Chapter #

MOBILE MULTICAST

Group communications in a wireless Internet

Elias C. Efstathiou and George C. Polyzos
Mobile Multimedia Laboratory, Athens University of Economics and Business

1. INTRODUCTION

During the late-1990s, the popularity of Internet applications like B2C (business-to-consumer) e-commerce and web-browsing led traditional network operators to seek ways which allowed their customers to access these new services. Today (2002), IP-based packet switching technology has finally matured enough to replace the circuit-switched backbones of telecommunication carriers. Convergence on the IP protocol is now happening and it will continue to happen throughout this decade. It is obvious that this one common protocol will ease the management of all wired and wireless data networks. IP will also fuel the deployment of novel applications. An important feature these applications will expect from an all-IP internetwork is full support for multipoint communications. Multipoint, or group, communications are best described by the term *multicast*, a term associated with network support for efficient data delivery to more than one interested recipients. Multicast's objective is to place the least amount of burden on network and end-host resources. Applications that could exploit this feature include conferencing, on-line games, software distribution, and others. Although multicast can be emulated by letting the data sources themselves send packet copies to all intended destinations, this "multi-unicast" solution offers no scalability as resources on both the source host and its local network would eventually be depleted.

Today, the many players involved in what is a rapid expansion of the Internet still place "all-IP" convergence several years away. All of these players are, in one way or another, focusing on Internet protocols, with 3G

licensees bearing much of the burden of bringing millions of new users and terminals to the Internet community. This worldwide deployment of the new 3G IP-based networks represents the Internet's biggest expansion yet. Meanwhile, digital broadcasters are also joining the IP bandwagon, with *DVB* (Digital Video Broadcasting – the digital TV standard) providing IP support and over 30 Mbps of shared bandwidth per DVB macro-cell. In addition, we must not forget all the older technologies that support IP over various link layers and which include Ethernet networks as well as all PSTN, ISDN, DSL and cable connections. Finally, all the wireless *IEEE 802.11* networks, also known as Wireless Local Area Networks or WLANs represent extremely important bridges between the older and newer technologies. The ease with which WLANs are deployed outweighs most of their limitations and, for many, WLANs represent the future in ubiquitous broadband wireless Internet access.

We can see, therefore, that although the "all-IP" goal has not yet been reached, it's clear that we are approaching it. All these converging, technologies were traditionally associated with particular services, e.g. bidirectional one-to-one audio (voice conversations) for cellular telephony, unidirectional one-to-many video (TV broadcasting) for DVB. The ideal all-IP internetwork must support each one of these services, integrate them, and "enhance" them with features such as higher interactivity and, most importantly, mobility and *tether-less* access. The requirement to support IP-based group communication is also dictated.

The IP suite was originally designed with one-to-one communications in mind. However, the advantages of one-to-many and many-to-many communications are numerous [10, 28] and we will not dwell on them further. Network technologies that provide "native broadcast" are better suited to support multipoint communications. Examples, such as the traditional Ethernet, IEEE 802.11 in all its flavors, and satellite, terrestrial or cable broadcast networks, allow for transmitted link-layer frames to be received efficiently by all hosts in a local subnet. Other network technologies, such as GSM and GPRS offer point-to-point links in the *user plane* [19], and only support broadcast in the *control plane.* With these cellular networks in particular, limited information can be delivered by *cell broadcasting*, but generally, users in a cell cannot engage in efficient group communication, although this is theoretically possible. Of course, there is absolutely no support in traditional GSM for inter-cell group communication. Finally, technologies like the PSTN, ISDN, and DSL networks are designed from the ground up to provide native support for point-to-point communications only.

IP currently (in both versions 4 and 6) supports efficient, local, network layer, many-to-many communications (i.e. multicast), wherever link level broadcast is provided. However, extending this support beyond a local sub-

net is non-trivial. Also, to support many-to-many communications in an internetwork topology that changes dynamically while IP hosts (and even entire subnets) are moving is even more challenging. Here we will focus on what we describe as *IP-based quasi-reliable mobile multicast*. IP multicast [10], which is the basis of our analysis, is a best effort multipoint communication protocol. Reliable, sequenced delivery extensions do exist [13]. However, applications don't always require this reliability, since they can adapt in their own more advanced ways. This, in particular, is the case with *streaming media* services (especially real-time media).

In an "*all-IP wireless Internet*," IP hosts should be able to change their network point of attachment with minimal disruption of ongoing communications. We may assume that one or more of the following may occur while an IP host is moving:

- the host may temporarily disconnect from the network;
- the host's IP address, the one used for packet routing, may change;
- in a wireless scenario, a horizontal or vertical (cross-technology) handoff may occur;
- handoffs may occur between cells belonging to different administrations.

Here we will offer our perspective on the issues involved in combining multicast capability with host mobility in an all-IP wireless environment. We will do this while assuming the existence of a fixed, wired routing infrastructure. Also, whenever we discuss IP-related protocols, such as IP multicast and Mobile IP, our focus is on version 4 of the IP protocol. We are aware that the move to IP version 6 is happening, albeit slower than expected, so IPv6 is taken into account when it helps with our analysis.

The remainder of this chapter is organized as follows. Section 2 shows how early Internet assumptions influenced present-day protocols and made it harder to solve our current mobility problems. Section 3 outlines the IP multicast protocol. Section 4 presents two IP mobility protocols, Mobile IP, the standard IETF macro-mobility protocol and Cellular IP, a less well-known IETF micro-mobility protocol. Section 5 will mention the usefulness of *transcoding* filters to the mobile multicast problem. Section 6 deals with the coexistence of mobility and IP multicast. In Section 7, we present our own perspective on the problem; we mention requirements for future mobile multicast protocols; we outline our own, Cellular-IP based, mobile multicast solution; and we describe in more detail the all-IP wireless Internet environment we envisage. Section 8 presents some concluding remarks.

2. BASIC ASSUMPTIONS INFLUENCING TCP/IP

The Internet was originally built to support packet based data communications among pairs of stationary IP hosts. This assumption influenced the design of most protocols in the IP suite [26]. Although the overall design of the Internet protocols is "layered," in practice, it is often found that interlayer dependencies exist and become apparent when attempting to port existing services to newer network technologies.

