Multicasting Streaming Media to Mobile Users

Ashutosh Dutta, Jasmine Chennikara, and Wai Chen, Telcordia Technologies Inc.
Onur Altintas, Toyota InfoTechnology Center
Henning Schulzrinne, Columbia University

ABSTRACT

Content distribution in general, and multicasting in particular, over a wired network to static hosts can be realized by placing proxies and gateways at several parts of the network. However, if the end hosts are mobile over heterogeneous wireless access networks, one needs to consider many operational issues such as network detection, handoff, join and leave latency, and desired level of quality of service, as well as caching and load balancing. This article surveys a set of protocols and technologies that offer multicast-based services for streaming multimedia in a mobile environment. It also brings forth some of the issues related to mobile content distribution in the wireless Internet that may be helpful during its deployment by application service providers.

INTRODUCTION

Lately, streaming real-time multimedia content over the Internet is gaining momentum in the communications, entertainment, music, automotive, and interactive game industries. Streaming applications include broadcasting multimedia content, multiparty conferences, collaborations, and multiplayer games. All of these applications also find use in a military context, including coordination, education, situation awareness, distributed simulation, and battlefield communication. Real-time streaming content (audio and video) is mostly an Real-Time Transport Protocol (RTP) [1] based application that has stringent delay and loss requirements. Mobility, on the other hand, affects the delay and transient loss for multimedia stream delivery to a great extent because of associated repeated handoffs. Thus, it becomes more challenging to maintain session continuity and provide proper quality of service (QoS).

In order to make efficient use of network bandwidth within the core of the network, IP multicasting is used in wide area networking. There are several proposed network layer schemes that provide native IP multicast routing over a wide area network, such as Protocol Independent Multicast (PIM), Multicast over Open Shortest Path First (MOSPF), Distance Vector Multicast Routing Protocol (DVMRP), Core Based Tree (CBT), and Border Gateway Multicast Protocol (BGMP). There can be several types of multicast models, such as one-to-many, many-to-many, and many-to-one. Examples of one-to-many applications include scheduled audio/video distribution, push media, file distribution, caching, and monitoring of stock prices. Multimedia conferencing, synchronized resources, concurrent processing, collaboration, distance learning, chat groups, distributed interactive simulations, multiplayer games, and jam sessions fall into the many-to-many category. Some of the many-to-one applications include resource discovery, data collection, auctions, polling, and accounting. Currently multicast is not widely deployed since there are many issues such as pricing, security, QoS, and maintenance of the router states in the core of the network (for a detailed discussion of these issues peculiar to wide area networking see [2]). On the other hand, local multicasting within a subnet becomes more attractive for mobile users experiencing intradomain handoffs because of its ease of deployment and ability to provide more flexible services such as localized advertisements, news broadcast, and location specific information.

This article is organized as follows. We provide some alternate proposals based on network and application layers that can build a multicasting content distribution network for both nonmobile and mobile users. We discuss some mobility components and highlight various issues involved in multicasting content distribution in mobile networks. We then conclude the article.

MULTICASTING STREAMING CONTENT OVER THE MOBILE INTERNET

Content distribution from a single source follows the one-to-many model. Most broadcasting

 0163-6804/03/$17.00 © 2003 IEEE

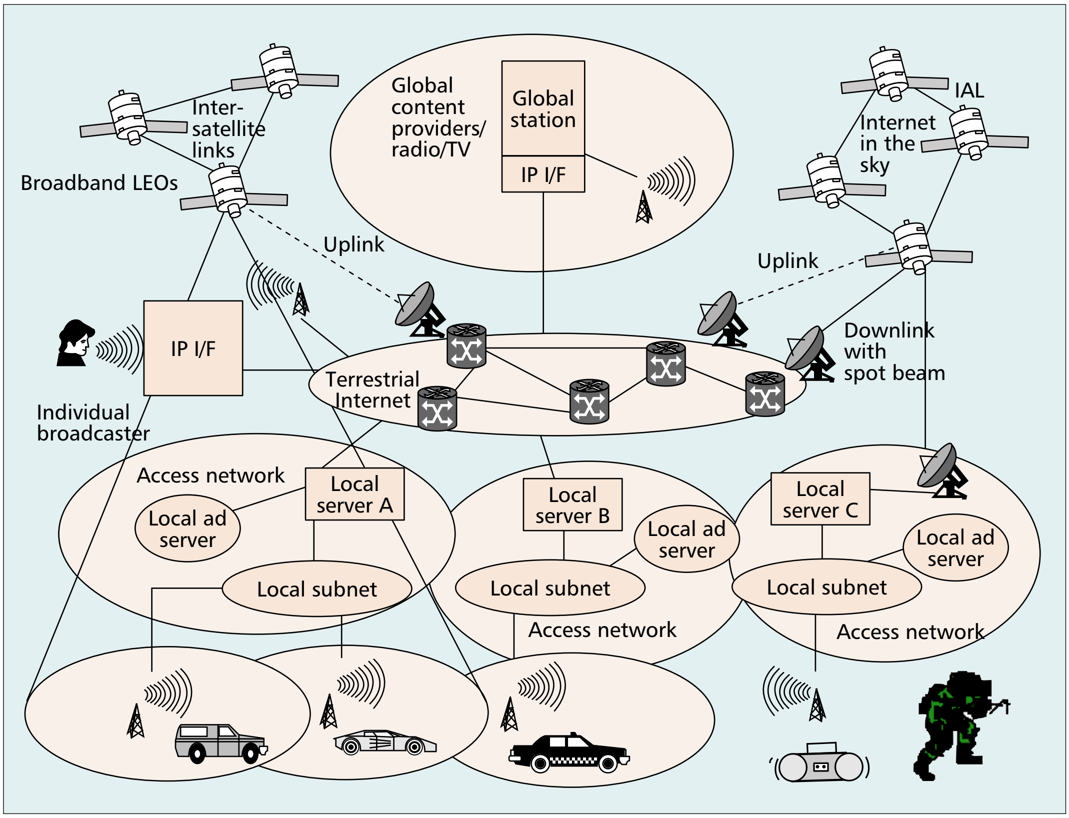

■ **Figure 1.** *A typical mobile content distribution network.*

By virtue of IP multicasting, IP packets are delivered from a single source to a group of receivers that are part of the same multicast group. Joining and advertisement of multicast groups is handled through IGMP and SAP, respectively

