Multicast: Wired to Wireless

Hrishikesh Gossain, Carlos de Morais Cordeiro, and Dharma P. Agrawal, University of Cincinnati

ABSTRACT

Recently there has been an increasing demand for applications like multiplayer online gaming, where players residing in different parts of the world participate in the same gaming session through the Internet. Multicasting could prove to be an efficient way of providing necessary services for these applications. Furthermore, with increasing popularity of handheld devices and mobile equipment, it is imperative to determine the best way to provide these services in a wireless environment. Due to very diverse requirements, it is necessary to investigate and discern the applicability of existing multicast protocols and qualify which is more suitable for which types of applications. This article provides a detailed description and comparison of IP-based wired and wireless multicast protocols. We hope that the discussion presented here will be helpful to application developers in selecting an appropriate multicast protocol for their specific needs.

INTRODUCTION

Imagine a scenario where a professor wants to conduct a real-time class with 50 students participating through the network. If the multimedia application for the conference employs unicasting, the professor's computer repeatedly sends out 50 audio streams to the students' computers. Unicasting wastes bandwidth because it sends 50 duplicate copies over the network, and causes a significant delay before the last student hears the professor. The audio stream could also flood every corner of the network and possibly bring the network down. Multicasting comes to the rescue by allowing the multicast host to send out only one copy of the information, and only those hosts that are part of that group receive it. In the class example, the professor's computer sends only one audio stream to the network, and only the targeted group of 50 students receive the stream. The information utilizes the minimum required network bandwidth and arrives at every student's computer without any noticeable delay.

This application is an example of the practical use of multicast in everyday life. The same is true for other applications like audio/videoconferencing, multiplayer online gaming, online/offline video distribution, news, and so on. As illustrated in Fig. 1, it is clear that even if there are only three receivers of a multimedia application, the bandwidth utilization between routers can be roughly reduced up to one-third if we use multicasting.

The concept of multicast was introduced by Steve Deering [1] in the late '80s followed by a widescale test during the Internet Engineering Task Force (IETF) "audiocast" in 1992. Adding multicast to the Internet does not alter the basic model of the network. Any host can send multicast data, but with a new type of address called a *host group address*. IPv4 has reserved class D addresses to support multicasting. A user can dynamically subscribe to the group to receive multicast traffic by informing a local router that it is interested in a particular multicast group. However, it is not necessary to belong to a group to send multicast. The delivery of multicast traffic in the Internet is accomplished by creating a multicast tree, with all its leaf nodes as recipients.

When providing multicasting to wireless hosts, the maintenance of this host grouping or the multicast delivery tree becomes a major issue. We have to focus not only on dynamic group membership, but also monitor the host movement as the mobile host moves to different cells or different service providers, which makes determining and maintaining the "optimal" multicast delivery tree very difficult. Reliability is another major issue in wireless, since a host may miss some packets, or may receive some duplicate packets due to host movement.

With this in mind, our goal here is to provide in-depth insight in the multicast research area, emphasizing its pros and cons. We first give an overview of wireline multicast followed by wireless multicast protocols. Note that our focus here is on solutions that make use of infrastructure provided by the wired network and not over infrastructureless ad hoc networks.

MULTICAST SUPPORT

Initial support of multicast in the Internet is done by adding multicast-capable routers (mrouters) and using *dedicated tunnels* to facilitate multicasting packets from one mrouter to another. The job of each mrouter is to encapsulate and decapsulate each multicast packet as a regular Internet Protocol (IP) packet and send it through the tunnel to

0163-6804/02/$17.00 © 2002 IEEE

another mrouter. This set of mrouters in the Internet is called MBone (multicast backbone) [2]. Presently, some of the existing Internet routers have been enhanced to support multicast, and there is no need to set up dedicated tunnels for them. This is called *native multicast*, and the Internet currently has a combination of both.

Multicast Group Membership

IP supports dynamic joining and leaving of a group by Internet Group Management Protocol (IGMP) [3]. A user wishing to join a multicast group sends an IGMP join message to its neighboring multicast router. If the multicast router is not a member of the group, it forwards the message upstream until it finds some router or source that is currently subscribing to the group. In a domain, if there are routers performing IP multicasting, one of them is elected as the multicast querier to control the multicasting in the domain. The multicast router also keeps track of current membership of any group by sending periodic messages to all hosts in the domain.

IGMP version 2 (IGMPv2) defines the procedure of selecting a multicast router for each LAN. A router with the lowest IP address is selected as the multicast querier. To reduce leave latency, each host is supposed to send a Leave Group Message to the multicast querier so that the multicast querier stops forwarding messages for that group if the host is the last one leaving the subscribed group in its domain. IGMPv3 adds the feature of selective group join and selective group leave.

Packet Forwarding

To provide Internet-wide delivery of multicast packets, it is necessary to develop a multicast forwarding algorithm, which is responsible for constructing a multicast delivery tree and forwarding of multicast packets. This section provides an overview of different algorithms employed by multicast routing protocols.

The simplest way of providing multicast in the Internet is through flooding. If a multicast router receives a multicast packet for the first time, it forwards this packet to all the outgoing interfaces except the one from which it receives. Although simple, this is very inefficient in terms of network bandwidth utilization. A better solution is to select a subset of Internet topology and build a spanning tree where a multicast router could forward multicast packets to all the interfaces that are part of a multicast tree except the source. Multicast forwarding algorithms can be classified into two categories: source-based and core-based.

