Multicast Routing Issues in Ad Hoc Networks*

Katia Obraczka Gene Tsudik

USC Information Sciences Institute
4676 Admiralty Way, Marina del Rey, CA 90292, USA
Voice (310)822-1511 Fax (310)823-6714 {*katia,gts*}*@isi.edu*

Abstract

The advent of ubiquitous computing and the proliferation of portable computing devices have raised the importance of mobile and wireless networking. At the same time, the popularity of group-oriented computing has grown tremendously. However, little has been accomplished to-date in bringing together the technologies for group-oriented communication and mobile networking. In particular, most modern wireless/mobile and ad hoc networks do not provide support for multicast communication. A major challenge lies in adapting multicast communication to environments where mobility is unlimited and outages/failures are frequent.

This paper motivates the need for new multicast routing protocols aimed specifically at fully-mobile (ad hoc) networks. Our premise is that, due to their inherent broadcast capability, wireless networks are well-suited for multicast communication. Unlike the evolution of routing in wired networks, we believe that—in ad hoc networks—it is more effective to treat multicast routing as a separate problem. The paper also identifies outstanding research issues pertaining to multicast routing in mobile and ad hoc networks, and discusses one possible approach to multicast routing and packet forwarding in ad hoc networks.

Subject Category: Mobility, wireless multimedia.
Keywords: Multicast, routing, ad hoc networks.

*This work is being supported in part by a grant from Hughes Research Laboratories (GP3044-97261).

1 Introduction

It is anticipated that a typical internetwork of the future will consist of a wired backbone, and a collection of fixed-infrastructure mobile, and fully mobile (or ad hoc) networks as depicted in Figure 1. Mobile hosts will be connected with the rest of the infrastructure through fixed switches (e.g., *base stations*), and ad hoc networks (AHNs) via satellites or fixed terrestrial switches.

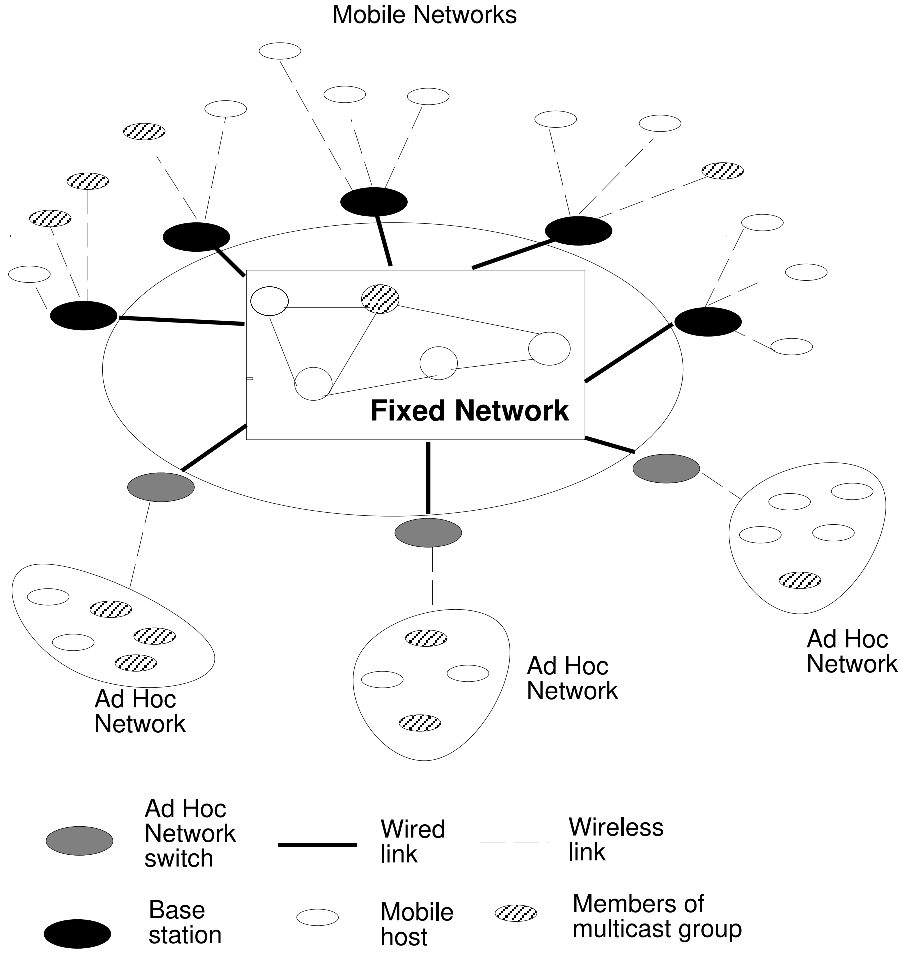

Figure 1: Internetwork of the future.

