

# MarconiNet: Overlay Mobile Content Distribution Network

*Ashutosh Dutta, Telcordia Technologies Inc.*

*Henning Schulzrinne, Columbia University*

## ABSTRACT

Flexible multimedia streaming such as advertisement insertion, location-based services, mobility, and wireless access are vital components that make existing Internet radio and TV networks more attractive to roaming users. All of these applications also provide added value to telematics and military usage, including coordination, education, situation awareness, distributed simulation, battlefield communication, and multiplayer games. While content distribution over a wired network can be realized by instituting proxies and gateways at several parts of the access network, to provide mobility over heterogeneous wireless access we need to consider many operational issues such as handoff, join and leave latency, and desired QoS level for mobile clients. This article discusses some application layer techniques that help build a flexible mobile content distribution overlay type network supporting a multitiered payment and security scheme. The proposed streaming network, called MarconiNet, is based on standard IETF protocols (e.g., SIP, SAP, SDP, RTSP, and RTP/RTCP) that help provide signaling, session announcement, session description, stream control, media delivery, and feedback control, respectively.

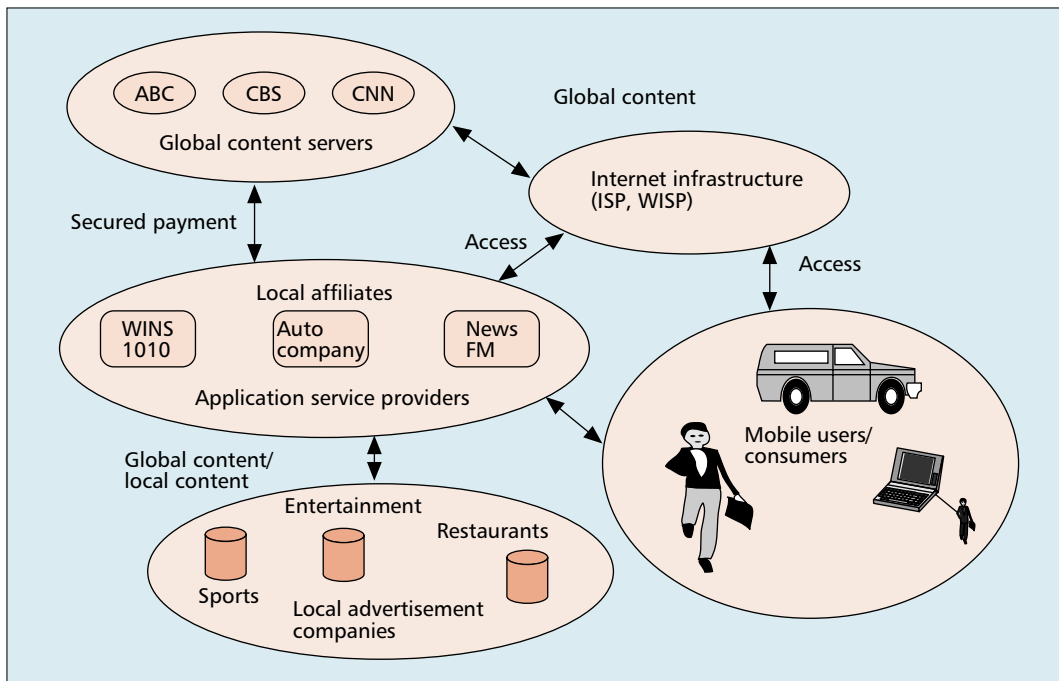
## INTRODUCTION

Streaming real-time multimedia content over the Internet is gaining momentum in the communications, entertainment, music, and interactive game industries. Streaming applications include IP telephony, broadcasting multimedia content, multiparty conferences, collaborations, and multiplayer games. These applications are far more demanding in terms of bandwidth, loss, latency, and reliability than traditional data applications, and are thus ideal drivers for the next-generation Internet.

Wireless and mobility are added dimensions to a multimedia streaming network. The wireless medium adds a constraint on the bandwidth requirement, while mobility affects the delay and transient loss for the multimedia stream delivery to a great extent, because of associated continuous handoff. A multicast infrastructure can alleviate many of the scalability issues with respect

to content distribution and will be desirable in third-/fourth-generation (3G/4G) wireless networks. However, in order to support flexible streaming services for a 4G type network, it is necessary to take care of several operational issues. These include distributing the bandwidth-intensive application to the end users over heterogeneous wireless access, maintaining continuity of the multimedia stream with minimum join/leave latencies, providing proper quality of service (QoS) to the mobiles as these move between subnets/domains. Performance can be optimized by using distributed local servers, and use of hierarchical and localized multicast. Localized multicasting becomes more attractive for the mobile users experiencing intradomain handoff because of its ease of deployment and its ability to provide more flexible services such as localized advertisement, news broadcast, and location-specific information in a bandwidth-limited wireless environment. Earlier attempts at next-generation radio and TV largely failed because they were closed systems. Current versions of Web radio or Web TV are not designed for large-scale multicast, cannot support low-latency constraints, and lack flexible programming and mobility support. More recently there has been an effort to introduce digital AM technology through a consortium called Digital Radio Mondiale (DRM) to provide added flexibility such as independent editorial control, coverage control, and efficient frequency planning.

MarconiNet [1] proposes a framework for an integrated streaming architecture to support multimedia applications such as IP telephony and broadcasting streaming content over the Internet. This article investigates some of the operational issues associated with supporting mobile content delivery and describes how MarconiNet can provide a viable business model for a 4G overlay network. MarconiNet supports flexible streaming multimedia by using several application layer multicasting techniques based on the standard Internet Engineering Task Force (IETF) protocols such as Session Initiation Protocol (SIP), Session Announcement Protocol (SAP), Session Description Protocol (SDP), Real-Time Transport/Transport Control Protocol (RTP/RTCP), and Real-Time Streaming Protocol (RTSP). SIP, SAP, and SDP, working



■ **Figure 1.** A business case scenario for MarconiNet.

There are several proposed schemes to provide native IP multicast routing over a wide area network. However, these mechanisms suffer from lack of wide deployment since there are many issues such as pricing, QoS, and maintenance of the router states in the core of the network.

in conjunction with each other, help set up, announce, and describe multimedia sessions, respectively. RTP/RTCP and RTSP help deliver and control multimedia traffic.

This article is organized as follows. The following section gives a description of related work. We then provide an overview of MarconiNet and discuss several architectural components. We discuss some of the operational issues associated with mobile content distribution in the context of MarconiNet, and next detail fast handoff techniques associated with seamless delivery of localized multicast stream. We describe some implementation details and results in an experimental testbed. The final section concludes the article.