In a source-based multicast forwarding algorithm, the multicast tree is constructed using the source of the multicast session as the root of the tree. There are three source-based multicast forwarding algorithms: Reverse Path Broadcasting (RPB), Truncated Reverse Path Broadcasting (TRPB), and Reverse Path Multicasting (RPM). In RPB, rather than creating a spanning tree for the complete Internet topology, a spanning tree is created for each (source, group) pair. If there are many potential sources for a group, a different spanning tree is constructed for each active (source, group) pair. RPB does not consider group membership; hence it forwards packets even if there are no members in the domain. TRPB addresses the issue of group membership by using IGMP, which allows a multicast router to determine if there is any member for a group in its domain. If not, it sends a prune message upstream to disassociate itself from the multicast delivery tree. TRPB solves some of the limitations of RPB by preventing forwarding of multicast packets to leaf routers, while packets are still being forwarded to non-leaf routers. RPM is an extension of both RPB and TRPB. In RPM, the spanning tree spans only to routers and subnetworks with group members. RPM allows a source-rooted multicast tree to get pruned so that a multicast packet gets forwarded only to those routers that could lead to its members.

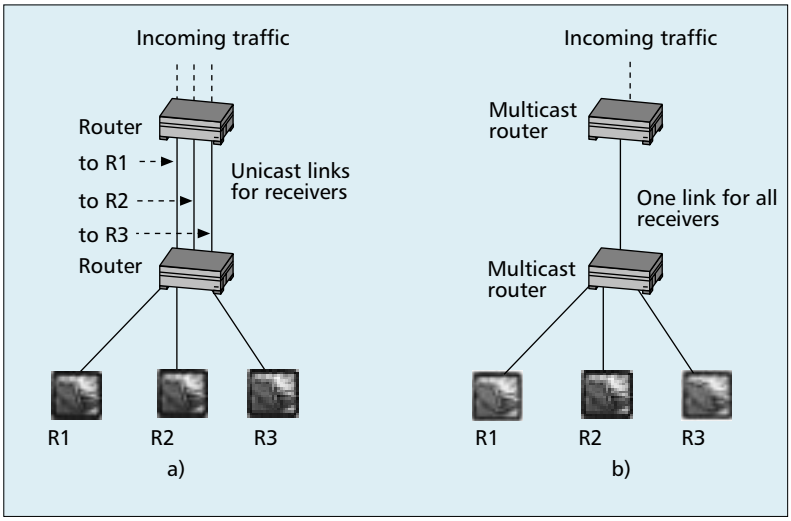

Figure 1. *Delivery of information using multiple unicasting and multicasting.*

Core Based Tree (CBT) is based on building a single shared multicast tree for each multicast group, with a set of multicast routers as a core for forwarding packets. All packets are forwarded through the core. The construction of the multicast tree is receiver-oriented. To join a multicast group, a user sends an explicit join toward the core of the group. This message is propagated upstream until it reaches one of the routers presently subscribed to that group. CBT efficiently handles the problem of scalability by building only one multicast tree for each group. However, the tree built may be suboptimal and may also lead to concentration of load on the core.

Multicast in the Internet

Layer Three (Routing) Protocols

Distance Vector Multicast Routing Protocol (DVMRP): DVMRP is based on a distance vector routing protocol. A recent version of DVMRP uses the RPM algorithm to forward multicast packets. If the time to live (TTL) and router interface permit, the first packet for every (source, group) is forwarded to the entire internetwork. If the leaf routers have no members presently belonging to that group, they send an explicit prune message back toward the source. Thus, a source-based multicast forwarding tree is created with all leaf routers as members for that

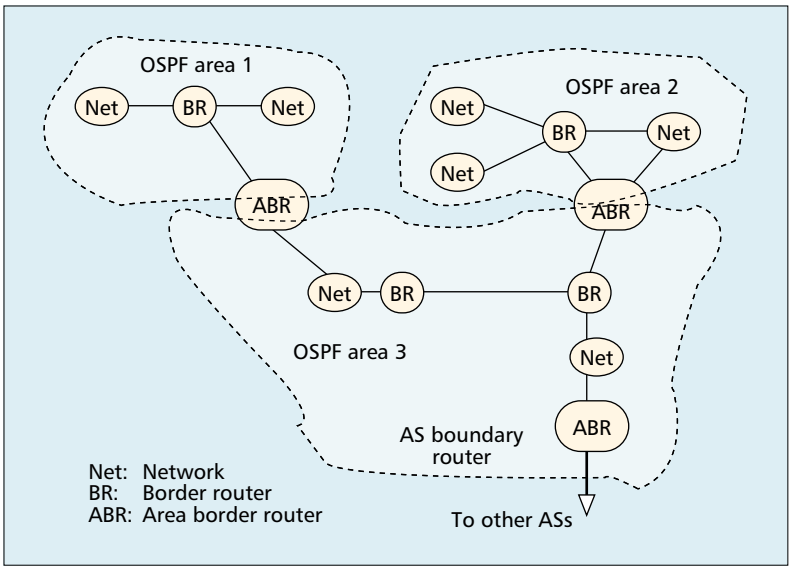

Figure 2. *Illustration of an autonomous system.*

group. DVMRP also supports a *graft* message through which a router that has previously sent a prune message may rejoin the group.