Regardless of the network environment, multicast communication is a very useful and efficient means of supporting group-oriented applications. This is especially the case in mobile/wireless environments where bandwidth is scarce and hosts have limited power. Example applications include audio- and video-conferencing as well as one-to-many data dissemination in critical situations such as disaster recovery or battlefield scenarios.

We argue that a global multicast solution for a future internetwork such as the one depicted in Figure 1 will have to consist of specialized solutions for each type of network and mechanisms for integrating these solutions. Existing multicast technology will be applied in the fixed portion of the network, while the research challenges are posed by i) multicast in fixed-infrastructure mobile networks as well as AHNs, and ii) inter-operation of different multicast mechanisms.

Multicast communication in mobile networks only very recently became a topic of active research. Although there has been some work in reliable message delivery in networks with mobile hosts [1], little has been done in the area of multicast routing for mobile environments. The main reason appears to be a popular belief that—similarly to the evolution of Internet routing—multicast routing in mobile networks will be built on top of the **unicast** routing infrastructure. For this reason, most research has focused on solving the unicast routing problem with mobile endpoints.

We maintain that, because of their broadcast capability, many types of mobile networks are better suited for multicast, rather than unicast, routing and, that it is more effective to solve the multicast routing problem separately.

Specifically, multicast routing and packet forwarding protocols in AHNs must emphasize the following:

Robustness versus Efficiency. Many multicast routing approaches rely on state in routers to keep track of multicast group members. This, coupled with the high volume of routing information exchanges and slow convergence make traditional multicast approaches untenable in highly dynamic AHNs composed of anemic (low-power, low storage capacity) hosts. Therefore, new techniques that stress rapid and robust delivery must be developed.

Active Adaptability. Hosts will migrate freely among ad hoc, fixed-infrastructure mobile, and wired networks. In order to adapt rapidly to infrastructure changes, an active networking approach can be employed: hosts can adapt in real-time by downloading appropriate multicast mechanisms.

Unlimited Mobility. Some existing multicast solutions are geared towards *discrete* mobility whereby periods of movement are interspersed with periods of rest. Others assume certain limits on direction, speed and number of simultaneously moving hosts. In contrast, we stress *universal*, unlimited mobility of *all* network components.

Integrated Multicast. Multicast solutions for AHNs will most likely differ substantially from those for fixed networks (one of the main reasons is the marked difference in transmission rates). In order to offer seamless and integrated multicast service, novel mechanisms must be developed for inter-operation of fixed and wireless multicast solutions.

In light of these requirements, this paper identifies the outstanding issues pertaining to multicast routing in fixed-infrastructure mobile and ad hoc networks. It also discusses an approach to multicast routing and packet forwarding for AHNs.

2 Mobile Networks and Mobility Patterns

The research issues raised by multicast routing in mobile environments are closely related to the type of mobile network in use. In this paper, we discuss multicast routing in the context of:

- Mobile networks with fixed infrastructure
- Ad hoc networks
- Internetworks comprising wired, fixed infrastructure, and ad hoc networks

Mobility patterns add another dimension to the problem of multicast routing in mobile networks. Table 1 illustrates sample scenarios for each combination of mobile infrastructure type (rows) and mobility patterns (columns). Mobility patterns are an important factor in the design of mobile multicast routing protocols. For instance, changes in topology occur much more frequently when mobility is continuous than when it is discrete.

	Mobility Type	
Infrastructure	Discrete	Continuous
Fixed	Traveling users	Wireless LAN
Ad Hoc	Impromptu conferencing	Battlefield

Table 1: Mobility patterns: sample scenarios.

3 Problem Scope

The main focus of this paper is on multicast routing in AHNs since we believe it to be a truly challenging and little-explored area. Research issues in AHNs are discussed in detail in Section 4. In contrast, multicast routing in fixed-infrastructure mobile networks presents only a few engineering obstacles which mainly have to do with the last hop delivery and membership tracking. This topic is addressed in Section 8.

This paper also discusses issues raised by providing an integrated multicast architecture for internets consisting of wired, fixed-infrastructure mobile, and AHNs. The goal is to provide seamless multicast service allowing multicast groups to span networks of different types. We address research issues in integrated multicast over mixed-media internets in Section 7.

4 Issues in Multicast Routing in AHNs

AHNs refer to (for the most part, wireless) networks where all network components are mobile. In general, there is no real distinction in an AHN between a host and a router since all network hosts can be endpoints as well as forwarders of traffic.

Most research in the area of routing for AHNs has concentrated on routing for unicast communication. Notable examples include: the Monarch project [11], the TORA algorithm [16], and the MMWS project [18]; they are reviewed in Section 9.1 below.