We will use TCP as an example of the dependencies, design limitations and assumptions that cause service disruption when a TCP application like Telnet or FTP is used in a mobile environment:

Interlayer dependency - Each TCP connection is identified by a unique pair of *sockets*, with each socket being a <*Host_IP_address*, *TCP_port*> pair. The TCP connection ID is, therefore, the 4-tuple <*Client_IP_addr*, *Client_TCP_Port*, *Server_IP_addr, Server_TCP_port*>. It is obvious that lower layer identifiers (IP addresses) are used as part of higher layer ones and that TCP's internal session state relies on these identifiers. If either end-host changes IP address, the TCP connection will break because, not only does TCP rely on lower layer identifiers, but also, it assumes that they will not change during the lifetime of a TCP connection [26]. So, in practice, a user (who for example has initiated a lengthy FTP file download) has no way of changing his or hers network point of attachment without needing to restart the FTP session. Users with ongoing Telnet sessions and WLAN/GPRS enabled laptops cannot handoff these TCP/IP connections to the GPRS network as they exit from a WLAN hotspot, since, most probably, the WLAN and GPRS networks will have different IP network prefixes and the mobile host will be assigned completely different IP addresses in each one.

Design limitations - "Classic" TCP (we refer here to IETF's RFC 793, including the 1988 Jacobson refinements) is designed to interpret packet losses as a sign of congestion along the router path between the two end-hosts. TCP normally adapts by lowering its rate and allowing router queues to drain [17]. However, if the packet loss was caused by the increased BER (Bit Error Rate) of a wireless link, lowering the rate only decreases TCP performance. Studies and simulations have shown that the connection's effective bandwidth is greatly reduced even with moderate increases in BER, mainly because of TCP's "slow-start" algorithm. The Internet community's underlying assumption for TCP was, of course, that the two TCP endpoints would be stationary and that they would normally communicate over wired links with extremely low BER [30]. With wireless error rates in the 10^{-5} and

10^{-4} range, even approaching 10^{-2} and 10^{-1} in extreme cases, this assumption needs reexamination. There are many TCP variations which address this issue, either by relying on associated protocols to explicitly deliver notifications about either congestion or loss events in addition to the simple TCP acknowledgment mechanism. One such notification mechanism is ECN, or *Explicit Congestion Notification*. Others try to recalibrate TCP timers and reconfigure the slow-start algorithm. Of course, there are also many solutions that try to increase reliability at the link-layer or attempt to add *Forward Error Correction* (FEC) information, usually combined with packet fragmentation and fragment reordering so as to minimize the effects of burst errors.

3. IP MULTICAST

In this section we present an outline of the basic IETF multipoint protocol, IP multicast. IP multicast is a many-to-many communication protocol. The *host group* service model defines its requirements: let *H* be the set of all IP hosts. Let E_G be a subset of H. Set E_G forms a *multicast group with group address G*, if and only if:

- members of H may join and leave E_G at any time;
- members of H can communicate unidirectionally with all members of E_G, using only identifier G. This identifier is also known as then *host group address*. This second requirement suggests that a host need not be a member of a multicast group in order to sent data to that particular group.

The definition above does not specify how to satisfy these requirements. They represent an idealized version of what group communication should be like and in practice the various implementations interpret these requirements liberally.

As IP hosts can belong to more than one group, mechanisms are needed in order to:

- associate hosts with groups;
- track group membership;
- route data to all group members [28].

Also ideally, IP multicast packet delivery should emulate the exactly-once semantics of packet delivery in a traditional (non-switched) Ethernet [2]. Since in an Ethernet all transmitted frames are received by all attached interfaces, exactly-once delivery comes for free. In an Ethernet, the only

additional mechanism needed is the mapping of multicast group identifiers to MAC addresses. In this way, Ethernet hosts know which frames to process and which to discard. (Switched Ethernets must also provide support for multicast frames.)

The IETF set aside all legal IPv4 class D addresses to be used as multicast group identifiers, underlining the fact that IP multicast is a network layer protocol. The special IPv4 class D address *224.0.0.1* always identifies, according to the specification, the link-local *all-hosts* multicast group, which all multicast IP hosts are required to be a part of [9]. The mechanisms needed to implement multicast packet delivery can be divided into *global* and *local* [28].

3.1 Global Mechanisms

In the multicast model, the burden of delivering multiple copies of packets falls on the network. The sender need only transmit one copy of each packet addressed to a multicast group identifier (a host group address). Multicast enabled routers will forward the packet as needed, replicating it onto more than one of their outgoing interfaces only when paths towards the destination group members diverge. The global multicast routing protocols deliver a group's packets to multicast routers that have expressed interest in receiving packets for the particular group. This interest, in turn, is triggered by hosts in the router's local IP subnets that declare their wish to join the specific group.

Proposed multicast routing protocols include the *Distance Vector Multicast Routing Protocol* (DVMRP), *Multicast Open Shortest Path First* (MOSPF), *Core Based Trees* (CBT), *Protocol Independent Multicast-Sparse Mode* (PIM-SM), and *Protocol Independent Multicast-Dense Mode* (PIM-DM). All these protocols attempt to build a multicast delivery tree of routers for each multicast group. DVMRP and MOSPF are less scalable (they construct one tree per sender per group) than CBT or PIM (they construct one tree per group, which senders share). CBT and PIM are also independent of the underlying unicast routing protocols used. PIM-SM is optimized for *sparse* multicast groups and PIM-DM for *dense* multicast groups. PIM works by choosing a *Rendezvous Point* (RP) when it constructs the multicast delivery tree for a group, where multicast senders can "meet" multicast receivers. With all protocols, there is an associated *graft* delay when a multicast router joins an existing tree because the tree has to be adjusted. All these protocols implicitly assume stationary hosts.

3.2 Local Mechanisms

The global protocols we described above are concerned with multicast senders, multicast routers and multicast groups, but not with the individual multicast listeners. These listener hosts are "hidden" behind their local multicast router. The router can be thought of as the interface between the local and the global mechanism [28]. Based on a local group management protocol, this router builds a list with all the different multicast groups its hosts have joined (on each one of the IP subnets it may serve). Only this aggregate list is exposed to the global mechanism.

The local group protocols are the *Internet Group Management Protocol* (IGMP) [12] for IPv4 and the *Multicast Listener Discovery* protocol (MLD – derived from IGMP version 2) [11] for IPv6. IGMP assumes the existence of link-level native broadcast (e.g. Ethernet) and is designed around the soft-state principle which traditionally leads to robust Internet protocols. IGMP works as follows: every *querying_period* the multicast router responsible for a local subnet sends out queries to the all-hosts group. Its objective is to refresh the group list it exposes to the global mechanism. All local listeners receive the query and respond (after a small *random_report_delay*) with a *group_report*, one for each of the groups they participate in. Each report is actually sent to the multicast address for the group reported so that all interested local listeners may learn of this fact.

The soft state principle in IGMP dictates that queries are router-initiated and that no explicit *leave_group* message is needed when a listener leaves a group: the router will discover it in the next query cycle. Extensions for unsolicited listener messages exist [12]. IGMP version 2 adds a *leave_group_message* and IGMP version 3 adds the ability to selectively join and leave multicast groups [1]. A good usage example of these extensions comes from the need to lower the *join delay* (the local equivalent of the global graft delay). A host joining a group not already present on the local link may send an unsolicited *group_report*. If not, it could very well wait up to *querying_period* + *random_report_delay* + *graft_delay* before multicast packets start arriving.