DVMRP treats MBone as a "flat" routing domain. With the growth of MBone, the routing table and periodic updates for every subnetwork became very large. To solve this problem, "hierarchical DVMRP" has been proposed, wherein MBone is divided into individual routing domains. Each domain runs its own multicast routing, while DVMRP is employed for interdomain routing.

Multicast Extensions to OSPF (MOSPF) — MOSPF is an extension of Open Shortest Path First (OSPF). With the help of the group membership obtained from IGMP and unicast routing information obtained from OSPF, MOSPF builds a multicast forwarding tree on demand for each (source, group) pair. MOSPF supports hierarchical routing where hosts are partitioned into *autonomous systems* (ASs) (Fig. 2). Within an AS, Interior Gateway Protocol (IGP) is used to distribute unicast routing information. Based on the address range, an AS is further split into OSPF areas that help border routers to identify every single node in the area. The concept of OSPF area is similar to subnetting in IP networks. Routing can be limited to a single OSPF or could cover multiple OSPFs as discussed below.

Intra-area routing using MOSPF protocol runs in a single OSPF area and supports multicast when the source and the multicast receivers are in the same OSPF area, or the entire AS is a single OSPF area. Each MOSPF router maintains a local group database that contains a list of directly attached group members. Each subnetwork consists of a designated router (DR), which is responsible for sending IGMP host membership queries and listens to the IGMP host membership reports. A DR propagates this group membership information to all other routers in the OSPF area by a group membership link state advertisement (LSA). Based on the router LSAs and network LSAs in the MOSPF link state database, a source-rooted shortest path tree is constructed. The source-rooted shortest path tree and the router's local group database information are used to build a forwarding cache at each router for each (source, group) pair. This forwarding cache is used to forward subsequent datagrams.

Inter-area routing using MOSPF is used when the datagram's source and destination are in different OSPF areas. Similar to intra-area routing, the forwarding cache is built from the local group database and the datagram shortest path trees. To support inter-area multicast, MOSPF uses a subset of the area's area border router (ABR) as *inter-area multicast forwarders*. They are responsible for forwarding group membership and multicast datagrams between different OSPF areas. These inter-area multicast forwarders summarize their attached area's group membership information to the backbone by originating new group membership LSAs. To permit forwarding of multicast traffic between areas, MOSPF introduces the concept of a *wild card multicast receiver*. This entity receives all multicast traffic generated in an area, regardless of the multicast group membership. For a nonbackbone area, an inter-area multicast forwarder works as a wild card multicast receiver so that all the multicast traffic can be forwarded to backbone and other nonbackbone areas. Thus, the backbone has the complete picture of group membership of different areas.

Protocol-Independent Multicast-Sparse Mode (PIM-SM) — PIM is designed to support multicast independent of any underlying unicast routing protocol and to ensure efficient Internet-wide multicast delivery. DVMRP and MOSPF are source-based multicast routing protocols and are meant for the environment where receivers are densely distributed. PIM distinguishes between dense mode (DM) and sparse mode (SM) multicast routing. PIM-DM is similar to DVMRP with some minor changes.

PIM-SM is designed for situations where receivers are sparsely distributed throughout the Internet. It is assumed that each receiver has to explicitly join a multicast tree if it wants to receive any multicast packet. It adopts some concepts from CBT by assigning a set of *rendezvous points* (RPs) for each group, where one works as the primary RP and is responsible for forwarding all packets destined for the multicast group. Each of the multicast domains selects a *designated router* (DR), which handles multicast group membership messages in its domain. If any host in its domain wants to join any group, it sends a join message to the DR. If the DR is not presently subscribed to that group, it looks for the RP of the group by a deterministic hash function search over the set of RPs and sends an explicit join message toward the RP. The DR and all the intermediate nonserving multicast routers create an (*.group) entry in the multicast routing table so that all future datagrams destined for that group should be forwarded to the DR. If a source wants to send a multicast packet, it has to first register itself with the RP by sending a PIM-SM-Register message toward the RP. The DR of this host encapsulates this message and forwards it toward the RP as a unicast

	Type	Protocol	Algorithm	Underlying unicast algorithm dependency?	Tree	Computation	Distribution tree
Layer three	Source-based	DVMRP	RPM	No	Unidirectional	On demand	Periodically flushed and rebuilt
		MOSPF	Dijkstra's	OSPF	Unidirectional	On demand and membership change	Changes with network topology
		PIM-DM	SPT	Yes	Unidirectional	On demand	Soft-state based
	Core-based	PIM-SM	SPT	Yes	Unidirectional	On demand	Soft-state based
		CBT	SPT	Yes	Bi-directional	On demand	Periodically refreshed

Table 1. *A comparison of multicast routing protocols.*

message. The RP then sends back a PIM-join message to the DR of the source.

Although PIM-SM is based on shared trees, it also allows a mechanism to provide a shortest path tree on behalf of the receivers. After joining the shared tree, if a receiver finds another optimal route to the source, it sends a join message toward the active source. After the source-based shortest path tree is constructed, the receiver sends a prune message to the RP and hence disassociates itself from the shared tree. Table 1 compares these layer 3 multicast routing protocols.