Multicast routing and packet forwarding in AHNs is a fairly unexplored area [18]. One recent exception is the Shared-Tree Wireless Network Multicast (ST-WIM) work at UCLA [3] which aims to adapt fixed-network multicast approaches (PIM Sparse Mode) [6] to AHNs. Our view is that, since fixed network multicast routing is based on state in routers (either hard or soft), it is fundamentally unsuitable for an AHN environment with unconstrained mobility.

In this context, *unconstrained mobility* implies the following:

- Host behavior completely independent of other hosts.
- No limit on host speed.
- No constraints on direction of movement.
- High probability of frequent, temporary network partitions.

We claim that all these factors no longer make it worthwhile for a mobile host to maintain any multicast-related state information other than its own. Furthermore, in many types of AHNs (e.g., where hosts are hand-held devices) **both storage capacity and power are severely limited**. This is yet another reason to avoid maintaining and exchanging multicast state.

Also, frequent changes in the topology make it difficult to apply clustering algorithms: heavy state maintenance and frequent elections of cluster leaders are too expensive for low-power hosts.

One other important consideration has to do with the mission of AHNs. AHNs are most often deployed in military (e.g., battlefield) and other emergency (e.g., disaster recovery) situations. In such critical environments **robustness** and **high quality-of-service** are of paramount concern. Consequently, multicast mechanisms (however attractive otherwise) that cannot provide the highest delivery guarantees are not appropriate.

This raises a fundamental issue:

> *if state-based multicast is not well-suited for AHNs, what is?*

5 A Potential Approach

One way of solving the multicast routing problem in the context of AHNs is to use a specialized form of *flooding* tailored towards unlimited mobility. We call it **hyper-flooding** since hosts may re-broadcast a packet. The prime emphasis of hyper-flooding is on **reliable delivery** and **minimal state retention**. Its salient features are:

- Maximum network diameter is assumed to be known.
- Maximum TTL (time-to-live) for a packet is known.
- Multicast packets carry hop counts as well as origination timestamps.
- Hop count is decremented each time a packet is re-broadcast (forwarded).
- Packets with hop count of zero or expired timestamps are not forwarded.
- Each host only keeps its own multicast group membership information.
- Every multicast packet is assigned a unique ID computed as a function of, at least: source address, current time, group ID, sequence number.
- Each host maintains a cache of IDs for recently processed multicast packets (state is only kept for packets with timestamps within max TTL).
- A packet that has already been forwarded can be re-forwarded by an intermediate host if that host has acquired at least one new neighbor since the last time the same packet has been forwarded.
- If a host is not aware of any new neighbors, the packet is not forwarded.

It is easy to see that hyper-flooding is not state-free. It requires each host to keep track of its current neighbors. However, neighbor awareness and discovery is something that many wireless/mobile technologies already include at either (or both) MAC and network layers. Thus, it is not a new requirement specific to multicast. Moreover, keeping track of neighbors is certainly much simpler than keeping track of group memberships for both neighbors and non-neighbors. The latter is exactly what hyper-flooding tries to avoid.

There is still the matter of keeping state related to the actual multicast packets. The purpose is to prune redundant broadcasts of the same multicast packet when no new neighbors are acquired in the interim. This technique is expected to be only partially effective since, in certain circumstances, redundant broadcasts will still occur. Consider, for example, what happens if a certain host's neighbor becomes unreachable for a short time and then re-appears. Now, if a re-broadcast of an old multicast packet is received by the host in question, it will be re-broadcast since, technically, a new neighbor has been acquired. While this is clearly suboptimal, little can be done to prevent it (barring having to maintain a snapshot of all neighbors for each cached multicast packet ID).

Whenever a host receives a previously seen multicast packet and, in the meantime, some new neighbors have been acquired, these neighbors must be identified in the packet when it is re-broadcast. If this is not done, broadcast storms can occur. Consider the situation in Figure 2. At time $T1$, packet $ID = x$ arrives. At time $T2$, a copy of the same packet x arrives. In the mean time, host 1 acquired a new neighbor, host 6. If host 1 were to re-broadcast packet x as is, each of its other neighbors (hosts 2-5) who have also acquired host 6 as a new neighbor, would re-broadcast the same packet. This would result in at least four unnecessary broadcasts. Instead, host 1 can inform its neighbors that packet x is being re-broadcast of the newly-acquired neighbor host 6. This new neighbor information will match the remaining neighbors' and will prevent further superfluous broadcasts.

