IP MULTICASTING FOR WIRELESS MOBILE HOSTS

George Xylomenos and George C. Polyzos

{xgeorge,polyzos}@cs.ucsd.edu

Computer Systems Laboratory

Department of Computer Science and Engineering

University of California, San Diego

La Jolla, CA 92093-0114

ABSTRACT

We consider the problem of efficient multicast support for mobile wireless hosts in TCP/IP networks. We summarize techniques supporting multicasting and mobility, along with their constraints and interactions, and explore architectural alternatives for solutions to the combined problem that remain compatible with the existing architecture. We propose a new mechanism for group management, optimized for point-to-point links. We also discuss three proposed multicast delivery mechanisms and compare them with respect to efficiency on wireless networks and impact on host software.

INTRODUCTION

While support for multicasting in the current Internet Protocol (IP) version (IPv4) has been evolving for more than five years, it is still regarded as experimental. However, the requirements for the next version of IP (IPv6) emphasize support for multicasting and encourage the replacement of broadcasting with multicasting [1][2] whenever possible. One reason for this interest is the ease of addressing services with a single multicast group identifier, thus enabling resource location [3] and distributed and replicated services. Another reason is the potential of multicasting for economizing on bandwidth, as datagrams to a multicast group are duplicated only when paths to their multiple destinations diverge. Bandwidth intensive services such as video distribution can become more cost effective due to this. In the meantime, the growth in wireless communications has attracted interest in the integration of wireless and wireline IP networks. A wide area wireless network allows devices to move without disrupting their communications. To achieve this, IP must be extended to transparently handle roaming hosts by hiding mobility from the transport service. Issues include not only continuous datagram delivery so that host connectivity is persistent, but also adaptation of hosts to visited networks, which could be achieved via multicast based resource discovery.

In the following we examine how multicasting and mobility can be combined in the IP world, by first presenting the relevant IP extensions and then examining some proposed solutions to the problems arising from their interactions, sep-

arating local from global, i.e. wide area, mechanisms. We are primarily concerned with the integration of such solutions with existing protocols and their efficiency in terms of utilizing the limited bandwidth and battery life of wireless hosts. We conclude with a service deployment plan.

IP MOBILITY

IP mobility support allows a *mobile host* (MH) to change its point of attachment to the network without losing connectivity, transparently to the transport layer [4]. Internet transport layer protocols (TCP/UDP) however assume that a host's address is fixed, so simply providing MHs with local addresses when attached to a new network cannot achieve transparency, as transport connections will have to be re-established. IP mobility extensions, and in particular the *Internet Mobile Host Protocol* (IMHP) [5], which we examine here, provide a mechanism for a MH to retain one address while roaming, called its *home address*, even though it connects to various wireless networks.

The problem to be solved is circumventing IP routing. A router receiving a datagram to a non local host forwards it based only on the network part of the destination address, thus keeping track of complete networks by their address prefixes. Datagrams are forwarded towards a router advertising reachability to their destination network. When this network is reached, the datagram is forwarded to the correct host by a local router with detailed knowledge of its attached hosts. With mobility, while a MH visits a remote network, called a *foreign network*, datagrams to the MH are still forwarded to the network indicated by its home address, which generally differs from the foreign network's address.

To solve this problem, a router on the home network, called the *home agent* (HA) and a router on the foreign network, called the *foreign agent* (FA), must cooperate. When the MH visits a foreign network, it locates the FA and *registers* with it, and then informs its HA of the FA currently serving it. Subsequently, the MH sends its datagrams via the FA, which forwards them normally as unicast IP routing ignores their source address. On the other hand, datagrams destined for the MH, are first delivered to the HA on the home network, which consults its tables to locate the FA serving the MH, and then *encapsulates* the datagrams inside new ones from the HA to the FA. The FA on receiving encapsulated datagrams, decapsulates and forwards them to

This research was supported in part by the UC MICRO program and by ARO FRI grant DAAH049510248

Fig. 1. (a) Triangle routing: datagrams from MH to CH are delivered directly, but datagrams from CH to MH must first pass through HA, which tunnels them to FA for delivery. (b) Sending multicasts: the MH tunnels multicasts to the HA which then forwards them as if they had originated in the home network, so that MR receives them normally.

the directly attached MHs. This technique, called *tunneling*, allows the MH to communicate continuously using its home address, despite its mobility, but it has two drawbacks. First, datagrams to the MH are always routed via its HA, resulting in suboptimal *triangle routing*, as shown in Figure 1.(a). Proposed extensions rectify this problem [5], but their applicability and performance are unclear. Second, there is no de-registration process, as the MH may lose connectivity with the FA at any time due to movement, so registrations must be repeated periodically, else the FA will delete the information on the visiting MH from its tables. The problems of when a MH should be assumed absent and how initial contact is established among the FA and the MH, are beyond the scope of this paper.