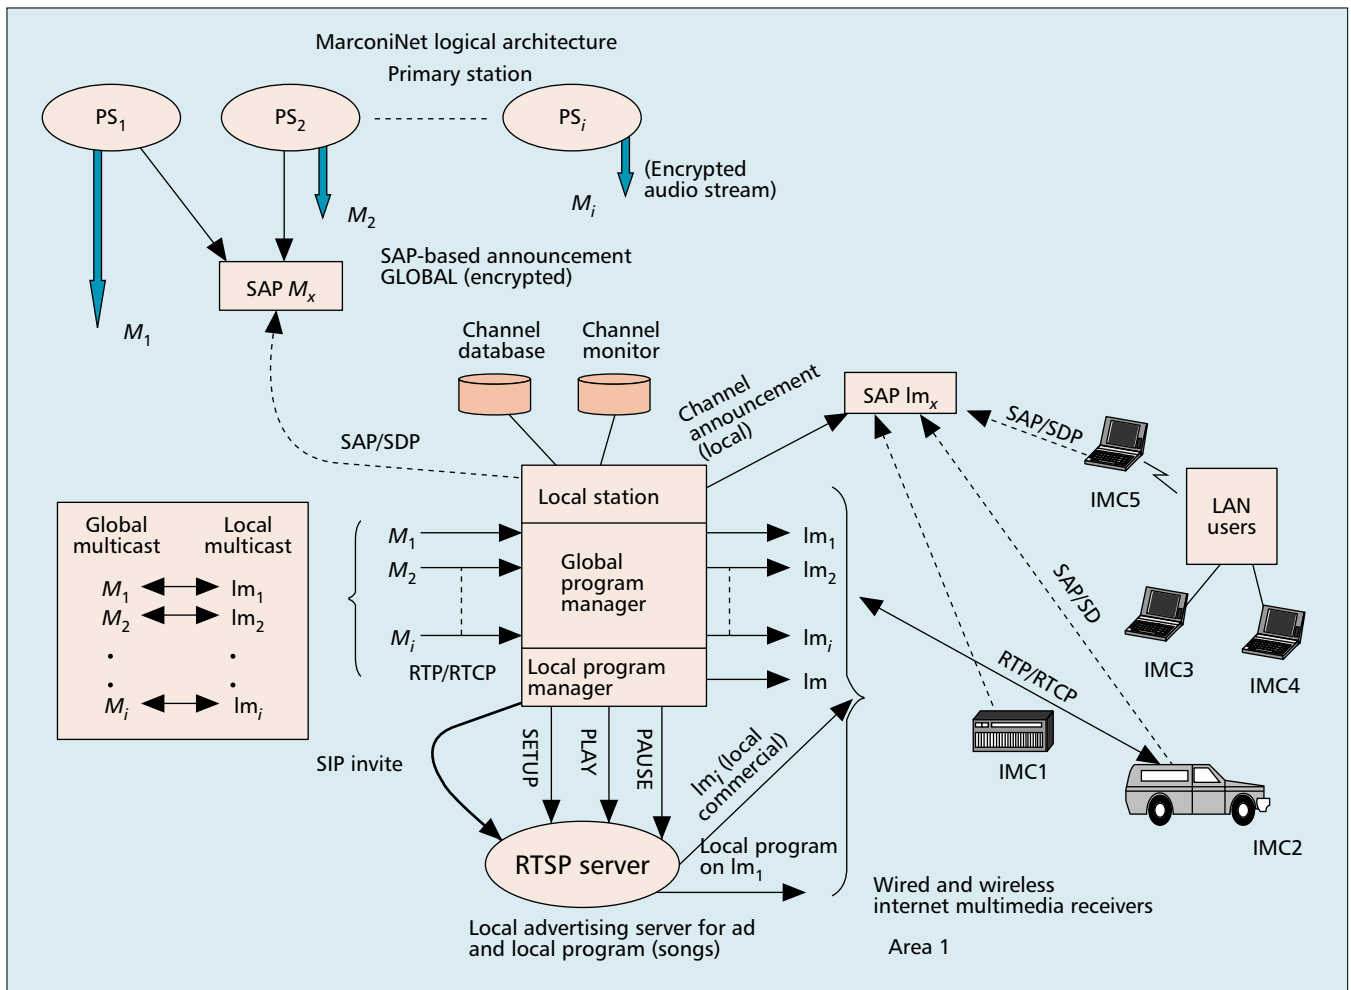
## RELATED WORK

There are several proposed schemes to provide native IP multicast routing over a wide area network such as Protocol Independent Multicast (PIM), Multicast Open Shortest Path First (MOSPF), Distance Vector Multicast Routing Protocol (DVMRP), Core Base Tree (CBT), and Border Gateway Multicast Protocol (BGMP). However, these mechanisms suffer from lack of wide deployment since there are many issues such as pricing, QoS, and maintenance of the router states in the core of the network as cited in [2]. Certain newly developed IETF-based protocols such as Source-Specific Multicast (SSM) and Explicit Multicast are suitable for broadcast type services. Source-Specific Multicast (SSM) addresses some of the issues, such as multicast address allocation, destination unawareness, interdomain routing, source advertisement, and connection state, that create a huge multicast forwarding table. SSM is ideal for Internet broadcast applications as it allows content providers to support services without requiring a unique IP

multicast address. Explicit multicast is useful for maintaining a large number of multicast groups when membership in each group is small.

Basso *et al.* [3] describe an overlay-type streaming network called PRISM, but it lacks mobility support. Almeroth *et al.* [4] describe multicast group behavior in the Internet by looking at Mbone's temporal statistics, although they do not take mobile users into consideration. Reference [5] describes many of the architectural issues associated with mobile hosts in a multicast environment. Reference [6] proposes a solution for fast handover with preregistration. In this solution there is an agent (Mobility Support Agent) in each subnet that invokes the join message on behalf of the impending mobile after being informed ahead of time. Mobile Multicast Proxy [7] uses a multicast proxy, where the proxy's clients do not themselves directly participate in the multicast tree; the multicast proxy participates in multicast tree formation for the groups of which its clients are members.

Mobile Multicast [8] provides a mobility scheme for multicast sessions for wide area networking and adopts a Mobile-IP-based approach to take care of mobility of multimedia traffic. It proposes to reduce the problem in bidirectional tunneling of many home agents (HAs) tunneling the same multicast packet to a foreign network. In this case one HA will be elected to tunnel multicast packets to a foreign network. In Mobicast [9], users continue to rejoin the multicast group in each visited network. This architecture adopts the domain foreign agent (DFA) concept to shield all mobility within the foreign domain from the main multicast delivery tree. In this scenario a DFA will send or receive the multicast traffic for a multicast group. When the MS receives the multicast traffic, the DFA will use a translated multicast address within its network to prevent multicast updates due to mobility.



■ Figure 2. MarconiNet's logical architecture.

Most recently, Packet Video in conjunction with DoCoMo has started providing wireless multicast streaming services to end users, but it 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 does not take the mobility of end clients into consideration. A few years ago iBEAM's product Activecast used to distribute streaming applications using geostationary Earth orbit (GEO) satellites, but it lacked flexible methods of advertisement insertion and mixing of local and global content, and did not take care of user mobility.

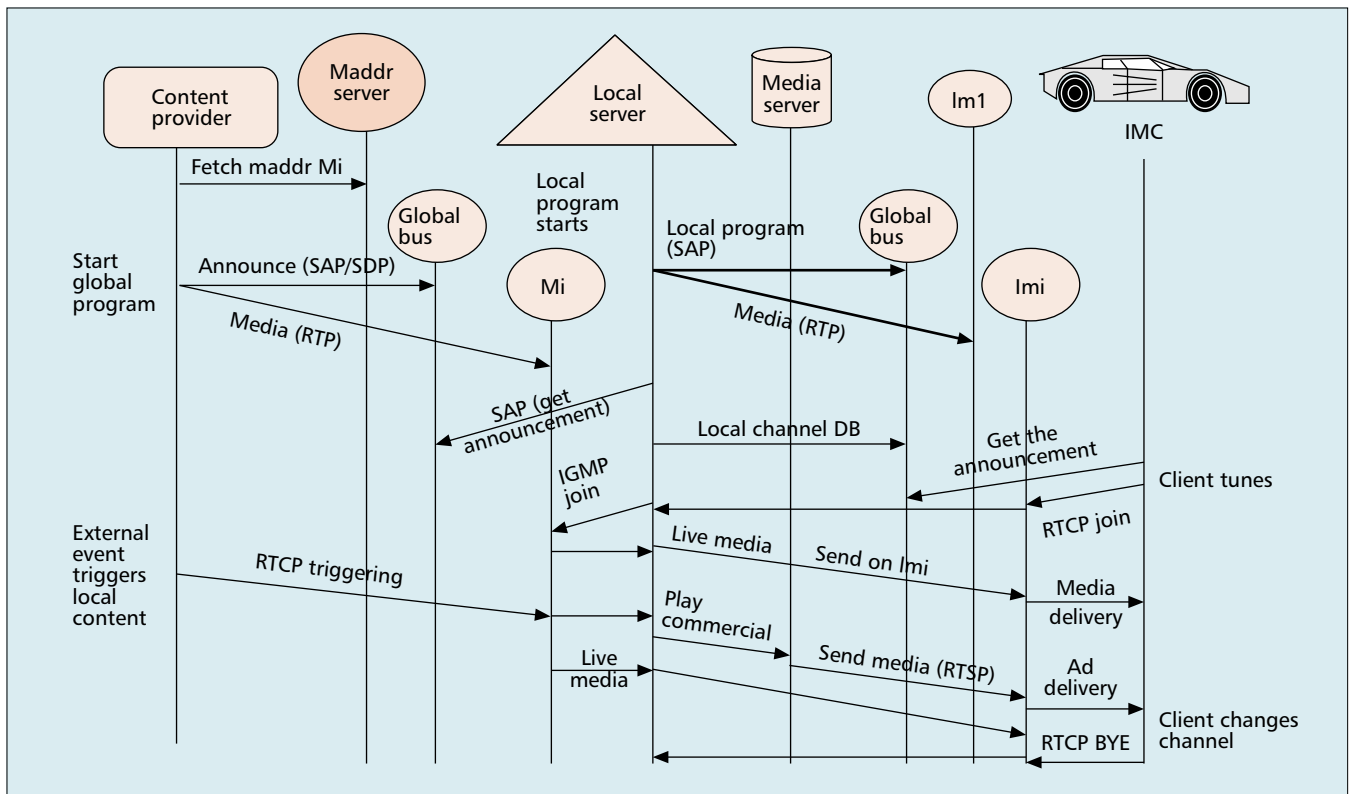
## MARCONINET OVERVIEW

MarconiNet offers flexible streaming wireless services over localized multicast while taking care of clients' mobility issues. Figure 1 presents business aspects of how MarconiNet can be helpful to end users, global broadcasting stations, local affiliates, and advertising companies. Figure 2 illustrates the logical architecture of MarconiNet, and Fig. 3 describes some of the protocol interaction.