4. IP MOBILITY PROTOCOLS

4.1 Mobile IP

The goal of *Mobile IP* (M-IP) [24] is to allow internetwork host mobility in a manner transparent to the transport layer. The objective is to leverage the investment on existing TCP/IP applications and avoid the need to

redesign new, mobility-aware ones. This mobility transparency can be achieved by assigning two different IP addresses to every *mobile host* (MH), a permanent one used by applications for identification, and another one, which may change, used for routing. The permanent IP address is called the MH's *home address*. If the MH changes its point of attachment and moves to a *foreign* IP subnet, a *Correspondent Host* (CH) may still deliver packets to the MH using the MH's home address as the destination. This is possible because a new entity at the MH's home subnet - a router known as the *Home Agent* (HA) - intercepts these packets and *tunnels* them (reroutes them using *IP-in-IP encapsulation* [23]) to the MH's current network address. To achieve this, HAs advertise reachability to the home network using standard IP routing protocols. This network could very well be a virtual one, i.e. have no physical instantiation: it could be represented by just a HA which keeps track of many MHs roaming around the Internet. In case the network is not a virtual one, the HA may use *proxy ARP* (ARP is the Address Resolution Protocol, the protocol for resolving IP addresses to link-layer MAC addresses in Ethernet and Ethernet-based LANs) in order to fool local network nodes into sending packets addressed to the MH to the HA instead.

The MH's current network address is known to the HA as a result of a registration procedure: when the MH first moves to the foreign subnet, it communicates back to its assigned HA its new *Care-of Address* (CoA). This new address is either a *co-located* CoA or a *Foreign Agent* (FA) CoA. In the co-located case, this new address is usually obtained through *DHCP* (Dynamic Host Configuration Protocol) on the foreign subnet. In the FA case, this address corresponds to a special router - the FA - that resides on the visited subnet. In both cases, this address is communicated back to the HA (either by the MH or the FA), where a new address *binding* with an associated lifetime is created, which binds the MH's home address with its current CoA.

In the co-located case, *reverse tunneling* (decapsulation of packets) is carried out by the MH, whereas in the FA case, reverse tunneling is the responsibility of the FA which will then forward each packet to the MH (FAs and MHs are normally assumed to have link-layer connectivity). In both cases, IP-in-IP from the HA will deliver a packet with an outer destination address matching the CoA and an inner destination address matching the MH's permanent home address.

The aforementioned mechanisms describe how CHs can always send packets to MHs. For the reverse path, the standard IP mechanism can be used: the MH will send packets addressed to the CH's address and place its home address in the source address field. This "trick" has no adverse effect since standard unicast IP routing normally depends only on the destination address (if no *ingress security filtering* is used).

The M-IP communication mechanism results in an inefficiency called *triangle routing* [28]. Extensions [25] to the original M-IP RFC allow CHs to make address bindings of their own which subsequently gives them the ability to discover the MH's new address and route packets directly to the MH bypassing its HA and home subnet. This feature can only be used if the CHs are also M-IP-aware. This particular feature has been slow in its adoption and will not really be used before IPv6 is deployed sufficiently, since it raises security concerns not easily solved within the confines of the M-IPv4 specification.

M-IP is built around the soft-state principle which dictates that all registrations with the HA (as well as all local registrations of MHs with FAs) should have an associated lifetime after which they expire. MHs need to refresh the bindings periodically. There is no explicit deregistration procedure. When an MH moves to yet another subnet or when it disconnects completely, the old FA (if one existed) and the HA in charge will soon learn of this fact. FAs maintain a list of all the visiting MHs they serve and HAs maintain a list containing the bindings of all the MHs for which they are responsible.

4.2 Cellular IP

M-IP was designed with "slow" macro-mobility in mind. The protocol incurs several delays, the most important one being the time MHs take to discover an FA in the foreign subnet, register with it, acquire a CoA, and register with their HA. This *agent discovery time* can be significant. If a moving MH traverses many *pico-cells*, each one controlled by a different FA (co-located with the cell's base station), the signaling required during every handoff may inhibit the normal reception of user data packets. In the case of a fast moving host, the so-called *mobility assumption* [3] may be violated: the total registration delay may be more than the time an MH spends inside a cell. As a result, no user packets will find the time to get rerouted to the MH.

Cellular IP (C-IP) [6] is a best effort *micro-mobility* protocol designed with campus-wide, partially-overlapping micro-cells and pico-cells in mind. Its objective is to provide seamless local mobility. C-IP can rely on M-IP for *macro-mobility* support (for example, an MH arriving at a campus internetwork will still register this fact with its M-IP HA). For intra-campus mobility and local handoff control, C-IP takes over and uses its own internal network of packet forwarders (which may be co-located with base stations). As far as the MH's HA is concerned the whole campus internetwork is controlled by only one FA (M-IP has to be used in FA mode for this to work).

A C-IP *gateway* can be co-located with M-IP's FA. This gateway is the interface between the M-IP routing infrastructure and the C-IP access

network (see Figure 1). Base stations within a C-IP network are fixed, interconnected IP forwarders with at least one wireless interface. They operate at layer 3 and can be connected via any number of layer 2 bridging nodes. To forward packets to MHs, they employ non-standard IP routing techniques: MHs are only identified by their M-IP home addresses and, unlike M-IP, no tunneling is employed [6]. The C-IP route learning mechanism greatly resembles the MAC bridge auto-learning mechanism and it follows the soft state principle: routes have to be refreshed periodically. This happens whenever MHs send regular user data packets towards the gateway. Therefore, there is no need for explicit signaling to notify about an MH's current location. In addition, there is no need for yet another type of end-host identifier since the simple and flat forwarding architecture, which resembles layer 2 bridging, uses the host's permanent home IP address as an identifier which is mapped to a particular port on each forwarder. These bindings are "auto-learned" whenever an MH sends regular data packets towards the gateway, by "snooping" and looking at the source IP address field in the IP packet. Although this IP-to-port mapping is not a very scalable solution, the size of real C-IP access networks is not expected to be unmanageable, since they will be limited to a certain geographical area and will serve a relatively small number of local, visiting or home, IP hosts.