IP MULTICASTING

IP multicasting is based on the *host group* concept: a dynamic set of hosts identified by a single address [6]. Hosts can *join* or *leave* groups at any time, to start or stop receiving datagrams sent to the group, while any host can send to a group. To deliver datagrams to a dynamic set of receivers we need to track group membership and route data to group members. Conceptually, we can split the required mechanisms into *local*, such as group membership management and local delivery, and *global*, such as routing from senders towards any interested networks.

Local Mechanisms

Local group membership is tracked using the *Internet Group Management Protocol* (IGMP) [6]. Each network supporting IP multicasting designates one router to periodically send queries for group membership in its local area, with attached hosts replying with the groups they want to participate in. The router then builds a list of groups whose messages should be distributed locally. The global mechanisms ensure datagram delivery towards the local router for these groups.

IGMP and the local delivery architecture were designed with broadcast based LANs in mind where *native* multicasting is available. Queries are multicast to an address to which all receivers are listening, and membership reports are sent to the multicast address for the group reported. The router and other group members listen to the group address, so that the router is informed of local group presence and other members suppress duplicate reports. Joining a group leads to an unsolicited report which is periodically repeated. If

no reports arrive for some time, the group is assumed absent, so leaving a group does not require explicit messages. Group members also send unsolicited reports to speed up delivery when first joining. Thus, for broadcast LANs one query and one report per group per querying period are required. Furthermore, since native multicast is available, the router simply records the presence of a group and multicasts any received datagrams.

When the router has to support a set of *point to point* (PtP) links, datagrams have to be separately unicast to each host. Thus, separate queries and reports are needed for each link and the router must record detailed membership information, such as a list of hosts per group or a list of groups per host, even though only the simple group list is required for global co-operation. Many wireless networks provide only PtP local links, and some proposals for combining multicasting and mobility in IP use *virtual* PtP links among router and receivers. Therefore, any improved local mechanisms for PtP links could be quite useful for the bandwidth limited wireless links.

Global Mechanisms

Using local mechanisms, routers learn which groups they must receive and how to deliver them locally. Global multicast delivery towards local routers requires cooperation among them. The most widespread routing mechanism is the *Distance Vector Multicast Routing Protocol* (DVMRP) [7]. In this algorithm each router keeps track of the first links in the best paths to datagram *sources*. Datagrams arriving from the first links in the best path to their sources are forwarded through all other links, while all other datagrams are discarded. Thus, datagrams are flooded over a tree composed of the best receiver to sender, or *reverse*, paths, with local routers receiving every group and forwarding locally only the present ones. DVMRP only tracks routes to *networks* to save routing table space, using a distance vector routing algorithm. As multicasting is not a required router feature, multicast routers communicate over non-multicast aware areas by setting up fixed *tunnels* among them, where multicast datagrams are encapsulated inside unicast datagrams at one tunnel endpoint and are decapsulated at the other. Tunnels are *virtual* links, so the collection of multicast aware areas connected by tunnels is a virtual network, known as the MBone. Datagram delivery scope is limited by the IP *time to live* (TTL) field, usually interpreted as a hop limit. Since virtual links look like a single hop, multicast routers attach TTL thresholds to tunnels to limit multicast delivery.

One alternative to DVMRP is the *Multicast Open Shortest Path First* (MOSPF) [8] protocol which uses a link state routing algorithm. In MOSPF routers flood topological information and group membership lists among them, so that each router has detailed knowledge of group membership. Datagrams arriving at a router are forwarded through all links leading towards the leaves of a shortest path tree from the sender network to all receiver networks. As all routers have the same network image, they compute the same trees.

Thus, datagrams are only propagated where needed, in contrast to DVMRP. Another alternative is the *Core Based Trees* (CBT) [9] protocol, which employs a single delivery tree for each group centered on an arbitrarily chosen *core* router. Local routers that want to receive a group contact its core so that a reverse shortest path tree from the core to all receiver networks is built. Datagrams to a group are initially sent towards its core, and when they reach any router in the tree they are forwarded over all tree links. Routing is normally less efficient than in the other proposals since shortest paths are used only on the tree, but a single tree per group simplifies tree management. CBT can employ any underlying unicast routing algorithm and makes routing decisions without considering the source address of datagrams. A last proposal, *Protocol Independent Multicast* (PIM) [10], combines core based and shortest path trees.