Following is a list of some of the important features supported by MarconiNet:

- Application layer multicasting to provide scalable live and on-demand streaming content for wider/global reach
- Automatic localized ad/commercial/content insertion into a global program; switching between global and local programs
- Reduction in join/leave latency by means of application layer triggering
- Faster registration, configuration, and application layer mobility
- Smooth and fast handoff for multicast streaming in a wireless environment, conserving bandwidth
- Real-time client (listener, viewer) index for global and local program; demographic information on users
- An adaptive QoS management scheme for mobile users

As illustrated in Fig. 2, there are four basic functional components in the MarconiNet architecture: global station (primary station), radio antenna server (RAS), which is the local station, and Internet multimedia clients (IMCs). A primary station acts like a multimedia streaming source. By introducing local servers one can have full control of the program being broadcast, such as transmitting global vs. local programs,



■ Figure 3. MarconiNet protocol interaction.

and introducing local advertisement. In each cell (subnet) there is at least one local server (RAS) that converts a globally scoped incoming multicast stream to a locally scoped one destined for local clients. To take care of dynamic load balancing and handoffs, multiple local servers can be placed within the same subnet or across the subnets in a domain. The commercial server is a streaming server and is usually collocated with the local affiliate. The local affiliate and commercial server can interact to provide stream control using protocols such as SIP and RTSP.

In this section we describe some of the basic mechanisms and key functional components associated with MarconiNet.

### CHANNEL ANNOUNCEMENT

A global streaming station (e.g., radio/TV station, individual broadcaster) denoted  $RSC_i/PS_i$  broadcasts its programs potentially to a global audience live on a unique multicast address,  $M_i$ , globally scoped and encrypted using RTP/UDP. These global stations send their session announcements using a subset of SDP parameters to a global multicast address,  $M_x$ , unencrypted. This common global multicast address provides a list of programs broadcast by primary stations (RSCs) all over the world. We have designed a Java-based interface called JS DR that is more customized than traditional SDR and provides a hierarchical searching functionality.

### CHANNEL MANAGEMENT

Channel management takes care of managing global and local programs. Each RAS gets a global encryption key it uses to tune to the global common multicast address  $M_x$  to get the listing of

channels and its contents. It decides to broadcast part of the list to the local domain based on demographics, and hence creates a local announcement database. The subset of channel descriptions announced by each radio station provides sufficient data for building a local channel database. This local announcement database contains the list of the supported channels, each with their appropriate attributes such as the name of the program, duration, type of content, and place of origin. The local station uses SAP to send this program index to a locally scoped common multicast address  $lm_x$  for announcement to local clients. The announcement on  $lm_x$  is not encrypted as the local station wants all of its clients to have access to what is relayed by it. The RAS also maintains a pair of multicast addresses for each channel. It keeps the mapping of the globally scoped multicast address on which the radio station sends its program and the locally scoped multicast address where it gets relayed. The RAS receives the audio stream on the global multicast address  $M_i$  and redirects it onto the local multicast address  $lm_i$  for the IMCs. Local programs are sent on a specific locally scoped multicast address  $lm_i$ . Reports generated by RTCP packets from the participating clients can be used for many statistical purposes such as billing, audio/video quality feedback, and membership information for a particular channel.

### CHANNEL TUNING

IMCs listen to the locally-scoped common multicast address  $lm_x$  to be aware of what is available using any tool based on SAP and SDP. According to the SAP specification, the local server updates the announcement information every

*The primary station collects the public keys from the local stations and adds these to its SEK distribution list upon their payment. This hierarchical-based security system will enable an ideal charging mechanism and payment model.*

several minutes or so. Each specific channel provides details of the program being broadcast.

#### LOCAL ADVERTISEMENT INSERTION

Localized advertisement/information insertion can be provisioned or event-driven, as in an emergency situation. Each global radio station knows the starting time and duration of a commercial break ahead of time, or has control over the time for the commercial break.

On receiving the signal for a commercial break via RTCP feedback, the management server at the local station requests the local media server via RTSP to play the local advertisement to a specific locally scoped multicast address  $lm_i$  corresponding to that station. It uses a set of RTSP commands like SETUP, PLAY, and PAUSE to control stream delivery of the commercial break. During this time, the management server stops forwarding the RTP stream from  $M_i$  to  $lm_i$  in the local domain. The local advertisement runs for a specific time conveyed to the server by means of RTCP reports. If a commercial break overlaps by several radio stations, the RTSP server can play several different local advertisements on different locally scoped multicast addresses.

#### CHANNEL MONITOR

The channel monitor mechanism provides an online estimate of the audience size tuning to a specific channel or listening to a particular advertisement at any instant. This is achieved by using RTCP feedback and is similar to many of the techniques addressed in [10]. For each channel diverted by the local station, an additional RTCP signaling channel is created with a different port. Each listener sends RTCP SDES packets to notify the local station (RAS) along with the client's demographic information and the respective channel to which it is subscribed. RAS maps each listener to the desired channel, which in turn increases the number of listeners for the particular channel. The listener-to-channel mapping is destroyed via RTCP BYE packets or the RTCP timeout feature. Doing this also decreases the number of listeners for the associated channel at the local station.

#### SECURITY

The proposed MarconiNet system presently offers several levels of security association between the global station, local affiliate, and end clients. The security model for a global multicast stream effectively prevents the IMCs, as well as nonpaying RASs (local affiliates) receiving the broadcast content. Thus, each radio/TV station (RSC) must maintain a secret key and encrypt all outgoing content so that only a ciphertext stream is transmitted. The basic strategy is to generate a symmetric encryption key at the station and securely distribute this key to a particular RAS or local server upon payment. The security association between the global station and local affiliate is taken care of as follows. Each local station generates its own PubKey/PrivKey pair. Each global station generates a secured encryption key (SEK) and begins transmitting encrypted audio content over SRTP. This SEK needs to be distributed to the participating local stations securely so that other local

stations that have not paid cannot obtain it. Public key technology is employed for this purpose. The local station submits its public key to the primary station along with its payment or certificates of authority (CAs). The local station (affiliate) in turn receives an integer ID from the global station that is later used to index the SEK distribution list. The primary station collects the public keys from the local stations and adds these to its SEK distribution list upon payment. This hierarchical security system will enable an ideal charging mechanism and payment model.

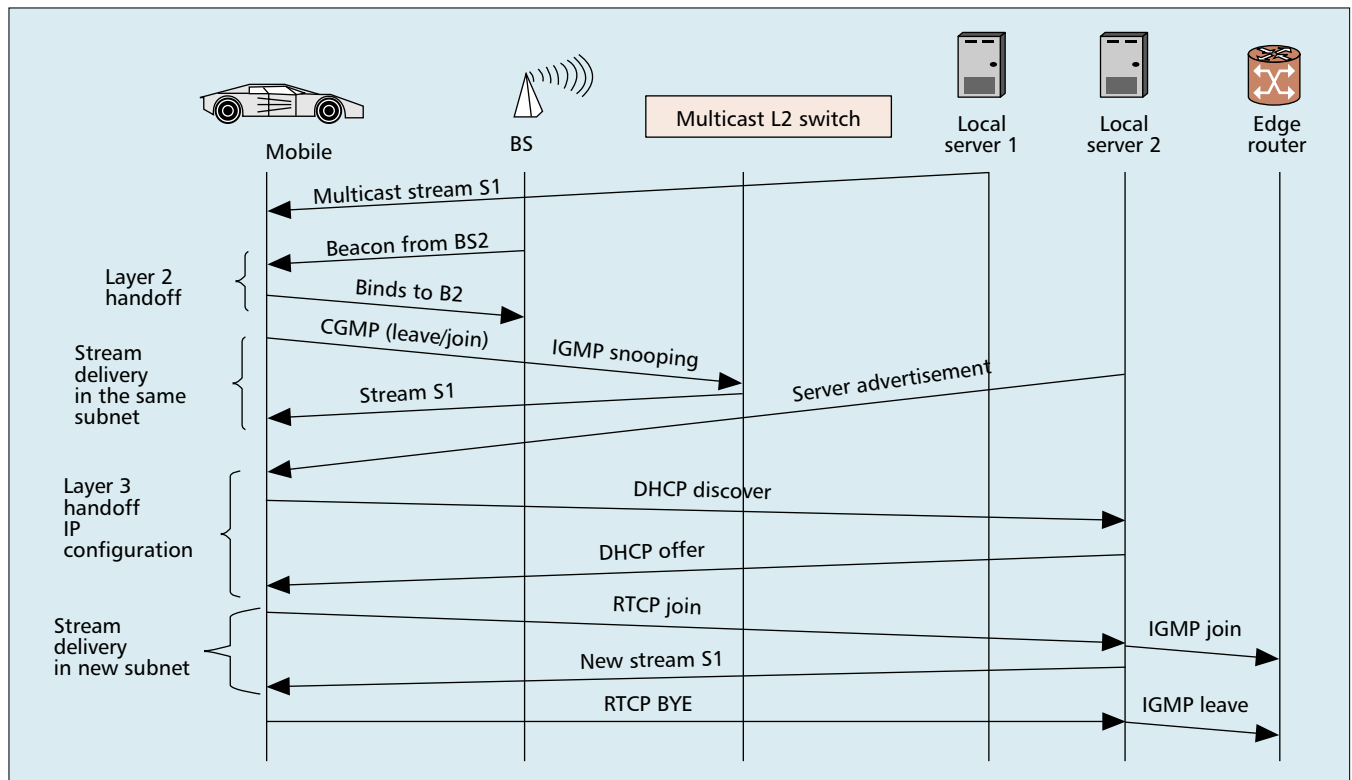
Encryption can also be extended to a local section as a second-level hierarchy and used for pay-per-listen programs announced to local multicast addresses. RTCP-based multicast group access control can be used, where an access token based on the host parameters can be sent as part of an RTCP join message before an end client is able to listen to multicast traffic similar to IGMP authentication mechanism. The security scheme also helps preserve sensitive information such as the secret keys (of RSCs and pay-per-listen channels), user accounts, and payment data. Advertising companies can be authenticated so that unauthorized companies cannot gain access to the local advertisement insertion system.