sources such as radio and TV networks follow this kind of model. Multicasting streaming content to end users over the Internet may include both mobile and nonmobile clients over wired and wireless media. Figure 1 illustrates a content distribution network, with multiple proxy servers, different kinds of sources, and several types of core and access networks, that offers flexible services to mobile users.

The subsections below discuss some of the related work applicable to both mobile and nonmobile users.

Multicasting to Nonmobile Users

By virtue of IP multicasting, IP packets are delivered from a single source to a group of receivers that are part of the same multicast group. Joining and advertisement of multicast groups is handled through Internet Group Management Protocol (IGMP) and Session Announcement Protocol (SAP), respectively, although some alternative application layer techniques are described in [3]. Multicast packets are generally routed along a single shared tree or multiple source-based spanning trees for efficient distribution. To support IP multicast, the network must maintain knowledge of its routing tables for the multicast routes as well as for the unicast routes. Traditional multicasting techniques do not handle large numbers of distinct multicast groups and do not provide a means to handle multicast when some routers may not be multicast-capable.

Explicit multicast (Xcast) being developed within the Internet Engineering Task Force (IETF) supports multicast groups when the membership in each group is small, unlike traditional multicast that supports a limited number of large multicast sessions. In Xcast, the sending node includes the IP addresses of all the members of the multicast group in the packet header. Intermediate routers use the header information to create unicast packets, encapsulating the multicast packet and forwarding it to the group members who are at the next hop. The routers then modify the packet header of the original multicast packet and remove the members who have been sent a copy of the multicast packet before relaying it to the next router. Implementation of this scheme requires modification to the sending and receiving nodes as well as to the intermediate routers.

An IETF developed protocol, Source Specific Multicast (SSM), addresses issues such as multicast address allocation, destination unawareness, interdomain routing, source advertisement, and connection state that creates a huge multicast forwarding table. In the SSM model, each multicast group is not only defined by a multicast address, but also by a sending or source node IP address. SSM does not require IP multicast address management since it does not need a unique multicast address for each group. Therefore, SSM is ideal for Internet broadcast applications, allowing content providers to support services without requiring a unique IP multicast address. However, this approach requires router modifications to handle multicast group identification based on both the source IP address and multicast group address.

There have been some advances in research with respect to supporting mobility for multicast users, specifically through Mobile IP. One such method is the bi-directional tunneling solution, which puts the multicasting burden on the home agent.