LOCAL MULTICASTING & MOBILE HOSTS

To co-operate with other routers, each multicast router tracks local group membership using IGMP, which was designed to complement local delivery mechanisms based on native multicast. If the MHs are attached to the router via point to point (PtP) links, either the physical ones of cellular telephone networks, or the virtual ones (tunnels) of some proposed schemes below, additional state is needed in the router beyond the list of present groups: either a list of hosts for each group or a list of groups for each host. As multicast datagrams have to be separately unicast over each PtP link, this state enables selective local multicast forwarding, instead of forwarding datagrams for all local groups over each bandwidth limited PtP link. We can employ this state to optimize IGMP by replacing the query/response mechanism. We propose that membership should be discovered by using explicit join/leave group messages sent by the MHs to the router. The router tracks group membership for each MH by listening to these messages, as long as the MH remains local. Thus, we replace the periodical IGMP reports containing complete membership information with state difference messages, in a variation of *header compression* [11]. If a MH belongs to n groups during its presence in the area, standard IGMP exchanges one query and n responses per query interval, while our scheme only requires n join and n leave messages regardless of membership duration. In addition, the explicit leave messages cause the router to stop forwarding multicast datagrams at once, rather than after the change is discovered during the next query interval. In terms of MH battery power, if groups are inactive the MH can switch to power saving mode without interruptions from IGMP.

For broadcast based wireless networks, datagrams have to be received by all MHs anyway, so native multicasting should be employed to minimize multicast delivery costs. However, standard IGMP may still be wasteful, as queries will have to be broadcast regularly even when no multicast receivers exist locally. Increasing the query interval reduces this management overhead but increases delivery overhead due to wasted transmissions after all MHs leave a group. By using

our proposed join/leave messages instead, we can use the more timely membership information to minimize delivery overhead. Management overhead is reduced by our method when few MHs are served by a router, group membership changes rarely and MHs receive distinct groups, while standard IGMP performs better when these conditions do not hold, where only a few messages are required to complete a query/response cycle. To use our method, additional state is required at the router, which should be weighed against the potential gains in bandwidth and power efficiency.

A network may even provide both PtP links and broadcast channels intended for common signaling. If the latter have spare capacity, it is possible to use them for native multicast and the standard IGMP algorithm. It is not clear whether this would optimize efficiency, as all MHs would have to receive such transmissions. Since native multicast saves bandwidth but consumes power at all MHs, a threshold on the ratio of group members to total population could be employed to decide among native multicast and multiple unicast for each group. To use such a threshold though, detailed membership information is required which can be discovered by our join/leave mechanism but not by standard IGMP.

MULTICASTING FROM MOBILE HOSTS

Unicast IP routing depends only on datagram destinations, so MHs can send datagrams from any point of attachment. DVMRP and MOSPF however route multicasts based also on the network part of a datagram's source. Specifically, multicasts from a MH are expected from the link used to reach its home network, but, if the MH has moved elsewhere, its datagrams will arrive from the link used to reach its current location. DVMRP drops such datagrams while MOSPF forwards them only towards the leaves of a distribution tree routed at the home network. In both cases, some destinations are not reached. Since CBT uses source independent distribution trees, routing depends only on a datagram's destination, enabling the MH to use normal routing mechanisms.

To solve routing problems, we can make multicasts originate from the current MH's network. We cannot use the FA's address as the source, as replies to multicasts would go to the FA, but we could assign a temporary local address to the MH instead. In this case, we would have to deal with the address shortage problems of the current IP version, and with the misdelivered replies to the MH's old temporary address after it leaves the local network, a problem preventing use of this solution even for the next version of IP. Another approach is to use complete addresses rather than network addresses for routing. This is viable for the few multicast areas currently in existence, but eventually routing table size will become a problem.

A practical approach is to circumvent routing by tunneling multicasts from the MH to the HA, which then forwards them as if they had originated locally, so that routers receive them from the expected links, as shown in Figure 1.(b). The HAs do not need to be multicast routers themselves, and

Fig. 2. (a) Home agent routing: MR delivers the datagrams to HA₁ and HA₂ which run IGMP. HA₁ uses two separate tunnels to MH₁ and MH₂, converging with a tunnel from HA₂ to MH₃. (b) Foreign agent routing: Multicasts are forwarded to the local multicast router (FA) for appropriate delivery.

since they must process encapsulated datagrams anyway to support mobility, they only need to be modified to recognize tunneled multicasts, while the FAs need no modifications. This approach leads to suboptimal triangle routing, as the HA is always used as an intermediate destination, but from the HA onward the standard multicast algorithms are used.