As the end client moves from one subnet/domain to another and associates itself with a different local server, it can use Protocol for Carrying Authentication to Network Access (PANA) with Diameter as the backend process to transfer the security association to the second server to which the client is moving.

#### PAYMENT MODEL

The logging mechanism helps advertisers decide and determine when and on which channels to place their commercials in order to maximize returns on investment. The higher level of advertising efficiency will drive more advertisers toward advertising on the Marconi server. This in turn will allow the Marconi server to pursue an increasing share of the Internet radio/TV advertising market by charging different prices for different segments of air time based on demand. The RTCP-based reporting mechanism enables Marconi server operators to track user listening trends and can provide to potential and existing advertisers detailed reports with users' listening statistics. RTCP is also ideally suited to allowing Marconi to gather user- and channel-specific information.

This architecture can support several kinds of payment models. Local stations (RASs) collect fees from local commercial companies to send advertisements during commercial breaks while relaying the global programs. Besides relaying the programs from the common global stations, the local stations may also relay some pay-per-listen programs, in which case it pays the global station depending on how many listeners are listening to a particular program. Every local station also broadcasts its local program to its local clients with segments of news or some other premium programs relayed from global stations. Two kinds of payment model can be designed: there may be users who like to listen to a program without being interrupted by advertisements; some would not really mind the local



■ **Figure 4** Handoff flow in MarconiNet.

advertisement. Channel monitoring, local advertisement insertion, logging mechanism, payment model, and security components associated with this architecture help provide a suitable business model to support mobile users in a 4G network.

## OPERATIONAL ISSUES WITH MARCONINET

This section describes how MarconiNet handles some of the operational issues associated with mobile content distribution.

### INTRADOMAIN MOBILITY COMPONENTS

In general, signaling and transport delays contribute to the latency of multimedia stream delivery during a node's movement. Operations 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 contribute to signaling delay, while transport delay dominates the component associated with mid-session mobility.

Latency associated with receiving continuous unicast/multicast stream from the same source while the client moves to the next cell consists of several components such as detection of a new cell/subnet/domain ( $\Delta_1$ ), address acquisition, and network configuration ( $\Delta_2$ ) using protocols such as DHCP, PPP, or Mobile IP, triggering of a multimedia stream to be delivered in the new subnet ( $\Delta_3$ ), and actual delivery of multimedia stream ( $\Delta_4$ ). While some of these factors are common to both unicast and multicast stream delivery (e.g., cell/subnet detection, IP parameter configuration), we concentrate mostly on delivery

schemes for multicast streaming traffic in the context of MarconiNet. For a mobile node in the MarconiNet environment, JOIN latency due to triggering, and mapping from a globally scoped multicast address to a locally scoped address at the RAS also add to latency during handoff.

Interdomain mobility includes additional steps for establishing security associations between the local servers in each domain, especially the steps involved in authentication and creating local security associations. When each adjacent server belongs to two different administrative domains, a combination of authentication, authorization, and accounting (AAA) protocols such as Diameter or Radius, and local authentication protocols such as Protocol for Carrying Authentication to Network Access (PANA) can be used for client authentication. Figure 4 shows several steps of the protocol sequence when a mobile is subjected to subnet handoff in an 802.11 environment using the MarconiNet scheme.

### MOVEMENT DETECTION

The process of discovering a new cell/subnet/domain can be realized in different layers. During a mobile'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 or other soft handoff technology. Using a layer 3 triggering mechanism similar to ICMP router advertisement, it can be determined if the client is in a different subnet. An application layer detection mechanism such as those described in [11, 12] can also be used if the client is involved in a real-time communication session. It may be possible to achieve faster handoff by

*In the lower hierarchy between local station and local clients, where the local station has the complete control, security is needed only for the pay-per-view type program. Security is not an issue if one is using an IMC set as a radio replacement.*

using layer 2 assisted triggering techniques, often discussed in IETF's Mobile IP working group.

### NETWORK CONFIGURATION

As a mobile moves between subnets, it needs to reconfigure itself by acquiring a new IP address, new default gateway, subnet mask, DNS server, and so on. The key requirements of the reconfiguration process include that it should not take more than a few hundred milliseconds to complete, and the DNS should get updated automatically to reflect the current name-to-address mapping and vice versa using dynamic DNS. As the client moves from one cell to another, where the new cell is another subnet, it will either obtain a new address from the DHCP/PPP server or use standard Mobile IP to obtain a new care-of address from the foreign agent. Each of these approaches will take a different amount of time ranging from a few hundred milliseconds to a few seconds to complete the network configuration. IPv6 provides a minimal delay for its stateless auto-configuration, amounting to on the order of  $\sim 1$  s.

### JOIN AND LEAVE LATENCY

The process of joining/leaving a specific multicast group while changing the cell or subnet is very similar to surfing TV/radio or flipping channels. Join and leave latency for multicast sessions during a mobile's movement between wireless cells or switching channels while within a cell contribute to the overall latency.

Since multicast communication is receiver-initiated, triggering techniques are very important for multimedia stream delivery. In order to maintain minimum loss and latency during a client's movement, it is desirable to minimize the handoff time and provide almost instantaneous flow of a multicast stream by adopting some novel triggering mechanism. Similarly, it may be required to avoid bandwidth waste due to the continuous flow of traffic associated with leave latency. A triggering technique for multimedia streams can be provided in several layers, including layer 2, layer 3, and the application layer.

In the traditional layer 3 method, triggering delay typically depends on IGMP query report after the node moves to a new subnet. A typical query interval for IGMP is by default 125 s, although this value is configurable in multicast routers. By using IGMP, a host will wait for 62.5 s on average in order to continue to receive multicast traffic after handover. Similarly, a typical leave latency once the host has moved to a new subnet is about 2 mins; that is, traffic still flows to the previous cell even after the client has moved out, thus wasting bandwidth in the previous cell. Thus, standard IGMP needs some modification in order to make it suitable for mobile users.

Layer 2 multicast comes into play for movement between cells within a subnet. In such a scenario triggering of a multimedia stream can be taken care of by Cisco Group Management Protocol (CGMP) or IGMP snooping. This layer 2 triggering mechanism also allows multicast traffic to be filtered out at the multicast switch to save valuable wireless bandwidth when there is no active participant in the adjacent cell.

The triggering technique for multicast multimedia streaming in the MarconiNet architecture takes advantage of triggering in all three layers to improve latency. Several layer 3 and 4 techniques, discussed below, help reduce the transient loss.

