INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

An Integrated Platform for Multicast, Reliable Multicast and Quality of Service Support in the Regional Registration Mobile-IP Environment

by

Hassan Omar

A dissertation submitted to the Graduate Faculty in Engineering in partial fulfillment of the requirements for the degree of Doctor of Philosophy, the City University of New York

2002

This manuscript has been read and accepted for the Graduate Faculty in Engineering in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

9/12/02
Date

Professor **Tarek Saadawi**
Chair of Examining Committee

9/12/2002
Date

Professor **Mumtaz Kassir**
Executive Officer

Professor **Myung Lee**
City University of New York

Professor **Umit Uyar**
City University of New York

Professor **Yi Sun**
City University of New York

Dr. **Donald DiMarzio**
Northrop Grumman

Supervisory Committee

The City University of New York

Abstract

An Integrated Platform for Multicast, Reliable Multicast and Quality of Service Support in the Regional Registration Mobile-IP Environment

by

Hassan Omar

Advisers: Professor T. Saadawi, Professor M. Lee

The standard based Mobile-IP protocol allows IP hosts to keep an active communication session, while moving between different networks. The regional registration Mobile-IP approach was introduced to reduce the overhead associated with the location registration process. The scalability features, associated with the regional registration Mobile-IP, qualify it as a deployment candidate to support mobility in IP networks.

A wide variety of applications, such as teleconferencing and video-on-demand type of services, will utilize IP multicast for traffic delivery. Providing an efficient system that supports IP multicast, reliable delivery, and differentiated quality of service is a challenge for systems providing mobility support.

Although few approaches were proposed to support such services in environments that support mobility, applicability to the Mobile-IP regional registration environment has not

been addressed yet. Since those proposed approaches may not work, or will not perform efficiently in the regional registration environment, providing an efficient solution to support multicast and reliable multicast delivery becomes a critical issue. In addition, the regional registration operation exhibits strong sensitivity to failure of mobility agents. To tolerate failure events, an efficient fault recovery scheme is needed, such that service disruption is minimized.

In this work we propose an integrated platform that supports IP multicast and reliable multicast in the Mobile-IPv4 regional environment. A robust procedure to recover from mobility agents failure events is described. In addition, we illustrate a platform extension to provide quality of service support based on the DiffServ model. The proposed schemes take advantage of the inherent characteristics of the regional registration environment, providing the intended functions with minimum additional overhead. A unique feature of this platform is that the support for those services, traditionally treated separately, is integrated with the control domain of the regional registration basic platform. This integration provides a very attractive set of features, such as the reuse of existing control data structures and control messages. The suggested integrated platform shows strong scalability features against the number of mobile nodes and mobility rates. Also, this platform tolerates the dynamics associated with the multicast group membership and the QoS requirements in an environment that supports mobility.

Acknowledgments

I would like to express my deepest appreciation to my advisers Professor Tarek Saadawi and Professor Myung Lee for their continuous support, encouragement and patience throughout my study. I am grateful for their guidance and advice on the technical research issues. I like to thank Professor Saadawi and Professor Lee for the aid in research through the different programs.

I would like to thank Professor Umit Uyar, Professor Myung Lee, Professor Yi Sun and Dr. Donald DiMarzio, members of my supervisory committee for taking the time to review this work and providing valuable comments. My sincere gratitude is to Professor Mumtaz Kassir, the Executive Officer of the Ph.D. program in Engineering for providing helpful support in the administrative issues.

My gratitude to the students at the computer communications lab, for the appreciated input during the meetings and the interesting technical discussions. I would like to thank the administrative staff, especially Victoria Ockay, for her assistance during my study in the Electrical Engineering Department.

Finally, I am eternally grateful to my parents and my family for their support and sincere encouragement.

Contents

List of Figures

1. Introduction

The advances in computing technology have produced miniaturized powerful computers easy to carry while roaming between different areas. Meanwhile, the growth of wireless communications has been rapid since the introduction of the cellular telephone. Current developments will make it possible for mobile users to transparently access communication networks from anywhere at any time. Such universal network connectivity offers great promises for mobile users to support different applications traditionally supported only in a stationary environment.

Existing IP network protocols were not designed with mobility support in mind, such that host mobility support involves network reconfiguration and the user is denied connectivity till the completion of the reconfiguration process. A true mobile user should not be required to consider any special action when moving from one service area to another. One of the most important problems, concerning the support of mobile users, is the routing scheme that allows mobile hosts to move seamlessly from one location to another. Any solution for this problem has to interwork with the existing Internet routing infrastructure. Another important aspect in the design of mobility systems is location management. Elements involved in supporting mobility should allow for the generation, processing and maintenance of location information. The location management task has to be accomplished while producing the minimum overhead, and without stressing network resources.

In the IPv4 mobility support [1] specified in the current IETF specifications, a readdressing function is implemented by the home agent node on the home network. All datagrams destined to the mobile node will be forwarded to the home network, where they are readdressed and forwarded to the foreign network where the mobile node is currently located. A foreign agent node on the visited network will provide the inverse readdressing function when the datagram is delivered to the foreign network, then it will deliver the datagram to the mobile node. This mechanism implies that the mobile node has to register its location information with the home network immediately after each mobility event.

The local registration Mobile-IP approach [2] was introduced to reduce the frequency by which the location registration with the remotely located home agent is needed. In this approach, foreign agents are strategically located on a hierarchical configuration. Mobility within the same region requires sending a local registration to the nearest regional foreign agent, while inter-regional mobility events should be followed by home registration. Since foreign agents receiving the local registration requests are distributed over the areas where the mobile node is likely to visit, the local registration approach aids in reducing the delay associated with the registration process.

Although the Mobile-IP allows a mobile node to receive unicast traffic at its new location, further support is needed to provide other services. Applications requiring multicast, reliable multicast and differentiated QoS are among the most demanded services by both fixed and mobile users. An IP-mobility platform will need to efficiently support those services, in order to ensure wide deployment. In addition, support for those

services is required to be compatible with existing network elements, to protect current investments. Accordingly, a platform that supports those services in an integrated manner represents a very attractive and desired solution.

Applications that need information to be replicated on multiple hosts in different locations, such as teleconferencing and broadcasting services, require network multicast support. An efficient multicast mechanism will have the objective of delivering a copy of the same information to multiple destinations, while keeping the number of duplicate messages and the distance traveled by those messages as minimum as possible. Since IP protocol is the most commonly used network protocol, there is a need for the research and implementation efforts to target the subject of integrating Mobile-IP and multicast.

At the current time, IP multicast is supported primarily using the Distant Vector Multicast Routing Protocol (DVMRP) [3] and is deployed in the different regions connected by the MBONE [4]. Providing an efficient system that support IP multicast, in an environment where the multicast group members frequently change their locations, is a challenge for systems providing mobility support. The mobility introduces two factors that affect the performance. First, multicast datagrams can be forwarded to the previous service area (and not to the current one) due to outdated information regarding the mobile node current location. Secondly, the delay associated with joining and leaving the different multicast groups can affect the user perception of the service. The resulting performance degradation can be considerable in a mobility environment, since the process of joining and leaving the multicast groups is needed upon each mobility event. Any proposed approach to support multicast in the regional registration model will need

to accommodate the fact that the complete information, regarding the location of a mobile node, is distributed over multiple agents and is not centralized in a single location.

The main objective of an efficient reliable multicast protocol is to ensure the delivery of datagrams to the recipients without unnecessary overhead. A retransmission scheme is needed to compensate for missing datagrams, and has to be designed to function properly with a finite amount of memory. The finite resources, available within the reliable multicast platform elements, dictate placing lower limits on the minimum receiving capacity of each group member. Accordingly, a receiver in violation of those limits may be denied the reliable service. Avoiding signaling implosion is another critical factor to be considered when designing a reliable delivery scheme. Supporting reliable delivery of multicast datagrams, in Mobile-IP networks, may necessitate the introduction of new elements and functions. Further, considerable additional signaling may be required to support this service. The mobility of the hosts provides another possible source for packet loss, pausing more challenges to the reliable multicast solution. Several protocols have been proposed to support reliable multicast delivery. Researchers have shown that the ACK-tree reliable multicast approach provides desired scalability features.

The local registration systems consider requirements and assumptions that may affect performance aspects of the Mobile-IP system such as fault tolerance. In an environment where service interruption may not be tolerated, or where those base stations are located in hazardous conditions, the problem of fault tolerance becomes of a particular interest. In this work, we will present the issues associated with fault tolerance in local registration Mobile-IP systems. We will suggest approaches to enhance the robustness of such

systems against both home agent and foreign agent failures. The regional registration approach relies on co-operation between the FAs on the different levels of the hierarchy, such that the failure of a FA will deny service to other downstream FAs. The scope of the failure is greatly associated with the location of the affected FA. A fault-tolerant Mobile-IP system should have the capability of restoring operation and of providing the support services to the MNs despite the failure of some of the elements supporting the system. An efficient fault-tolerant system will provide a resilient system with minimum overhead.

To support QoS over IP networks, two architectures have been suggested. The Integrated Services architecture (IntServ) [5] uses a signaling mechanism to allocate per-flow resources over the network elements. Due to the maintenance overhead, associated with the per-flow states, the IntServ model does not provide attractive scalability features and thus limiting its implementation to medium size networks. The Differentiated Services (DiffServ) [6] model is positioned as a simple and scalable solution that pushes the complexity of the per-flow processing to the edges of the IP networks. In the core of the DiffServ network, the differentiated service is provided per traffic aggregate.

The DiffServ architecture was designed with no particular support for neither mobility nor multicast. In a mobile environment, resources associated with mobile nodes need to be allocated and released according to the mobility of the hosts. A home agent and multiple foreign agents will support the forwarding of the traffic between the mobile node and the corresponding node. Accordingly, multiple DiffServ domains will need to be involved in the QoS provisioning process. In addition to the base Mobile-IP [1], the route optimization [7] and the bi-directional tunnel [8] cases need to be supported in the

proposed platform. Supporting different multicast groups, with constantly changing memberships and multiple classes of service, presents additional challenges to the QoS support scheme. If the DSCP (Differentiated Service Code Point) value is copied from the original datagram to all transmitted copies at a multicast branching point, several problems can affect the performance of the system [9].

In this work, we will describe an integrated platform that supports multicast, reliable multicast and QoS services in the regional registration Mobile-IP environment. Capitalizing on the inherent features of the environment, those services will be supported with minimum additional overhead that leads to a corresponding better performance. To support the multicast service, a Local Multicast Service Provider (LMSP) is dynamically selected to allow flexibility and better performance in multicast datagrams forwarding. This approach can be tuned to satisfy a particular performance requirement. The reliable multicast delivery is supported using an ACK based tree, taking advantage of the hierarchical characteristics in the environment. To tolerate the failure of a FA, a service recovery scheme is proposed to accommodate both unicast and multicast services. A DiffServ scheme is integrated in this platform, allowing the support of QoS in the regional registration environment.

The objective of our proposed integrated platform is to offer a highly efficient mobility support environment, which allows wider deployment for Mobile-IP systems. Support for advanced services is provided, while keeping the system scalable and efficient.

2. The Mobile-IP Protocol

The Mobile-IP is an extension to the IP protocol that allows mobile nodes to continue in receiving datagrams wherever they happen to be attached to the Internet. The Mobile-IP protocol describes the additional control messages that are needed to allow involved nodes to manage their IP routing tables. The Mobile-IP protocol as described in [1], and the associated specifications, provide for a system that allows IP hosts [10] to move between different sub-networks without the need to tear down the established transport layer sessions. The mobility of the IP hosts is supported by two agents, the Home Agent (HA) and the Foreign Agent (FA). The Home Agent keeps track of the current location of its Mobile Nodes (MNs) and tries to keep the Correspondent Nodes (CNs), which are communicating with those MNs, updated with the current location information. The FA forwards packets to and from the MNs currently located in its service area. The Mobile-IP with the route optimization extension [7] describes a location information management scheme, where correspondent nodes can cache the MNs binding information. To ensure information consistency, both the HA and the FA need to co-operate to keep the CNs updated with the current MNs location information.

A mobile node is assigned a long-term IP address on its home network. This address is treated similarly to that permanent IP address associated with a stationary host. When away from its home network, the mobile node is associated with a care-of address that reflects the mobile node's current point of attachment. The MN may acquire the IP address through the DHCP protocol [70] or through another scheme. Hosts interested in

sending datagrams to the mobile node will use the home address as the destination address. The same address will be used by the mobile node as its source IP address for upstream datagrams. Such arrangement allows the mobile node to appear, to the correspondent node, as if it was a stationary host. Correspondent nodes do not need to change the destination address of the datagram upon mobility events associated with the MN. Figure 2.1 illustrates the elements involved in supporting the Mobile-IP protocol.

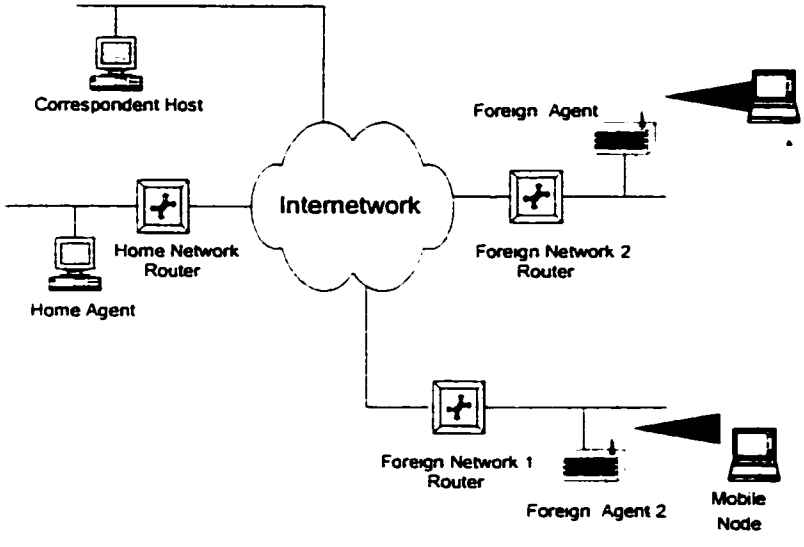

Figure 2.1 Elements supporting host mobility in the Mobile-IP environment

A deviation from this base Mobile-IP mechanism, which provides mobility transparency, is represented by the route optimization approach where the current location information of the mobile node is made available to the correspondent hosts. The route optimization requires the location information, at the CN, to be updated regularly.

