Mobile Multicast

Thomas C. Schmidt
HAW Hamburg, Germany

Matthias Wählisch
FHTW Berlin, Germany

INTRODUCTION

The submission of datagrams from a node to a group of receivers must be seen as a powerful extension of the Internet routing layer (Deering, 1989). Multicast group communication forms an integral building block of a wide variety of applications, ranging from public content distribution and streaming over voice and videoconferencing, collaborative environments, and gaming up to the self-organization of distributed systems. Its support by network layer multicast will be needed, whenever globally distributed, scalable, serverless, or instantaneous communication is required. As broadband media delivery more and more emerges to be a typical mass scenario, scalability and bandwidth efficiency of multicast routing continuously gains relevance. Internet multicasting will be of particular importance to mobile environments, where users commonly share frequency bands of limited capacity. The rapidly increasing mobile reception of 'infotainment' streams may soon require a wide deployment of mobile multicast services.

The fundamental approach to deal with mobility in the next-generation Internet is stated in the Mobile IPv6 RFC (Johnson, Perkins, & Arkko, 2004). Multicast has only roughly been treated therein, but raises quite distinctive aspects within a mobility-aware Internet infrastructure. On the one hand multicast routing itself supports dynamic route configuration, as members may join and leave ongoing group communication over time. On the other hand 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.

Mobile multicast has been a concern for about ten years (Xylomenos & Polyzos, 1997) and led to innumerous proposals for solutions. 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, Kellil, Lach, Bouabdallah, & Bettahar, 2004a). In multimedia conference scenarios each member commonly operates as receiver and as sender for multicast-based group communication. In addition, real-time communication such as voice or video over IP places a severe temporal requirement on mobility protocols: seamless handover scenarios need to limit disruptions or delay to less than 100 milliseconds. Jitter disturbances are not to exceed 50 milliseconds. Note that 100 milliseconds is about the duration of a spoken syllable in real-time audio.

While multicast routing itself has been proposed to support mobility on the Internet (Helmy, 2000), consensus on an efficient, robust, and widely deployable scheme of multicast for mobile hosts is still lacking (Schmidt & Wählisch, 2005a). In this review we will summarize the state of the art in current work to multicast to Mobile IPv6 networks. The principle conceptual problems are discussed and analyzed. Propositions for improvement and possible directions to further proceed towards a mobile multicast are presented.

BACKGROUND

Multicast mobility must be considered a generic term that subsumes a collection of quite distinct functions. First, multicast communication divides into any source multicast (ASM) (Deering, 1989) and source-specific multicast (SSM) (Bhattacharyya, 2003; Holbrook & Cain, 2005). Second, the roles of senders and receivers are asymmetric and need distinction. Both individually may be mobile. Their interaction is facilitated by a multicast routing function—such as DVMRP (Waitzman, Partridge, & Deering, 1988), PIM-SM/SSM (Estrin et al., 1998) or CBT (Ballardie, 1997)—and the IPv6 multicast listener discovery protocol (Deering, Fenner, & Haberman, 1999).

Any multicast mobility solution must account for all of these functional blocks. It should enable seamless continuity of multicast sessions when moving from one IPv6 subnet to another. It should preserve the multicast nature of packet distribution and approximate optimal routing. It should support per-flow handover for multicast traffic, as properties and designations of flows may be of individual kind.

Multicast routing dynamically adapts to session topologies, which then may change under mobility. However, routing convergence arrives at a time scale of seconds, even minutes and is far too slow to support seamless handovers for interactive or real-time media sessions. The actual temporal behavior strongly depends on the routing protocol in use and on the geometry of the current distribution tree. A mobility scheme that arranges for adjustments—that is, partial changes or full reconstruction—of multicast trees is forced to make provision for time buffers sufficient for protocol convergence. Special attention is needed with a possible rapid movement of the mobile node, as this may occur at much higher rates than compatible with protocol convergence.

IP layer multicast packet distribution is an unreliable service, which is bound to connectionless transport protocols. Packet loss thus will not be handled in a predetermined fashion. Mobile multicast handovers should not cause significant packet drops. Due to statelessness the bi-casting of multicast flows does not cause foreseeable degradations of the transport layer.

Group addresses in general are location transparent, even though there are proposals to embed unicast prefixes or rendezvous point addresses (Savola & Haberman, 2004). Source addresses contributing to a multicast session are interpreted by the routing infrastructure and by receiver applications, which frequently are source address aware. Multicast therefore inherits the mobility address duality problem for source addresses, being a logical node identifier (HoA) on the one hand and a topological locator (CoA) on the other.

Multicast sources in general operate decoupled from their receivers in the following sense: a multicast source submits data to a group of unknown receivers, thus operating without any feedback channel. It neither has means to inquire on properties of its delivery trees, nor will it be able to learn about the state of its receivers. In the event of an inter-tree handover, a mobile multicast source therefore is vulnerable to losing receivers without taking notice.

MULTICAST LISTENER MOBILITY

A mobile multicast listener entering a new IP subnet may simply re-subscribe to its previously received groups. It thereby will encounter one of the following conditions. In the new network the requested multicast service may be supported, but the multicast groups under subscription may not be forwarded to it. The current distribution trees for the desired groups may reside at large routing distance. It may as well occur that some or all groups under subscription of the mobile node are received by one or several local group members at the instance of arrival and that multicast streams natively flow.

The temporal behavior of the multicast handover will largely depend on the conditions met. The mobile node may employ predictive or reactive accelerating schemes to reduce performance degradation. For a detailed analysis, see Schmidt and Wählisch (2005b).

ASM SOURCE MOBILITY

A node submitting data to an Any Source Multicast group defines the root of either a shared or a source-specific delivery tree. Beside root location forwarding along, this delivery tree will be bound to a topological network address due to reverse path forwarding (RPF) checks. A mobile multicast source moving away is *solely* enabled to either inject data into a previously established delivery tree by using its previous topologically correct source address, or to (re-)define a multicast distribution tree compliant to its new location. In pursuing the latter the mobile sender will have to proceed without control of the new tree construction due to decoupling of sender and receivers.

Conforming to address transparency and temporal handover constraints will be the key problems for any route optimizing mobility solution. Additional issues arrive from possible packet loss and from multicast scoping. A mobile source away from home must attend scoping restrictions which arise from its home and its visited location (Johnson et al., 2004).

SSM SOURCE MOBILITY

Fundamentally Source Specific Multicast has been designed for changeless addresses of multicast senders. Source addresses in client subscription to SSM groups are directly used for route identification. Any SSM subscriber is thus forced to know the topological address of its group contributors. SSM source identification invalidates when source addresses change under mobility.

Consequently, source mobility for SSM packet distribution introduces a significant conceptual complexity in addition to the problems of mobile ASM. As a listener is subscribed to an (S,G) channel membership and as routers have established an (S,G)-state shortest path tree rooted at source S, any change of source addresses under mobility requests for state updates at all routers and all receivers. A moving source would have to update its change of CoA with all listeners, which subsequently had to newly subscribe and initiate corresponding source-specific trees. As the principle multicast decoupling of a sender from its receivers likewise holds for SSM, the need for client update turns into a severe problem.