### KEY EXCHANGE

As discussed in the overview section, there are two levels of hierarchy in the distribution of a multimedia stream. Encryption and authentication between local affiliates and global stations were discussed earlier. Secured key exchange can be provided by using a variation of Group Key Management Protocol (GKMP) that generates and maintains symmetric keys for the members of a multicast group. This key can be used for encryption and group authentication. Each multicast group has a dedicated *group controller* responsible for managing the group keys. Most recently, the MMUSIC working group of IETF is working on adding extensions to SDP to support key exchange. Any session announcement can be protected by using the authentication header defined in SAP. However, Multicast Key Management Protocol (MKMP) can provide this key distribution in a more scalable way.

In the lower hierarchy between a local station and local clients, where the local station has complete control, security is needed only for pay-per-view type programs. Security is not an issue if one uses an IMC set as a radio replacement. In a pay-per-view scenario media can still be encrypted using Secured RTP (SRTP), and the keys can be distributed using Internet Key Exchange (IKE). IKE can work in conjunction with PANA-based protocols to provide local-client-based authentication. The MMUSIC working group within IETF is also considering several ways to protect this key during distribution by using different methodologies such as S/MIME and Multimedia Internet Keying (MIKEY).

### QoS MANAGEMENT

It is desirable for a mobile that is part of a multicast communication to be able to maintain the same QoS as it moves across subnets within a domain, and connect to a different server. In order to minimize disruption it is important that the local server is made aware of the bandwidth requirement of the client's current application proactively. We propose an approach combining IETF-based differentiated services (DiffServ) and application layer triggering based on RTCP feedback using RTCP Receiver Record (RR) and SDES. Our approach uses DiffServ by instituting policy control [13] in the edge routers/servers, but relies on the RTCP feedback report to maintain QoS parameters for the multicast streaming traffic. As part of its RTCP feedback, each mobile keeps on advertising its QoS parameters (e.g., delay, jitter) on the locally scoped multicast address using SAP. This information gets advertised in the globally scoped multicast address (M) where both adjacent servers can determine the QoS parameters for each mobile about to move. If a mobile listening to a specific multicast address happens to be the first client in the adjacent cell, it needs to make a reservation a priori to receive the desired QoS for the streaming traffic as soon as it moves in

there. In the testbed, Linux's traffic control (tc) mechanism was used to set up the traffic priority destined for a particular multicast address in one of its interfaces. While within a cell the client can also negotiate with the Linux server to control the traffic rate by making additional real-time requests. Protocols such as General Internet Messaging Protocol for Signaling (GIMPS) discussed in the NSIS working group within IETF can also be used to provide QoS information of the client to the upstream server.

### SUPPORTING HETEROGENEOUS ACCESS

MarconiNet is designed to support receivers over a variety of access technologies such as 802.11b, code-division multiple access (CDMA), General Packet Radio Service (GPRS), and 802.3, supporting a variety of speeds (11 Mb/s, 384 kb/s, 144 kb/s, and 10 Mb/s, respectively). A receiver-based adaptation scheme such as layered encoding and multiple multicast groups for one session [3] can be used to take care of receivers with different QoS requirements. Mobility between different types of access networks needs to take into account the delay factor associated with layer 2 detection and IP address configuration. As an example, a handover between a LAN-based 802.11 network to a WAN-based CDMA1XRTT network takes about 15 s, about 4 s in reverse. Multi-interface mobility management can help reduce the transient data loss due to time associated with IP address acquisition.

### LOAD BALANCING

In a content distribution network, multiple content servers may need to coordinate among each other to transfer the desired multimedia content to the mobile. If one particular server gets heavily loaded the adjacent server needs to be able to direct the multimedia content to a mobile that is part of a multicast stream. RTCP statistics provide feedback about the status of the server with which the client is communicating at any particular point of time. If a client experiences congestion because of server overload, it can switch back to the secondary server configured in the client during its IP address acquisition process. A server that is currently serving a client can also help direct the multimedia stream from a less overloaded server.

### MOVEMENT BETWEEN DIVERSE NETWORKS

Multicast support between autonomous systems spanning over multiple service providers is still not readily available because routers may not have inherent multicast support. There are several solutions based on user- or network-level applications such as UDP Multicast Tunneling Protocol (UMTP) and Automatic Multicast Tunneling (AMT) that help connect these multicast-enabled islands. Simple UDP servers can also be used to connect autonomous systems that are multicast-enabled but whose core network does not have underlying multicast support. As an alternative to a wired core network, companies such as Inktomi and Coolcast are already providing multicast services through satellite to reach a wide range of users. Similarly, in order to support movement from a multicast-enabled network to a non-multicast-enabled network, the

tunnels have to be set up on the end clients proactively in cooperation with the gateway proxies. A typical technique to support mobile users between such diverse networks (e.g., multicast and non-multicast) has been proposed in [15], where the proxy helps set up the tunnels between content servers and the end client as it prepares to move to a non-multicast-enabled network.

## FAST HANDOFF TECHNIQUES

Figure 5 shows a typical address assignment where  $I_a$ ,  $I_b$ , and  $I_c$  are the globally known subnets connected to the primary interfaces of the servers, and  $i_a$ ,  $i_b$ , and  $i_c$  are the local subnets connected to the secondary interfaces of the servers and could be local. According to the MarconiNet model, a server receives a multicast stream through its global interface and redirects it out through a local interface for the local clients in each cell. In this particular picture, S1, S2, ..., S5 are local stations (servers) connected to upstream routers. Each server (with the exception of S2 and S3) is connected to a different subnet and has a separate interface. S2 and S3 are connected to the same subnet via a multicast switch that takes care of layer 2 multicast. Policy-based mapping can be implemented in the adjacent local station/antenna servers to ensure that there is an identical one-to-one mapping between global and local addresses at each local station.

The following subsections provide several methods associated with faster stream delivery in the MarconiNet environment.

### POST-REGISTRATION

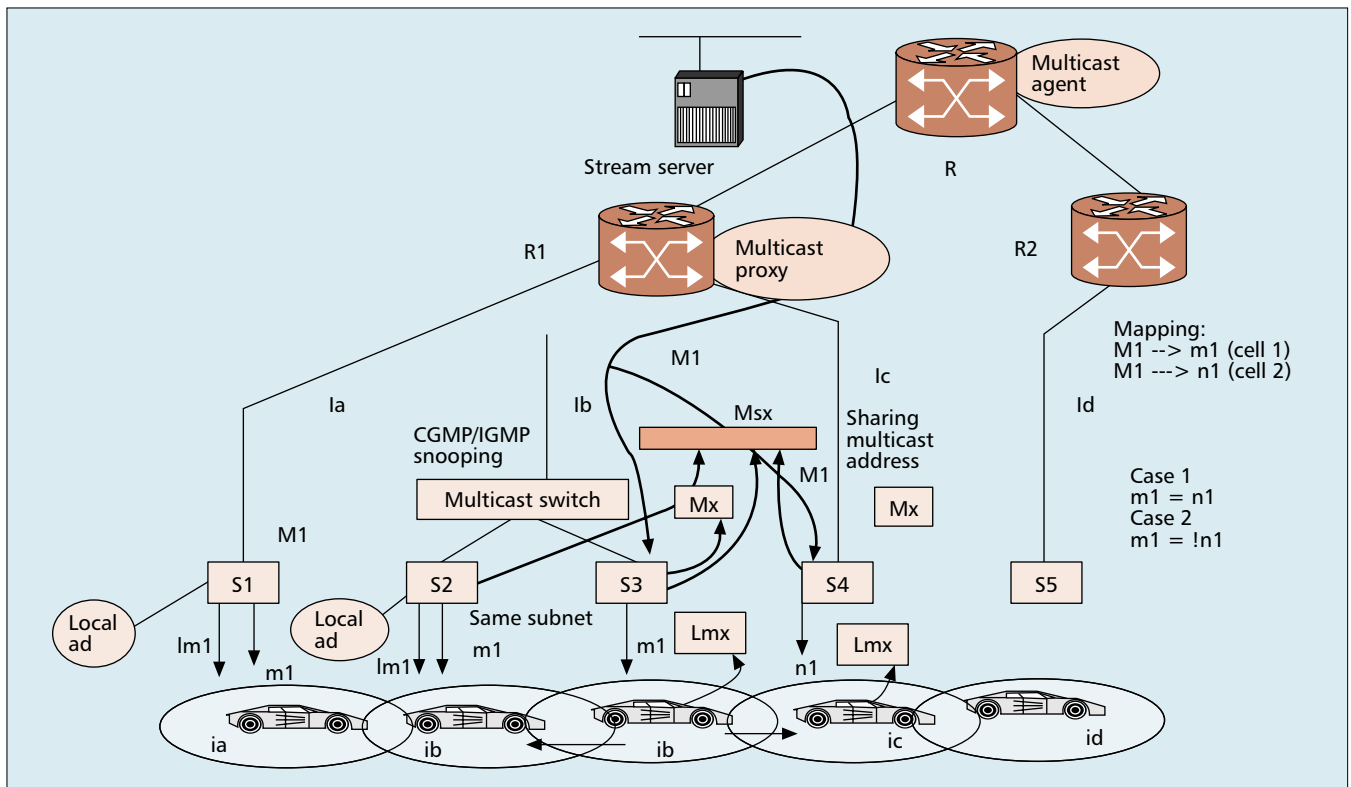
In a post-registration scenario, when client moves to a new subnet it obtains a new IP address and then sends a join query via IGMP. In some cases IGMP could be modified to provide an aggregate group report to expedite join latency. A similar technique can also be applied here. However, in the MarconiNet environment we propose a new mechanism where the client uses an application layer triggering mechanism based on RTCP in order to facilitate the join and leave. Triggering at the lower hierarchy is accomplished by RTCP feedback, but the local server may trigger the multicast flow from the upstream router using IGMP. Using an RTCP-based triggering technique offers a solution at the user space and thus is most suitable where the end client's kernel does not need to be multicast-enabled. Besides, this method does not depend on the join/leave latency.