The IP mobility as described earlier is intended for IPv4 environment. Capitalizing on features supported in the IPv6 [11], different specification describes the Mobile-IP support in the IPv6 environment [12].

2.1 The Base and the Route Optimization Mobile-IP: Overview

In the Base Mobile IP Protocol, when packets generated by a correspondent node are destined to a mobile node, the transmitted packets from the CN travel to the home agent which tunnels the packets and forwards them to the current care-of address (e.g., Foreign Agent) of the mobile node. The foreign agent de-tunnels the packets, and then forwards them to the mobile node. When the mobile node sends packets to a correspondent node, it sends them directly through the foreign agent. This sequence causes a triangle routing problem in the CN-MN direction, leading to performance degradation. The binding cache concept was introduced in the Routing Optimization IETF draft [7] to resolve the triangular routing problem. Figure 2.2 illustrates the triangle route path associated with the base Mobile-IP environment. One advantage of the base Mobile-IP is that it is transparent to the correspondent node, such that no modification is required in order to send and receive datagrams from a mobile node. The route optimization scheme requires the correspondent node to be enhanced with the binding cache support. Correspondent node with no support for the binding cache will revert to the base Mobile-IP mode of operation.

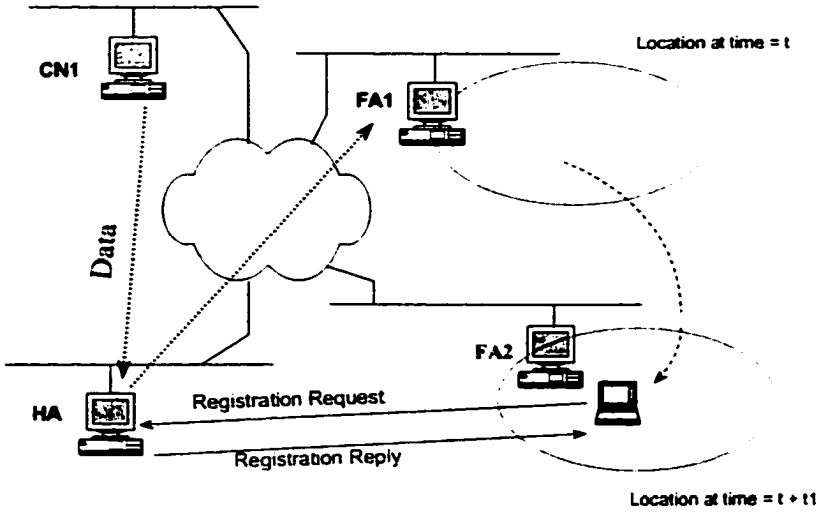

Figure 2.2 Triangular routing and registration messages in base Mobile-IP

The binding cache approach relies on updating the correspondent node with the new location information such that it can start sending future datagrams directly to the care-of address, bypassing the home agent as an intermediate node. The HA and the FA are to co-operate to provide the correspondent nodes with the updated location information. To support this capability, the HA and the FA need to detect and correct the situation where a CN is using an outdated binding information. Accordingly, the Binding Update and Binding Warning messages were introduced in [7]. Considering that the generation of those messages can be triggered by receiving an untunneled datagram, and not only upon mobility events, a robust mechanism is needed to limit the generation rate of those messages. Appendix A illustrates an adaptive approach that we propose to limit the number of those control messages. The suggested scheme adapts to the network

conditions to select an appropriate limiting rate for binding messages, maximizing the use of the available bandwidth and without introducing significant overhead delay. This features prevent the HA and the FAs from unnecessary generating large number of signaling messages, thus saving the scarce resources of the network.

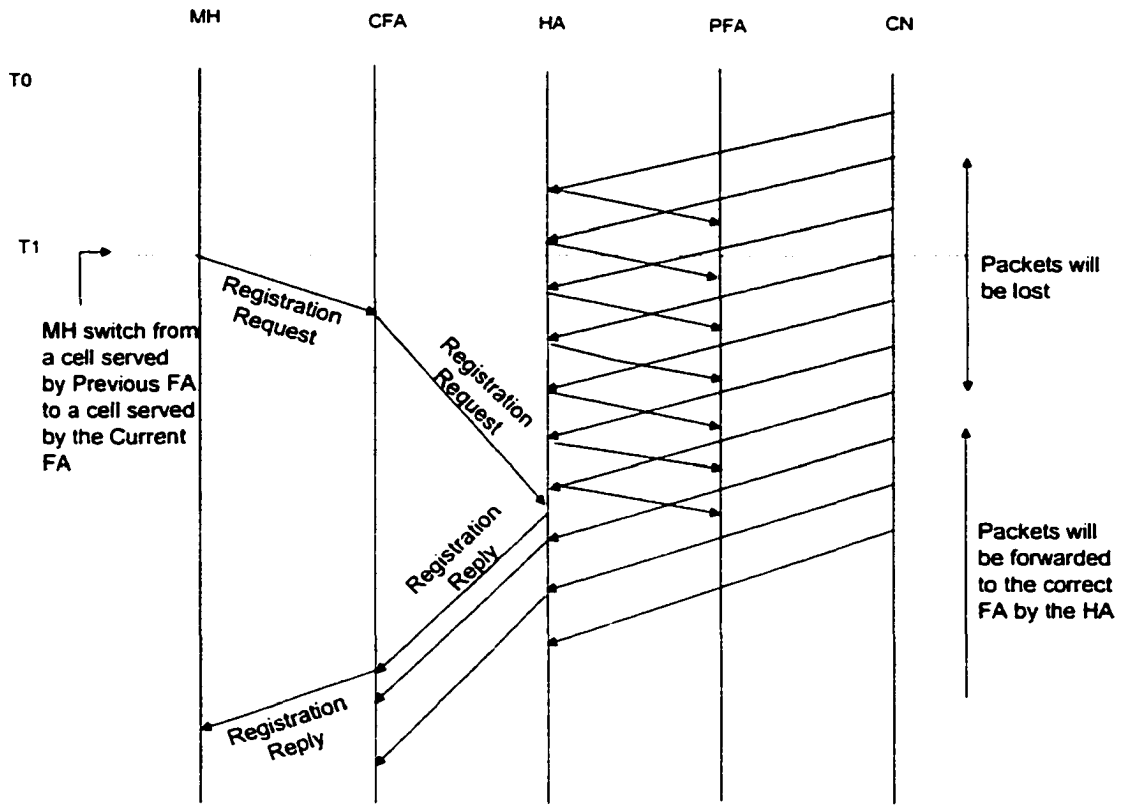

Figure 2.3 Control messages exchanged in Base Mobile-IP

Figures 2.3 and 2.4 illustrate the control messages associated with the Base Mobile-IP and the Mobile-IP with the route optimization extension, respectively. The route optimization scheme relies on support from the previous FA (PFA), which is the FA where the FA was located before the last mobility event.

In the base Mobile-IP, the control messages are limited to the Registration Request and the Registration replies messages. Datagrams lost are those intercepted by the HA while using the outdated binding information, before receiving the new registration request.

Datagrams lost in the case of route optimization are those generated by the CN before receiving the binding update from the HA. Datagrams intercepted by the previous FA before receiving the previous FA extension, informing about the new location of the mobile node, are also lost.

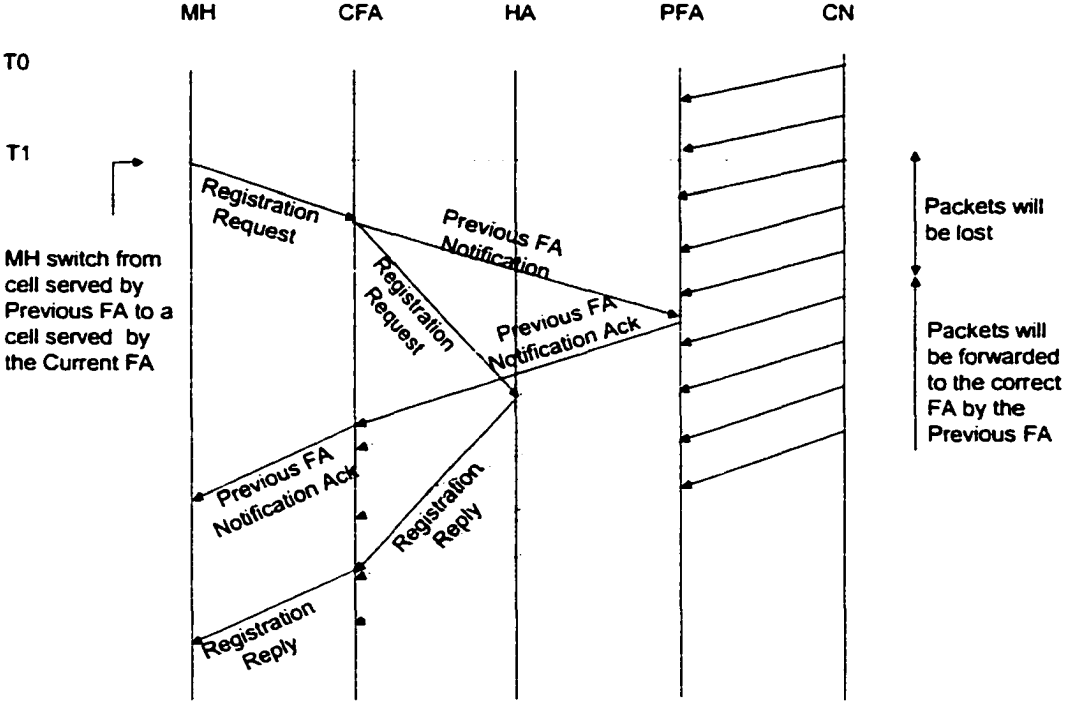

Figure 2.4 Control messages exchanged in Mobile-IP with the Route Optimization extension

Datagrams generated by the MN, and destined to a CN, carry the MN's home address as its IP source address. This address is not topologically correct, considering the network

address of the foreign network where the MN is currently residing. Firewalls placed on the edges of networks can detect this situation, and may interpret it as a spoofing attempt [13].

In addition to flagging this event as a security threat, the firewall will discard those datagrams. A possible solution is presented in [14], where datagrams towards the CNs are tunneled from the FA to the HA. This solution comes with the expense of non-optimal route for traffic generated by the MN.

The Mobile-IP fields used in the registration request and registration reply messages are illustrated in figure 2.5 and 2.6 respectively. Mobile-IP messages use the UDP [72] protocol. The extension part of the registration request allows the MN to ask for optional services or treatment.

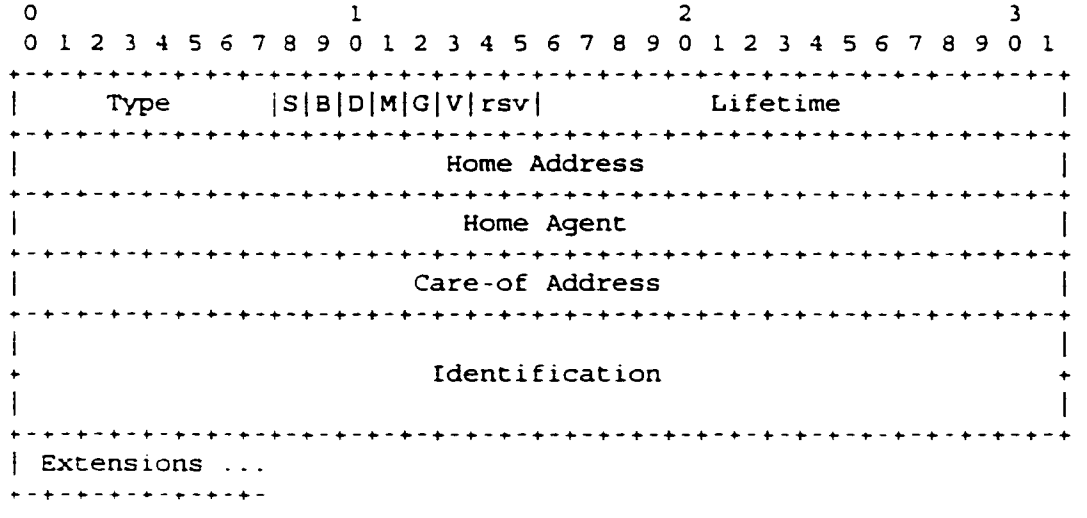

Figure 2.5 Fields in the Mobile-IP registration request message

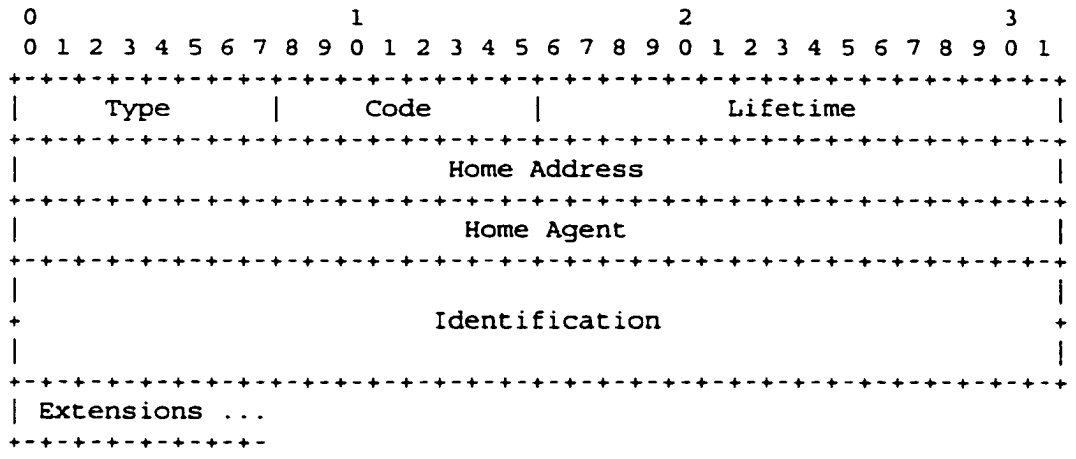

Figure 2.6 Fields in the Mobile-IP registration reply message

2.2 Tunneling and Encapsulation

The basic idea of tunneling is to encapsulate a datagram within an additional header that hides the details of the original datagram information. This external header entitles the encapsulated datagram to a different treatment throughout the networks between the two ends of the tunnel.