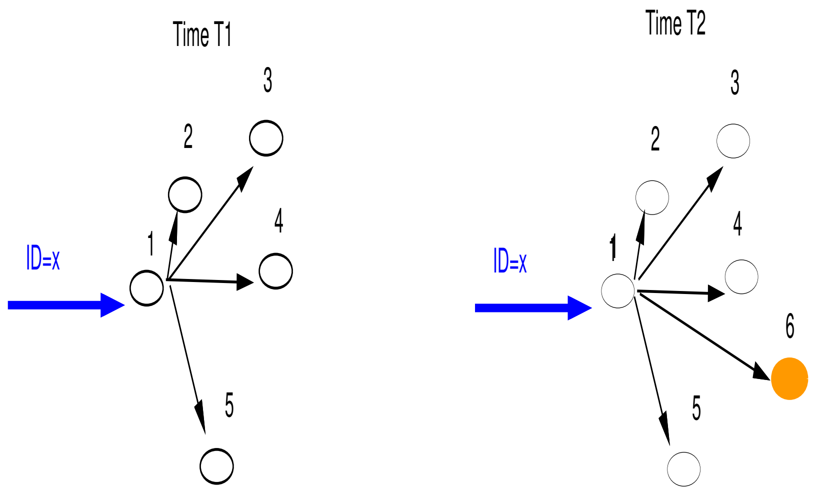

Figure 2: Pruning unnecessary re-broadcasts.

Alternatively, suppose that a host acquires no new neighbors in the interval between receiving two successive copies of the same multicast packet. Ideally, this should mean that there is no need to re-broadcast the packet the second time. However, it is possible that one or more of the host's neighbors have themselves acquired new neighbors. An example is illustrated in Figure 3. Hosts 3 and 5 acquire new neighbors 6 and 7, respectively. Host 1, however, is not aware of hosts 3 and 5 new neighbors and does not re-broadcast the packet. The end-result is potential unreliability since packets may not reach certain hosts (note that hosts 6 and 7 may get packet x from some other neighbors).

Thus, there is a clear tradeoff between robustness and efficiency: while re-broadcasting a packet will lead to improved reliability, it also involves transmission of unnecessary duplicates, and, consequently, a waste of both bandwidth and power.

It is easy to see that hyper-flooding is a fairly simple and intuitive mechanism. On the other hand, it is not, in and of itself, a complete solution to AHN multicast routing. Optimization of hyper-flooding algorithms to balance robustness and resource consumption is an important area for further study. An important component thereof will be the use of simulation in order to investigate the performance and general behavior of hyper-

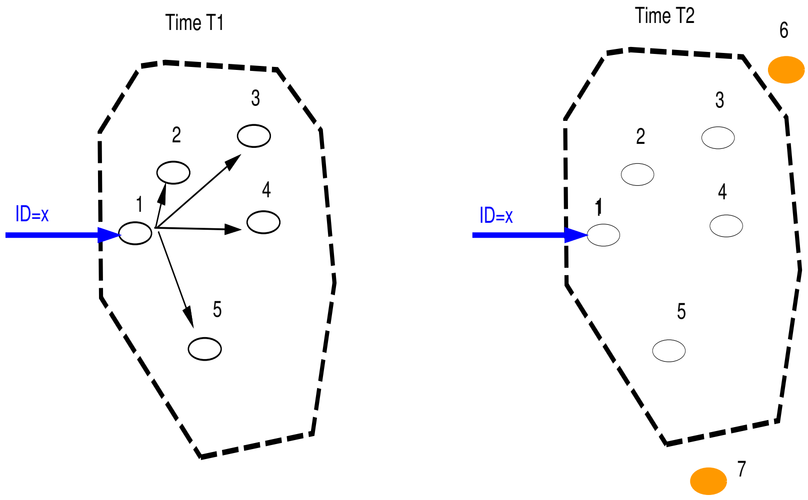

Figure 3: Incomplete coverage.

flooding in different AHN settings. Simulation is key to studying the algorithm for different types of multicast group organizations (e.g., sparse or dense groups, groups that span large geographic areas or have most members concentrated in a limited area), mobility and user access patterns. It will also facilitate the evaluation of different hyper-flooding flavors and tradeoffs (e.g., among the amount of state kept by hosts, protocol robustness and efficiency).

6 Accommodating Different AHNs

We claim that there is no single multicast routing solution for all types of AHNs. (Due to the sheer diversity of AHNs, no one approach can be a *panacea*.) We anticipate that relatively stable AHNs with few host failures, outages and infrequent movements will lend themselves to approaches different from those best-suited for highly dynamic and volatile AHNs. Consequently, one of the challenges is in determining the best multicast approach when faced with a specific AHN configuration and parameters.

More concretely, we envisage a suite of solutions each geared towards (and nearly optimal in the context of) a specific AHN type and endpoint mobility pattern. This is very much in line with multicast development in the fixed networks such as the Internet. The two modes of Protocol-Independent Multicast (PIM) [6] (sparse and dense modes) are a case in point.

