankowski, Multicast tree structure and the power law, *IEEE Transact. Information Theory* 52(4) (2006) 1508–1521.
2. T. Aura, Cryptographically generated addresses (CGA), RFC 3972, IETF (2005).
3. S. Bhattacharyya, An overview of source-specific multicast (SSM), RFC 3569, IETF (2003).
4. R.C. Chalmers and K.C. Almeroth, On the topology of multicast trees, *IEEE/ACM Trans. Netw.* 11(1) (2003) 153–165.
5. R.-S. Chang and Y.-S. Yen, A multicast routing protocol with dynamic tree adjustment for mobile IPv6, *Journ. Information Science and Engineering* 20 (2004) 1109–1124.
6. J.C.I. Chuang and M.A. Sirbu, Pricing multicast communication: A cost-based approach, *Telecommunication Systems* 17(3) (2001) 281–297. Presented at the INET'98, Geneva, Switzerland, July 1998.
7. S.E. Deering, Host extensions for IP multicasting, RFC 1112, IETF (1989).
8. B. Fenner, M. Handley, H. Holbrook and I. Kouvelas, Protocol independent multicast—sparse mode (PIM-SM): protocol specification (Revised), Internet Draft—Work in Progress 12, IETF (2006).
9. R. Govindan and H. Tangmunarunkit, Heuristics for internet map discovery, in: *Procedures. IEEE INFOCOM 2000*, Vol. 3, IEEE Computer Society, Tel Aviv, Israel (2000) pp. 1371–1380.
10. O. Heckmann, M. Piringier, 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*, ACM Press, New York, NY, USA (2003) pp. 28–32.
11. H. Holbrook and B. Cain, Source-specific multicast for ip, internet Draft—Work in Progress 07, IETF. <ftp.rfc-editor.org/in-notes/internet-drafts/draft-ietf-ssm-arch-07.txt> (2005).
12. M. Janic and P. Van Mieghem, On properties of multicast routing trees, *Int. J. Commun. Syst.* 19(1) (2006) 95–114.
13. C. Jelger and T. Noel, Supporting mobile SSM sources for IPv6 (MSSMSv6), Internet Draft—Work in Progress (expired) 00, individual (2002).
14. D.B. Johnson, C. Perkins and J. Arkko, Mobility Support in IPv6, RFC 3775, IETF (2004).
15. M. Kellil, I. Romdhani, H.-Y. Lach, A. Bouabdallah and H. Bettahar, Multicast receiver and sender access control and its applicability to mobile IP environments: A survey, *IEEE Comm. Surveys & Tutorials* 7(2) (2005) 46–70.
16. R. Koodli, Fast handovers for mobile IPv6, RFC 4068, IETF (2005).
17. 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 (2003). <http://atnac2003.atcrc.com/ORALS/KURUP.pdf>
18. H. Lee, S. Han and J. Hong, Efficient mechanism for source mobility in source specific multicast, in: K. Kawahara and I. Chong, Eds., *Proceedings of ICOIN2006*, Vol. 3961 of *LNCS*, SpringerVerlag, Berlin, Heidelberg (in press) (2006).
19. D. Magoni, nem: A software for network topology analysis and modeling, in: *Proceedings of the 10th IEEE Symposium on Modeling, Analysis and Simulation of Computer & Telecomm. Systems (MASCOTS'02)*, IEEE Computer Society, Fort Worth, Texas, USA (2002) pp. 364–371.
20. D. Magoni and J.-J. Pansiot, Internet topology modeler based on map sampling, in: *Proceedings of the 7th IEEE Symposium on Computers and Communications*, IEEE Computer Society, Taormina, Italy (2002) pp. 1021–1027.
21. MaxMind LLC (2006), GeoIP, <http://www.maxmind.com>.
22. A. Medina, A. Lakhina, I. Matta and J. Byers, 'BRITe: Boston university Representative Internet Topology gEnerator,' <http://www.cs.bu.edu/brite/> (2005).
23. A. O'Neill, (2002) Mobility Management and IP Multicast, Internet Draft—Work in Progress (expired) 01, IETF.
24. G. Phillips, S. Shenker and H. Tangmunarunkit, Scaling of multicast trees: Comments on the chuang-sirbu scaling law, in: *SIGCOMM '99: Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ACM Press, New York, NY, USA (1999) pp. 41–51.
25. I. Romdhani, H. Bettahar and A. Bouabdallah, Transparent handover for mobile multicast sources, in: P. Lorenz and P. Dini, Eds., *Proceedings of the IEEE ICN'06*, IEEE Press (2006).
26. I. Romdhani, M. Kellil, H.-Y. Lach, A. Bouabdallah and H. Bettahar, IP mobile multicast: challenges and solutions, *IEEE Comm. Surveys & Tutorials* 6(1) (2004) 18–41.
27. SCAN project: Internet Maps. SCAN+Lucent map (2005), <http://www.isi.edu/scan/mercator/maps.html>.
28. T.C. Schmidt and M. Wählisch, Extending SSM to MIPv6—problems, solutions and improvements, *Computational Methods in Science and Technology* 11(2) (2005a) 147–152. Selected Papers from TERENA Networking Conference, Poznań (2005a).

29. T.C. Schmidt and M. Wählisch, Multicast mobility in MIPv6: Problem Statement, IRTF Internet Draft—Work in Progress 00, MobOpts, (2005b).
30. 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) (2005c) 123–142.
31. T.C. Schmidt and M. Wählisch, Seamless multicast handover in a hierarchical mobile IPv6 environment (M-HMIPv6), Internet Draft—Work in Progress 04, individual (2005d). <ftp://ftp.rfc-editor.org/in-notes/internet-drafts/draft-schmidt-waehlich-mhmipv6-04.txt>
32. J.S. Silva, S. Duarte, E. Monteiro and F. Boavida (2003) MNet—A new multicast approach for the future Internet, in: *Proceedings of the 10th Intern. Conference on Telecommunications*, Vol. 1 (2005d) pp. 340–347.
33. H. Soliman, C. Castelluccia, K. Malki and L. Bellier, Hierarchical mobile IPv6 mobility management (HMIPv6), RFC 4140, IETF (2005).
34. W.R. Stevens, *TCP/IP Illustrated, The Protocols*, Vol. 1. Addison Wesley, Reading, MA (1994).
35. D. Thaler, Supporting mobile SSM sources for IPv6, in: proceedings of ietf meeting, individual (2001). www.ietf.org/proceedings/01dec/slides/magma-2.pdf
36. P. Van Mieghem, *Performance Analysis of Communication Networks and Systems* Cambridge University Press, Cambridge, (2006).
37. P. Van Mieghem, G. Hooghiemstra and R. van der Hofstad, On the efficiency of multicast, *IEEE/ACM Trans. Netw.* 9(6) (2001) 719–732.
38. A. Varga et al. The OMNeT++ discrete event simulation system (2005). <http://www.omnetpp.org>.
39. G. Xylomenos and G.C. Polyzos, IP multicast for mobile hosts, *IEEE Comm. Mag.* 35(1) (1997) 54–58.