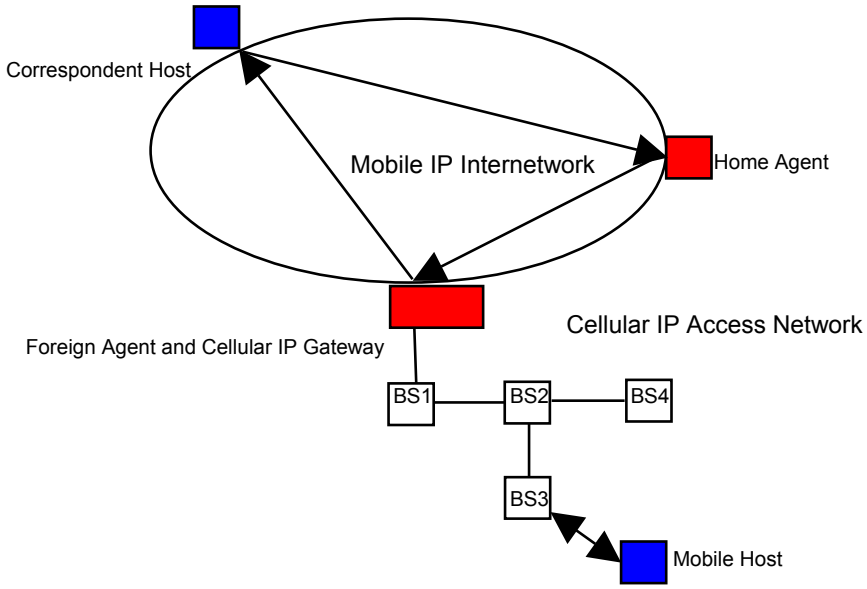

Figure #-1. Cellular IP Access Network

Certain key notions at the heart of C-IP are based on existing mechanisms used in cellular telephony networks. These include the notions of *active* and *idle* nodes and the notion of *paging* [6]. To minimize signaling overhead, MHs only update their location information when they are active, i.e. engaged in communication. Since C-IP is as connectionless as IP is, MHs can only be deemed idle after a certain period of inactivity. When they are idle, the fixed base station infrastructure has no knowledge of their whereabouts if the soft state routes have also expired. When new incoming packets have to be delivered to an MH, paging is used to discover its exact cell: the C-IP forwarders flood all their outgoing interfaces (unless a cached route entry exists), except the one they received the packet from. C-IP supports smoother handoffs by allowing MHs to send *route_update* packets as soon as they detect that they have moved to a different cell.

5. TRANSCODING FILTERS

There exist software components, collectively called *media filters* or *transcoders*, which, given an input media stream, can produce a different, but similar, output media stream. Their purpose is to help maintain a level of *User-Perceived Quality of Service* [14] as the media stream traverses many internetwork links of varying bandwidth. These filters can be divided into *smart* and *simple* [14]. Smart filters do not have to decode the media stream (e.g. an MPEG-2 video stream) to a raw format and then re-encode it. Rather, they operate on the encoded stream directly. Smart filters differ from simple ones in that they require more processing power, but, generally, because they take advantage of a previous encoding stage they are also faster and produce better perceived output than their simpler counterparts.

Filters can be used to reduce the bit rate of a media stream before it is injected into a bandwidth limited link, such as a wireless link [22]. Filters can support multicast for wireless communications (many believe they are absolutely necessary) by allowing wireless receivers that cannot receive the source stream at its original bit rate to receive a "cut-down" version of it. Usually, this means lower resolution, fewer frames per second and smaller frame-size. It can also mean lower color-depth, less smooth movement and the tweaking of a variety of parameters that depend on the media encoding format. This sort of *flexibility* represents a subset of the requirements for fully programmable and *reconfigurable networks* at all layers. Technologists are planning ahead for these "software" networks, whose appearance will change many of today's accepted networking principles. The vision also involves dynamic reconfiguration of radio modulation parameters, dynamic spectrum allocation and dynamic protocol downloading and configuration.

On the same front, dynamic encodings like *layered coding* and *multiresolution layered coding* [15] can be used to separate a media stream into more than one streams, each carrying, progressively, more information. Listeners can receive as many of the streams their incoming bandwidth allows. This can be used in conjunction with filtering. A proposal also suggests using different multicast groups to transmit each one of these substreams [20]. Listeners may then "tune into" as many groups as possible.

Filters become very useful to multimedia communications in wireless environments. A natural location to place such filters is at the boundary between the wired and the wireless part of the internetwork. There, they can modify the passing stream based on actual measurements of the conditions on the wireless link. In a multicast scenario, different receivers may require different filters. In the case where many receivers at the same multicast sub-tree request the same transcoding function, the relevant filters can propagate upstream (towards the sub-tree root - their nearest common ancestor - closer

to the stream source), where they can combine into one and serve all of the sub-tree's requesting nodes [21]. In the general case, an *active* or *configurable* IP router (one that allows third-party code uploading and execution) is a natural location for these mobile filters to reside in.

6. COMBINING IP MULTICAST AND MOBILITY

The IP multicast protocol was designed to bring multipoint communication capability to the Internet. At the same time, Mobile IP and Cellular IP were designed to allow transport-layer-transparent mobility. Trying to integrate the protocols into a coherent IP framework for multicast mobility exposes several fundamental issues. Existing approaches [16, 18, 24, 27] offer some functionality, but further simulation and deployment is needed before the scalability of each proposed solution can be judged.

6.1 "Fixed Host" Assumption – IP Addressing Issues

We already mentioned that the "fixed host" assumption influenced the design of several protocols in the IP suite. When considering mobility, we need to take into account that mobile devices can be very different from fixed hosts connected to an Ethernet network. Some basic differences are listed below:

- limited battery life dictates that unnecessary operations should be avoided, so, protocols that rely on constant monitoring of network traffic are impractical;
- protocols that assume high bandwidth, low latency connectivity may become inoperable;
- protocols that assume low BER and no disconnections are faced with a hostile wireless environment;
- the notion of handing-off the connection to a different cell, neighboring or overlaid, is non-existent in the wired world;
- protocols that rely on link-local broadcast capability are not always easy to port to a wireless environment because the techniques used in several wireless networks cannot be easily modified to support link-level multicast. GSM's time division multiple access, CDMA's power control, and the fact that, in most current cellular systems, the base station is not an IP-layer router but a lower layer entity, make such attempts difficult.

IPv4 address shortage is also a problem now. For example, GPRS operators usually have to rely on *Network Address Translation* (NAT) and are

forced to assign private IP addresses (usually over PPP) to the mobile devices that register with their networks. All the usual NAT limitations affect IP multicasting although work-arounds do exist [31]. Also, some cellular operators have their own interpretation of 2.5G and 3G multicasting (simple cell broadcast usually) which is more limited than the IP-based multicasting envisaged for an all-IP wireless Internet.

6.2 IETF Mobile IP Multicast Support

The current IETF proposed standard for Mobile IPv4, RFC 3220 [24], devotes no more than a single page to multicast packet routing in conjunction with Mobile IP. Two methods are mentioned, which are referred to by [8] as *Remote Subscription* (MIP-RS) and *Bi-directional Tunneled Multicast* (MIP-BT). Both methods are only relevant when MHs are visiting a foreign subnet. We only describe them with M-IP operating in FA mode, since this is the preferred mode of operation both for wireless communications and for interoperability with C-IP.