Large, Widely Distributed AHN: An Example

It appears highly unlikely that vanilla hyper-flooding will work well in populous and geographically distributed AHNs as it is primarily geared towards small and very dynamic AHNs. On the other hand, an AHN comprised of a large number of hosts spanning large geographical areas would lend itself to some kind of a hierarchy. This is especially applicable in military scenarios which are inherently hierarchical. If hosts in an AHN are partitioned into clusters as, for example, in BBN's MMWS (see Section 9.1), a two-tiered approach to multicast can be envisioned, as shown in Figure 4:

- The AHN is partitioned into clusters however hosts within a cluster are not necessarily one-hop away (i.e., a cluster is still a multi-hop AHN).
- Intra-cluster mobility is arbitrary but inter-cluster migrations are relatively rare.
- Each cluster is assigned (or elects) a clusterhead.
- Intra-cluster (tier 1) multicast is performed with hyper-flooding or a variation thereof.
- Inter-cluster (tier 2) multicast is accomplished by traditional multicasting, or even unicasting, among clusterheads. Clusterheads maintain state (binary state, not per host) regarding multicast group membership of their constituent hosts.

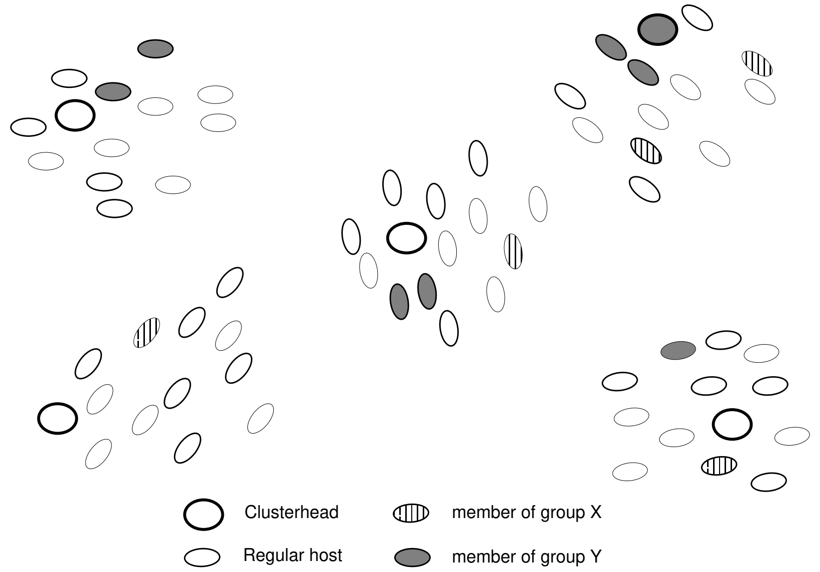

Figure 4: 2-tiered Multicast.

Although, in the end, only *hard-core* simulations and other types of experiments can confirm our intuition, the above approach appears both natural and workable.

7 Integration and Seamless Operation

Once the appropriate multicast solutions for both mobile and ad hoc networks are identified, analyzed and evaluated, the next challenge is their integration with (fixed) Internet multicast mechanisms to produce an integrated multicast architecture. The goal, of course, is to provide a seamless multicast service whereby a single multicast group can span all network types (fixed, mobile, AHN). This would let a given host to partake in multicast communication regardless of the currently underlying network type. Specifically, this integration translates into:

- Design of mobile $\leftrightarrow$ wired multicast gateway.

- Design of ad-hoc $\leftrightarrow$ wired multicast gateway.
- Mechanisms for active, on-the-fly switching among different multicast routing mechanisms as a host changes the network type it is part of.

We anticipate the last two to be truly challenging. There is little or no experience in the research community in interconnecting different multicast algorithms and protocols. (Albeit, some proposals have been recently floated in the IETF [8, 20, 13].) We expect the difference in transmission speeds (between AHNs and wired networks) to be one of the greatest obstacles. One important consequence thereof is the need to buffer multicast packets in fixed $\rightarrow$ ad-hoc gateways.

Furthermore, in order to provide connectivity for freely roaming hosts, mechanisms for active, real time adaptation to the underlying environment can be employed. This is a difficult undertaking since **limited capabilities of a typical mobile host conflict with the need to carry, and switch among, multiple multicast mechanisms.** Our preliminary analysis leads us to consider a somewhat unorthodox solution:

> *A mobile host carries only one multicast mechanism at a time which corresponds to the type of network the host is currently part of. As the underlying environment changes, new mechanisms are downloaded dynamically.*

There are several outstanding research issues raised by this "active" approach:

- How does a host detect when it needs to download a new routing mechanism, i.e., how does it detect changes in the underlying environment that should trigger the download and activation of a new multicast routing protocol?
- How does a host determine where to obtain a new multicast routing protocol? Also, what strategy is used to designate *carrier* hosts, i.e., those carrying multiple protocols?