Mobile-IP requires each home agent and foreign agent to support tunneling datagrams using IP-in-IP encapsulation [15]. Minimal encapsulation [16] and GRE encapsulation [17] are alternate encapsulation methods that may optionally be supported by mobility agents and mobile nodes. Encapsulation is used also in the case of bi-directional tunneling [14], where datagrams are tunneled from the FA to the HA. To encapsulate an IP datagram using IP-in-IP encapsulation, an outer IP header is inserted before the

datagram's existing IP header. The outer IP header source address and destination address identify the end-points of the tunnel. The inner IP header source address and destination address identify the original sender and the ultimate recipient of the datagram respectively. The inner IP header is not changed by encapsulator except to decrement the TTL, and remains unchanged during its delivery to the tunnel exit point. Figure 2.7 illustrates the encapsulation process as implemented by a HA. On receiving a datagram from the CN, the HA will query its local table and identify the current care-of address. This address is then used as the destination address in the external header, while the home agent address is used as the source address. With the exception of the case when bi-directional tunneling is used, datagrams generated from the MNs are not tunneled.

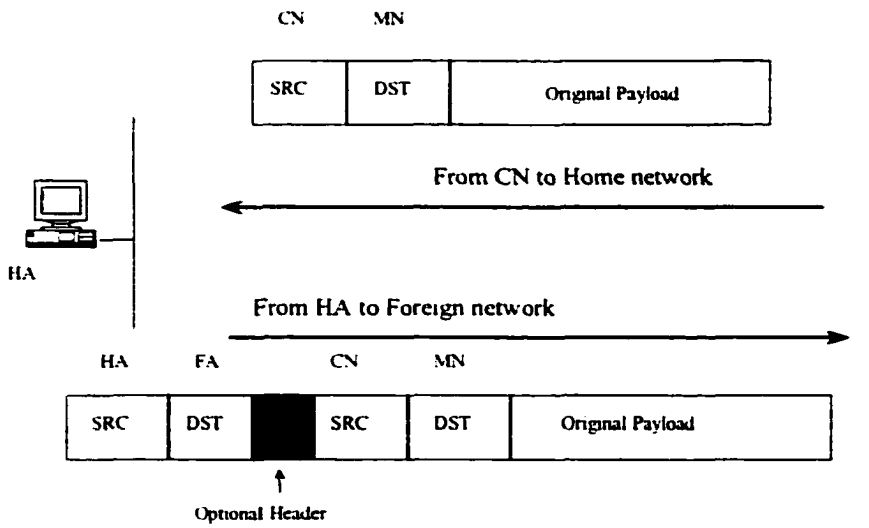

Figure 2.7 Datagram encapsulation and tunneling

When using encapsulation, care should be considered to ensure that the new external header allows the datagram to be treated as intended. IP header fields, similar to the TOS (Type of Service) field, may carry a new value after the encapsulation process. This can affect QoS treatment, if there is an intermediate node (or nodes) that uses this field to select a class of service. The same concern is valid also if other information in the header, like the transport protocol type, are being used for the same purpose.

3. Regional Registration Mobile-IP

Both the Base Mobile-IP and the Route Optimization schemes provide considerable amount of signaling messages to support mobility. Therefore, the regional registration approach [2] was introduced to address this concern.

3.1 The Regional Registration Concept

The basic idea behind applying the regional registration approach [2, 18] to the Mobile-IP is to allow the mobile node to send registration requests to a regional FA, that tracks regional movements but does not forward the requests to the HA. FAs are arranged hierarchically in a regional topology, and the MN is allowed to move from one serving area to another in the same regional topology without the need to bind location information at the HA.

3.2 Architecture and Operation

In this work we will consider local registration Mobile-IP systems were the FAs are strategically located on a hierarchy. For short, we will call such system HLRM-IP (Hierarchical Local Registration M-IP). In an HLRM-IP system, the MN is trying to minimize the amount of location information tracking required to maintain its

connectivity by identifying the smallest region for which the mobile node has not traversed any regional boundary. To accomplish this functionality, each ancestral FA considers the MN to be registered at the downstream FA. The FA on the top of the regional domain is referred to as the Gateway FA (GFA) or as the root FA. The FA advertisement contains the complete regional hierarchy of FAs supporting the local registration, along the path starting from this FA to the GFA. When moving to a new service area, the MN examines the hierarchical list of FAs in the new FA advertisement in search for the nearest common ancestor between the care-of-addresses of the new and previous service areas. The local registration request generated by the MN is then relayed from the FA currently serving the MN to the next higher level of the hierarchy and towards the common ancestor FA. In this way, each FA in the hierarchy between the MN and the gateway FA will be able to maintain a binding for the MN. The registration reply follows the same path but from the gateway FA to the leaf FA direction, allowing the intermediate FAs to examine the status of the registration request and update the binding accordingly.

The operation of the system may best be described by an example. Consider figure 3.1, the environment includes one HA and two gateway FAs, FA1 and FA2. The nodes named FAx are FAs supporting the HLRM-IP, while nodes tagged as Rx are regular routers supporting no FA functionality. Each FA announces the upstream lineage of the hierarchy that this FA is located on. For example, FA4 will announce the lineage FA4/FA3/FA1 while FA11 will announce FA11/FA7/FA5/FA3/FA1. We will consider the case where a MN is moving from home network to FA19, visiting the intermediate

agents FA4, FA8, FA11 and FA14. It is clear that it is only when MN moves to FA19 (step 4), then the MN will need to send the Registration Request to the HA. As long as the MN is moving within the area served by the same gateway FA (FA1 is the gateway FA for the FAs visited in steps 0 through 3), the HA does not need to be involved in the MN registration.

In the case of regional registration, the decapsulation and the reencapsulation occur at each level of the hierarchy until the datagram reaches the last tunnel endpoint. The actual decapsulation does not need to occur at each level of the hierarchy, and can be substituted by changing the source and destination IP addresses of the encapsulating IP header.

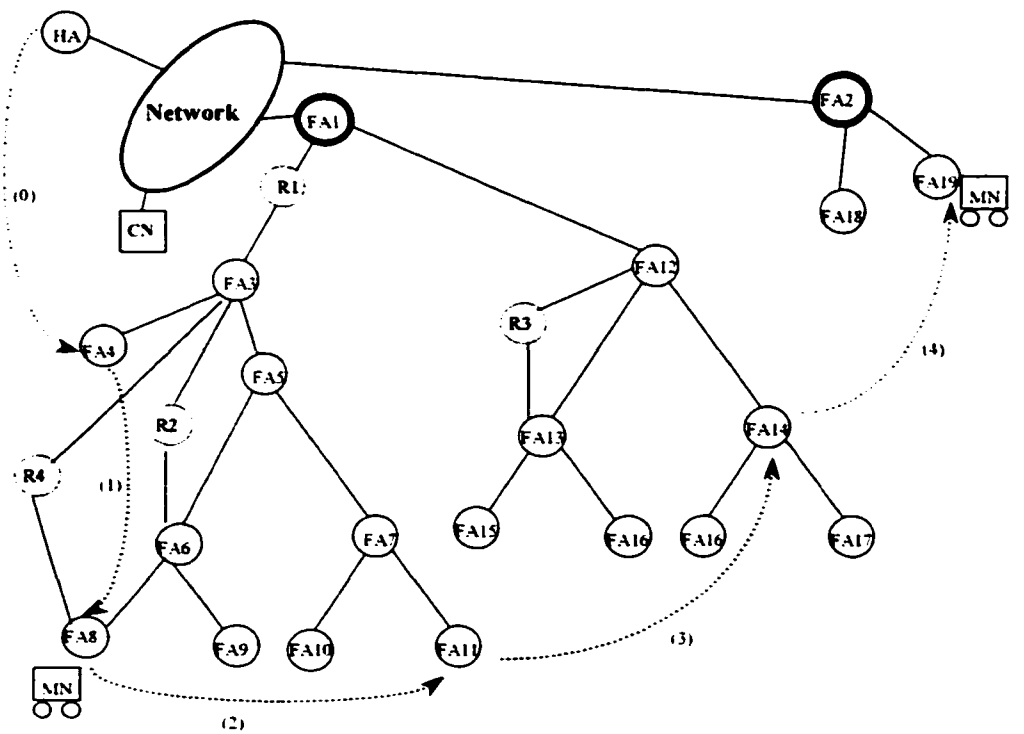

Figure 3.1 Environment that supports regional registration

3.3 Regional Registration Mobile-IP: Issues

The regional registration Mobile-IP provides undeniable advantage from the point of view of reducing the delay associated with signaling needed to support mobility. On the other hand, some of the features associated with this environment make it more sensitive to failure of FAs. The failure of a single FA can have a wide scope that causes service to be denied to a large number of MNs and FAs. Also, some proposed approaches to support multicast will not work in this environment. This may be the situation since they rely on the assumption that the HA has exact information regarding the MN's location, or that the location information is centralized on a single FA. This is not the case in the regional registration environment, where location information is distributed over the FAs constituting the lineage ending with the leaf FA. On the other hand, the regional registration environment exhibits some features that can be capitalized on to provide an efficient support for advanced services such as multicast and reliable multicast.

The research work related to local registration environment has focused mainly on the hand-off process to achieve optimum solution regarding datagram loss. In [19, 20], the authors provide a number of extensions to provide faster handoffs and to avoid the triangle routing. In [21] an extension that addresses reduced handoff latency and key management problems is provided.

To the best of our knowledge, little researches have been done to investigate the support for services like multicast, reliable multicast and DiffServ in the regional registration Mobile-IP environment. The work presented here is novel in that it considers

the integration of those advanced services support with the regional domain mobility management platform, leading to an efficient and scalable solution.

4. Multicast and the Regional Mobile-IP Environment

The multicast service allows information to be simultaneously distributed to different sites, without the need to send a separate copy of the same datagram using unicast. Applications requiring network multicast support are gaining increasing popularity. In this work, we will describe the issues and limitations associated with supporting multicast in the local registration Mobile-IP systems utilizing hierarchical organization. In particular, three different schemes that support multicast in this environment will be described. Those mechanisms will adapt to the architecture and requirement for the local registration mechanism and will take advantage of its inherent characteristics to provide better performance. We will point to the performance aspects where one approach outperforms the other.

4.1 IP Multicast

Applications that need information to be replicated on multiple hosts in different locations such as teleconferencing, broadcasting services and redundancy using checkpointing are in need for network multicast support. An efficient multicast scheme will have the objective of delivering a copy of the same information to multiple destinations, while keeping the number of needed duplicate messages and the distance

traveled by those messages as minimum as possible. The emerging and accelerated growth of portable computers and wireless communications suggests that mobile hosts are highly expected to be among the sources and the recipients of such multicast traffic. Civil and military applications are attracted to systems that provide service for a majority of mobile nodes with support from few stationary hosts. Since IP protocol is the most commonly used networking protocol, it is expected that research and implementation efforts will be directed to provide multicast support for IP hosts complying with the Mobile-IP protocol.

At the current time, IP multicast is supported primarily using the Distant Vector Multicast Routing Protocol (DVMRP) [3] and is deployed in the different regions connected by the MBONE [4]. Another routing mechanism, the Multicast Open Shortest Path First (MOSPF) protocol [23], is based on a link state algorithm. While the DVMRP constructs a delivery tree for each source in each multicast group, the Core Based Trees (CBT) approach [24] considers a single delivery tree for each multicast group. The Protocol-Independent Multicast (PIM) [25] uses an architecture that is independent of the employed protocol for unicast routing.

In addition to addresses referring to individual hosts, IP also accommodates addresses that refer to group of hosts on one or more networks referred to as multicast addresses. Multicasting protocols need to support sending a packet from a source to the members of a multicast group. Two tasks have to be implemented to support multicast. First, a process should be implemented to track group memberships and to enable hosts to join

and leave multicast groups. A second scheme is needed to route multicast datagrams, from the source to the recipients, within the network.

The Internet Group Management Protocol (IGMP) [26] is used by hosts and routers to exchange multicast membership information over a local LAN. To join a group, a host sends an IGMP report message that can be intercepted by the multicast router on the LAN. To maintain a valid current list of active group addresses, a host claiming membership in a group should respond with a report message to an IGMP query message that is transmitted periodically by the multicast router.

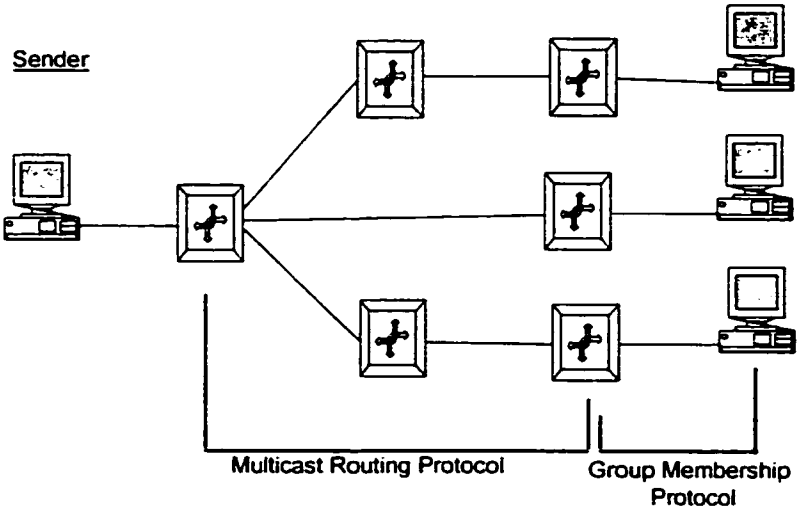

Figure 4.1 Positioning of IGMP and multicast routing protocols

A basic principle in multicast routing is that network routers must be able to use multicast addresses to direct traffic towards the intended receivers in the multicast group.