Remote subscription - This may be used only when a multicast router is present in the visited subnet (this router may be co-located with the FA). The MH can use IGMP to subscribe to any number of groups using this router.

Bi-directional tunneled multicast - Another option is for the MH to setup a bi-directional tunnel to its HA (this HA must also be a multicast router). The HA will join groups on the MH's behalf. The tunnel is used to send IGMP messages and receive multicast packets. In this case, double encapsulation is required: first, the HA has to encapsulate the multicast packet inside another packet addressed to the MH's home address, and then, encapsulate once more as specified by Mobile IP. This means that the MH must be able to decapsulate the multicast packet even if it uses an FA for the standard Mobile IP decapsulation procedure.

Both approaches have their disadvantages. MIP-RS requires that the MH re-subscribes to a potentially large number of multicast groups after every subnet move. This delay will cause packet losses for the MH and will require rearrangements of the multicast tree with each move. MIP-RS assumes that a multicast router will exist in the visited subnet. Also, MIP-RS generally causes *get-ahead* and *lag-behind* effects (terms borrowed from [6]). These may happen when MHs register with a new multicast router and re-subscribe to multicast groups. Because of the complex nature of a multicast router tree, some edge routers may lag behind others in their reception of multicast packets. One solution to the get-ahead problem is for the host to accept the loss of some packets and rely on higher layer protocols to recover from this loss. More sophisticated solutions involve buffering at the multicast routers. The lag-behind problem may be solved locally at the host, where higher

layer protocols can discard already received packets, or at the router if somehow the router is signaled to drop packets before forwarding them to its internal subnet (which is important if the subnet is a wireless, bandwidth limited one).

With MIP-BT, the multicast tree does not have to be rearranged, since the MH's HA will be a stationary multicast receiver. Also, if the HA buffers multicast packets, no get-ahead or lag-behind problems will exist for the MH. MIP-BT, however, is associated with the three following unwanted phenomena:

- If an HA supports many MHs, all visiting the same foreign subnet and all subscribing to the same groups, then multiple point-to-point channels will have to be setup between the same HA and the foreign subnet, each one carrying the same information. The duplicate packets will strain the (potentially wireless and bandwidth limited) visited subnet.
- If, somehow, the HA detects this and sends only one packet copy per foreign subnet, then a problem known as *tunnel convergence* (see figure 2) may still occur if *many* HAs have MHs, all visiting the same foreign subnet and all subscribing to the same group. This will also lead to unnecessary packet duplication.
- Another way for unwanted duplicate packets to appear on the foreign subnet is if a host, local to that subnet, already subscribes to one or more of the groups some MHs wish to subscribe to. Since MHs will receive multicast packets encapsulated in unicast packets addressed to their home address (see MIP-BT description), it would not be easy for the FA or some multicast router in the foreign subnet to detect this duplication and stop the unnecessary transmission.

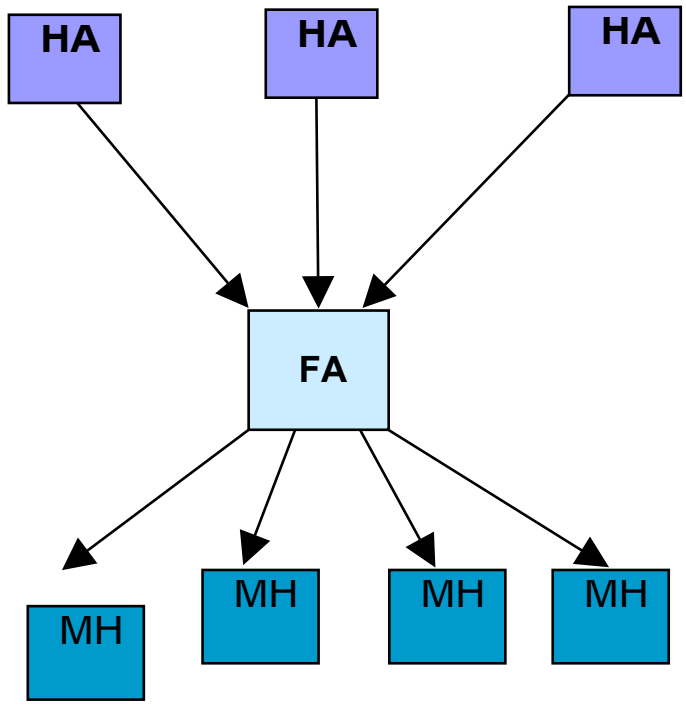

Figure #-2. The Tunnel Convergence Problem

6.3 Extensions to the IETF Approach

The following examples are alternatives or extensions to the basic Mobile IP and IP multicast interoperability approaches proposed by the IETF. Each one of these examples improves on the basic mechanisms but we believe that further refinement is needed before these solutions can be widely deployed.

Mobile Multicast (MoM) *Protocol*

MoM [8, 16] is based on MIP-BT and its key extension is the use of a *Designated Multicast Service Provider* (DMSP). DMSPs attempt to solve the tunnel convergence problem (see section 6.2). A DMSP for a given multicast group is an HA chosen by the visited subnet's FA out of the many HAs that forward packets for the specific group to the visited subnet. MoM

supports choosing more than one DMSP for redundancy and it also supports *DMSP-handoff*, which is necessary when a DMSP has no more MHs of its own in the visited subnet that require packet tunneling.

MoM-specific algorithms are executed every time:

- MHs arrive at a foreign subnet;
- MHs depart from a foreign subnet;
- MH registrations with the FA time out;
- unicast or multicast packets from the HAs arrive at the FA.

The subnet's FA keeps track of HAs, MHs, and multicast groups, so it always has enough information in order to choose a DMSP and to perform DMSP-handoffs whenever required. MoM is the most cited alternative to MIP-BT and MIP-RS.

Mobile Multicast with Routing Optimization (MMROP)

MMROP [18] is based on MIP-RS and its key extension is the introduction of the *Mobility Agent* (MA) entity, which attempts to solve the get-ahead problem due to handoffs (see section 6.2). This is done to "ensure routing efficiency and no packet losses from roaming" [18]. MAs are FAs that route missing packets (via tunneling) to neighboring subnets. MMROP works as follows: let FA1 and FA2 be an MH's old and new foreign agents respectively. Let's assume the MH was subscribed to group G through FA1. The MH will attempt to resubscribe to group G through FA2, at which point FA2 will start buffering packets. Upon joining the new subnet, the MH will look at the sequence numbers of the packets for G and decide whether or not it should ask for cached packets from FA2. (MMROP assumes packets are somehow numbered.) If FA2 cannot supply these packets, it will request FA1 to continue transmitting packets to the MH through a tunnel between FA1 and FA2, until FA2 and the MH are synchronized, at which point FA2 will start delivering packets to the MH through its own multicast subscription.