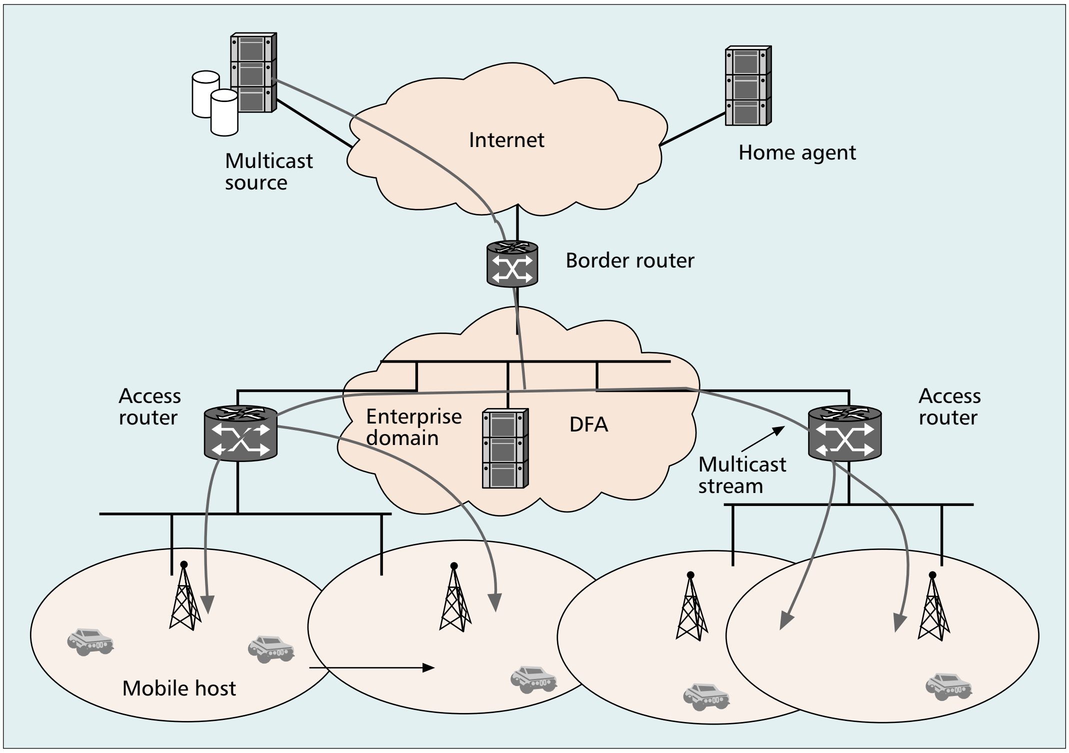

Figure 2. *Mobility support for multicast in Mobicast.*

MULTICASTING TO MOBILE USERS

There have been some advances in research on supporting mobility for multicast users, specifically through Mobile IP. One such method is the bidirectional tunneling solution, which puts the multicasting burden on the home agent (HA). In this case, a user wanting to join a certain multicast group joins the group through the user's HA using IGMP. When the user moves to a foreign network, the HA is responsible for tunneling multicast packets to the user. However, when a single HA or multiple HAs have users in the same multicast group visiting the same foreign network, tunneling multiple multicast packets to the foreign network from one or more HAs is inefficient. In order to avoid duplication of multicast packets being tunneled to foreign networks, one proposed solution is remote subscription. In this approach, a user will join the desired multicast group in each visited network through the foreign agent (FA). However, this requires that after each handoff the user rejoin the multicast group, and the multicast trees used to route multicast packets be updated to track the multicast group members. In order to limit the tree updates and duplication of multicast packets, proxy or agent-based solutions have been proposed. For example, in the Mobicast solution [4], users continue to rejoin the multicast group in each visited network. This architecture introduces the domain foreign agent (DFA) concept to hide all mobility within the foreign domain from the main multicast delivery tree. In this scenario the DFA will send or receive the multicast traffic to a multicast group. When the mobile host (MH) is receiving multicast traffic, the DFA uses a translated multicast address within its network to prevent multicast updates due to mobility. Figure 2 shows a practical application of how Mobicast can be used to distribute content to mobile users using the DFA approach.

Mobile Multicast (MOM) [5] provides a mobility scheme for multicast multimedia sessions for wide area networking and adopts a Mobile IP-based approach. It proposes to reduce the problem in bidirectional tunneling of delivering multiple copies of the same multicast packet to a foreign network. In this solution, one HA will be elected to tunnel multicast packets to a foreign network. Range-based MOM [6] takes the MOM approach one step further and elects a multicast agent close to the FA to tunnel multicast packets to the foreign network.

Mysore and Bhargavan [7] provide a scheme to take care of the loss of transient data for mobile hosts by assigning location-independent unique multicast addresses to each MH, so the MH will experience less transient data loss. However, this scheme does not address the mobility of localized multicast sessions.

The Mobile Multicast Proxy [8] approach proposes that the proxy's clients do not directly participate in the multicast tree. The multicast proxy participates in the multicast tree formation for the groups to which its clients belong. In this case, a multicast proxy performs a similar function to a designated router. However, the multicast proxy can be outside the member's subnet and can forward multicast messages to its receivers using unicast, multicast, or a limited scope broadcast. It also provides for primary and secondary multicast proxies, where a secondary proxy is located closer to the mobile clients, but both primary and secondary proxies can communicate using unicast or multicast tunnels.

Also, there are few commercial content distri-

Protocols	Functionality
PIM	Protocol Independent Multicast
MOSPF	Multicast extension to OSPF
DVMRP	Distance Vector Multicast Routing
ICAP	Content Distribution Protocol
CBT	Shared Multicast Tree
BGMP	Routing between Domains
IGMP	Host-to-router JOIN protocol
CGMP	Layer 2 Join Protocol (Cisco)
DRCP	Faster IP address acquisition
UMTP	Multicast Tunneling for UDP
AMT	Automatic Multicast Tunnel
RTSP	Real-Time Streaming Protocol

Table 1. *Protocol galaxy.*

bution networks that use multicasting as their core technology. Most recently, Packet Video, in conjunction with DoCoMo, has started providing wireless multicast streaming services to end users, but has not taken into account the subnet mobility factor. Akamai and Network Appliance have developed a new protocol called Internet Content Adaptation Protocol (ICAP) that enables communication between edge content devices (e.g., Web caches and Internet content delivery servers) and application servers that modify content. But this does not support multicast yet, and has not included mobility of the end clients in its consideration. Companies such as Inktomi and Coolcast are already providing such multicast services. These services are supported via satellite to reach a wide range of nonmobile users.

In addition, iBEAM's infrastructure and its product Activecast distribute streaming media over the Internet using geostationary Earth orbit (GEO) satellites for content distribution, without addressing user mobility or providing flexible methods of advertisement insertion for mixing of local and global content.

While many of the multicast content distribution approaches discussed above provide network layer solutions, there are a few architectures that provide application layer multicast techniques. These do not depend on underlying IP multicast support, but can use it where it is available. In these architectures the clients can use an overlay network to disseminate the information. Application layer multicast trades off ease of deployment, flexible access control, and simplified configuration at a cost of higher data traffic than in network-layer multicast. Scattercast [9] overlays broadcasting architecture on top of the Internet. It introduces a set of network agents called Scattercast proxies (SCXs) that may connect with each other using unicast connections. CoopNet [10] combines infrastructure-based CDN and peer-to-peer systems. CoopNet alleviates the server overloading problem by having clients cooperate with each other to distribute content. For on-demand content, the server redirects clients to others who may have previously cached it; for live streaming, a distribution tree rooted at the server is formed with each of its clients as its members. MarconiNet [11] proposes an integrated streaming architecture to support multimedia applications such as IP telephony and broadcasting streaming content over the Internet using both wired and wireless access. It provides an application layer approach using a real-time feedback mechanism based on Real-Time Control Protocol (RTCP), Session Description Protocol (SDP), SAP, Real-Time Streaming Protocol (RTSP), and Session Initiation Protocol (SIP). MarconiNet relies on local servers in the access network, and uses application-layer multicasting and triggering techniques to provide flexible streaming services such as localized advertisement, local/global channel management, and fast handoff for the mobiles with the desired level of QoS. Table 1 shows a taxonomy of the protocols used in this article.