Distributed Core Multicast (DCM) — DCM [4] is a multicast routing protocol intended for use within a large single Internet domain network with a very large number of multicast groups for a relatively small number of receivers. Such a case occurs whenever multicast addresses are allocated to mobile hosts, as a mechanism to manage Internet host mobility. For such networks, existing dense or sparse mode multicast routing algorithms do not scale well with the number of multicast groups. It is an extension of the CBT approach and employs several core routers called *distributed core routers* (DCRs).

Layer Four (Transport) Protocols — The IP multicast service extends traditional best effort datagram delivery service to efficient multipoint communication. Similar to unreliable IP unicast service, multicast packets may be dropped or may arrive out of order, therefore causing deferred services to higher layers.

Recently, researchers have demonstrated multicasting real-time data, such as audio and video, over the Internet using the MBone. Since most real-time applications can tolerate some data loss but not the delay associated with retransmissions, they either accept some loss of data or use forward error correction (FEC). However, the main objective of these unreliable multicast protocols is to guarantee quality of service in terms of reduced end-to-end delay. On the other hand, although these applications accept losses and are thus well suited to IP multicast, a number of other applications (e.g., software distribution, multiplayer games) require reliable delivery of packets. Many reliable protocols have been developed to support this kind of applications such as Scalable Reliable Multicast (SRM), Reliable Multicast Transport Protocol (RMTP), Log-Based Reliable Multicast (LBRM), and Scoped Hybrid Automatic Repeat Request with FEC (SHARQFEC); details can be found in [5].

A simple approach to recover from packet losses is to make the sender retransmit the lost packets to individual receivers. However, such a *sender-centric* retransmission strategy does not scale well. In a large-scale multicast session, the probability that a given packet is lost by any one receiver is rather high; thus, the sender may end up retransmitting every packet. Moreover, if a packet were lost near the sender in the multicast tree, most receivers would not receive that packet. This leads to repair traffic that is proportional to session size and could be prohibitive. We may avoid this problem by allowing the sender to multicast the retransmission. However, as recent MBone studies indicate, many packet losses are not correlated, and different receivers may experience different loss rates. This leads to the *repair-locality problem* wherein repair traffic is not localized to its desired receivers, but rather multicast to all. When the retransmission is multicast, receivers may end up receiving many "unwanted" packets (duplicate packets) in the repair traffic. Sender-centric retransmission schemes also suffer from the well-known *implosion problem* in which the sender is potentially flooded by control traffic (request packets, negative acknowledgments from receivers, status requests, etc.).

The common consensus on addressing these problems now seems to be to effectively delegate the responsibility for recovery to the receivers (*receiver-oriented* approach). SRM, perhaps the most popular scheme for reliable multicast, allows receivers to multicast requests to the entire group. Any receiver with requested packets can multicast it. With clever use of randomized timers and suppression, SRM effectively solves the implosion problem. Unfortunately, by local and hierarchical scoping (grouping of receivers) [6], SRM can alleviate but not eliminate the repair locality problem.

Tree-structured protocols such as RMTP and LBRM solve the implosion and repair locality problems by imposing a logical tree structure to the multicast session. Specialized receivers located at the root of the subtrees of the logical tree receive requests and initiate retransmission only to their own children in the tree. These protocols work without any router support, but need specialized receivers. There are others protocols, such as Pragmatic General Multicast Protocol (PGMP) and Large-Scale Multicast Scheme (LMS), which propose to modify routers so that repair packets

Multicasting for wireless mobile hosts in an IP network is a challenging task. The addition of mobility in the host group model implies that the multicast-forwarding algorithm ought to focus not only on the issue of dynamic group membership, but also on the host location.

can be localized to the most effective region. Special routers need to be widely deployed to enhance the effectiveness of these protocols.

Another approach to solving both the implosion and repair locality problems is to use FEC. This involves splitting the original packet stream into groups of B packets, called *blocks*. For each block, h FEC encoded packets are generated for suitably chosen h. Receivers can recover the original block by receiving any B packets out of the $B + h$ ones. When combined with an appropriate automatic repeat request (ARQ) technique, FEC incurs very low network delay [7]. The FEC approach can be employed in any of the protocols we have described, such as SRM and tree-based protocols. An example is the SHARQFEC protocol that combines hierarchical scoping and hybrid FEC/ARQ. It breaks the entire multicast group into hierarchically nested scopes and designates a receiver within each scope as the *zone closest receiver* (ZCR). After receiving each data block, each ZCR proactively multicasts FEC packets to the receivers in its scope. Since proactive FEC packets and SRM-style suppression can much reduce repair and request traffic, a scope can contain many more receivers than a subtree in tree-based protocols.

To summarize, scoping and employing FEC together with suppression techniques lead to reliable and scalable multicast. However, two main issues remain unresolved: how many FEC packets need to be transmitted in a scope, and which protocol is suitable to schedule the transmission of these FEC packets. Most existing schemes force the sender to multicast as many FEC packets as needed by the lossy receiver and is effective only if all the receivers within a scope have similar loss rates. Unfortunately, in most existing protocols scopes are not defined by the loss rates of receivers. In [8], scopes are defined to be receivers within a certain hop count or TTL, and hence are based on relative physical locations.