Constraint Tree Migration Scheme (CTMS)

CTMS [7] is an attempt to design a new global multicast routing protocol that would improve on CBT [5] when it comes to highly dynamic multicast configurations, such as those found when multicast listeners are mobile. CTMS "automatically [migrates multicast trees] to better ones, while maintaining the QoS guarantees specified by mobile users" [7]. CTMS uses fewer resources per multicast tree and, as a result, packet losses due to

reconfigurations and join delays are reduced. CTMS is a good alternative to existing mobile routing protocols but its adoption will be difficult, considering most multicast routers still run DVMRP.

Multicast Scheme for Wireless Networks (MobiCast)

MobiCast [27] is based on MIP-RS and its key extension is the introduction of the *Domain Foreign Agent* (DFA) which serves many small adjacent wireless cells. A hierarchy is introduced, with small cells being organized into one *Dynamic Virtual Macrocell* (DVM). Micromobility is thus hidden from the global multicast mechanism, which does not require re-configuring when handoffs occur within the same DVM.

6.4 IGMP Mobility Support and IGMP Assumptions

IGMP was designed with Ethernet networks in mind. Its basic functionality assumes link-level native multicast. Also, its soft-state timers require that multicast listeners repeatedly announce all the multicast groups to which they are subscribed. This happens every time the subnet's designated multicast router issues a query. We present IGMP's two main problems with respect to mobility support:

IGMP is not suitable for point-to-point links – If the local multicast router is not only connected to an Ethernet subnet, but also has interfaces that connect to *point-to-point* links, then, IGMP queries have to be issued to each one of these interfaces [28]. The IGMP replies will not be heard by all other participants unless the router specifically multi-unicasts them to every one of the point-to-point links that it supports. This increases delay, data traffic and state information needed at the router. Instead of just the group list which the global multicast mechanism requires (see section 3), the multicast router must record per-host information [28]. Indeed, as we mentioned, many cellular networks currently offer only point-to-point links for user data. The PPP protocol that is usually used on these networks (as well as on most Internet home connection technologies) to support IP packet transfer does not have the same semantics nor does it use the same techniques as the shared Ethernet protocol.

IGMP is not suitable for mobile hosts – We already mentioned (section 6.1) that mobile hosts don't have the luxury to monitor network traffic constantly. That would place a burden on their battery and processing power. Also, mobile hosts should not be forced to keep resending unnecessary information if this can be avoided. The IGMP soft state timers, although they are simple to implement and they contribute to IGMP's robustness, force the hosts to keep repeating the same data for as long as each one of their group

subscriptions is active. A solution proposed in [29] suggests using explicit *join_group* and *leave_group* messages which would require from hosts (which are only multicast listeners) to send significantly fewer packets. This scheme may interoperate with "traditional" IGMP.

7. NEW PERSPECTIVES ON MOBILE MULTICAST

7.1 Multicast Semantics and Mobility

An extension to the issues raised in section 6.1 is that multicast semantics need to be reexamined in the presence of mobility. We present a simple example based on ideas raised on [8] that exposes this problem. Let's assume X and Y are two Ethernet-based IP subnets. Let's also assume that MHs with home addresses in subnet X are designated Xi and that MHs with home addresses in subnet Y are designated Yj. Some Xi MHs are visiting subnet Y and some Yj MHs are visiting subnet X. If an IP multicast packet addressed to the link-local all-hosts group 224.0.0.1 is directed towards subnet X, then there are three possibilities, according to [8], about what should happen:

- the packet should only be delivered to MHs in subnet X, regardless of whether they belong to the Xi or Yj set;
- the packet should only be delivered to all MHs in subnet X that belong to the Xi set;
- the packet should only be delivered to all MHs belonging to the Xi set, irrespective of their current location.

Obviously, there exist techniques for each one of the three possibilities. But which one is semantically correct? There is no right or wrong answer. This depends on the service protocol that originated the link-local all-hosts multicast. We refer to three service examples, each one assuming a different interpretation:

- an advertisement for public network printers that are present on a specific floor;
- an advertisement for available high-quality color photo printers, to be used as part of subnet X's core business;
- an administrator's advertisement, describing the new authentication procedure for subnet X's SMTP server.

Currently, there is no clear IPv4 mechanism that would help mark the

advertisement packet and allow the mobility protocol to make an informed decision. IPv6 can, however, differentiate between *link local*, *site local* and *organizational local* multicast scopes.

7.2 Mobile Multicast Requirements

Some general issues that should affect all mobile multicast solutions are the following:

Significant vs non-significant moves - Let MH be a host subscribed to a set of multicast groups. If the move of MH to a new subnet causes the subnet's multicast router to subscribe to new multicast groups, then the move is said to be *significant*. Otherwise, if due to existing subscriptions packets addressed to the MH's set of groups are already being transmitted on that subnet, then the move is *non-significant*. Mobile multicast protocols will have to differentiate quickly between these two types of moves. Ideally, non-significant moves must have no effect at all on the global mechanisms. All the necessary information to identify non-significant moves can be found within the local subnet and at the subnet's multicast router. In addition, significant moves should appear identical to non-significant moves from a user's perspective.

Multicast packet buffering – Although, in theory, IP multicast is a best effort protocol, in practice, if mobile multicast schemes are to work efficiently, packet buffering should happen at the IP multicast layer. For example, with MIP-BT, the HA may buffer packets before tunneling them to the MH. This is necessary in order to achieve smooth handover when the MH moves to a new subnet and reestablishes the bi-directional tunnel. With MIP-RS the situation is more complicated. Depending on how MIP-RS is used, both the FA and a local multicast router on the visited subnet are candidates for buffering packets. The main problem with buffering is the following: buffer packets until when? In a wireless environment, with significant and non-significant moves, disconnections due to handoffs, disconnections due to physical layer problems and disconnections due to user intent, it will not be easy to judge how long buffering should go on. If multiple entities buffer simultaneously complexity increases. Soft-state timers must adapt based on a number of parameters and on input from both lower and higher protocol layers.

Mobile subnets - Ships, planes, trains, and even cars can be thought of as mobile subnets, each one with several local mobile hosts. Dealing with these as one logical entity will greatly assist routing protocols, ease tunnel convergence problems and minimize state information kept throughout the internetwork.

Roaming - The problem of global roaming between different

administrations is very difficult to solve, even for point-to-point communications (the same applies to vertical handoffs). Even if mobile multicast routing protocols are simulated and tested, true mobile multipoint communications will ultimately rely on sophisticated authentication mechanisms and pricing schemes.