### PRE-REGISTRATION METHOD

Pre-registration has an advantage over post-registration since a multimedia stream gets triggered ahead of time. It helps reduce the join latency for an impending client at the expense of extra stream flow in the adjacent cell for a certain time. Two kinds of pre-registration have been designed in the MarconiNet environment.

**Pre-Registering a Multicast Address Announcer** — For each neighboring station sharing an overlapping area with another station there is a multicast announcement (address) associated that can be preprovisioned. Each

*MarconiNet is designed to support receivers over a variety of access technologies such as 802.11b, CDMA, GPRS, and 802.3, supporting a variety of speeds (11 Mb/s, 384 kb/s, 144 kb/s, and 10 Mb/s, respectively).*



■ Figure 5. The MarconiNet fast handoff scenario.

local station can point to find out the program subscribed to (group address) by the impending mobile host. Just before a mobile node leaves the cell, a local policy decision (e.g., signal-to-noise ratio) will trigger an RTCP message to the local announcement address; the server in turn announces that to the sharing multicast addresses to which the neighboring stations listen in the global space. The neighboring stations (servers) look up the multicast address and check with their own databases to see the group association of this multicast stream.

In the absence of this association, the server sends an IGMP message to the upstream router and passes the stream to the local cells using a locally scoped multicast address, even before the client has moved to the new cell. As soon as the client moves into the next cell it still receives same stream without any interruption. Similarly the client sends an RTCP BYE to the server as it moves away from the previous server.

**Pre-Registration with Multicast Proxy Agent** — Another approach is to deploy proxy agents in each subnet. These proxy agents join the upstream multicast tree on behalf of the servers even before the clients move into the cell. The neighboring proxy servers then listen to a common multicast address to figure out the impending host's subscribed multicast address. In this case the multicast proxy sends the IGMP query messages beforehand on behalf of the local servers. This is a typical case where layer 3 multicasting has been used for the proxy servers. Similarly, a multicast proxy agent within each upstream router can help forward the global

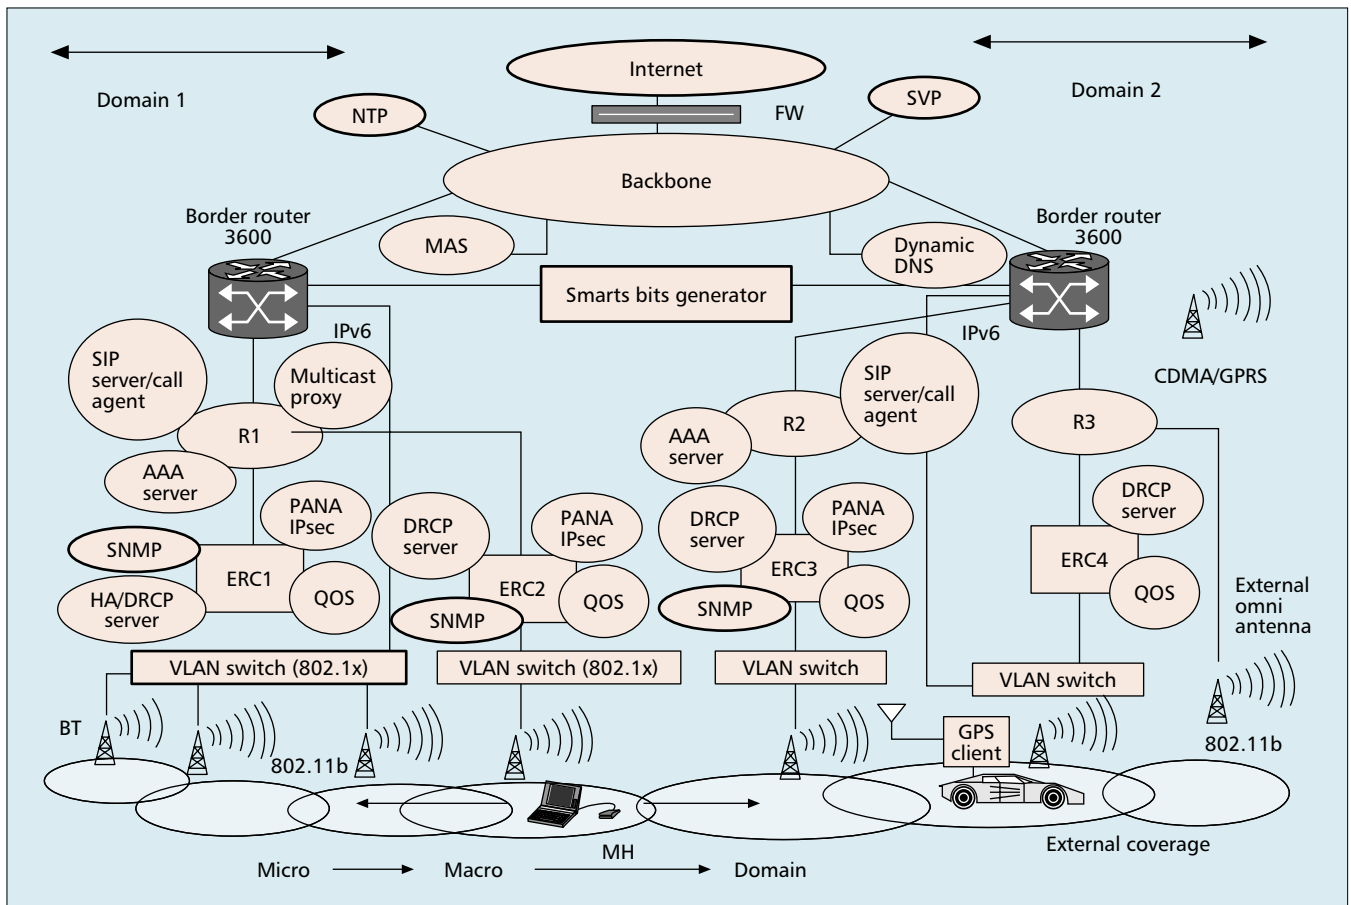
stream to the respective global multicast addresses (e.g., for areas where these clients are about to move) in each subnet for a specific period of time determined by the client's entry to the cell. Thus, each neighboring server can receive the stream regardless of whether the mobile node is moving into that cell or not.

In either case, as soon as the mobile node moves into the new cell it sends a trigger signal to indicate that the mobile node has moved in. Upon receipt of this signal a subnet proxy agent leaves the tree or a multicast proxy agent stops forwarding the traffic.

#### IN REGISTRATION

Group membership information can also be passed during the client's registration to the new network. During a node's movement between subnets it can send the request for the particular multicast address in its DHCP Discover option/PPP server option message about the local multicast address it was listening to in the previous subnet. During the process of obtaining the IP address from the DHCP server, the client can send the locally scoped multicast address to the server. Thus, the server can join the desired group during the time the client is in the process of getting configured.

There are trade-offs associated with each of these techniques. Pre-registration helps reduce the join latency at the cost of bandwidth in the previous cell/subnet. The in-registration process provides a better solution by striking a balance between both. In a situation where a client is subjected to zero-connected interruption state, end system buffering can be applied similar to satellite radio systems such as Sirius and XM.



■ Figure 6. MarconiNet realization in an integrated testbed.