Multicast in Wireless Mobile Environments

Multicasting for wireless mobile hosts in an IP network is a challenging task. The addition of mobility in the host group model implies that the multicast forwarding algorithm ought to focus not only on the issue of dynamic group membership, but also on host location. Reconstructing a multicast tree every time a host moves to a new domain could result in substantial overhead. We first introduce the issues involved in providing multicast in a mobile environment and then discuss details of proposed multicast routing protocols.

Issues in Mobile Environments

Multicast Forwarding Algorithm — Multicast forwarding algorithms have been designed considering the static nature of the hosts. A source-based protocol like DVMRP creates problems if a host's point of attachment changes. DVMRP forwards multicast packets only if it receives them on the correct interface. This is a serious drawback in a mobile environment. When a mobile host (MH) moves to a new domain, its interface to the multicast router changes, resulting in the packets being dropped.

Dense or Sparse Mode Protocols — Multicast routing protocols are designed for dense mode or sparse mode. The dense mode protocols are useful if there are a lot of users and bandwidth availabilitiy is not a problem, whereas sparse mode protocols are meant for widely distributed users with limited bandwidth. This assumption may not be true in the mobile environment where, due to movement of the MH, user density patterns may change dynamically.

QoS Provisioning — It is difficult to do resource reservation in a multipoint communication, and the situation becomes more complex for wireless mobile environments. Quality of service (QoS)-based wireless multicasting is an open issue, and research is being carried out to address associated problems. There is a general feeling that future Internet access for wireless networks may also be IP-based. The IETF has proposed Mobile IP [9] as a possible solution for host mobility. The solution for implementing mobile multicast over the Internet may have to be some variant of Mobile IP. In this section we first describe Mobile-IP-based multicasting protocols followed by protocols that provide reliable mobile multicasting.

Proposals for Multicast over MIP

The IETF has proposed two approaches to provide multicast over Mobile IP. They are called *remote subscription* and *bidirectional tunneling*. This section gives an overview of these proposed methods.

Remote Subscription — In remote subscription an MH resubscribes to the multicast group each time it moves to a new foreign network. It is the simplest way of providing multicast through Mobile IP. There is no special encapsulation needed, and it works well with basic Mobile IP. However, this approach is not suitable for highly mobile users since frequent resubscription in each foreign network may lead to lost packets. Moreover, frequent reconfiguration of the multicast delivery tree may result in substantial control overhead.

Bidirectional Tunneling — In this scheme MHs send and receive multicast packets by way of their home agents (HAs), using unicast Mobile IP tunnels from their HAs. This requires the HA of an MH to be a multicast router. All subscription of the MH is done through the HA. Both the above two methods address the basic issue of providing multicast over Mobile IP, but not the issues when multicast services are extended to Mobile IP.

Multicast Support Using Mobile IP (MoM) — The bidirectional tunneling solution for multicast over Mobile IP creates an interesting situation when many MHs, belonging to different HAs, move to the same foreign agent (FA). According to bidirectional tunneling, each of the respective HAs creates a separate tunnel to the FA so that multicast packets to their respective MHs can be forwarded. If these MHs were subscribed to the same group, all of the tunnels

from different HAs to the FA would carry the same multicast packet (Fig. 3), resulting in packet duplication. This is called the tunnel convergence problem. MoM [10] addresses this by selecting only one HA among the given set of HAs. The selected HA among the given set of HAs is called designated multicast service provider (DMSP) and solves the problem of tunnel convergence, but it may result in packet loss if the MH belonging to the currently serving DMSP moves out. An FA comes to know about the movement of the MH only after its lifetime expires, while an HA comes to know about the movement of the MH as soon as it receives a register message for a different care-of address (CoA). As a result, the HA stops sending packets to its old interface, thinking that the MH is no longer in the previous CoA. So, within the handoff duration and the MH's lifetime, the FA will not receive any packets for that group; hence, there could be packet loss. Another variation of DMSP addresses this issue of packet loss during handoff and suggests selecting more than one DMSP (less than three) at any time, so even if one HA serving as DMSP stops forwarding, the FA could still receive packets from another DMSP. But this comes at the cost of packet duplication.

An enhancement of MoM called Range-Based MoM (RBMoM) provides a trade-off between the shortest delivery path and the frequency of multicast tree reconfiguration. It selects a router, called a multicast HA (MHA), which is responsible for tunneling multicast packets to the FA to which the MH is currently subscribed. The MHA can only serve MHs that are roaming around foreign networks and are within its service range. If an MH is out of service range, MHA handoff will occur. Initially, the MHA of na MH is set to its HA. Every MH can have only one MHA, which changes dynamically with the location of the MH, whereas the HA of an MH never changes. This protocol requires that each MHA be a multicast group member.

Multicast for Mobility Protocols (MMP) — The focus of this protocol [11] is to provide fast and efficient handoffs for MHs in foreign networks and enable location-independent addressing. MMP combines Mobile IP and CBT where the former controls communication up to the foreign network, and the latter manages movement of hosts inside them. It assumes the foreign domain to form a hierarchy of multicast supporting routers. Similar to the concept of an FA, a base station acting as a multicast router transmits periodic beacons, including one multicast CoA. Upon acquiring a CoA, the MH sends a registration message to the base station, which triggers a multicast tree join and transmits a CBT join request to the core. The core takes care of relaying the registration request to the HA of the MH by replacing the CoA to its own address, thus hiding the multicast part of the protocol and acting as a sole FA. In a domain consisting of a hierarchy of multicast routers, the border router can be selected as the core of the network. The drawback of the protocol is that it assumes a large-scale deployment of multicast-capable routers in each domain. Also, it is not protocol-independent, which is a major limitation considering the diversity of the Internet.