Routers will interact with each other to exchange information about neighboring routers. Figure 4.1 illustrates the domains of IGMP and multicast routing protocols.

Among the available approaches to support IP multicast routing, the following schemes are popular:

Multicast Open Shortest Path First (MOSPF): MOSPF [23] is an extension to the OSPF and is designed to operate within a single autonomous system. Periodically, each router floods information about local group membership to all other routers in its area. The result is that all routers in the area are able to know the location of all group members for each multicast group. Each router calculates the shortest-path spanning tree from a source to all networks containing members of a multicast group. This calculation is performed only on demand, and thus reducing the associated overhead.

Distance-Vector Multicast Routing Protocol (DVMRP): A distant-vector multicast routing protocol [3] that has been implemented as an extension to RIP. A datagram from a multicast source is initially propagated downstream to all other multicast routers that forward only the required groups locally. With DVMRP each router keeps track of the best routes to the source of multicast datagrams. A multicast datagram received via the best path to its source is forwarded, otherwise it is discarded. DVMRP produces source-specific shortest path trees.

Core Based Trees (CBT): CBT [24] utilizes a single tree to forward datagrams for each group instead of one tree per multicast source. The delivery tree is centered around a core router for each multicast group. This tree is shared by all sources for this specific

multicast group. The delivery tree will extend from the core to the multicast routers. The CBT provides a scalable and flexible approach to support multicast.

Protocol Independent Multicast (PIM): The MOSPF and the DVMRP are extensions of existing unicast routing protocols. In most cases, the multicast protocol is designed to be efficient when there is a relatively high concentration of multicast members. To provide a more general solution to multicast routing, PIM [25] has been developed. PIM is designed to extract needed routing information from any unicast routing protocol and may work across multiple Autonomous systems (AS). PIM defines two modes of operation, the dense-mode and the sparse mode, to accommodate different requirements. For a multicast group, one router is designated as a rendezvous point similar in function to that of the core of the CBT.

4.2 Multicast in the Mobile Environment

A basic design aspect, in supporting the delivery of multicast traffic to mobile nodes, is specifying which element to join the multicast tree associated with the group requested by the MN. The IETF Mobile-IP [1] proposed two approaches to provide the multicast support. The first approach utilizes the current FA serving the MN to join the multicast group, while the HA is considered by the second approach for the same purpose.

Another aspect of the multicast support in Mobile-IP systems is ensuring interworking with the existing multicast infrastructure. In this context a Mobile-IP system supporting

multicast, using either the HA or the FA subscription method, should be capable of utilizing the underlying multicast routing mechanisms mentioned earlier.

In [29], the author provides an overview of different approaches to support multicast in the mobility environment. A multicast support scheme for mobile hosts, based on the CBT approach, is presented in [30]. The MobiCast scheme [31] introduces a new element, the domain foreign agent, to support multicast in the mobile environment while minimizing packet loss due to mobility.

4.2.1 Foreign Agent Subscription

In this approach, the MN will use the membership report to express interest in receiving the datagrams associated with a specific multicast group. Accordingly, the FA will join this group using a multicast protocol and will become a leaf on the delivery tree. This approach requires the FA visited by the MN to be willing to provide a multicast service for the visitors.

The main advantage of the FA subscription approach is optimal routing. The FA can interact with the multicast routing protocol and attach itself to the multicast delivery tree, providing the optimum available routing from the multicast source to the MNs. On the other hand the approach relies on the assumption that the visited FAs are willing to tolerate the overhead associated with providing the multicast service to the MNs. The visited FAs may provide the FA service but may not welcome the support of multicast service on their local networks.

4.2.2 HA Subscription

This option specifies that the multicast group subscription has to be accomplished through the home agent. When the mobile host is visiting a FA, a bi-directional tunnel to the HA is created. A MN will send membership reports to the HA through the established tunnel, then the HA will join the intended multicast group. Multicast datagrams will be tunneled from the HA to the FA, where they are delivered to the MNs. In [27, 28], an approach similar to that of the bi-directional tunnel is used to support mobility. The proposed approach relies on link-level multicast packets generated by the FA at different foreign networks to forward the multicast traffic to the MNs.

The advantage of the HA subscription scheme is that multicasting is completely transparent to the FAs that the MNs may visit. Only the HA needs to join the delivery tree of the multicast group. One of the disadvantages of this approach is the resulting sub-optimal routing since multicast packets have to be forwarded to the FA through the HA. Another drawback is the behavior described by the tunnel convergence problem. A FA, serving multiple MNs interested in receiving the multicast of a particular group, will receive a copy of every multicast packet from each home agent forwarding traffic to its MN. This will lead to multiple copies of the same datagram to be delivered to the FA. The authors in [27] suggested a solution for this problem where the FA selects only one HA to deliver the multicast traffic associated with a particular group. In this work, we will describe a protocol that allows a single HA to forward multicast traffic to multiple FAs, in contrast to serving only one FA as in the traditional HA subscription approach.

Figure 4.2 illustrates the HA and FA positioning difference in both the HA and FA approaches.

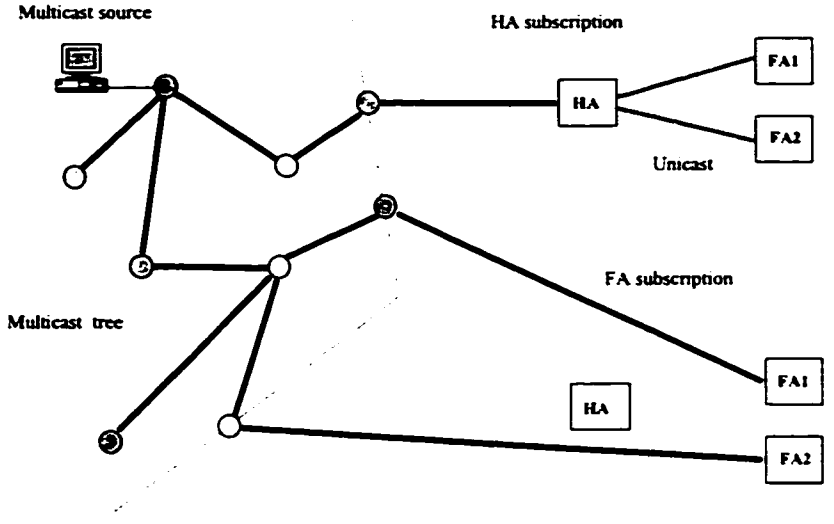

Figure 4.2 FA and HA subscription to support multicast

4.2.3 Multicasting from Mobile Nodes

In the previous two sub-sections we presented the basic ideas behind two possible approaches to support multicast in the Mobile-IP environment. The previous discussion considered the MNs as recipients for the multicast traffic. Although both of the presented approaches can be used to support the delivery of multicast traffic originating from the MNs, we will select the HA subscription approach for this purpose considering the following point. If the MN is expected to be characterized by a high mobility rate then

supporting the MN, as a multicast traffic source using the FA subscription approach, will involve reconfiguring the multicast delivery tree upon each move. This situation will result in a considerable overhead that is proportional to the mobility rate of the multicast source. In the case of a stationary source and that of the FA subscription scheme, moving to a FA that does not provide multicast service will deny service to this single MN. On the other hand, if the MN is the multicast source then this movement will cause the multicast service to be denied for all recipients. The objective of providing a robust and stable system suggests using the HA subscription approach that comes with the cost of sub optimal routing through the HA.

4.2.4 Multicast Support Objectives

An important topic to discuss is the aspects that represent an efficient multicast approach for the Mobile-IP environment. The objective of a multicast mechanism, in either a stationary or a mobile environment, is to deliver the multicast traffic to its recipients with minimum overhead. We can summarize the features required in a multicast support for mobility systems as follows:

- Deliverability of multicast datagrams. An efficient mechanism should minimize the number of datagrams not received by the destined MNs due to mobility events.
- Reduce the delay associated with multicast packet delivery and provide fairness in packets delivery time.

- Transparency to existing support elements. For example a protocol that requires a code to be uploaded to the correspondent nodes is less attractive than the one that requires modification on the HA only.
- Reduce the number of control messages needed to provide the multicast support.
- Minimize the total cost associated with transporting the multicast packets.
- Eliminating or reducing the number of elements representing a single point of failure or a bottleneck.
- Fair packets distribution over different links.
- Reduce the number of elements needed for the multicast support.

4.3 Proposed Solutions

In this subsection we present the details associated with our proposed schemes [33] to support multicast on local registration system. In those approaches, the MN will express his interest in receiving the multicast traffic of a specific group by sending a membership report. The MN will send this membership report using a unicast datagram to the current FA. The final destination of this report differs according to the approach considered. This report may accommodate a request to receive traffic for one or multiple multicast groups. In this subsection we will describe a mechanism to utilize the FA subscription in local registration system, where the gateway FA joins multicast groups on behalf of the MNs. A modified scheme that allows an intermediate FA to join the multicast group will be

presented, providing a more distributed approach. To accommodate deployment situations where transparency to FAs is a requirement, the HA subscription approach is presented.

4.3.1 Approach A: Gateway FA Subscription

In this scheme, the gateway FA joins the multicast groups requested by MNs currently served by FAs on its hierarchy. The gateway FA receives summary reports, from the FAs on lower levels of the hierarchy, regarding the multicast groups to be considered.

Group subscription: A MN, wishing to receive a multicast group datagrams, sends a membership report to the current FA requesting to receive the traffic associated with this specific group. On receiving the report, the current FA will consult its local membership table that includes entries for all groups requested by serviced MNs. If the FA determines that it is not receiving the datagrams for this newly requested group, it will send a summarized report to the FA on the higher level. The summarized membership report contains multiple entries, one for each requested group, covering all groups requested by the MNs residing within the FA's service area. No report is to be generated if the FA determined that it is already receiving the multicast traffic for the requested group. The summarized report will be forwarded upward along the hierarchy towards the gateway FA.

An intermediate FA may receive multiple reports from different FAs on the lower hierarchical level regarding the same multicast group. Also the received report may be

concerning a group that the FA has already subscribed in. In this case, the FA will update, summarize or suppress the forwarded report to the upper level while updating the entries in its local database as needed. On receiving the membership report, each intermediate FA will create an entry indicating the interest of the lower FAs in receiving the multicast traffic associated with particular groups. This process will result in the gateway FA receiving a single report, from each of the FAs in the downstream level, listing the groups that the FAs need to receive traffic from. Accordingly, the gateway FA will join all of those multicast groups using the underlying IP multicast infrastructure.

Multicast data delivery: Upon receiving a multicast datagram, the gateway FA consults its local tables and verifies if there is any of the visitors MNs is a member of this multicast group. If one or more members exist then the multicast datagram will be delivered locally. The table is queried also for FAs on lower levels interested in this group. One copy of the datagram is forwarded to each of those FAs. Those same steps will be repeated on each of the intermediate FAs till the last FA that is interested in this group member receives the multicast datagram.

Handling mobility: When a MN moves from one serving area to another, the MN will send a local registration request to the common ancestor FA associated with both the previous and the new FAs. Intermediate FAs between the common FA and the new current FA need to add the multicast support entry associated with the MN, while FAs between the common FA and the previous FA need to remove the corresponding entries. When the MN moves to the new FA, in addition to sending the local registration request, it will generate a membership report to the current FA to be forwarded with a final

destination pointing to the common FA. A point to consider here is that the common FA does not need to forward the report upward any further; it only needs to update its entry to point to the FA on the lower level on the new lineage instead of that of the previous one.

To remove the multicast entries on FAs between the common FA and the previous FA, more than one possible mechanism can be used. The simplest one is to rely on a future membership report from the lower FA that carries the most updated information. This approach comes with the cost that multicast traffic will continue to be unnecessarily delivered to the old FA for some time. If conserving bandwidth on the hierarchy is a major concern, the common FA may send a query for group membership. Issuing this query is triggered by the common FA receiving the local registration request. The common FA will then extract the identification information for the old FA, and send a membership query for this particular FA. This query will trigger the old FA to send its report upward and eventually the intermediate Foreign Agents will incorporate the membership into the summarized report. The common FA will end with a single updated summary report from the lineage associated with the old FA.

Our proposed mechanism may be best described by an example as follows:

Consider a Hierarchical local registration system as illustrated in figure 4.3 and MN1 moving from the area served by FA23 to that served by FA21. MN1 will listen to the advertisements containing the complete regional hierarchy of FAs supporting the local registration.

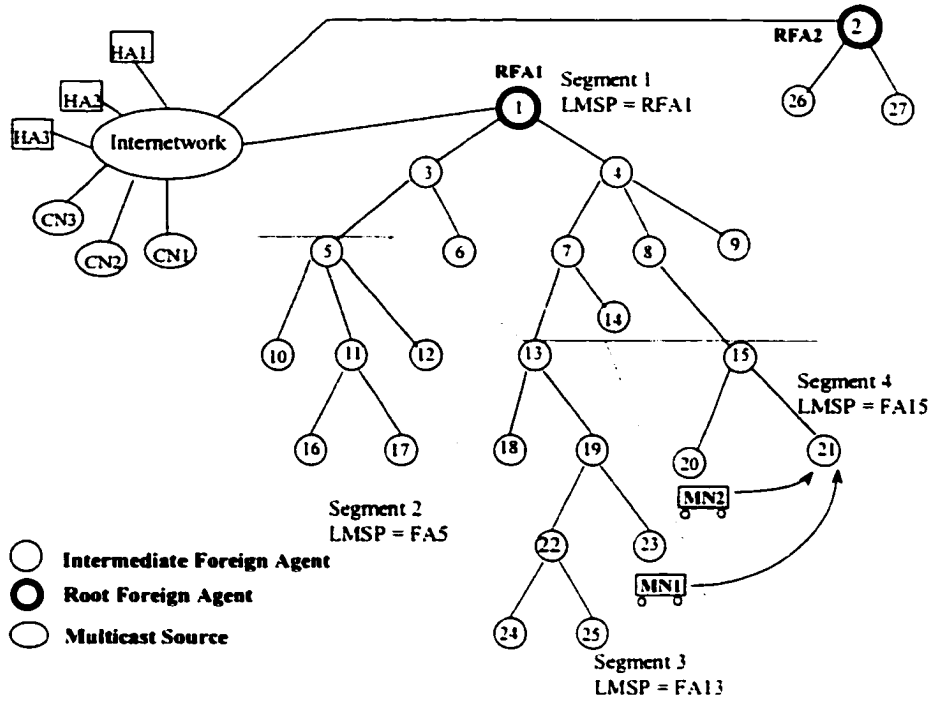

Figure 4.3 Multicast support in local registration Mobile-IP

MN1 examines the hierarchical list of FAs in the new FA's advertisement in search for the nearest common ancestor to the care-of-addresses at the new and previous service areas. MN1 will send the local registration request to the common ancestor FA (FA4). Assume that while MN1 is residing in the area served by FA21, it decided to start receiving the multicast traffic associated with the multicast group G1. MN1 will send a membership report to FA21. This message will be forwarded upward from the current FA (FA21), till it reaches a FA currently receiving the traffic of G1. In our example, no such FA exists, and the report will continue its trip till received by the gateway FA (FA1). The intermediate FAs (FA15, FA8 and FA4) will create an entry for multicast group G1.