MarconiNet relies on local servers in the access network and uses application-layer multicasting and triggering techniques to provide flexible streaming services such as localized advertisement, local/global channel management, and fast handoff for the mobiles with the desired level of QoS.

Deployment Issues for Mobile Content Distribution via Multicasting

Diot *et al.* [2] provide a survey of deployment issues for IP multicast service and architecture. But they do not discuss details of issues related to content distribution for mobile networks. Here we focus on some of the issues related to content distribution in mobile networks.

A. Movement between Diverse Networks

Native multicast support among autonomous systems spanning multiple service providers is still not easily available because the core routers may not have multicast support. In order to extend the multicast-based applications to a network where multicast is not supported, there need to be changes in the design approach. There are several solutions based on user level or network level application available by using IETF protocols such as UDP Tunneling Multicast Protocol (UMTP), and Automatic Multicast without Explicit Tunnel (AMT) to connect these multicast-enabled islands. These enable multicast applications to span over networks and reach end users. These multicast applications can also be distributed via broadband low Earth orbit (LEO) satellite systems as part of several spot beams.

When the mobile is moving from a multicast-enabled network to a non-multicast-enabled network, UMTP or AMT tunnels must be set up proactively between end clients and the intermediary node supporting multicast routing. This approach can be realized easily by installing gateway proxies between the boundaries of the two networks and activating the UMTP tunnels when the client moves to the new access network.

Figure 3 shows how AMT or UMTP tunnels are set up as the mobile moves between multicast and non-multicast-enabled networks. More details of how mobile users move between such

In general, when a node moves, signaling and transport delays contribute to the latency of multimedia stream delivery associated with any two-party or multi-party communication session.

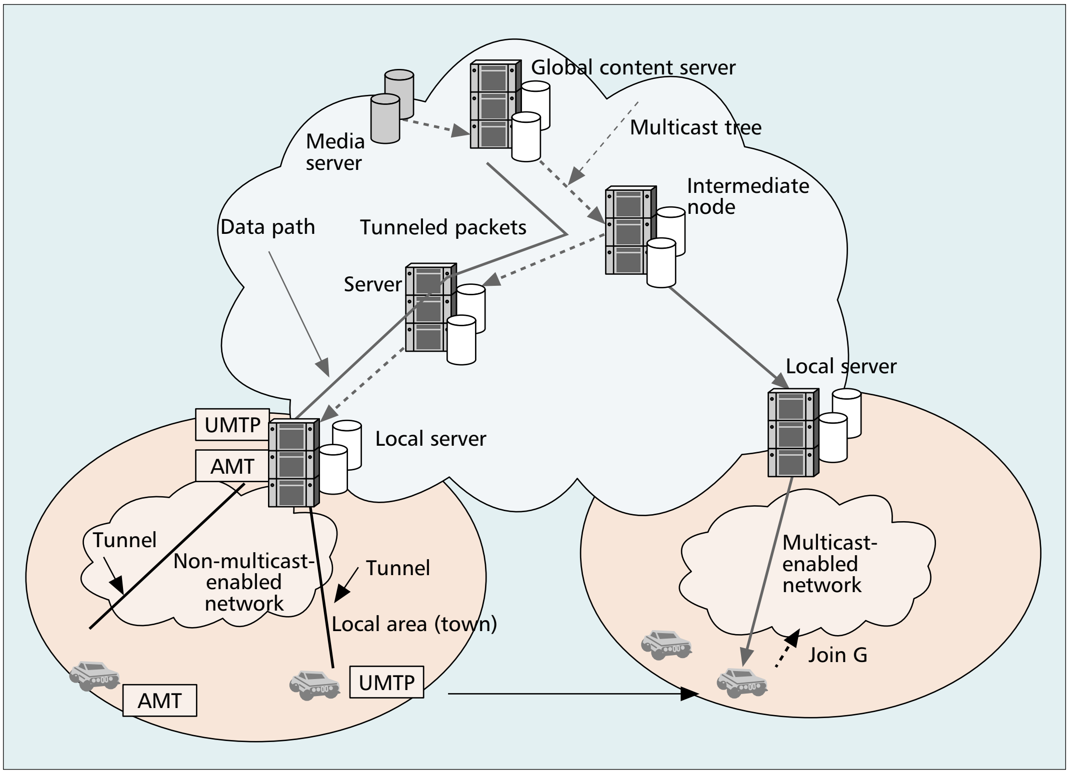

■ **Figure 3.** *Movement between diverse networks.*

diverse networks (e.g., multicast and non-multicast) are given in [11]. Figure 4 shows an example of how tunneling techniques can be used in the MarconiNet environment to supplement the lack of multicast support within the core of the network. The network elements in each of these multicast enabled islands in Fig. 4 are the radio station client (RSC), radio antenna server (RAS), and mobile clients. M and lm are globally scoped multicast address and locally scoped multicast address, respectively. A UDP server tunnels multicast packets between these islands.

Intradomain Mobility Components

In general, when a node moves, signaling and transport delays contribute to the latency of multimedia stream delivery associated with any two-party or multiparty communication session. Signaling, such as registering with a new server, notifying the communicating party of the mobile node's new contact address, and inviting another user to a streaming session, generally constitutes signaling delay; while transport delay dominates the delay component associated with mid-session mobility.