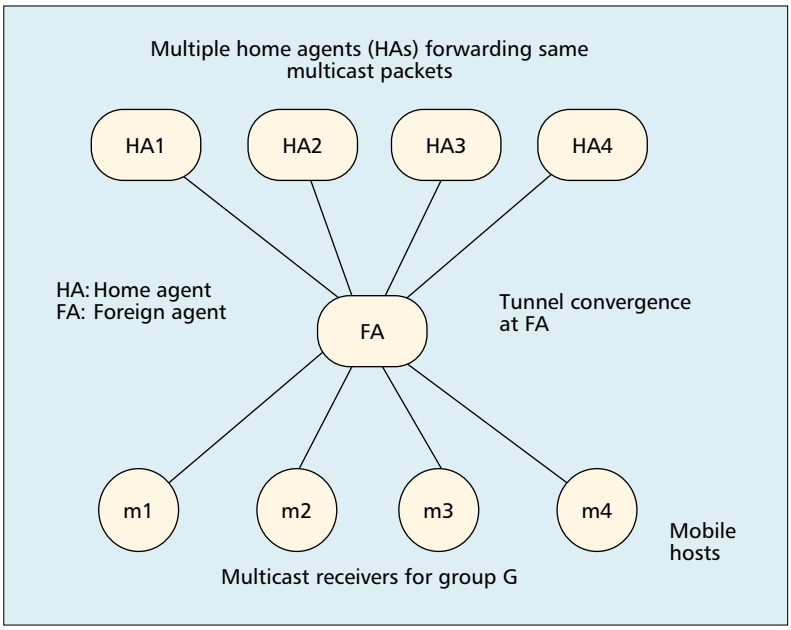

Figure 3. *The tunnel convergence problem in multicasting.*

MobiCast — MobiCast [12] is designed for an internetwork environment with small wireless cells, with many cells grouped together and served by domain FAs (DFAs). DFAs are multicast FAs and are meant to isolate the mobility of the MH from the main multicast delivery tree. This hierarchical mobility management approach isolates the mobility of the MHs from the main multicast delivery tree.

A framework to handle multicast source movement over Mobile IP is proposed in [13]. In case an MH is serving as both source and recipient of a multicast group, during the movement to a new FA it creates a bidirectional from its new FA if needed. In the meantime, the FA initiates a multicast tree join in case it is not presently subscribed to the multicast group. Once it starts receiving multicast packets through the tree, it disassociates its reverse tunnel to the HA and only keeps the forward tunnel to send multicast packets.

Reliable Wireless Multicast Protocols — Reliable multicasting has become an important issue. Here we discuss two different multicast protocols designed for reliable delivery of messages.

RMDP Protocol — A *Reliable Multicast Data Distribution Protocol* (RMDP) presented in [14] was originally implemented for use on the MBone. It relies on the use of FEC and ARQ information to provide reliable transfer. Redundant information is inserted into the FEC, often enabling a receiver to reconstruct the original packet. In the event that such information is not enough, an ARQ is sent to the multicast source, which, in turn, multicasts the requested packet to all receivers. In RMDP, a data object to be transmitted is a *file*, identified by a unique name, say its uniform resource locator (URL). The file