7.3 Cellular IP and Mobile Multicast

In this section we present a model that integrates local multicast mechanisms to Cellular IP (C-IP). By using C-IP in conjunction with Mobile IP (M-IP) in a hierarchical manner similar to the MobiCast scheme (see section 6.3) we outline a mobile multicast scheme based on C-IP and M-IP interoperability ideas developed within the IETF. In our scheme, a MobiCast DFA is a C-IP gateway and a MobiCast DVM is a C-IP subnet. Introducing multicast support to C-IP is relatively straightforward considering that the basic C-IP forwarding mechanism is simple. However, a real deployment would be necessary in order to test the scalability of our proposed architecture.

We chose C-IP and a method based on MIP-RS because we took into account not only the current evolution of the Mobile IP and Cellular IP specifications but also the real network configurations that people deploy. These include campus-wide 802.11 internet-works, UMTS cells and the future DVB-T macro-cells. It is our position that the MIP-BT based tunneling schemes (although friendlier to current multicast routing protocols) are simply not scalable enough. The current *Content Delivery Network* (CDN) trends strengthen our belief that content and services need to be pushed as much as possible to the edges of the Internet and that hosts should first try to exploit whatever resources they have available in their immediate environment (i.e. follow the *locality* principle) and only when this fails should hosts try to access more distant resources.

As we have already mentioned, the way Cellular IP forwarders view the MH address space is “flat”. Inside a Cellular IP access network, the MHs use their home address as identifiers. Since the relatively simple mapping inside a forwarder maps these IP identifiers to forwarder ports, we can safely say that multicast (Class D) IP addresses would not appear inherently different from unicast addresses, at least as far as a potential mapping is concerned.

In addition, C-IP has keep-alive mechanisms that look similar to IGMP’s keep-alive soft-state mechanisms. These mechanisms serve similar purposes: C-IP maintains the mappings that concern active hosts as long as these hosts send data packets or the special C-IP route_update packets. On the other hand, IGMP responses to router queries serve to maintain the subscription to a specific multicast group. It is theoretically possible to adapt the C-IP route

update mechanism to supplant IGMP's operational semantics. In practice, to achieve this, MHs that are subscribed or wish to subscribe to a particular multicast group can send C-IP route_update packets, but instead of using their own IP address in the source IP field, they can use the multicast group address instead. These packets will not only update the forwarders' mappings, but also, when they reach the C-IP gateway (which we assume to be a M-IP FA and a multicast router as well) they can cause the gateway to subscribe to the particular group using the global multicast mechanisms. Of course, if the gateway is receiving the group already, there is no need to graft to the particular multicast tree again. In this way, we replace IGMP by native C-IP mechanisms.

One may say that we view multicast groups as virtual C-IP hosts. In order to complete the picture, C-IP forwarders should be able to handle *IP address-to-multiple port mappings*. Also, MHs should be able to use their IP stack in a rather unconventional way (placing the multicast address in the source field can be considered "non-standard"). A result of all the above is that packet duplication inside the C-IP access network due to the multicast transmission will only happen when paths towards receivers diverge, because of the way the C-IP forwarders update their mappings. This is a basic requirement for efficient multicast protocols. The following figure (figure 3) depicts a multicast transmission source (sending data to multicast "channel" 224.1.2.3) somewhere on the Internet and 3 MHs inside a C-IP access network that are subscribed to the 224.1.2.3 group. Until they reach base station/C-IP forwarder BS2, there is no need for the multicast data packets to be duplicated.

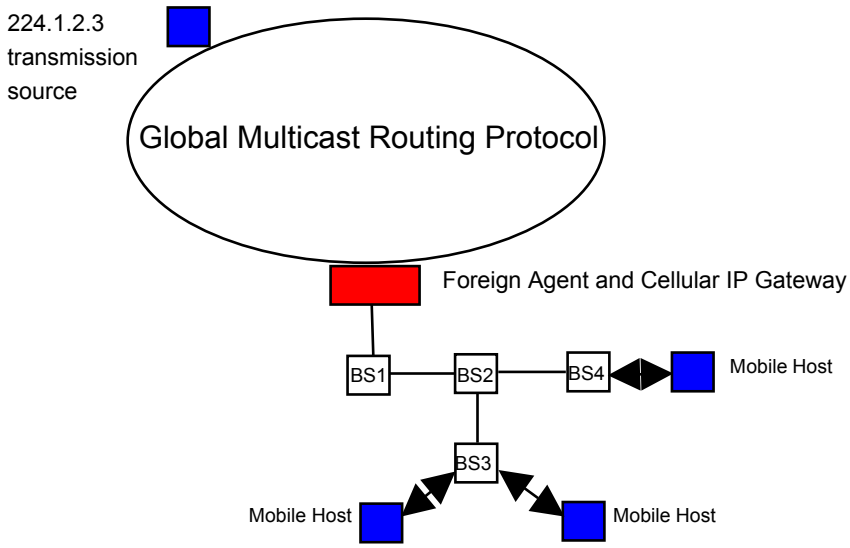

Figure #-3. Combining Cellular IP and Mobile Multicast

7.4 The All-IP Wireless Internet

Having completed our discussion on the basic mobile and multicast protocols and their interoperability issues, we describe the all-IP wireless internetwork that we envisage will support improved versions of all the aforementioned protocols. It will be made up of many, superimposed, cellular technologies. If we ignore the satellite component, which is usually the least cost-effective way to create wireless cells, we still have many promising technologies with which to build this hierarchical cell structure. DVB-T may be used for metropolitan size (1–100 Km) macro-cells, 3G systems, such as UMTS, may be used for neighborhood size cells, and the 802.11 variations, wherever available, for local micro-cells and pico-cells. Although DVB-T is unidirectional, extensions do exist that provide bidirectional functionality. Even without these extensions, 3G networks are perfectly suited to provide the necessary return channel [4].

DVB-T, with 5-30 Mbps of shared bandwidth and excellent support for mobility within a DVB-T macro-cell, could become the technology of choice for delivering IP multicast traffic, bypassing most of the problems we described in previous sections. Going down the cellular hierarchy, we would

have overlapping 3G cells (running versions of Mobile IP) and, then, localized 802.11 networks (running versions of Cellular IP). This structure needs to be augmented with umbrella cells that can handle traffic for fast moving receivers in trains or in cars. Mobile subnets could also be supported by installing Cellular IP gateways on ships and on trains, which would support micro-mobility within the moving subnets.

Devices with multiple interfaces are starting to appear. Common offerings include support for 802.11, combined with either DVB or GPRS. As long as IP and Mobile IP are accepted standards we can expect that all future devices and networks will support them. For unicast applications, alternatives to TCP will be used most of the time and improved variations of TCP will provide backwards compatibility. Transcoding functions and filter mobility protocols will be standardized and they will be put to use by all wireless network providers.