MULTICASTING TO MOBILE HOSTS

Home Agent Routing

A simple approach for multicast reception on MHs is to let the HA handle routing, by executing IGMP and delivering multicasts to its MHs as if they were on the home network. When a MH is not at home, datagrams may be delivered by tunneling them through the FA, with IGMP membership reports from the MH being unicast to the HA, as shown in Figure 2.(a), so that HA and MH communicate via two virtual PtP links. As discussed earlier, for PtP links, per MH information must be kept in the local router (HA), so IGMP can be modified to use explicit join/leave messages, thus optimizing transmissions over the wireless part of the virtual PtP link. To implement this scheme we only need to extend a similar proposed mechanism that tunnels local *broadcasts* by encapsulating datagrams twice: the outer header, addressed to the FA, is striped by the FA which delivers locally the encapsulated broadcast, addressed to the MH. The MH strips the inner header to uncover the broadcast datagram itself. Broadcast tunneling is activated by a flag on the registration messages to the HA. Exactly the same mechanism can be used for multicasts, with a flag indicating that the HA should both run IGMP and forward multicasts.

This approach interoperates with existing networks since multicast routing is transparent to the FA, while the MH and the HA that need to be modified are generally under the same administrative control. Thus, the MH will receive multicasts even on foreign networks that do not support multicasting. In addition, the modifications needed are minor extensions of existing mechanisms. On the other hand, resource utilization is inefficient even if the IGMP optimizations are employed. First, suboptimal triangle routing is used. Second, with the virtual PtP links datagrams are unicast separately to each MH over multiple tunnels, even when the wireless network supports native multicast. Third, multiple tunnels from separate HAs are used to deliver the same group to a wireless network, leading to the *tunnel convergence problem* [12]. Whatever the source of multiple tunnels, duplication occurs at the bandwidth constrained wireless link.

Fig. 3. Combined routing: (a) MH₃ first reported group membership to FA and a tunnel was set up from HA₂ to FA which delivers datagrams locally. (b) After MH₃ leaves FA's network, the tunnel from HA₂ is torn down, and a new one is set up from HA₁ when another MH reports group membership.

Foreign Agent Routing

When the FA supports multicast routing, the existing IP multicast model can be used for the wireless network. The FA executes IGMP, receives datagrams, and forwards them to the MHs as shown in Figure 2.(b). Depending on the network, PtP and/or broadcast links may be available, so the earlier discussion for local multicast applies. Since global multicasting is concerned with forwarding multicasts to complete networks, the FA can hide the home addresses of the MHs. Implementation of this scheme is the same as implementation of IP multicast in general, i.e. executing IGMP and arranging for fixed DVMRP tunnels to the MBone. Actually, any model that separates local and global mechanisms as described above, including MOSPF and CBT, can use this scheme to accommodate MHs, without incurring any additional overhead for multicast management and delivery.

By separating global from local mechanisms, the FA can employ group management and delivery mechanisms optimized for its specific network and enforce local administrative policies concerning multicasting when delivery tradeoffs exist. The HA does not need to support multicasting and tunnel convergence is avoided. The drawback of this scheme is that the wireless network provider may not want to provide multicasting, either because it consumes precious bandwidth, or because it is an experimental and evolving service. Although the latter is a reasonable concern, the former is self defying, as the MHs can always set up tunnels from their home networks to receive multicasts, thus multiplying overhead. A more practical consideration is ensuring that a MH supports the optimized schemes that a wireless network may employ. For multicast delivery, when the wireless network offers both PtP and broadcast links, the FA could employ encapsulation to switch to unicast mode when desired, so the MH should be prepared to handle this case too. For IGMP operation, the choice between normal and join/leave mode can be made at registration time, with the MH indicating what it supports and the FA making the choice, so the FA should be prepared to handle both options.

Combined Routing

A third approach to multicast reception combines tunneling from the HA with local multicast service from the FA [12]. The FA executes IGMP locally and sets up unique tunnels on demand for all required groups, originating at the HA of the first MH that joined each group, as shown in Figure 3.(a). Thus, global routing is performed by the FA and some HAs

	Change Scale	Changed Entities	Protocol Overhead	Delivery Overhead	Multicast Routing	Local Operation	Local Network
Home Agent	Minor	HA,MH	Yes	Yes	Suboptimal	Suboptimal	Home
Foreign Agent	Minor	FA	No	No	Optimal	Optimal	Foreign
Combined	Major	HA,FA,MH	Yes	No	Suboptimal	Optimal	Both

TABLE I
COMPARISON OF MULTICAST RECEPTION TECHNIQUES

together, while local delivery and management is performed by the FA only. A HA whose MHs have all left the foreign network will tear down the tunnel and inform the FA, who sets up a new tunnel after a new membership report arrives, or immediately if join/leave messages are used, as shown in Figure 3.(b). The HA must be a multicast router, and it must notify the FA before tearing down a tunnel, else inactive tunnels will not be distinguished from disconnected ones.