8 Mobility with Fixed Infrastructure

Mobile networks with fixed infrastructure, or cellular networks, consist of stationary switches and mobile endpoints. Every switch, or *base station*, is assigned a geographic area, or *cell* and is responsible for connecting endpoints in that cell to the wired portion of the network.

Intuitively, this kind of environment presents much less of a challenge for multicast than AHNs. We claim that mobility in fixed infrastructure raise engineering, rather than research, issues. The fixed portion of the network including the base stations will use the existing multicast routing infrastructure. Multicast routing between the base stations and the mobile endpoints is what needs to be addressed. There are two possible scenarios: one where mobile hosts are receivers in the multicast group, i.e., mobile endpoints are at the leaves of the multicast distribution tree, and the other—where the multicast sender is mobile, i.e., mobile host is at the root of the multicast tree.

8.1 Mobile Receivers

This case is quite straightforward since mobility is encountered only at the last hop. Existing protocols like Mobile IP (IPv4 and IPv6) [17, 10] can be used provided that they are extended to handle multicast communication. (We should point out that in the mobile IP world, mobile endpoints are not constantly on the move like in cellular networks; typically, mobile users take their laptop when they travel, plug it in for the time they are away and move on.) For instance, a base station (Foreign Agent or FA in Mobile IP terminology) needs to join every multicast group which has members currently in its cell. By not refreshing its entry in the immediate upstream multicast router, a base station leaves a multicast group when the corresponding members move to another cell. Leaving a multicast group should be done gracefully so that data is not lost during hand-off. However, only a limited amount of buffering can be accomplished by a base station to accommodate continuous movement patterns, where a mobile host goes from one cell to the next. In the case of a host leaving the network and subsequently re-appearing somewhere else, recovering from data loss needs to be addressed by higher-level protocols.

8.2 Mobile Senders

Mobile senders raise some issues with respect to addressing. The main concern is whether a mobile host should keep a single constant address or change its address each time it moves. If the mobile sender's address does not change, traffic originated at the sender will have to be tunneled to the sender's home (specifically, to its home agent) and from there, it will actually be multicast to the group. In Mobile IP and cellular networks (e.g., PCS), senders change address depending on their location. Data can thus be multicast directly. Thus, every time a sender moves, it creates an appearance of data coming from a different host. In case of source-rooted multicast distribution trees, every time a multicast sender moves, the corresponding multicast tree needs to be re-built anew. This is probably acceptable in discrete moving patterns. However, when a host is continuously moving, building the multicast tree every time is inefficient. On the other hand, in cellular networks, typically only the first hop changes. Therefore, instead of building source-specific multicast trees, routing using group-specific trees (e.g., core-based tree multicast routing [2]) will likely work well.

Changing sender's addresses may also raise higher-level issues such as how to piece together data sent by the same host while moving. Other concerns include group security and membership. (We are assuming that, as in IP multicast, senders do not have to be members of a multicast group to which they are sending data.)

9 Related Work

In this section we review previous and on-going work in related areas, including: unicast and multicast routing in AHNs, multicast routing in fixed-infrastructure mobile networks, and integrated multicast in wired internetworks.

9.1 Unicast Routing in AHNs

Most research in the area of AHN routing has focused on unicast communication. This is probably motivated by the evolution of Internet routing, where multicast routing service was built on top of the existing unicast routing infrastructure. In this section, we review several research efforts in mobile unicast for AHNs.

The routing algorithm proposed in the Monarch project [11] is based on source routing (similar to that used in IP options). Each mobile endpoint maintains a route cache where they keep previously learned source routes. When an endpoint needs to send a packet, it performs a route cache lookup for the corresponding destination. If there is no route known for that destination, the source may attempt to discover one using the route discovery protocol. Route discovery floods *route requests* packets containing the requesting host ID and the route's destination. Successful route requests result in *route replies* containing the requested source route.

The Wireless Routing Protocol (WRP) [14] is a path-finding based routing algorithm target at reducing temporary loops and routing table updates. The latter feature is particularly attractive in mobile network environments, where topology changes are quite frequent.

The Temporally-Ordered Routing Algorithm (TORA) [16] is another recently developed routing solution for AHNs. TORA's approach is based on building a directed acyclic graph (DAG) rooted at the destination. Besides route creation, TORA also implements route maintenance and route deletion in the case of partitions.

The two-level hierarchical unicast routing algorithm for AHNs presented in [5] uses a *spine* for intra-cluster routing and a link-state algorithm for routing among clusters. For each cluster, a spine is computed using a minimum connected dominating set (MCDS) algorithm. Shortest-path routes are computed between the spine and other cluster members. For shortest-path route computation, spine nodes gather the AHN's topology information and keep track of topology changes.