## TESTBED PROTOTYPE AND RESULTS

Figure 6 shows a comprehensive multimedia integrated testbed, where the MarconiNet architecture has been realized supporting mobile content distribution features over several types of access technologies such as 802.11b, GPRS, and CDMA1XRTT. This multimedia testbed provides both indoor and outdoor mobility functionality, and emulates a wide-area wireless Internet by deploying a cross-traffic generator and a delay simulator. A mobility scenario involving three layers (e.g., layer 2, network layer, and application layer) was used for the experiment. In the testbed the edge router controllers (e.g., ERC1, ERC2, ERC3, ERC4) are Linux systems that also act as local station servers. As part of the initial experiment many multimedia streaming applications were demonstrated using some of the Mbone tools such as sdr, rat, vic, and wb, and MarconiNet software modules for the RAS, RSC, and IMC. UMTF tunneling techniques have been used to demonstrate the streaming application in the event of unavailability of multicast routers in some parts of the network. Experiments were conducted for cell, subnet, and domain mobility. Several measurements were taken including the time for layer 2 movement detection, IP address acquisition, and join/leave latency using network layer IGMP and application layer RTCP. CGMP implementation of Cisco's layer 2 switch was used to control the multicast stream during a mobile's movement

between cells within a subnet. Tools like tcpdump and ethereal were used to collect the data. Other functional aspects of MarconiNet, such as QoS management based on real-time feedback, several location-based services, and localized advertisement insertion, were also realized in the testbed. A typical layer 2 detection delay depends on the beacon interval; layer 3 IP address acquisition delay depends on the client's movement type (e.g., LAN to WAN, WAN to LAN) as it involves different protocols (e.g., DHCP, PPP) to obtain the IP address. DHCP without ARP takes about 4 s for IP address acquisition, DRCP takes about 500 ms, and PPP takes about 10 s to obtain an IP address. Inter-domain mobility and local authentication adds additional delay due to protocols such as AAA and PANA. IGMP/RTCP join/leave latency are independent of the IP address acquisition technique. A proxy-based expedited JOIN reduces the join latency from 60 s to almost zero. It was found that during subnet movement in a regular 802.11 network a client does not send a IGMP/RTCP leave message when it moves out of a subnet. This contributes to undesirable leave latency, but it helps the client only if it moves back and forth between subnets within a specified time. A proxy-based IGMP/RTCP Leave Group message can be sent as soon as the client's IP address changes. Figures 7 and 8 show the protocol sequence and advantage of a proxy-based handoff scheme that reduces join and leave latency. Figure 9 shows

some of the prototype components of MarconiNet, primarily the local station manager, Internet client, and channel monitor, showing online statistics of the tuned-in audience.

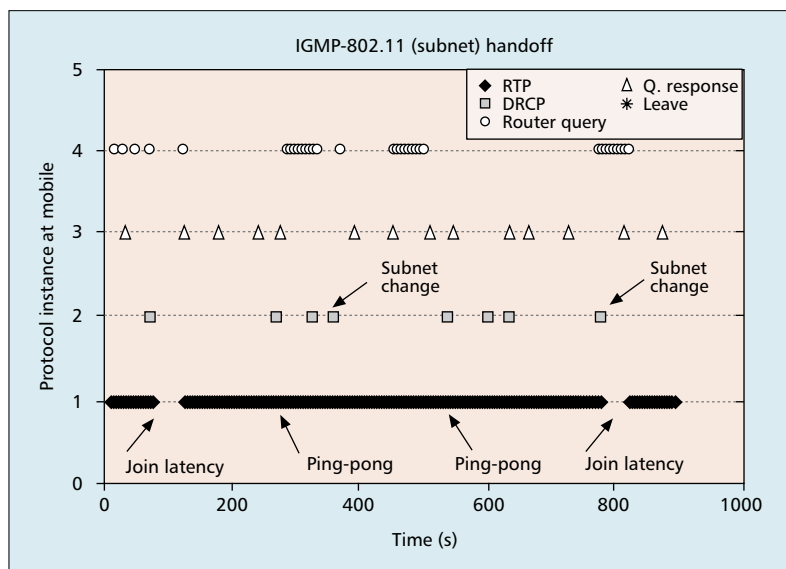
## CONCLUSIONS

In this article we present a flexible streaming content distribution architecture called MarconiNet. Several features, issues, and implementation details are discussed, and results of MarconiNet architecture implementation in a multimedia testbed are presented. Here we have focused on the operational techniques needed to build a business model for mobile users in a 4G-based overlay network. These include secured content distribution, a hierarchical payment scheme, faster join and leave latency for mobile clients using application layer techniques, timely

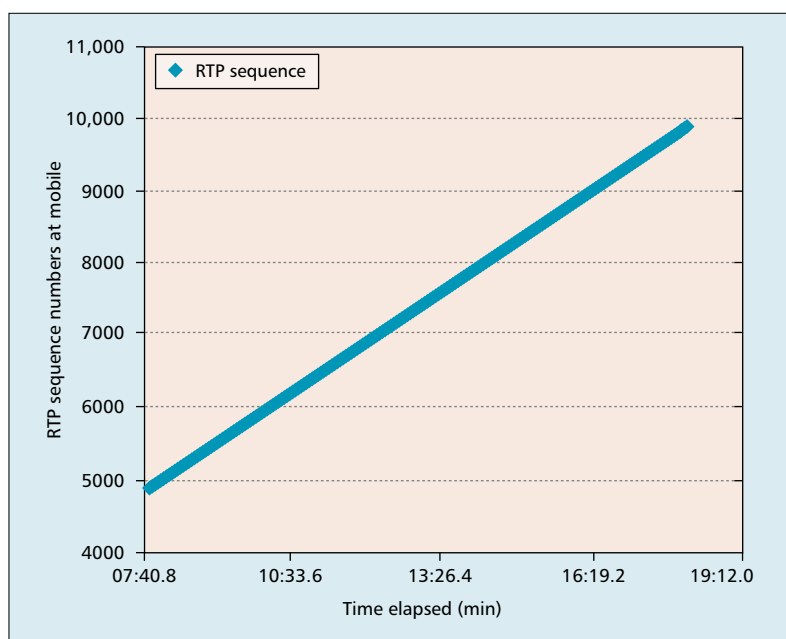
delivery of multimedia content for moving clients, and ensuring the desired QoS in a wireless mobile multicast environment. Deployment of such a system can help provide flexible streaming services to support mobile commerce in a 3G network and beyond.

## REFERENCES

- [1] A. Dutta and H. Schulzrinne, "Marconinet — An Architecture for Internet Radio and TV Networks," *Proc. NOSSDAV*, Basking Ridge, NJ, June 1999, pp. 241–45.
- [2] C. Diot et al., "Deployment Issues for the IP Multicast Service and architecture," *IEEE Network*, vol. 14, Jan. 2000.
- [3] A. Basso, C. Cranor, and S. Sibal, "Prism: An IP-Based Architecture for Broadband Access to TV and Other Streaming Network," *Proc. NOSSDAV*, Chapel Hill, NC, June 2000, pp. 11–17.
- [4] K. Almeroth and M. Ammar, "Multicast Group Behavior in the Internet's Multicast Backbone (MBone)," *IEEE Commun. Mag.*, vol. 35, no. 6, June 1997.
- [5] G. Xylomenos and G. C. Polyzos, "IP Multicast for Mobile Hosts," *IEEE Commun. Mag.*, vol. 35, no. 1, Jan. 1997.
- [6] J.-L. C. Wu, "An IP Mobility Support Architecture for 4GW Wireless Infrastructure," *PCC*, Nov. 1999.
- [7] A. McAuley et al., "Mobile Multicast Proxy," *IEEE MILCOM*, Atlantic City, NJ, Nov. 1999.
- [8] C. Williamson et al., "Performance Evaluation of the MoM Mobile Multicast Protocol," *ACM MONET J.*, vol. 3, Aug. 1998, pp. 189–201.
- [9] C. L. Tan and S. Pink, "Mobicast: A Multicast Scheme for Wireless Networks," *Mobile Networks and Apps.*, vol. 5, Jan. 2000, pp. 259–71.
- [10] T. Friedman and D. F. Towsley, "Multicast Session Membership Size Estimation," *Proc. IEEE INFOCOM*, New York, NY, Mar. 1999.
- [11] J.-O. Vatn and G. C. Maguire, "The Effect of Using Co-located Care-of Addresses on Macro Handover Latency," *14th Nordic Tele-traffic Seminar*, Tech. Univ. of Denmark, Lyngby, Aug. 1998.
- [12] A. McAuley et al., "Dynamic Registration and Configuration Protocol (DRCP)," IETF Internet draft, July 2000, work in progress.
- [13] E. Magna, E. Izkue, and J. Villadangos, "Review of Traffic Scheduler Features on General Purpose Platforms," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 31, Apr. 2001, pp. 50–79.
- [14] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-Driven Layered Multicast," *SIGCOMM Symp. on Commun. Architectures and Protocols*, Palo Alto, CA, Aug. 1996, pp. 117–30.
- [15] J. Chennikara et al., "Application Layer Multicast for Mobile Users in Diverse Networks," *Proc. GLOBECOM*, Taipei, Taiwan, Nov. 2002.