Assume that MN2 (which is currently visiting FA20) wishes to be a member of both G1 and G2, then it will send a membership report for the two groups. FA15 will receive reports for G1 and G2 from FA20 and those reports will be compared against the existing local multicast group entries (in our case, this will be G1). Accordingly FA15 will have G1 and G2 in his summary report. By the time the gateway FA receives the report, all the intermediate FAs will have information about the multicast traffic that needs to be forwarded down the hierarchy and which FAs to forward that traffic to.

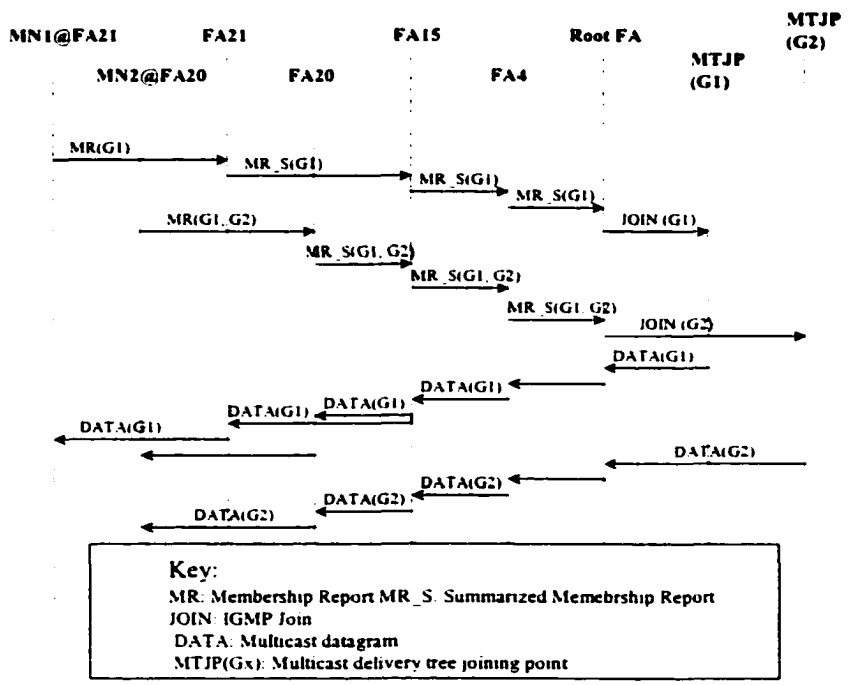

Figure 4.4 Messages associated with the GFA subscription scheme

When MN2 moves from FA20 to FA 21, it will send the local registration request to FA15. In addition, a membership report concerning G1 and G2 will be sent to FA 21.

FA21 will send a summary report upward for G1 and G2. On receiving the report, FA15 will not forward the report any further, and will adjust its local table to point to FA21 as interesting in receiving traffic for G1 and G2. Since the entries for G1 and G2 on FA20 and the pointer to FA20 on FA15 are not needed any more, FA15 will extract the previous FA identification (FA21) from the local registration request and issue a membership query towards FA20. FA20 will broadcast the query on its network. Eventually, FA20 will report no requested groups and FA15 will remove the entry pointing to FA20.

4.3.2 Approach B: Intermediate FA Subscription

It is clear that the gateway FA subscription mode has the advantage of supporting multicast on the local registration system, while requiring only the support of the gateway FA in joining the multicast delivery trees. On the other hand, this approach represents a single point of failure represented by the gateway FA. In addition, since multicast traffic destined for different FAs has to flow on the hierarchy passing through the gateway and the intermediate FAs, the approach does not provide an optimum route. This is in contrast to the optimal route resulting when the current FA is to join the multicast tree directly. Considering those observations, we will suggest a modified approach that will provide a more efficient routing leading to less average packet delay and avoid having the gateway FA as a single point of failure. Those benefits come with a reasonable cost represented by larger number of FAs required to join the multicast trees. The modified approach will

intelligently select those FAs to minimize the associated group join/leave overhead. The basic idea of the approach is to distribute the burden of joining the multicast groups among different strategically selected FAs, those FAs will be referred to as Local Multicast Service Providers (LMSP). In the following discussion, we will propose two possible approaches to select the LMSPs.

LMSP selection based on the regional domain topology:

In this simple approach, each regional domain is logically divided into multiple segments (or service domains), where each segment includes a number of neighboring FAs. A FA is assigned to each segment to be the local LMSP and is responsible for joining the multicast groups requested by the FAs associated with this particular segment. This mechanism may be used to avoid adding excessive delay to datagrams received by FAs on lower levels on the hierarchy, by assigning segments and FAs in a fairly distributed manner over the hierarchy.

LMSP selection based on MNs density:

This scheme aims at reducing the overhead associated with the join/leave process, by adapting to the number of MNs currently served by the FAs. To accomplish this task, the FAs on the bottom of the hierarchy send their group reports upward including the count of MNs interested in particular multicast group. Each FA will keep track of the number of MNs, and as soon as the number associated with a multicast group exceeds a particular threshold then this FA is triggered to act as an LMSP. The responsibility of the LMSP is to provide the multicast service for the FAs in its service domain, where the service domain is defined as those FAs on lower levels not serviced by other LMSPs.

Accordingly, the LMSP will have to join the different multicast groups requested by FAs residing in its domain. The Intermediate FA subscription differs from the gateway FA approach in the following areas:

Group subscription: A MN will send its membership report as usual to the current FA. In the case of LMSP selection based on MNs density, the report forwarded by FA will contain an additional entry that records the total number of MNs wishing to receive the traffic of the different multicast groups. The report will be forwarded upward till received by a FA that identifies itself as an LMSP for the requested group. The LMSP will intercept the report, suppress the entry associated with the group or groups it assumes service for and forward the modified report upward.

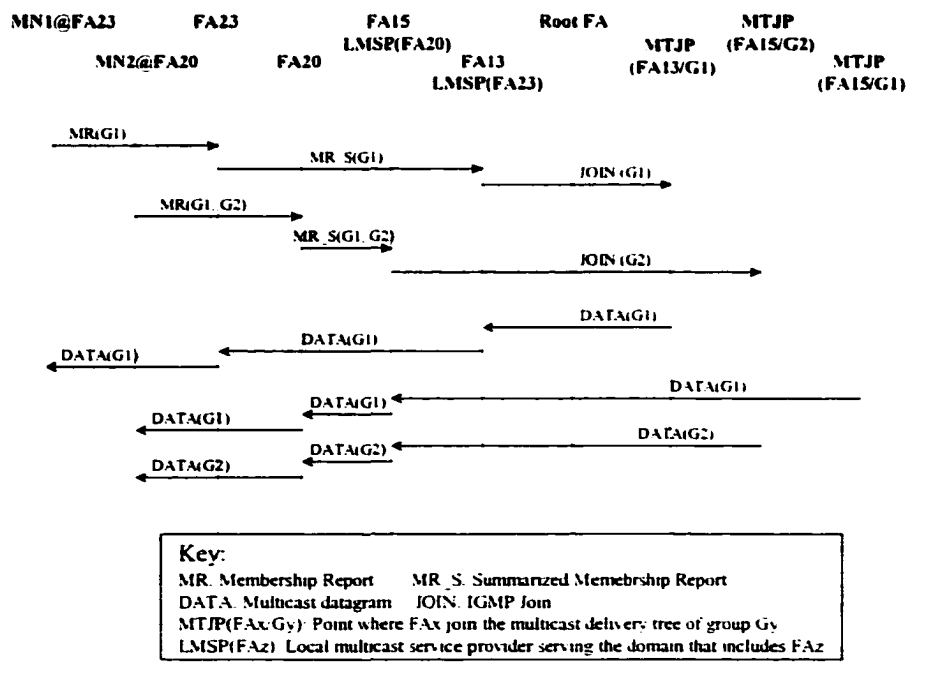

Figure 4.5 Control and data messages for the LMSP approach

Multicast data delivery: After joining the multicast group, the LMSP will start receiving the multicast datagrams. The LMSP will consult its table and forward the datagrams downward towards the interested FAs. Eventually the datagrams will not be forwarded anymore when received by the last interested FA in the service domain.

Handling mobility: A movement of MN can be classified as intra_service-domain or inter_service-domain. In contrast to the gateway FA approach, where the common ancestor FA between previous and new FAs is guaranteed to have an entry for the multicast group requested by the MN, interdomain movement may necessitate the new LMSP to join the requested multicast group. A mechanism to remove the old entries and pointers, similar to that mentioned in the gateway FA subscription approach, will work in this case also.

Considering figure 4.3, and assuming that applying the segmented hierarchy as the LMSP selection policy resulted four segments. When MN1 is located with FA23, its membership report requesting G1 will be forwarded to FA13 triggering it to join G1. When MN1 moves to FA23, located in a domain served by FA15 as the LMSP, a new report will be forwarded to that LMSP. FA4 will trigger FA23 to issue a membership query to update entries associated with the old binding. The corresponding control messages flow is illustrated in figure 4.5.

4.3.3 HA Subscription in Local Registration

Agents providing the FA service in areas visited by MNs are not necessarily willing to provide, or capable of supporting, the multicast service if requested. This situation suggests the need for a multicast support scheme that expects limited support from the FA and relies on the HA in the costly process of joining the requested multicast groups. Since transparency to FAs is one of the key advantages of the HA subscription, care has been considered to avoid unnecessary processing of packets by intermediate FAs.

Approach C: HA subscription

In this mechanism, the HA will join the delivery tree associated with the multicast group requested by its MN. A basic characteristic of this approach is that FA does not need to provide support for the multicast service, which can be useful when the MN is moving to an area where FAs are not capable of acting as LMSPs.

Group subscription: The MN will send a unicast membership report to the gateway FA through the current FA. The current FA will keep information regarding the current visitors, their HAs and the subscription in multicast groups if any. The current FA will process the information, associated with MNs interested in receiving multicast traffic of a particular group, and send a report to the gateway FA. The report generated by the FA contains information regarding the multicast group and the HA supporting the MN on this FA. It is expected that the gateway FA will receive multiple reports from different FAs requesting to receive the multicast traffic associated with the same multicast group through different HAs. The GFA consider those HAs as candidate HA_MSP (HA Multicast Service Provider) for this group, and should select only one HA to forward the multicast traffic.

If the gateway FA receives only one report or multiple reports for the same HA claiming membership to the same group, then this HA will be selected as the HA_MSP. If multiple HAs candidates exist, then a criteria will be applied to select the HA_MSP. Since the HA_MSP is servicing now multiple FAs, in contrast of servicing one FA as in the case of the MoM protocol, care should be considered when deciding the criteria used for HA_MSP selection. Another important process happening at that time is that the gateway FA processes all the received reports and accordingly constructs a database with all FAs interested in receiving multicast traffic and the associated multicast groups. After the HA_MSP selection for the particular group, negative caching [32] is to be used to suppress the transmission from other HAs not selected as the HA_MSP.

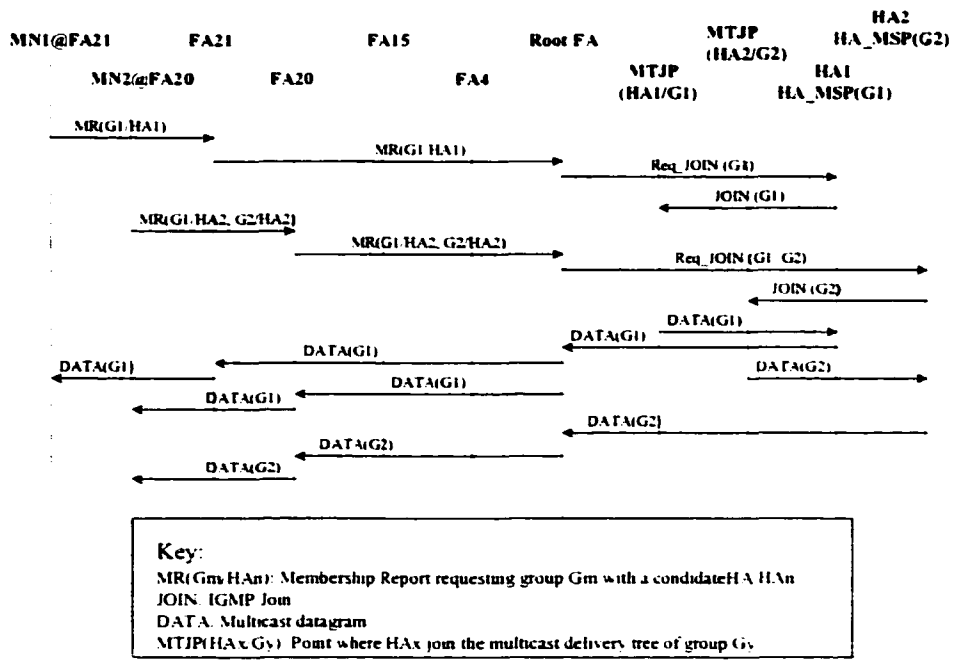

Figure 4.6 Control and Data messages for approach "C"

Multicast data delivery: The datagrams forwarded from the HA will be encapsulated and destined to the MN, then re-encapsulated again with the destination of the gateway FA. The HA will send one copy of the multicast datagram to the gateway FA even if there are more than one MN currently have location binding for this gateway FA. On receiving the forwarded multicast packet, the gateway FA will decapsulate the datagram and forward it to the FA in the lower hierarchical level as indicated in its local table.