The latency associated with receiving a continuous unicast or multicast stream from a single source while the client moves to the next cell consists of several components, such as detection of a new cell, subnet, or domain; address acquisition, network configuration; triggering of a multimedia stream to be delivered in the new subnet; and actual delivery of the multimedia stream.

While some of these factors are common to both unicast and multicast stream delivery (e.g., cell or subnet detection, IP parameter configuration), this article mostly focuses on delivery schemes for multicast streaming traffic.

In a typical mobile content distribution network that uses multicast technology, latency mostly consists of the actual handoff time and join interval. Handoff delay includes the time to detect that the mobile node is in a new cell, subnet, or domain, the time for obtaining an IP address from a DHCP (or PPP server if it is moving between subnets), as well as time for some triggering mechanism (e.g., RTCP or IGMP-based) that will help initiate the multimedia flow to the MH's new location. Figure 5 shows the flow with a standard sequence of events for micro (cell) and macro (subnet) mobility for multicast streaming. For interdomain mobility, there are other factors, such as profile verification due to accounting, authentication, and authorization (AAA), that contribute to the delay preceding media delivery.

Network Movement Detection

Discovery of a new cell, subnet, or domain can be realized in different layers. During an MH's handoff, first movement detection takes place in layer 2, where the client decides to switch over to a new base station based on the signal strength of the received beacon. In the case of code-division multiple access (CDMA), soft handoff is initiated so that the client can listen to both base stations, receive both streams, and then decide which will be accepted from the mixed signal. As soon as it switches over to a new base station and layer 2 handoff completes, the client needs to figure out if it is in a new subnet or domain altogether. Using a layer 3 triggering mechanism (e.g., router advertisement or ICMP advertisement in Mobile IP), it can be determined if the client is in a different subnet.

An application layer detection mechanism such as server advertisement can be used as well if the client is involved in a real-time communication session. However, it may be faster to achieve handoff notification using layer 2 mechanisms.

CLIENT CONFIGURATION

The client configuration process helps configure a client in the new network with the IP address and other parameters such as Domain Name Service (DNS) server. As the client moves from one cell to another, if the new cell is in another subnet, it will either obtain a new address from the DHCP server or use the standard Mobile IP approach to obtain a new care-of address from the FA while keeping its original home IP address unchanged. In subnet movement, the typical time for acquiring a DHCP address can be on the order of 5–15 s, although there are various other alternatives, such as Dynamic Rapid Configuration Protocol (DRCP), DHCP without ARP, and auto-configuration in DHCPv6, that provide faster IP address acquisition. DRCP reduces the IP address acquisition time to a few hundred milliseconds. DHCPv6 provides similar measurement in stateless auto-configuration mode, while DHCP without ARP option takes about a second to configure the IP address. PPP for wide area network roaming takes up to 15 s before a handshake is complete and an IP address is assigned. However, in some of the micromobility approaches such as cellular IP or HAWAII, clients may not need to change their IP addresses during their movement.

JOIN/LEAVE LATENCY

The process of joining or leaving a specific multicast group while changing the cell or subnet can be treated as equivalent to surfing a TV/radio channel or flipping the channels. Since multicast communication is receiver-initiated, triggering techniques are very important for multimedia stream delivery. The triggering technique of a multimedia stream can be implemented in several layers such as layer 2, layer 3, and the application layer.

Traditionally, configuration time to obtain the network parameters and IGMP join/leave latencies contribute to the transient data loss or waste of bandwidth. In order to maintain minimum loss and latency during the client's movement, it is desirable to minimize handoff time and provide almost instantaneous flow of the multicast stream by adopting some novel triggering mechanism. Similarly, it may be required to avoid the waste of bandwidth associated with the leave latency.