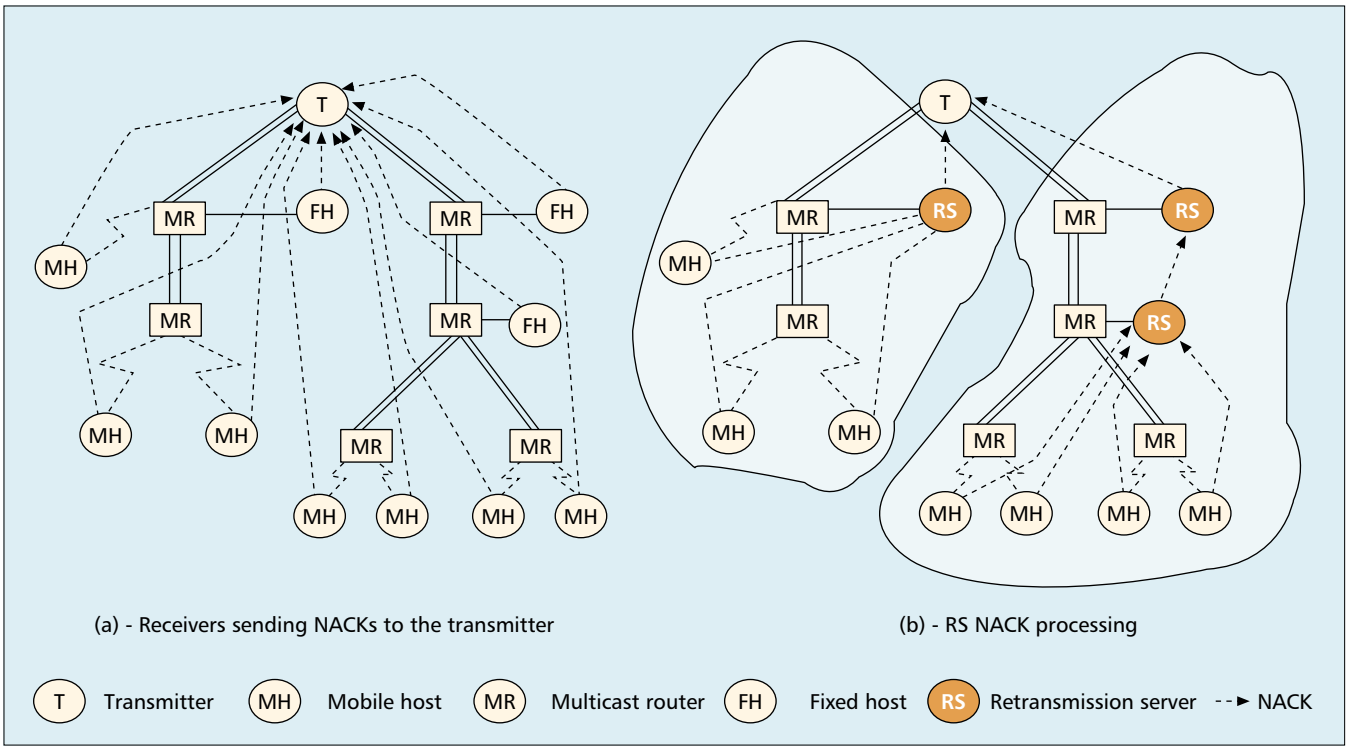

Figure 4. *RM2 support for reliable multicast sessions.*

has a finite size, and is split into packets of s bytes each. RMDP uses an (n, k) encoder with $n \gg k$ to generate packets for transmission, and assumes the existence of a multicast network that provides unreliable but efficient delivery of data packets.

RMDP does offer scalability and efficiency when used in commonly reliable media. On the other hand, one of the main drawbacks of RMDP is that data encoding/decoding is done through software, resulting in processing overhead and therefore performance degradation. For resource-limited receivers, decoding cost is of concern. Furthermore, in the presence of highly unreliable wireless media, errors typically occur in burst, causing the protocol to generate a large amount of ARQ packets, which triggers a substantial amount of retransmission packets multicast to all receivers. Clearly, RMDP's retransmission scheme based on ARQ packets does not seek to conserve network resources. This problem could be minimized if a hierarchal protocol is employed.

RM2 Protocol — RM2 [15] is a reliable multicast protocol that can be used for both wired and wireless environments. Furthermore, RM2 guarantees sequential packet delivery with no packet loss to all its multicast members. RM2 relies on the Internet Group Management Protocol (IGMP) [3] to manage group membership, and on the IETF's Mobile IP to support user mobility through a CoA. RM2 is a hierarchical protocol that divides a multicast tree into subtrees where subcasting within these smaller regions is applied using a tree of retransmission servers (RSs) with each RS having a retransmission subcast address shared by its members and which may be dynamically configured using the IETF's Multicast Address Dynamic Client Allocation Protocol (MADCAP). In order to guarantee end-to-end reliability, the receivers are required to send negative acknowledgments (NACKs), pointing out the packets to be retransmitted. In other words, RM2 implements selective packet retransmission. Figures 4a and 4b illustrate the concept of RSs and how they are applied to reduce network overload. RM2 assigns RS functionality to fixed hosts selected on a network topology basis. In RM2, RSs may be specially configured to manage retransmission for fixed/mobile users or a mixture of them. In [15] it is shown that this subcasting scheme greatly improves the scalability and efficiency of the protocol.

One important factor that determines the efficiency of RM2 is its retransmission algorithm. While existing reliable multicast protocols adopt a static packet retransmission scheme — unicast or multicast — often leading to performance degradation due to wasted bandwidth, RM2 adopts a retransmission algorithm that dynamically switches between unicast and multicast modes to save network and wireless resources. This algorithm has been developed taking into consideration the network topology and the number of fixed and mobile users. As a result, precise conditions and rules are defined taking into account the dynamics of the network state. Thus, RM2 implements a protocol designed to cope with the scarce bandwidth available and scalability, but suffering from cache management. As packets are received by the RSs, they are stored in cache for future retransmission requests. An effective cache management algorithm needs to be developed in order to determine how long a packet should be stored in the cache and the cache's optimum size. Table 2 compares various wireless multicast routing protocols.

	Mobility protocol	Optimal routing	Reliability	Packet redundancy	Multicast protocol dependency	Join and graft delays
Remote subscription		Yes	No	No	Independent	Yes
Bidirectional tunneling		No	No	Yes	Independent	No
MoM	Mobile IP	No	No	Minimal	Independent	No
MMP		No	No	Minimal	CBT	No
Mobicast		Yes	No	Minimal	Independent	Yes
RMDP		No	Yes	Yes	Independent	Yes
RM2		Yes	Yes	No	Independent	Yes

Table 2. *A comparison of IP-based wireless multicast routing protocols.*

Conclusions and Future Directions

Multicast is a field in which there is no *one-size-fits-all* protocol that can optimally serve the needs of all types of multicast applications. Both wired and wireless multicast proposals have been designed to cope with specific applications types, which often lead to unexpected behavior when applied to unfamiliar environments. Since the impact of multicast spans numerous areas, analyzing and indicating the suitability of a protocol is very hard, specially when considering both wired and wireless multicast. This article is an attempt to give an overview of current research in wired and wireless multicast field and show that this area is rapidly growing and evolving, and there are still many challenges that need to be addressed. Future directions in wireless multicast need to consider QoS, security, and so on. Questions to be addressed are whether multicast ought to be application-oriented, and how to integrate the wireless multicast infrastructure into Mobile IP. Furthermore, a detailed investigation is desirable for both unreliable and reliable environments.