Handling mobility: In this approach, multicast datagrams forwarded from the HA are forwarded over the same route as that of unicast traffic. The standard local registration scheme requires the MN to send a local registration request when moving to a new area within the same hierarchy. Accordingly, no extra messages are required to continue in delivering the multicast traffic to the MN in its new area. The mobility of the MN affects the system in the case when the MN is moving to an area served by a new gateway FA. If it happened that this MN is the only host that the HA_MSP has binding for, then in this case the gateway FA has to select a new HA_MSP. Several HA_MSP selection policies can be considered, where an efficient policy is the one that leads to a less number of HA_MSP switching events and provides routing that is near optimum as possible.

Considering our example in figure 4.3, and that MN3, MN4 and MN5 belong to HA1, HA2 and HA3 respectively. MN3 is interested in joining groups G1 and G2, MN4 in G1 and MN5 in G1 and G3. MN3, MN4 and MN5 are located with FA20, FA21 and FA16 respectively.

The MN will listen to the advertisements containing the complete regional hierarchy of FAs supporting the local registration. The MN examines the hierarchical list of FAs in

the new FA's advertisement. The MN will send the local registration request to the common ancestor FA. Assume that while MN is residing in this area, it decided to start receiving the multicast traffic associated with the multicast group G1. The MN will send a membership report to the HA using a unicast datagram. Accordingly, the HA will join the delivery tree associated with group G1. The current FA will keep information regarding the current visitors, their HAs and the subscription in multicast group if any. The current FA will process the information associated with MNs interested in receiving multicast traffic of a particular group and accordingly select a HA to be recommended to the gateway FA as a candidate HA_MSP.

Considering figure 4.3, FA20 will send a report asking to receive the multicast traffic of G1 and G2 through HA1, while FA21 will ask for the traffic of G1 through HA2. Similarly, FA16 will ask for the traffic associated with G3 through HA3. The gateway FA will construct the HA_MSP pools for multicast groups G1, G2 and G3 as follows:

$$HA\_MSP\_Pool(G1) = \{HA1, HA2, HA3\}$$

$$HA\_MSP\_Pool(G2) = \{HA1\}$$

$$HA\_MSP\_Pool(G3) = \{HA3\}$$

Clearly, the gateway FA (FA1) has one option when selecting the HA_MSP for G2 and G3. When it comes to G1, the FA1 has to select one of the candidates. Considering nearest to gateway FA policy, HA1 will be selected as the HA_MSP for G1.

After the HA_MSP selection for the particular group, negative caching is to be used to suppress the transmission from other HAs not selected as the HA_MSP (HA2 and HA3). HA1will now start forwarding G1multicast traffic to the gateway FA.

4.3.4 Membership Reports

In all of the proposed subscription schemes, the membership report is embedded in the local registration request, where the report is treated as an extension. Accordingly, multiple groups can be requested in the same registration message. Also, the MN can specify different subscription schemes to different groups.

A possible message structure is illustrated in figure 4.7, where the registration request message is extended with the report for two multicast groups. The base registration request message is appended by the appropriate extension for local registration. The multicast extension is included in the registration request several times, according to the number of groups requested. To indicate the required multicast subscription scheme, we consider the usage of two bits of the reserved bits as follows:

00 GFA subscription

01 LMSP subscription

10 HA subscription

In the GFA and the HA subscription schemes, if there is need to join a new group then the membership report will propagate to the root FA. On the intermediate FAs, entries will be set indicating the presence of a receiver on a particular interface.

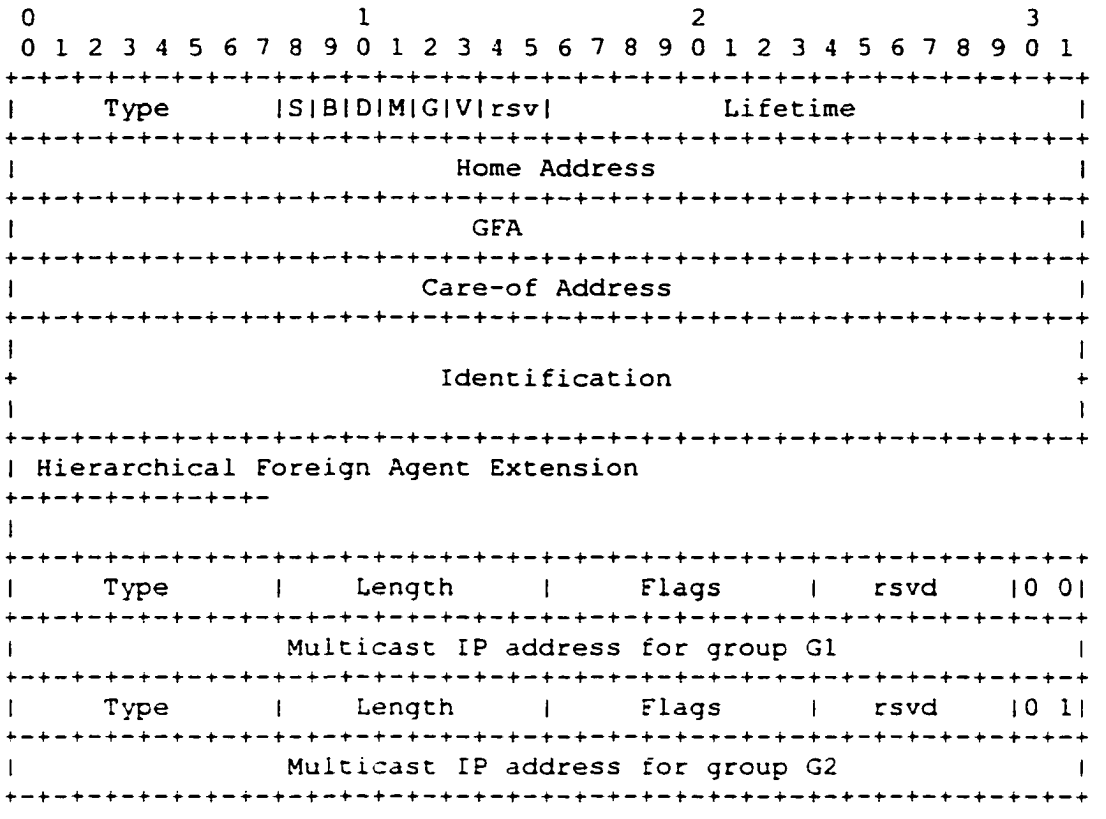

Figure 4.7 Format for registration request with the multicast report

4.4 Performance

In this section, we address performance issues associated with the multicast support. We will start with a comparison between the proposed approaches considering different multicast aspects. Next, we will use simulation to illustrate how the performance will vary under different conditions with the different schemes.

4.4.1 Performance Aspects of the Proposed Schemes

A. Route optimality and delay

The route considered by multicast datagrams, from the source to final destinations, directly affects the delay experienced by those datagrams. The delay associated with datagrams delivery is composed of two components. The first delay component is the one between the multicast source and the element joining the multicast delivery tree, where this element may be the root FA, the LMSP FA or the HA according to the considered approach. The second component, the local delivery delay, is associated with the delivery delay from the element joining the multicast tree to the recipients MNs.

The LMSP FA approach provides the most efficient routing, capitalizing on an optimal path from the multicast source to the LMSP. The root FA subscription comes second where datagrams travel the optimized path only between the multicast source and the root FA. When considering the HA subscription, the triangle routing will result in a path that is less than optimum. Considering the local delivery delay, the intermediate FA scheme provides best results, while the root FA and the HA approaches will provide comparable results to each other.

B. Duplicate packets

When using the HA approach, a scheme is needed to tackle the tunnel convergence problem. This is implemented using the HA_MSP selection policy. The root FA subscription approach does not rely on HAs to forward multicast datagrams. Instead, the

root FA collects information regarding the needed groups and joins the corresponding delivery trees, eliminating duplicate packets from the multicast tree. On the other hand, in the intermediate FA approach, datagrams are not duplicated within a region served by a LMSP when using the topology based LMSP selection scheme. In the density based LMSP selection approach, a MN may be serviced by multiple LMSPs when subscribing to multiple groups, reducing the number of duplicate packets.

C. Total datagrams delivery cost

The previous two multicast performance aspects can be considered together to provide an important descriptor for the system efficiency. The delivery cost can be evaluated as the aggregation of the individual delivery cost for each multicast datagram received by MNs. A smaller value for this descriptor indicates an overall shorter path considered by datagrams and less copies of the datagrams needed to complete the multicast service. Formulas used to evaluate the total delivery cost (T_Cost) for each of our proposed approaches can be expressed as follows:

- **HA subscription:**

$$T\_Cost = \sum_{HA\_MSP_i, i=1}^{n} \sum_{Gk(HA\_MSP_i), k=1}^{m} \sum_{PxGk, x=1}^{y} \left(D_{MTJP(Gx) \leftrightarrow HA} + D_{HA \leftrightarrow RFA} + \sum_{FA_l(Gx), l=1}^{f} D_{RFA \leftrightarrow CFA} \right) \quad (4.1)$$

Where:

HA_MSPi: is the ith HA_MSP providing service to the system.

N: is the number of HAs selected as HA_MSPs and is upper bounded by the number of requested groups by the root FA.

$G_k(HA\_MSP_i)$: is the kth multicast group that HA_MSP_i is responsible for forwarding its traffic.

m: is the number of multicast group assigned to HA_MSP_i and is upper bounded by the total number of multicast groups requested by MNs that belong to this HA.

P_xG_k: is the x*th* datagram that belongs to the multicast traffic of G_k.

y: total number of packets to be delivered associated with G_k.

f: is the number of FAs to receive the multicast of a particular group when using HA subscription.

$D_{MTJP(Gx) \leftrightarrow HA}$:
Delay between the HA and the point where the HA will be attached to the delivery tree of G_x

$D_{HA \leftrightarrow RFA}$:
Delay between the HA and the Root FA.

$D_{RFA \leftrightarrow CFA}$:
Delay between the Root FA and the current FA. Since both of those two elements are on the hierarchy, the value of this delay component is expected to be relatively less than that between the HA and the root FA.

- **Root FA subscription**:

$$T\_Cost = \sum_{G_k(RFA),\, k=1}^{m} \; \sum_{P_xG_k,\, x=1}^{y} (D_{MTJP(Gx) \leftrightarrow RFA} + \sum_{C_qP_xG_k,\, q=1}^{r} D_{FA\_E1 \leftrightarrow FA\_E2}) \tag{4.2}$$

$D_{MTJP(G_x) \leftrightarrow RFA}$:

Delay between the Root FA and the point where the Root FA will be attached to the delivery tree of G_x.

$D_{FA\_E1 \leftrightarrow FA\_E2}$:

Corresponds to the delay experienced by a single copy of the datagram between points E1 and E2. No additional copies of the same datagram are needed between E1 and E2.

Cq:

Number of copies of the same datagram needed to complete the delivery. Each copy will travel from E1 to E2 before the need to duplicate another copy.

r:

Number of copies needed, from a particular datagram at the different branching points.

- **Intermediate FA:**

$$T\_Cost = \sum_{regions} \sum_{LMSP_i,\ i=1}^{n} \sum_{G_k(LMSP_i),\ k=1}^{m} \sum_{P_iG_k,\ x=1}^{y} (D_{MTJP(G_x) \leftrightarrow LMSP_i} + \sum_{C_qP_iG_k,\ q=1}^{r} D_{FA\_E1 \leftrightarrow FA\_E2}) \quad (4.3)$$

$D_{MTJP(G_x) \leftrightarrow LMSP}$:

Delay between the LMSP and the point where the LMSP will be attached to the delivery tree of G_x

Regions:

Number of regions (service domains) in the hierarchical topology. In the case of the topology based, this is a predetermined value and n is equal to 1 for a single LMSP per region. For the density based, the number of regions is one and n is the total number of LMSPs.

D. Deliverability

Multicast packets are not delivered to recipients when a mobile node moves to a new service area, and elements supporting the multicast delivery have not yet adjusted the local tables to correctly deliver the datagrams to the new location. The number of

undelivered datagrams, due to the mobility, depends on the amount of time needed to complete the process of updating the tables. Another source of undelivered packets is specific to the HA subscription approach. This packet loss is observed since that the primary HA_MSP may stop forwarding datagrams when none of its MNs is currently residing in the hierarchy served by the root FA. As described earlier, the HA subscription scheme will recover from this situation by instructing the secondary HA_MSP to start forwarding multicast. In general, the root FA subscriber approach provides the least number of packets dropped due to the fact that the root FA will remain most of the time attached to the different multicast delivery trees, minimizing service disruption due to group joining delay. The intermediate FAs approach comes second, since it is possible that the MN will move to a new FA served by an LMSP that is not attached to the requested group delivery tree. Accordingly, additional delay may occur during recovery while the LMSP is joining the requested group. HA subscription approach provides larger number of undelivered datagrams due to two factors. The first is that the delay from the root FA to the secondary HA_MSP is expected to be considerably high relative to delay values between nodes on the hierarchy. This will cause a larger number of undelivered datagrams during HA_MSP switching. The second factor is that if the MN is moving to a new FA that is not receiving the multicast traffic of the requested group, the new membership report has to travel all the way up to the root FA.

E. Transparency and fault tolerance

The HA approach centralizes the multicast support functionality in the HA, while limited support is required also from the root FA to support the HA_MSP selection. No support

for multicast functions is expected from the intermediate FA. On the other hand, the root FA subscription approach relies on the root FA in supporting the service. Also support from intermediate FA is needed to process the membership reports. In this approach, the root FA represents a single point of failure. A point to consider here is that the root FA provides the multicast service and local registration support for all the FAs in lower hierarchical levels. The failure of the root FA will cause denial of unicast service to MNs in lower levels, in addition to disrupting the multicast service. The intermediate FA approach requires more intermediate FAs (LMSP) to support the multicast service. On the other hand, it is the most robust approach among the proposed ones such that the failure of the root FA will affect only FAs in the domain for which the root acts as an LMSP. A detailed study of tolerating the failure of FAs, in local registration Mobile-IP systems, will be presented in a later chapter.