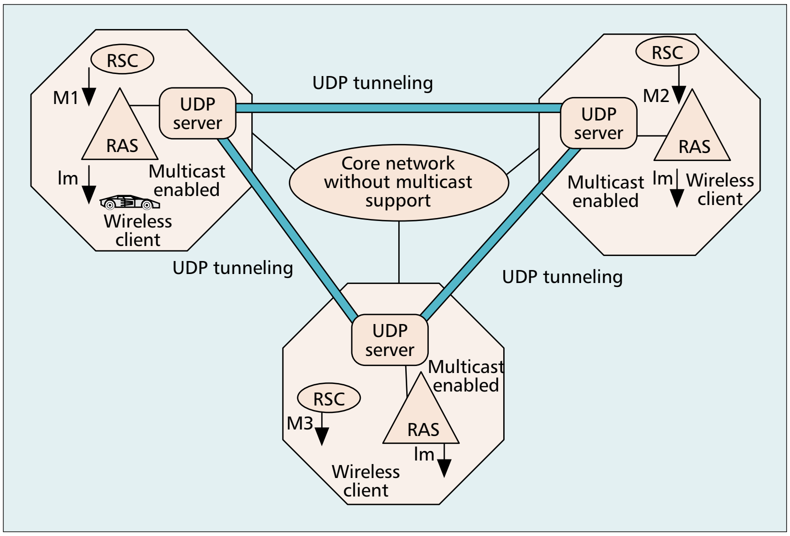

■ **Figure 4.** *Interconnection of a multicast-enabled network.*

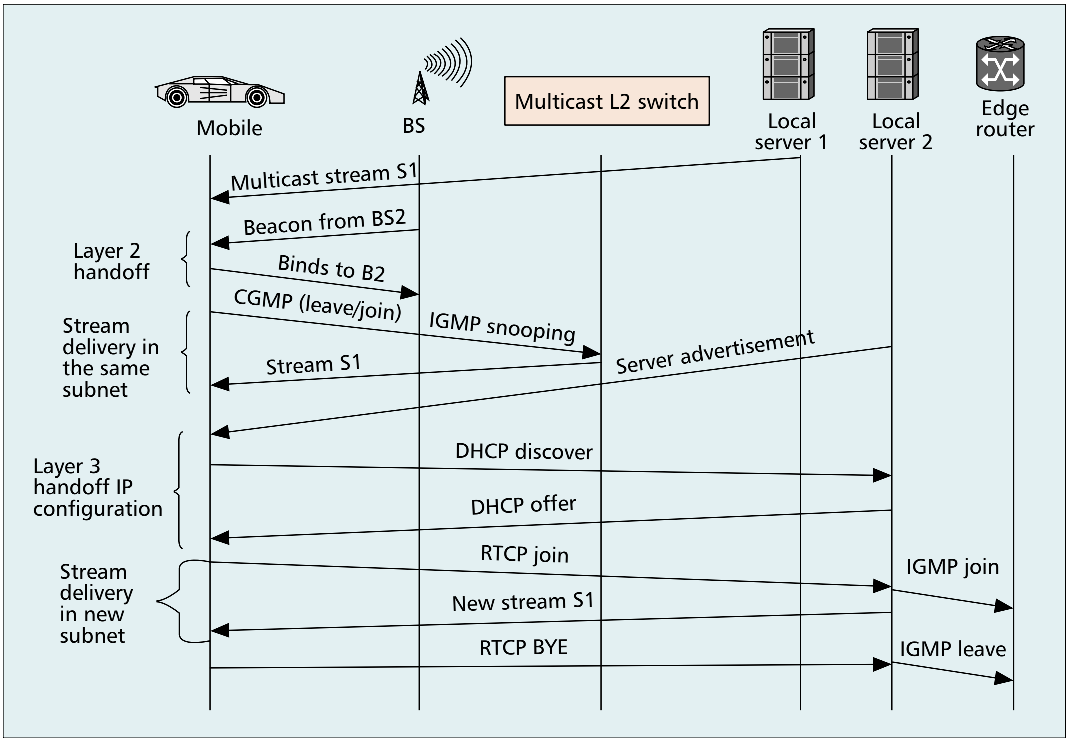

■ **Figure 5.** *Multicast stream handoff flow.*

For certain applications there may be a large number of users in the same multicast group, and this may become a problem if the backbone network and access network are not multicast-enabled because duplication of packets and multicast streams will occur.

Some previous work discusses the group join/leave behavior in the Internet, the effect of channel surfing, and mobility of a multicast stream. Almeroth *et al.* [12] describe the multicast group behavior in the Internet, and cite some results on surfing by looking at Mbone statistics. It shows that within a time interval of 2 min, a user leaving one session either joins another session or becomes inactive. Although this is very similar to a mobility model, where the user leaves one group and rejoins the same group in the next cell, the study does not take into account the mobility of the users and the associated parameters. However, Varshney and Chatterji *et al.* [13] describe many of the architectural issues associated with MHs in a multicast environment.

Also, Wu *et al.* [14], propose ways to handle fast delivery of the multicast stream when the end hosts are mobile within a domain. It proposes a solution of handover with preregistration in order to provide fast handoff for the multicast streams while moving between subnets. This is accomplished by sending a unicast datagram to the neighboring station about the multicast address to which the client is subscribed, so the neighboring station will have joined the multicast tree even before the client moves into the neighboring cell. This solution assumes there is an agent (mobility support agent) in each subnet that invokes the join message.

Layer 2 multicast-based triggering comes into the picture when the adjacent cell to which the client moves belongs to the same subnet. Here it does not have to spend time obtaining a new IP address; rather, handoff is handled at layer 2 with the help of coordination between the adjacent access points. Several methods such as IGMP snooping or Cisco Group Management Protocol (CGMP) can be used to take care of layer 2 handoff for multicast streams.

If the destination cell belongs to the same subnet and both cells are served by the same local server, multicast stream flows to both cells in the absence of any multicast switch. Although this helps reduce triggering time, it will contribute to waste of bandwidth if there is no active participant in the adjacent cell. Typically, as soon as the change of base station is detected, CGMP Join/Leave is triggered, thus starting the flow of a multicast stream or keeping it from flowing to the next cell.

While using a layer 3 triggering method, triggering delay consists of an IGMP query report after the node moves to the new cell to be part of the same multicast tree. A typical query interval for IGMP is by default 125 s, although this value is configurable in multicast routers. In order to avoid flooding the LAN with IGMP messages, this value cannot be made very small. Reference [14] shows that by using IGMP, a host will wait for 65 s on average in order to continue to receive the multicast traffic after a handover. This is because IGMP was not designed for roaming clients in a wireless environment. Similarly, typical leave latency once the host has moved to a new subnet is about 2 min; traffic would still flow to the previous cell even after the client has moved out. Kaur *et al.* [15] propose modification to IGMP for faster join and leave latency.

Similarly, application layer triggering techniques can be used when clients are involved in a real-time communication involving RTP/UDP. MarconiNet [3] implements all three kinds of triggering techniques to support cell and subnet mobility using the RTCP feedback mechanism. However, it is to be noted that performance of application layer triggering techniques will depend on the processing power of the end hosts.