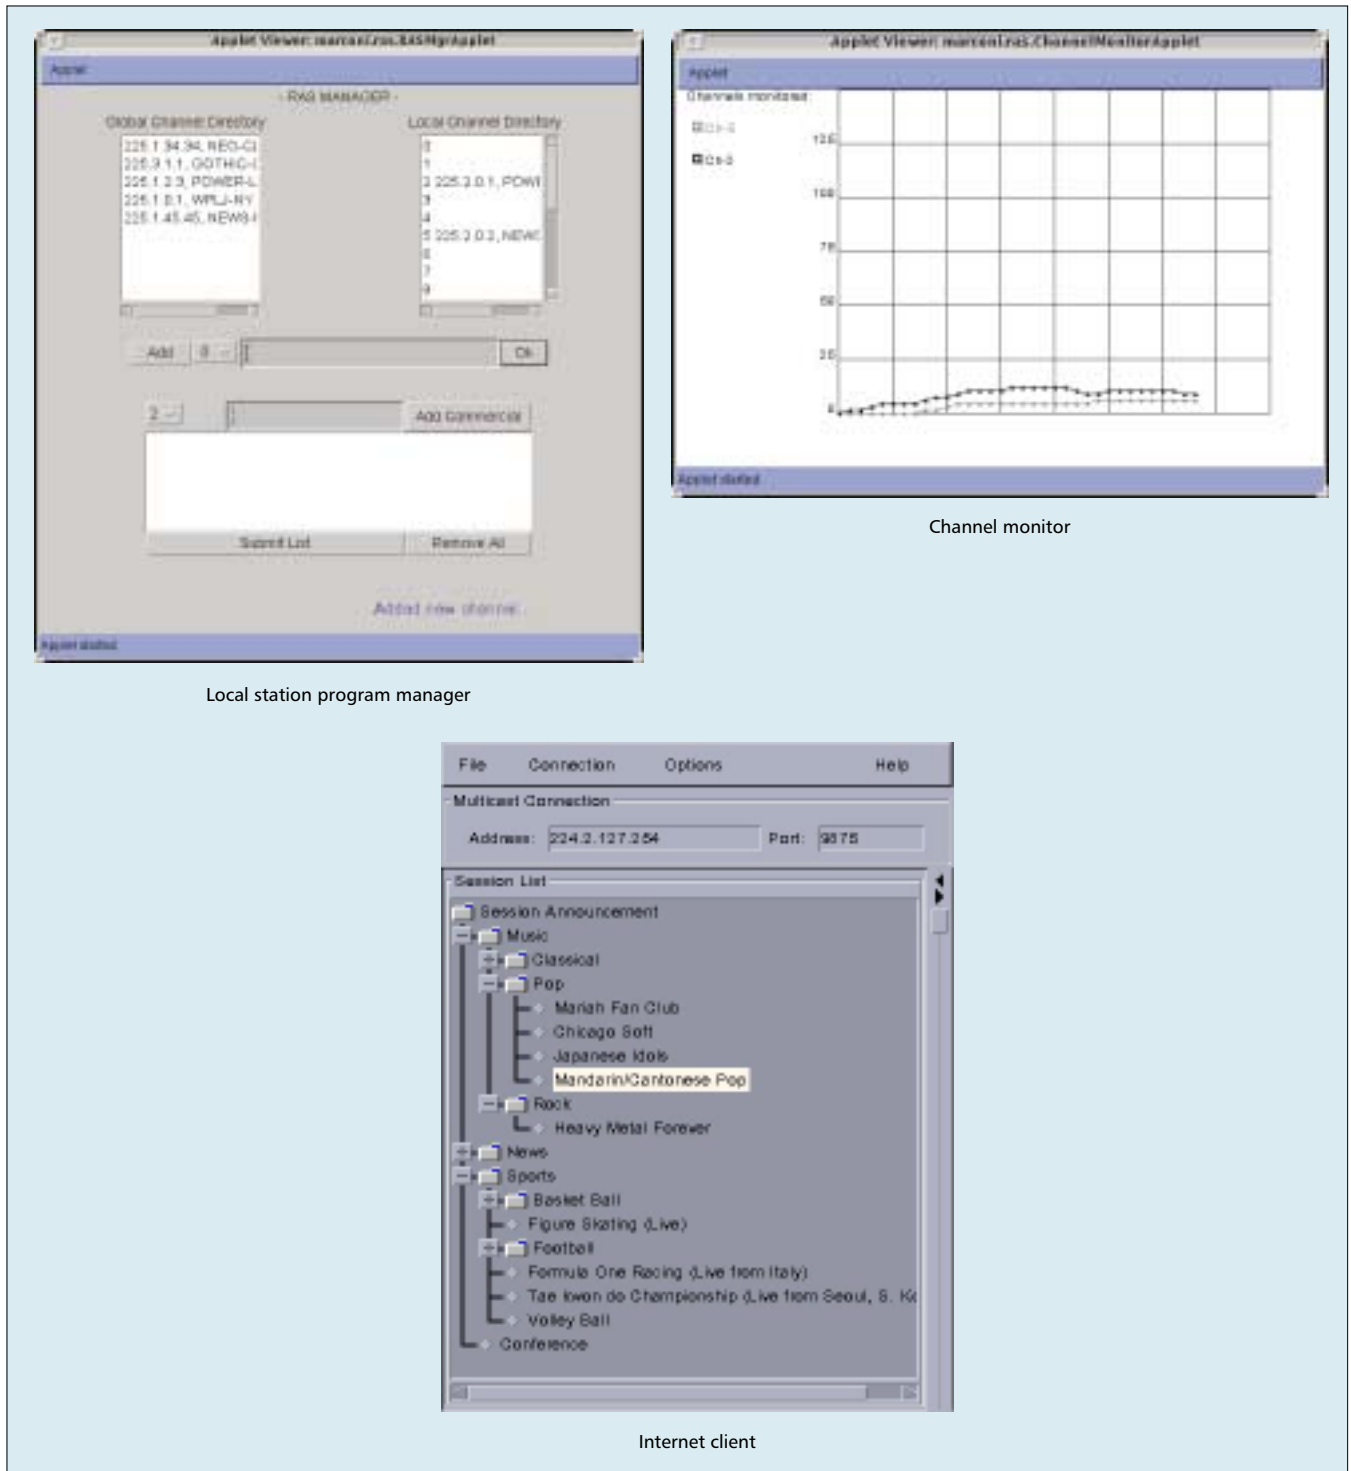
■ Figure 7. Protocol flow for IGMP handoff.



■ Figure 8. Proxy-based handoff.

## BIOGRAPHIES

ASHUTOSH DUTTA [SM] (adutta@research.telcordia.com) is currently a research scientist in Telcordia Technology's Internet Network Research Laboratory with an emphasis on mobile networking and middleware applications for wireless Internet. 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 network, 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 has been active in IETF's SIP, MMUSIC, and Mobile IP working groups, and served on the organizing committee of MOBICOM 2000. He has been a frequent speaker at VoIP related trade conferences. He was co-project leader of a DARPA-funded Airborne Communication Node project, and technical lead for integrated mobility management for the project. He received the 2000 and 2002 Telcordia CEO Award and SAIC's ESTC 2002 best paper award in the Information and Technology category. He has a B.S.E.E. (1985) from National Institute of Technology, Rourkela, India, an M.S. in computer science (1989) from New Jersey Institute of Technology, and a professional engineering degree from Columbia University. Currently he is pursuing his Ph.D. part-time at Columbia University.



■ **Figure 9.** MarconiNet prototype components.

HENNING SCHULZRINNE 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 at Amherst. He was a member of technical staff at AT&T Bell Laboratories, Murray Hill, New Jersey, and an associate department head at GMD-Fokus (Berlin) before joining the Computer Science and Electrical Engineering Departments at 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 the *Journal of Communications and Networks*, an editor of *IEEE/ACM Transactions on Network-*

*ing*, and a former editor of *IEEE Internet Computing Magazine* and *IEEE Transactions on Image Processing*. He is a member of the Board of Governors of the IEEE Communications Society and the ACM SIGCOMM Executive Committee, and former chair of the IEEE Communications Society Technical Committees on Computer Communications and the Internet. He was also a member of the Internet Architecture Board. 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. He serves as chief scientist for SIPquest Inc. and chief scientific advisor for Ubiquity Software Corporation.