4.4.2 Simulation Results

To illustrate the behavior of the proposed schemes, we considered discrete events simulation to evaluate aspects of the multicast support performance. The objective of this simulation is to highlight some the characteristics of the different schemes in an environment with different conditions, where one approach may outperform the other.

Hierarchical network model: The Network environment used for simulation has a hierarchy that has one root FA, five HAs and 6 levels of intermediate FAs. The simulation environment includes 350 MNs moving among 30 FAs. Each MN can request to receive multicast traffic from up to 5 groups from 20 available groups. The delay over

links between two FAs in the hierarchy is set to one msec. The delay between any FA and a routing element is set to 2 msec, and the delay over wireless link is set to 2.5 msec. The delay between a HA and the root FA ranges from 10 to 12 msec. The delay figures used here represent an aggregation for all delay components, such as propagation and processing. Different values for the delay should not affect the correctness of the approach. The only restriction is that the delay from the MN to the HA should be larger than that from the MN to the root FA, which is the expected topology and the motivation behind the local registration approach. The multicast traffic is generated from four CNs. Simulation time of 1500 minutes was considered, with uniformly random distributed mobility rate between 1 and 5 moves/minute. The case of the Base M-IP is considered in this simulation. The objective of demonstrating the simulation results mentioned here is to verify the concept correctness, and to illustrate the characteristics and the behavior of the different schemes under different conditions.

A. Deliverability

Figure 4.8 shows the percentage of packets not received with the number of MNs. As a general observation, the larger the number of MNs the less the number of undelivered datagrams. This observation comes from the fact that with the larger number of MNs, it is more probable that a MN moving to a new service area will find the new FA or a nearby FA already subscribed to the requested group. With increasing number of MNs, the density based approach will optimize the location of the LMSPs such that it will

outperform the topology based approach. This is the reason behind the density based approach outperforming the topology based selection scheme.

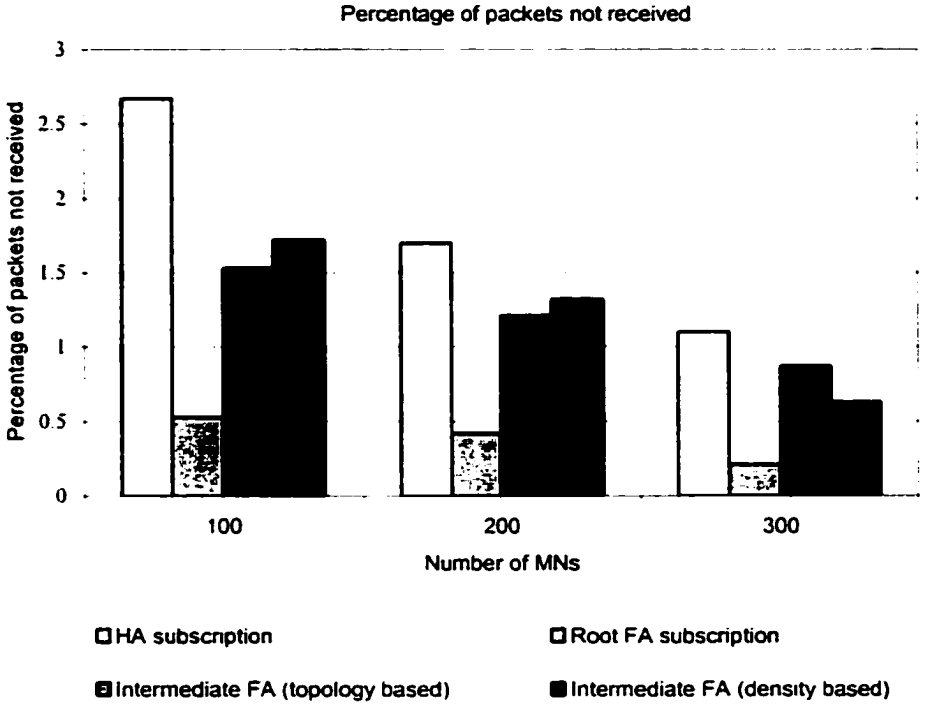

Figure 4.8 Multicast datagrams dropped with number of MNs

B. Delivery cost

To illustrate the effect of the density distribution of MNs on the delivery cost, these performance aspects are evaluated over different MNs distributions as shown in figure 4.9. A higher percentage of MNs in clusters is an indication that more MNs will be able to find a nearby FA that is already a subscriber in the multicast group requested by MNs. A value of 20% of MNs in clusters indicates that 20 % of the total number of MNs

considered in the simulation are forming clusters (with immediate neighbors), where each cluster includes at least 10% of the total number of MNs. The total delivery cost is evaluated as mentioned earlier. The normalized cost refers to the value of the cost relative to a reference value. This reference value is considered as the delivery cost of the topology based approach at the lower clustering percentage.

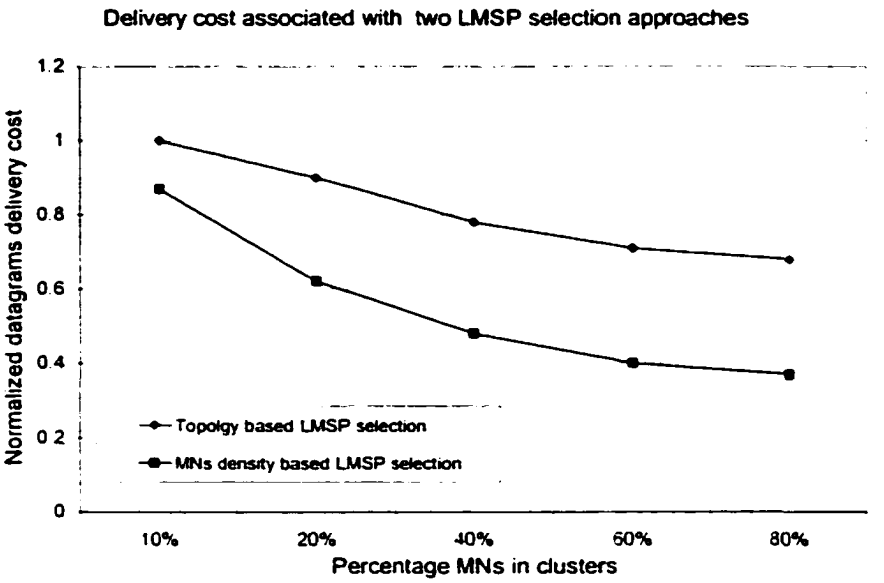

Figure 4.9 Delivery cost with MNs distribution

For this simulation, the topology scheme partitions the hierarchy into seven segments (LMSP service domains), while the density based assigns an LMSP for a group when the number of interested MNs exceeds a threshold of 50. Recalling that the total delivery cost takes both delay and number of datagrams copies into consideration and that the density based provides relatively higher delay for lower clustering percentage, the cost of the density based is still better than that of the topology based. The reason is that with lower

density, less number of FAs will be selected as LMSPs and accordingly less number of replicated datagrams will be generated. With higher number of MNs in clusters, the density based scheme will optimize the number and the location of the LMSPs providing lower delivery cost. Careful selection of the number of MNs threshold, used in the density based scheme, is critical. Statistical information regarding the expected number of MNs, their distribution and the number of multicast groups requested can be of a great help in selecting this threshold. The same comment is true regarding the topology based, when selecting the number and the location of the LMSPs.

C. Multicast delivery delay

Figure 4.10 shows how the number of MNs affects the delay experienced by multicast datagrams, considering the two proposed approaches for LMSP selection. We will consider the delay associated with optimal routing, where the current FA joins the multicast group, as our reference. A MN receiving 25% or more of the datagrams delayed with more than 15% extra of the delay experienced with optimal routing will be counted as unsatisfied MN. For the density based approach and a smaller number of MNs, limited number of LMSPs will exist leading to larger delay for traffic on the hierarchy. More MNs will lead to larger number of LMSPs and will provide less delay. Considering the topology based approach, the number of LMSPs is fixed and is not affected by number of MNs, but delay will slightly decrease due to higher probability that the new LMSP will be already a member in the requested group.

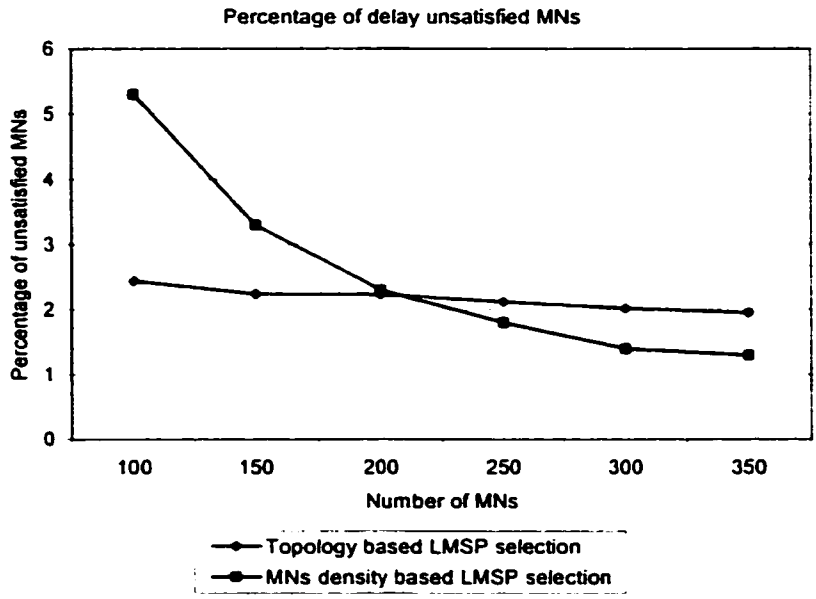

Figure 4.10 Delay and number of MNs

An observation regarding the density based approach is that it provides relatively larger delay compared to the topology based scheme for smaller number of MNs. This behavior of the density based approach can be enhanced by imposing an upper limit on the number of hops that a membership can travel upstream without finding an LMSP. If the hop count threshold is exceeded without satisfying the density of MNs required to assign an LMSP, the first available FA is selected as the LMSP bypassing the density constraint.

It is expected that applications using multicast will emphasize one performance aspect, while the elements in the architecture supporting the service may relax the constraint on this aspect and enforce it on another. This situation suggests that the proposed different schemes can be considered individually, or can co-operate to provide a better performance, taking advantage of the distinguished characteristics of each scheme.

The intermediate HA subscription scheme provides a promising approach because of its robustness, its better routing and the absence of a single point of failure. In the case where support from FAs is not expected, the HA subscription provides an efficient solution. In an environment where the root FA and the associated links have the needed capacity and where failure rates are not relatively high, the root FA subscription is an available option. For smaller number of MNs, the delay associated with datagrams delivery can be relatively large. This is the situation in figure 4.10 for number of MNs less than 150 MNs.

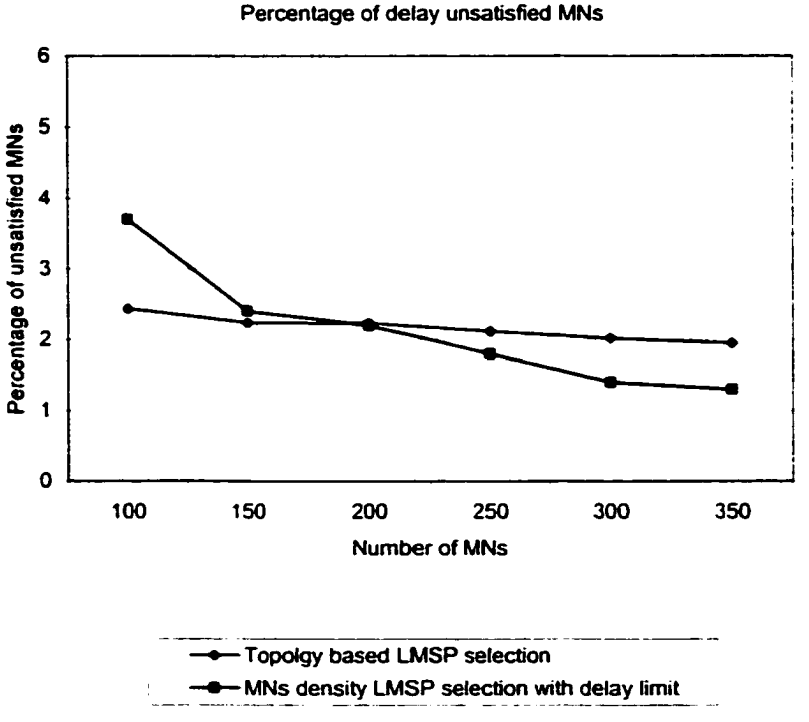

Figure 4.11 Delay and number of MNs (using report propagation limit)

A membership report can propagate upstream in the density based mode, without being processed by an LMSP. Limiting the number of hops traveled by the report can be used to place an upper limit on the delay associated with local delivery. The effect of applying the delay bound is apparent in the delay values in figure 4.11, where the report has to be processed within two hops away from the source. This points to the flexibility of the density based approach regarding the LMSP selection. The value for this bound threshold can be configured manually or can be downloaded to the different FAs as needed. A third option, is that the maximum number of hops can be treated as a quality of service descriptor, such that reports for MNs with higher class of service may propagate for less number of hops.

5. Reliable Multicast Support in the Regional Registration Mobile-IP

5.1 Reliable Multicast

A basic design aspect, in supporting the delivery of multicast traffic to mobile nodes, is identifying the element that is to join the multicast tree associated with the requested group. In [1], two approaches were proposed to provide multicast support. The FA subscription approach utilizes the current FA (CFA) serving the MN to join the multicast group, while the HA joins the requested groups on behalf of its MNs in the HA subscription scheme.