8. CONCLUSIONS

We examined the multipoint communications problem assuming an all-IP wireless Internet with mobile hosts, a fixed network infrastructure, and a best effort network layer. We focused on reusing IETF protocols where possible. IP multicast and Mobile IP were obvious choices. We showed how existing transcoding techniques may be used. However, we saw that the proposed IETF IP multicast and Mobile IP interoperability solutions are not perfect and that they require extensions. We presented a number of additions and alternatives to the basic scheme. We mentioned an approach involving Cellular IP and we offered our perspective on additional multicast mobility issues. Finally, we described the all-IP wireless Internet we envisage.

Still a lot of functionality needs to be added to the basic IP multicast and Mobile IP offerings before infrastructure-less networks, strong reliability and security are also supported. Ultimately, global roaming and pricing agreements will be necessary to complete this ideal vision of mobile multicast in an all-IP wireless Internet.

REFERENCES

[1] D. Agrawal, C. Cordeiro, H. Gossain, “Multicast: Wired to Wireless,” *IEEE Communications Magazine*, 40(6):116-123, June 2002.

[2] A. Acharya, A. Bakre, and B.R. Badrinath, “IP Multicast Extensions for Mobile Internetworking,” *Proceedings of 1996 IEEE INFOCOM*, pp. 67-74, San Francisco, CA, March 1996.

[3] A. Acharya and B.R. Badrinath, "A framework for delivering multicast messages in networks with mobile hosts," *ACM/Baltzer Journal of Mobile Networks and Applications*, 1(2):199-219, October 1996.

[4] Ad hoc Group DVB-UMTS, "The Convergence of Broadcast & Telecomms Platforms," *Public Document, www.dvb.org*, March 2001.

[5] A. Ballardie, J. Crowcroft, and P. Francis, "Core based Trees (CBT) – An architecture for scalable inter-domain multicast routing," *Computer Communications Review*, 23(4):85-95, October 1993. (Proceedings of the ACM SIGCOMM'93.)

[6] A.T. Campbell, J. Gomez, and A.G. Valko, "An Overview of Cellular IP," *IEEE WCNC*, New Orleans, September 1999.

[7] K. Chen, N. Huang, and B. Li, "CTMS: A novel constrained tree migration scheme for multicast services in generic wireless systems," *IEEE JSAC.,* 19:1998-2014, October 2001.

[8] V. Chikarmane, C.L. Williamson, R.B. Bunt, W. Mackrell, "Multicast Support for Mobile Hosts Using Mobile IP: Design Issues and Proposed Architecture," *ACM/Baltzer Mobile Networks and Applications*, 3(4):365-379, Jan. 1999.

[9] S. Deering, "Host Extensions for IP Multicasting," *RFC 1112*, August 1989.

[10] S. Deering, "Multicast Routing in a Datagram Internetwork," *Ph.D. Thesis, Department of Computer Science, Stanford University*, 1991.

[11] S. Deering, W. Fenner, and B. Haberman, "Multicast Listener Discovery for IPv6," *RFC 2710*, October 1999.

[12] W. Fenner, "Internet Group Management Protocol, Version 2," *RFC 2236*, November 1997.

[13] S. Floyd, V. Jacobson, and S. McCanne, "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing," *IEEE/ACM Transactions on Networking*, 1995.

[14] G. Gardikis, E. Pallis, A. Kourtis, "Beyond 3G: A Multi-Service Broadband Wireless Network with Bandwidth Optimisation," *IST Project "Multi-Services Management Wireless Network with Bandwidth Optimisation," Public Document*, *www.openmux.com-/mambo/Public/Beyond 3G.pdf*, 2000.

[15] J.K. Han and G.C. Polyzos, "Multi-Resolution Layered Coding for Real-Time Image Transmission: Architectural and Error Control Considerations," *Real-Time Imaging*, 4(4):275-298, Academic Press, August 1998.

[16] T. Harrison, C.L. Williamson, W.L. Mackrell, and R.B. Bunt, "Mobile Multicast (MoM) Protocol: Multicast Support for Mobile Hosts," *Proceedings ACM MOBICOM*, Budapest, Hungary, September 1997.

[17] V. Jacobson, "Congestion Avoidance and Control," *Proceedings ACM SIGCOMM*, CA, USA, August 1988.

[18] J. Lai, W. Liao, M. Jiang, and C. Ke, "Mobile Multicast with Routing Optimization for Recipient Mobility," *Proceedings IEEE ICC2001,* pp. 1340 - 1344, June 2001.

[19] Y. Lin, "A Multicast Mechanism for Mobile Networks," *IEEE Communications Letters*, 5(11), November 2001.

[20] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-driven Layered Multicast," *Proceedings ACM SIGCOMM*, Stanford, CA, August 1996.

[21] J.C. Pasquale, G.C. Polyzos, E.W. Anderson, and V.P. Kompella, "Filter Propagation in Dissemination Trees: Trading off Bandwidth with Processing in Continuous Media Networks," *NOSSDAV, D. Lecture Notes in Computer Science 846*, Shepherd et al., eds., Springer-Verlag, Berlin, Germany, pp. 259-268, 1994.

[22] J.C. Pasquale, G.C. Polyzos, E.W. Anderson, and V.P. Kompella, "The Multimedia Multicast Channel," *Internetworking: Research and Experience*, 4:151-162, 1994.

[23] C. Perkins, "IP Encapsulation within IP," *RFC 2003*, October 1996.

[24] C. Perkins, editor, "IP Mobility Support for IPv4," *RFC 3220*, January 2002.

[25] C. Perkins and D.B. Johnson, "Route Optimization in Mobile IP," *Internet Draft draft-ietf-mobileip-optim-11*, September 2001.

[26] A.C. Snoeren, H. Balakrishnan, and M.F. Kaashoek, "Reconsidering Internet Mobility," *Proceedings HotOS-VIII*, May 2001.

[27] C.L. Tan and S. Pink, "MobiCast: A Multicast Scheme for Wireless Networks," *ACM MONET*, 5(4):259-271, 2000.

[28] G. Xylomenos and G.C. Polyzos, "IP Multicast for Mobile Hosts," *IEEE Communications*, 35(1):54-58, January 1997.

[29] G. Xylomenos and G.C. Polyzos, "IP Multicast Group Management for Point-to-Point Local Distribution," *Computer Communications,* 21(18):1645-1654, Elsevier Science, December 1998.

[30] G. Xylomenos, G.C. Polyzos, P. Mahonen, and M. Saaranen, "TCP Performance Issues over Wireless Links," *IEEE Communications*, 39(4):52-59, April 2001.

[31] "How Does Multicast NAT Work on Cisco Routers?" *Technical Note*, www.cisco.com/warp/public/105/multicast_nat.html, 2001.