The Mobile Multimedia Wireless Network (MMWS) project provides link- and network-layer communication services to support real-time, distributed multimedia applications in large-scale, wireless networks. It assumes is a hierarchical network architecture, where endpoints are grouped in cells around cellheads or switches. These switches form clusters, each of which functions as a multihop packet radio network. Clusters, in turn, can form higher-level clusters. MMWS's routing is of link state variety. For control overhead scalability, a hierarchical approach is employed whereby routing details are confined to a given level in the hierarchy. Thus, route optimality is traded off for scalability.

In summary, existing unicast routing algorithms for AHNs were not designed with multicast extensions in mind. Therefore, they do not naturally lend themselves to multicast routing solutions. For example, the scalability properties of TORA's destination-rooted DAG or Monarch's source-routing approaches are questionable for multicast routing in AHNs.

9.2 Fixed-Infrastructure Mobile Networks

Support for Host Mobility.

As an alternative to Mobile IP [17, 10], Mobility Support using Multicasting in IP (MSM-IP) [15] supports host mobility using IP multicast. Since IP multicast already addresses the problem of location-independent host addressing, MSM-IP makes use of the IP multicast infrastructure for addressing and routing packets to mobile hosts. The basic approach is to assign a unique multicast address to a mobile, and as the host migrates it registers with the current network's multicast router in order to receive packets.

Multicast for Mobile Hosts.

As mentioned in Section 8 above, the issues raised by multicast routing in mobile networks with fixed infrastructure are confined to the first and/or last hops depending on whether the source and/or receiver(s) are mobile. The work described in [4] uses Mobile IP tunnels to solve the first/last hop problem associated with endpoint mobility in multicast delivery. To handle multicast source mobility, the mobile source establishes a tunnel to its home agent. In the case of receiver mobility, the home agent tunnels multicast packets to mobile receivers through the corresponding foreign agents.

9.3 Multicast Routing in AHNs

Little work has been done in the area of multicast routing in AHNs. The only current research effort in the field is the Shared-Tree Wireless Network Multicast (ST-WIM) algorithm [3]. The ST-WIM approach is based on adapting PIM's sparse-mode algorithm [6] to AHNs. As stated above, our view is that, since fixed network multicast routing is based on state in routers (either hard or soft), it is fundamentally unsuitable for an AHN environment with unconstrained mobility. This is confirmed by ST-WIM's own simulation results, which show that the performance of both hard- and soft-state multicast tree maintenance mechanisms degrade rapidly with increased mobility. (Host mobility was simulated by having hosts roam randomly with a preset average speed.)

9.4 Other Related Work

On a different front, the need for integrated multicast service in fixed (wired) internetworks has been recently recognized. Consequently, interoperability mechanisms to allow current multicast routing protocols to coexist in an internetwork have been proposed. For example, PIM-SM to DVMRP gateways [8] allow the interconnection of PIM-SM domains to a DVMRP backbone. Multicast routing protocol interoperability issues are discussed in [20, 13]. We intend to leverage this work in our integration efforts.

Another related topic is the interconnection of AHNs and fixed networks. To this end, the Wireless Internet Gateways (WINGS) [9] project at UC Santa Cruz developed wireless IP routers for connecting AHNs with the fixed IP Internet. Each WING router runs WIRP (an extension of RIP and RIPv2) that effectively handles frequent topology changes. We note that multicast is not addressed by WINGS.

10 Summary and Research Directions

There is significant work to be done in the area of multicast routing for AHNs and in providing a seamless integrated multicast routing architecture across mixed-media internets. We envision this research being carried out as follows:

- **Exploration of the design space of multicast routing algorithms suitable for AHNs**. In this phase, simulations will be useful in examining the various tradeoffs and alternatives. This would allow different multicast routing protocols (including hyper-flooding and state-based schemes) to be contrasted by varying a number of parameters such as traffic patterns and group membership dynamics. Moreover, different variations of hyper-flooding can be evaluated for different types of AHN environments and endpoint mobility patterns.

  Simulations can conducted in the context of the Network Simulator (ns) [12], originally developed at LBNL, and currently being extended as part of the VINT project [7] at ISI. Another candidate simulator is UCLA's Maisie network simulator [19]. One of the challenges raised by simulating mobile networks in general, and AHNs in particular, is defining "typical" AHN behavior, endpoint mobility patterns, multicast group membership dynamics, and end-user workloads. We anticipate that establishing these simulation parameters will be an important contribution.

  Another significant contribution of a thorough simulation study will be a multicast routing evaluation framework for AHNs, which will include a set of performance metrics, e.g.:

  - Robustness measured by the total number of packets that were not received by all group members.
  - Efficiency measured by the total number of duplicate packets received.
  - State retention by hosts.
  - How long it takes for a multicast packet to reach all group members.