Scalability

Scalability often refers to the ability to support large groups without possible degradation of quality. Scalability is an important concern while designing a mobile content distribution network. For certain applications there may be a large number of users in the same multicast group, which may become a problem if the backbone and access networks are not multicast-enabled because duplication of packets and multicast streams will occur. For location-specific applications it is more likely that there will be many distinct multicast groups (i.e., one for each local area). Management of multicast addresses for the groups is an important issue. Both time to live (TTL) scoped and administratively scoped address management may be considered to avoid multicast address management or overlapping problems, thus increasing the scalability factor.

Load Balancing

In a content distribution network, there may be multiple local content servers that will need to coordinate among each other to transfer multimedia content to the mobile client as it moves between cells. When one particular server is heavily loaded, the adjacent server needs to be able to direct the multimedia content to the mobile client that is part of the multicast stream. An alternate server can be selected based on the location of a Global Positioning System (GPS)-equipped mobile. Participating clients can also take part in delivering cached on-demand or live content to other neighboring clients in the case of server overload.

Quality of Service

Maintaining the same QoS as the mobile traverses subnets while being part of the same group is an important issue address. When a different server is serving the client after handoff, it is important that the server be made aware of the bandwidth requirement of the client ahead of time. Thus, the server can institute a proper bandwidth control mechanism at either its interface or the downstream router based on a request from the impending end client. A differentiated services (DiffServ)-based mechanism can be used to handle QoS for intradomain mobility.

Location Awareness

If a user is to be provided location-specific information, the user's geographical location must be known to the streaming source or the multicast proxy filtering the information to the user so that the local information can be communicated to the media server that will filter the local news and traffic information accordingly. If IP address assignment is location-specific, the network provider can supply the media servers with a database of IP address pools for different locations or access networks. In a mobile environment several types of triangulation scheme (i.e., 802.11 triangulation) can be utilized to exactly pinpoint the location of a mobile. Based on its coordinates a GPS equipped terminal can easily locate the nearest server that can serve it. Recent work within IETF's Geopriv working group on associating IP address with the GPS coordinates as part of a DHCP option will help achieve this.

CONTENT CACHING

Caching is a process by which streaming content gets dynamically replicated closer to the users to provide better quality. For an on-demand streaming session, sources use multicast to reduce bandwidth usage in the network, but this may introduce delay for the clients until multicasting starts. Hierarchical caching lowers the latency and bandwidth usage of streaming media being delivered to the client. Thus, a hierarchical regional cache server will be able to reduce this initial playout delay at clients by sending the prefix of a requested stream while waiting to get the multicast stream. Prefix caching at the proxies augments fast triggering techniques when a mobile moves from one cell to another and in the process gets served by a new multicast proxy. Using an RTSP server can provide a typical streaming cache proxy. However, there are still issues related to caching such as transfer loss, transformation loss, cache coherency, access accounting, authorization, and copy protection.

SECURITY

While multicasting multimedia stream, a source may want to encrypt a specific stream (audio or video) based on the type of program and nature of the audience. Secured content distribution may be desirable in many cases to promote a proper business model between the content providers, local affiliate, and client in a commercial environment. Group key management and media encryption are two very important factors for mobile multicast. Because of the one-to-many nature of the application, key distribution can be made possible by adopting a centralized key management architecture complemented by SAP-based security association. Encryption for media can be provided at different layers such as 802.11-based WEP , IPSec, and Secured RTP in transport layer. Since a mobile client is subjected to heterogeneous access (e.g., 802.11, CDMA1XRTT) and IP address change as a result of repeated handoff, it becomes difficult to maintain layer 2 and 3 based security association. In a mobile multicast environment secured RTP (SRTP) is preferred over standard IPSec encryption as it avoids the encapsulation overhead and tunnel setup time associated with IPSec.

CONCLUSION

In this article we have used multicast technology as one of the approaches to build a content distribution network for mobile users in the wireless Internet. We have surveyed related techniques and architectures that can be used to provide multicast steaming media to clients in a highly mobile environment. Many of the issues involved in building a content distribution network for mobile users using multicast technology have been discussed. Successful deployment of mobile content distribution using a combination of application layer and network layer multicast technology will usher in a new era for mobile commerce.

An alternate server can be selected based on the location of the mobile that is GPS-equipped. Participating clients can also take part in delivering cached on-demand or live content to other neighboring clients in case of server overload.

REFERENCES

[1] H. Schulzrinne *et al.*, "RTP: A Transport Protocol for Real-Time Applications," RFC 1889, IETF, Jan. 1996.
[2] C. Diot and Brian Levine *et al.*, "Deployment Issues for the IP Multicast Service and Architecture," *IEEE Network*, Mar. 2001.
[3] A. Dutta and H. Schulzrinne, "A Streaming Architecture for Next Generation Internet," *ICC*, Helsinki, June 2001.
[4] C. L. Tan and S. Pink, "Mobicast: a Multicast Scheme for Wireless Networks," *Mobile Nets. and Apps.*, 1999.
[5] T. G Harrison, C. L Williamson, and W. L Mackrell, "Mobile Multicast (MoM) Protocol: Multicast Support for Mobile Hosts," *Proc. Mobicom '97*.
[6] C. R Lin and K.-M. Wang, "Mobile Multicast Support in IP Networks," *IEICE Trans. Commun.*, vol. E-84-B.
[7] J. Mysore and V. Bharghavan, "A New Multicasting-based Architecture for Internet Host Mobility," *Proc. Mobicom '97*.
[8] A. J. McAuley *et al.*, "Mobile Multicast Proxy," *MILCOM '99*.
[9] Y. Chawathe, "Scattercast: An Architecture for Internet Broadcast Distribution as an Infrastructure Service," Ph.D. dissertation, UC Berkeley, Dec. 2000.
[10] V. Padmanabhan *et al.*, " Distributing Streaming Media Content Using Cooperative Networking," *NOSSDAV 2002*, May 12–14, Miami, FL.
[11] J. Chennikara *et al.*, "Application Layer Multicast for Mobile Users in Diverse Networks," *IEEE GLOBECOM 2002*.
[12] K. Almeroth and M. Ammar "Multicast Group Behaviour in the Internet's Multicast Backbone," *IEEE Commun. Mag.*, June 1997.
[13] U. Varshney and S. Chatterji, "Architectural Issues to support Multicasting over Wireless and Mobile Networks," WCNC, 1998.
[14] J. Wu *et al.*, "An IP Mobility Support for 4GW wireless Infrastructure," *PCC Workshop '99*, Nov. 1999.
[15] S. Kaur *et al.*, "Multicast Support for Mobile IP Using a Modified IGMP," *Proc. IEEE WCNC '99*, Sept. 1999.