Acknowledgment

This work has been supported by the Ohio Board of Regents, Doctoral Enhancements Funds, and the National Science Foundation under Grant CCR-0013361.

References

[1] S. Deering, "Multicast Routing in a Datagram Network," Ph.D. dissertation, Stanford Univ., 1991.
[2] H. Eriksson, "MBONE: The Multicast Backbone," *Commun. ACM*, vol. 37, no. 8, Aug. 1994, pp 54–60.
[3] W. Fenner, "Internet Group Management Protocol, Version 2," Internet draft, Apr. 1996.
[4] L. Blazevic, and J. Boudec, "Distributed Core Multicast (DCM): A Routing Protocol for Many Small Groups with Application to Mobile IP Telephony, IETF draft, June 1999.
[5] Reliable Multicast Links: http://research.ivv.nasa.gov/rmp/links.html.
[6] P. Sharma *et al.*, "Scalable Session Messages in SRM Using Self-Configuration," USC tech. rep., July 1998.
[7] J. Nonnenmacher *et al.*, "How Bad Is Reliable Multicast Without Local Recovery," *Proc. IEEE INFOCOM*, Mar. 1998, pp. 972–79.
[8] S. Floyd *et al.*, "A Reliable Multicast Framework for Lightweight Sessions and Application Level Framing," ext. rep., Sept. 1995.
[9] C. Perkins, "IP Mobility Support," RFC 2002.
[10] V. Chikarmane *et al.*, "Multicast Support Using Mobile IP: Design Issues and Proposed Architectures," *ACM/Baltzer J. Mobile Net. App.*, vol. 3, no. 4, 1998, pp. 365–79.
[11] A. Mihailovic, M. Shabeer, and A. H. Aghvami, "Sparse Mode Multicast as a Mobility Solution for Internet Campus Networks," *Proc. PIMRC '99*, Osaka, Japan, Sept. 1999.
[12] C. L. Tan, and S. Pink, "Mobicast: A Multicast Scheme for Wireless Networks," *Mobile Net. App. 5*, 4, Dec. 2000, pp. 259–71.
[13] H. Gossain, S. Kamat, and D. P. Agrawal, "A Framework for Handling Multicast Source Movement over Mobile-IP," to appear, *IEEE ICC*, May 2002.
[14] L. Rizzo, and L. Vicisano, "RMDP: A FEC-Based Reliable Multicast Protocol for Wireless Environments," *Mobile Comp. Commun. Rev.*, vol. 2, no. 2, 1998.
[15] D. H. Sadok *et al.*, "A Reliable Subcasting Protocol for Wireless Environments," *2nd Int'l. Conf. Mobile Wire. Commun. Net.*, Paris, France, May 2000.

Biographies

Hrishikesh Gossain (hgossain@ececs.uc.edu) is a graduate student and research assistant at the Department of ECECS, University of Cincinnati. He received his B.E. in electronics engineering from Motilal Nehru Regional Engineering College, Allahabad, India, in 1998, where he was undergraduate Gold Medallist. He is co-inventor of four patents filed in QoS and e-media. His research interests include multicast, mobility management in the wireless environment, Mobile IP, QoS, and multimedia communications. He previously worked at Nortel Networks in Richardson, Texas.

Carols de Morais Cordeiro (cordeicm@ececs.uc.edu) received his B.Sc. and M.Sc. in computer science in 1998 and 2000, respectively, from the Federal University of Pernambuco, Brazil. He is presently pursuing a Ph.D. degree in computer science and engineering at the University of Cincinnati where he is a research assistant in the Center of Distributed and Mobile Computing. His research interests and papers are mostly in the area of wireless and mobile communications including ad hoc networks, multicast, TCP over wireless, and short-range radio communication such as Bluetooth. He has prior work experience with IBM in San Jose, California.

Dharma P. Agrawal [F] (dpa@ececs.uc.edu) is the Ohio Board of Regents (OBR) Distinguished Professor of Computer Science and Computer Engineering and founding director of the Research Center for Distributed and Mobile Computing in the Department of Electrical and Computer Engineering and Computer Science, University of Cincinnati, Ohio. He has been a faculty member at Wayne State University, Detroit, Michigan (1977–1982) and North Carolina State University, Raleigh (1982–1998). His current research interests include wireless and mobile networks, distributed processing, and scheduling techniques. He is an editor for *Journal of Parallel and Distributed Systems* and *International Journal of High Speed Computing*. He has served as an editor of *IEEE Computer* and *IEEE Transactions on Computers*. He has been program chair and general chair for numerous international conferences and meetings. He was selected for the Third Millennium Medal by the IEEE for his outstanding contributions.