  The evaluation framework must also include a set of variables that form the design space of mobile multicast routing protocols. The parameters that can drive the simulations include: endpoint mobility patterns, AHN mobility dynamics, multicast group membership dynamics and coverage as well as traffic patterns reflecting end-user workload. A parallel goal is to determine boundary conditions, i.e., conditions under which one routing approach ceases, and another starts, being optimal. The next step is to identify and piece together the sub-components of the various routing approaches that will form a comprehensive multicast routing framework for AHNs.

- **Development of the multicast gateway for inter-connecting wired network multicast with AHN-based multicast**. The major challenge here lies not just in the interconnection of multiple disparate multicast mechanisms but in dealing with different transmission rates and reliability/reachability factors. This stage also needs to encompass the development of mechanisms for **active integrated multicast** to allow on-the-fly switching among different multicast protocols as a host moves from one network type to another.

- **Prototyping and trial deployment**. The prototype of the AHN multicast routing framework will have to be evaluated through live experiments. These, in turn, will require an AHN testbed that will include mobility and workload parameters listed above. The prototype will consist of the AHN multicast routing framework and the mechanisms for its inter-operation with wired multicast routing. The AHN testbed will need to accommodate both wired and fixed infrastructure mobile networks. This mixed-media network testbed will enable the evaluation of the integrated multicast service in terms of functionality and seamlessness.

References

[1] A. Acharya and B.R. Badrinath. A framework for delivering multicast messages in networks with mobile hosts. *In ACM/Baltzer Journal of Mobile Networks and Applications*, 1(II):199–219, 1996.

[2] A.J. Ballardi, P.F. Francis, and J. Crowcroft. Core based trees (cbt). *Proc. of the ACM SIGCOMM'93*, August 1993.

[3] C.C. Chiang, M. Gerla, and L. Zhang. Shared tree wireless network multicast. *Proceedings of the IEEE 6th International Conference on Computer Communications and Networks (ICCCN'97)*, September 1997.

[4] V. Chikarmane, C.L. Williamson, and R.B. Bunt. Providing multicast for mobile hosts in a tcp/ip internetwork. *Proceedings of the ACM Mobicom97*, September 1997.

[5] B. Das and V. Bharghavan. Routing in ad-hoc networks using a spine. *Proceedings of the IEEE 6th International Conference on Computer Communications and Networks (ICCCN'97)*, September 1997.

[6] PIM development group. Protocol independent multicast. Available from http://netweb.usc.edu/pim/, November 1997.

[7] D. Estrin et. al. Virtual internet testbed (vint): Methods and system. White Paper, 1996.

[8] D. Estrin et. al. Pim border router specification for connecting pim-sm domains to a dvmrp backbone. Internet draft draft-ietf-idmr-PIM-SM-spec-09.ps, October 1996.

[9] J.J. Garcia-Luna-Aceves, C.L. Fullmer, E. Madruga, D. Beyer, and T. Frivold. Wireless internet gateways (wings).

[10] D. Johnson and C. Perkins. Mobility support in IPv6. Work in progress, draft-ietf-mobileip-ipv6-02.txt, November 1996.

[11] D.B. Johnson and D.A. Maltz. Dynamic source routing in ad hoc wireless networks. *Mobile Computing, edited by Tomasz Imielinski and Hank Korth, Chapter 5, Kluwer Academic Publishers*, pages 153–181, 1996.

[12] S. McCanne. ns - LBNL network simulator. Available from http://www-nrg.ee.lbl.gov/ns/.

[13] D. Meyer. Some issues for an inter-domain multicast routing protocols. Internet draft draft-ietf-mboned-issues-03.txt, November 1997.

[14] S. Murthy and J.J Garcia-Luna-Aceves. An efficient routing protocol for wireless networks. *ACM Mobile Networks and Application Journal, Special Issue on Routing in Mobile Communications Networks*, 1(2), 1996.

[15] J. Mysore and V. Bharghavan. A new multicasting-based architecture for internet host mobility. *Proceedings of the ACM Mobicom97*, September 1997.

[16] V.D. Park and M.S. Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. *Proceedings of the IEEE Infocom 1997*, April 1997.

[17] C. Perkins. Ip mobility support. Internet Request for Comments RFC 2002.

[18] S. Ramanathan and M. Steenstrup. A survey of routing techniques for mobile communications networks. *ACM/Baltzer MONET*, pages 89–104, October 1996.

[19] Maisie Project Team. Maisie parallel simulation home page. Available from http://may.cs.ucla.edu/projects/maisie/, November 1997.

[20] D. Thaler. Interoperability rules for multicast routing protocols. Internet draft draft-thaler-interop-00.ps, November 1996.