BIOGRAPHIES

ASHUTOSH DUTTA [SM] (adutta@telcordia.com) is currently a research scientist in Telcordia Technology's Internet Network Research Laboratory, Morristown, New Jersey. For the past 15 years he has dealt with a variety of high-speed networks and computer systems, and has been responsible for designing and implementing many enterprise networks, and wireless and mobile computing related projects. Prior to joining Telcordia Technologies, he was director of Central Research Facilities at Columbia University, from 1989 to 1997. His research interests include session control protocols, streaming multimedia, wireless multicast, and mobile wireless Internet. He was a recipient of 2000 and 2002 Telcordia CEO Awards and winner of SAIC's ESTC 2002 best paper award in information and technology. He has been co-project leader of a DARPA- funded Airborne Communication Node project, and technical lead for integrated mobility management for the project. He has a B.S. in electrical engineering (1985) from NIT Rourkela, India, an M.S. in computer science (1989) from the New Jersey Institute of Technology, and a professional engineering degree in electrical engineering from Columbia University. He is currently pursuing his Ph.D. degree at Columbia University related to MarconiNet and streaming multimedia. He is also a member of ACM and is currently serving as secretary of the IEEE Princeton and Central Jersey section.

JASMINE CHENNIKARA received a B.S. degree in electrical engineering from Manhattan College, New York, in 1994. She received her M.S. and Ph.D. in electrical engineering from Polytechnic University, Brooklyn, New York, in 1996 and 1999, respectively. Since 1999 she has been working at Tel-

Successful deployment of mobile content distribution using a combination of application layer and network layer multicast technology will usher a new era for mobile commerce.

cordia Technologies, Morristown, New Jersey, as a research scientist. Over the past few years she has worked on QoS and traffic prioritization issues for high-speed networks. She has also done research on call admission, and traffic modeling issues for satellite and wireless networks. More recently, she has worked in the areas of multicasting, dynamic virtual private networking, resource management, and information assurance for wireless networks. She was a recipient of the Telcordia CEO award in 2000 and 2002. She was also a member of the team that won the 2001 SAIC Executive Science Technology Committee Publication Award in information and communications technology.

ONUR ALTINTAS [M] is a corporate manager in the R&D Group of Toyota InfoTechnology Center, United States. From 1995 to 1999 he was a research scientist at Ultra High Speed Network and Computer Technology Labs (UNCL), Tokyo. Before joining Toyota InfoTechnology Center in 2001, he was with Toyota Motor Corporation. He received his B.S. (1987) and M.S. (1990) degrees from Middle East Technical University, Ankara, Turkey, and his Ph.D. (1995) degree from the University of Tokyo, Japan; all in electrical engineering. He is a member of ACM and MAA.

WAI CHEN [SM] is program director of the Ubiquitous Mobile Computing and Networking Research Group in Telcordia's Applied Research. The group conducts research in automotive networking, mobile networking middleware, and mobile ad hoc networks, among others. He has been project manager for an automotive networking project with Toyota ITC since 2000. He has also been a principal investigator of a U.S. Army Research Laboratory project on mobile ad hoc networks since 2001. He received a B.S.E.E. degree from Zhejiang University, Zhejiang, China (1982); and M.S.E.E. (1984), M.Phil. (1986), and Ph.D. degree (1989) in electrical engineering from Columbia University, New York. He is a member of Sigma Xi, the American Mathematical Society (AMS), and the Mathematical Association of America (MAA).

HENNING SCHULZRINNE [F] received his undergraduate degree in economics and electrical engineering from Darmstadt University of Technology, Germany, his M.S.E.E. degree as a Fulbright scholar from the University of Cincinnati, Ohio, and his Ph.D. degree from the University of Massachusetts, Amherst. He was a member of technical staff at AT&T Bell Laboratories, Murray Hill, and an associate department head at GMD-Fokus (Berlin), before joining the faculty of the Computer Science and Electrical Engineering Departments of Columbia University, New York. His research interests encompass real-time multimedia network services in the Internet, and modeling and performance evaluation. He is a division editor of *Journal of Communications and Networks*, an editor of *IEEE/ACM Transactions on Networking*, and a former editor of *IEEE Internet Computing* and *IEEE Transactions on Image Processing*. He is a member of the ACM SIGCOMM Executive Committee, former chair of the IEEE Communications Society Technical Committees on Computer Communications and the Internet, and has been technical program chair of Global Internet, INFOCOM, NOSSDAV, and IPtel. He was also a member of the Internet Architecture Board (IAB). Protocols co-developed by him are now Internet standards, used by almost all Internet telephony and multimedia applications. His research interests include Internet multimedia systems, quality of service, and performance evaluation.