Besides allowing local optimizations, this model also enables multicasting without the FA being a multicast router, as the HAs are responsible for tunneling multicasts to the FA. On the other hand, suboptimal triangle routing is used and tunnel management overhead is repeatedly incurred when tunnels are set up and torn down. It is unclear how the HA will learn whether any of its MHs require the groups that it is asked to tunnel, and whether any are still members of this group while the tunnel exists: either the HA would have to trust the separately controlled FA, or an extra handshaking protocol will have to be devised to let the HA know of the membership status of its MHs, so that detailed membership information will have to travel beyond the FA. Finally, since both FA and HA have to be modified to use a non standard protocol for tunnel maintenance, it may be preferable to simply support standard multicast routing at the FA itself.

Comparison of Approaches

A point summary of the approaches described above for multicast reception appears on Table I, examining how easily they can be integrated with existing IP mechanisms and how efficient they are. Regarding interoperability, the *Change Scale* and *Changed Entities* criteria show the extent and location of required host software modifications. Regarding performance, *Protocol Overhead* and *Delivery Overhead* examine overhead beyond the standard IP mechanisms, *Multicast Routing* compares each approach with standard multicast routing, and *Local Operation* examines whether local IGMP and delivery optimizations can be transparently employed. HA routing is easy to implement transparently as it is limited to hosts under the same administrative control, but suffers from tunneling and routing overhead. FA routing requires multicast support at the FA, but does not involve any other overhead. Combined routing is harder to implement as it is non standard and stands between the other two approaches in performance. Last, *Local Network* shows which network's limited TTL, i.e. local, multicasts will be received by the MH, and depends on the host responsible for local

delivery. Considering the foreign network as local enables the use of multicasting for resource discovery purposes [3].

CONCLUSION

We have discussed IP multicast mechanisms that interoperate with existing protocols without sacrificing efficiency, by tailoring protocols to wireless network needs. Locally, the join/leave model can be used whenever the network technology is based on PtP links. Globally, we should avoid combined routing as it is closer to the worst rather than the best solutions. As long as IP multicasting support is limited, home agent routing can be employed to support multicasting by modifying hosts under the same control, but as multicasting spreads it will be preferable to use the more efficient foreign agent routing. By enhancing the draft IETF mobility standard (based on IMHP [5]) to support both methods, as well as a means of switching among them, we can get the best of both worlds at once.

REFERENCES

- [1] F. Kastenholtz and C. Partridge, "Technical criteria for choosing IP: The next generation (IPng)", December 1994, RFC 1726.
- [2] B. Carpenter, "IPng white paper on transition and other considerations", August 1994, RFC 1671.
- [3] P. Bhagwat, C. Perkins, and S.K. Tripathi, "Transparent resource discovery for mobile computers", in *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, 1995, pp. 116–118.
- [4] J. Ioannidis, D. Duchamp, and G.Q. Maguire Jr., "IP based protocols for mobile internetworking", *Computer Communication Review*, vol. 21, no. 4, pp. 235–245, September 1991.
- [5] A. Myles, D.B. Johnson, and C. Perkins, "A mobile host protocol supporting route optimization and authentication", *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 5, pp. 839–849, June 1995.
- [6] S. Deering, "Host extensions for IP multicasting", August 1989, RFC 1112.
- [7] S. Deering, C. Partridge, and D. Waitzman, "Distance vector multicast routing protocol", November 1988, RFC 1075.
- [8] J. Moy, "Multicast routing extensions for OSPF", *Communications of the ACM*, vol. 37, no. 8, pp. 61–66, August 1994.
- [9] A. Ballardie, J. Crowcroft, and P. Francis, "Core based trees (CBT) — An architecture for scalable inter-domain multicast routing", *Computer Communications Review*, vol. 23, no. 4, pp. 85–95, October 1993.
- [10] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C. Liu, and L. Wei, "An architecture for wide-area multicast routing", *Computer Communications Review*, vol. 24, no. 4, pp. 126–135, October 1994.
- [11] V. Jacobson, "Compressing TCP/IP headers for low-speed serial links", Internet Request For Comments, February 1990, RFC 1144.
- [12] V. Chikarmane, R. Bunt, and C. Williamson, "Mobile IP-based multicast as a service for mobile hosts", in *Proceedings of the 2nd IEEE International Workshop on Services in Distributed and Networked Environments*, 1995, pp. 11–18.