A reliable multicast protocol has the objective of ensuring message delivery to all members in the multicast group. In a mobile environment, the reliable delivery needs to extend its service to cover the case when MNs continuously change their locations during the multicast session. To support reliable multicast, two main approaches were proposed and implemented in different studies. The first approach considers sending redundant data such that the frequency of unrecovered loss events is reduced. This approach is useful for delay sensitive application where the delay associated with retransmission requests and responses cannot be tolerated, and when bandwidth is available to carry the redundant data. The second approach relies on loss detection and retransmission mechanisms. Different schemes were proposed and implemented representing the second approach and can be classified as follows:

ACK-based Protocols: In this approach, the source of the multicast datagrams is responsible for collecting acknowledgments from all receivers. The datagrams have to be kept in cache till the corresponding ACKs are received from all members.

Receiver-Driven Protocols: Instead of sending ACKs, in this approach the receivers send NACKs (Negative Acknowledgement) to the source when retransmission is needed. To avoid the problem of NACK-implosion when multiple receivers send NACKs at the same time, a NACK-implosion avoidance scheme should be used.

Tree-Based Protocols: In this approach, and in contrast to the previous two schemes where the receivers are not organized into any specific structure, receivers are grouped into local regions. Receivers in each region send their ACKs upward towards the gateway of the tree. ACKs can be aggregated along the tree, thus avoiding the ACK implosion problem. The merits of the ACK tree approach, and a comparison of different available approaches of reliable multicast protocols can be found in [34, 35] respectively. The TRAM protocol [36] uses a tree based scalable approach to support bulk data transfer. In [37], the reliability in multicast environment and issues associated with large-scale networks are discussed. In [38], criteria for reviewing reliable multicast protocols are provided.

Several protocols to support reliable multicast in a mobile environment have been proposed. In [39], a special node called supervisor host is introduced to relief the mobile support station from the overhead of supporting reliable multicast. A single supervisor host is positioned to provide its services, such as routing to mobile nodes and flow control, to multiple mobile support stations. In [40], a flooding mechanism is used to

deliver datagrams only to the multi-hop radio network parts that are affected by topology changes. A protocol that utilizes a logical ring and a token passing scheme, to pass the sequence number information needed for cache management, is proposed in [41].

Different from those mentioned approaches, our proposed protocol does not require additional elements for the reliable multicast support. Taking advantage of the inherent local registration features, the logical retransmission tree is constructed to overlay on the local multicast tree. Further, multicast reliable delivery and FAs fault tolerance services are integrated in a single platform, minimizing the overhead and maximizing the efficiency of existing resources. The measurement based observation in [42] shows that the majority of multicast datagram losses occur away from the core network. Such observation suggests that reliable multicast support is needed, in particular, at the network edges. Accordingly, the implementation of our proposed protocol, close to receivers in the access network, provides robustness and efficiency to the end-to-end reliable multicast service.

5.2 Design Aspects of the Proposed Scheme

The reliable multicast delivery scheme proposed in this paper is designed considering the following points:

- The objective is to provide a reliable multicast delivery service with minimal additional overhead.

- The service shall be integrated with the multicast and the failure recovery services, leading to more efficiency in using existing signaling messages and functions.
- In the case of retransmission requests, the platform will retrieve the datagrams from the nearest FA that still has a copy of the requested datagrams.
- It is highly probable that the external protocol, running outside of the local registration domain, will apply a rate-based flow control policy. The presence of low rate receivers may result in reducing the transmission rate of the multicast source, or in denying the reliable service to those slow receivers. The LMSP (Local Multicast Service Provider) implements its own flow control scheme that can accommodate, to the extent possible, the retransmissions and the lower rates associated with mobile hosts. An objective of this scheme is to reduce the negative effect on the multicast source transmission rate.
- It is possible that a MN may have the reliable multicast service, as part of its QoS profile, with specific minimum rate guarantee. The proposed platform should provide the capability of supporting higher forwarding rate, to preferred nodes, even in the presence of slower receivers.

5.3 Operation of the Proposed Protocol

Our proposed protocol for reliable multicast support [43] is based on the ACK tree approach, and considers a multicast support scheme based on the FA subscription similar

to the one proposed in section 4.3.2. To provide the reliable multicast delivery service, two basic functions are needed to be implemented; multicast datagrams forwarding and the flow control. The proposed platform implements internal schemes to support both MNs mobility and tolerance for FAs failures as illustrated in the following subsections.

5.3.1 Delivery of Multicast Messages

Based on the received membership reports, as described in section 4.3, the LMSP will join the requested multicast groups. Also, the LMSP will engage in a reliable multicast session with the external reliable protocol. The multicast source will give each transmitted datagram a unique sequence number. The LMSP will use this external sequence number to detect any missing datagrams. In such case, the LMSP will utilize the underlying mechanism to indicate a request for retransmission. This can be implemented through a sequence number and a bitmap mechanism similar to that mentioned in [44]. Each FA will keep track of the identification of FAs requesting the multicast traffic of a particular group. We will refer to those entries as the immediate receivers set. This set includes also those MNs which are currently residing in this FA service area and are interested in the same multicast group.

A received multicast datagram will be associated with two processes on the LMSP. First, the LMSP will inspect the sequence number in the received datagram, and generate an ACK or retransmission request if needed towards the underlying multicast protocol. The LMSP and the intermediate FAs will co-operate in locally retrieving lost datagrams,

such that a minimum number of retransmission requests are propagated to the external multicast source, avoiding slowing the transmission rate. The LMSP will keep a copy of this datagram in cache till receiving an ACK from all members in the immediate receivers set. The cache management process will remove packets with no corresponding ACK, using a FIFO approach, if the cache space is needed to accommodate new packets.

The second activity is associated with the datagram forwarding process, using the local delivery protocol. The LMSP will generate a local sequence number for the datagram, while keeping a mapping between the local and the external sequence numbers. Although the local sequence number can match that of the external one, in some situations the LMSP may elect to manipulate the forwarded data to optimize performance using segmentation. Using the local sequence number is not necessary for the operation of the protocol, and is introduced to provide flexibility to the LMSP to implement its own policy. For the remaining of this work, we will assume that the local sequence number is equal to the external sequence number.

At this point, the LMSP is responsible of reliably delivering the datagram to the local MNs currently residing in its service area. The LMSP will query its local table for entries pointing to elements receiving the multicast traffic for this group. Also, the LMSP will update the data structure to keep track of the sequence number of datagrams delivered to and acknowledged by each of those recipients. Upon receiving the datagram on its trip downstream on the hierarchy, each of the intermediate FA with interest in this group will record the received sequence number, and implement the flow control scheme. Those steps will be repeated at each FA traversed by the datagram.

5.3.2 Rate Adaptation

In an environment where elements may communicate using a combination of wired and wireless links, rate adaptation is a highly needed feature. In this work, we will use a scheme that adapts its rate of transmission according to the receiving capacity indicated by recipients. The FA will start forwarding traffic with a preconfigured initial forwarding rate. If the recipient FA is not detecting any datagram loss, it will acknowledge all the datagrams. The FA will increase the forwarding rate, using a rate increase function, as long as the ACK messages received from the FAs indicate no congestion. When the rate of retransmission requests received from FAs downstream exceeds a threshold, the FA receiving those requests will send an upstream congestion indication. On receiving the indication, the FA will reduce its forwarding rate.

In our proposed scheme, the FA will adapt the forwarding rate on each interface to the capacity associated with each FA member in the immediate receivers set. For each interface, the FA will keep track of the optimum rate for each supported multicast group. The objective of this approach is to provide those FAs, capable of receiving higher rates, with the opportunity to enjoy their available resources. Since the FA will be serving local MNs in addition to the FAs, the FA will adapt to a single rate associated with the MNs. The benefits of such scheme can be observed as long as there is available free cache on the FA. If one of the receivers is not able to handle the minimum transmission rate from the FA, it will be denied the reliable service unless different treatment is specified in the MN's profile.

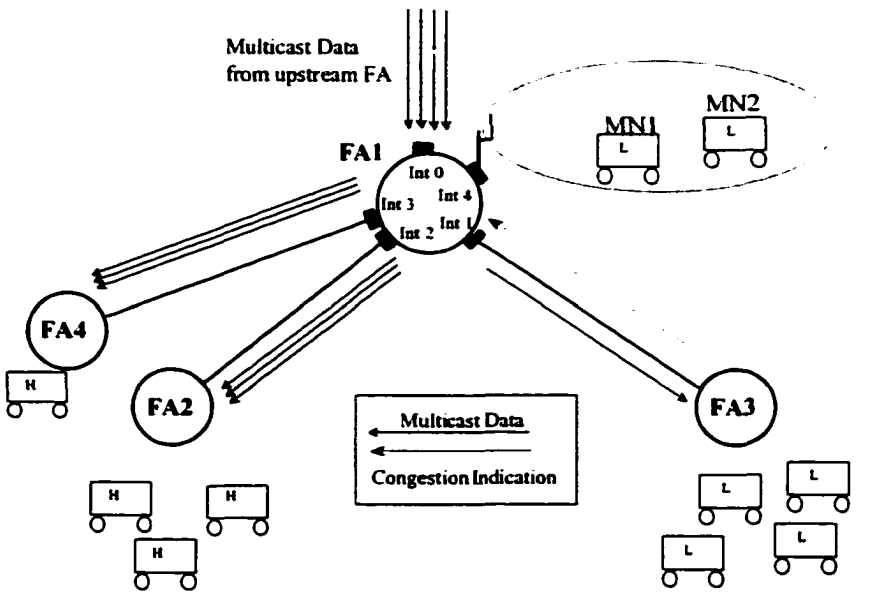

Figure 5.1 Support for different adaptive forwarding rates on multiple interfaces

Figure 5.1 illustrates a situation where MNs with different receiving rates exist in the network. FA1 is forwarding G1 traffic to FA2, FA3 and FA4. FA1 and FA3 are serving MNs with low rate interfaces, while FA2 and FA4 are supporting MNs with higher rates. FA1 will start increasing the forwarding rate until it starts receiving congestion indication from FA3. FA1 will continue forwarding to FA2 and FA4 with the higher rate as long as local resources allow this process to continue. The support for such feature allows MNs to continue receiving traffic with bursty characteristics, similar to some multimedia application traffic, with less effect from other slower receivers.

This arrangement can be supported without keeping multiple copies of the same datagram for each interface. The multicast datagrams are kept in a shared storage space, while a data structure associated with each group and interface combination is managed to keep track of the reliable multicast support states. This approach is different from

another simpler one where all downstream interfaces are treated as one interface. In such simple approach, multiple forwarding rates are not supported, and responding to a retransmission request involves broadcasting the requested datagrams over all interfaces with the same rate. Our proposed scheme provides more efficient retransmission with minimum additional overhead in the size of the management data structure.

Figure 5.2 illustrates a conceptual data structure for such support. The support for the reliable multicast does not introduce new data structure; instead it expands on those structures needed for the unreliable multicast service. For example, the unreliable multicast support requires keeping information of the groups and the associated interfaces. The reliable multicast will introduce additional information related to the last transmitted and acknowledged sequence numbers, congestion indications and current transmission rates.

To adapt the forwarding rate, the FA will calculate the amount of time to delay a newly arriving datagram. The data structure also will include the moving average of the arrival rate, used for identifying drop priority among different receivers. If the FA is acting as a LMSP, then additional states are kept to support the external reliable multicast session. The structure can be extended to support multiple rates per wireless interface for the same group, to support different QoS in the same service area, if the wireless transport system allows such support.

On receiving a multicast datagram, the FA will start a process for the datagram handling based on the multicast group ID and the associated sequence number. In

addition to forwarding datagram on the corresponding interfaces, the current forwarding rate associated with each interface is verified to avoid exceeding the adaptive rate selected for this interface. This adaptive rate is updated according to the ACKs and the congestion indications received from the downstream FAs.

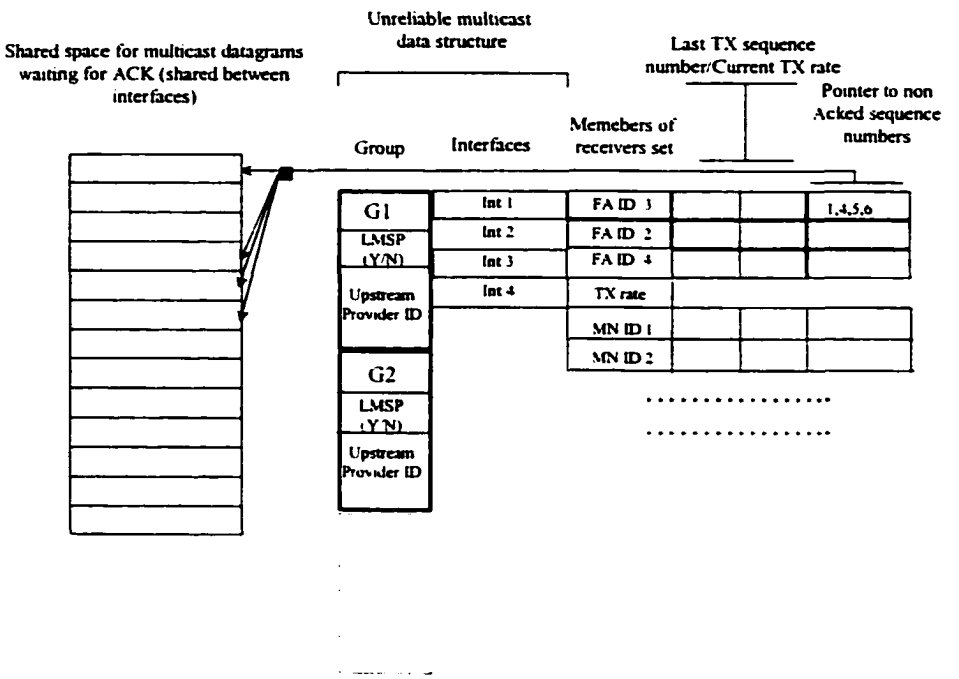

Figure 5.2 Data structures at the FA for multicast with reliable delivery support

When a FA receives a multicast datagram, it starts a handling process for forwarding on the appropriate interfaces with the required rates. Figure 5.3 illustrates a conceptual process flow on an intermediate FA to handle an incoming multicast datagram.

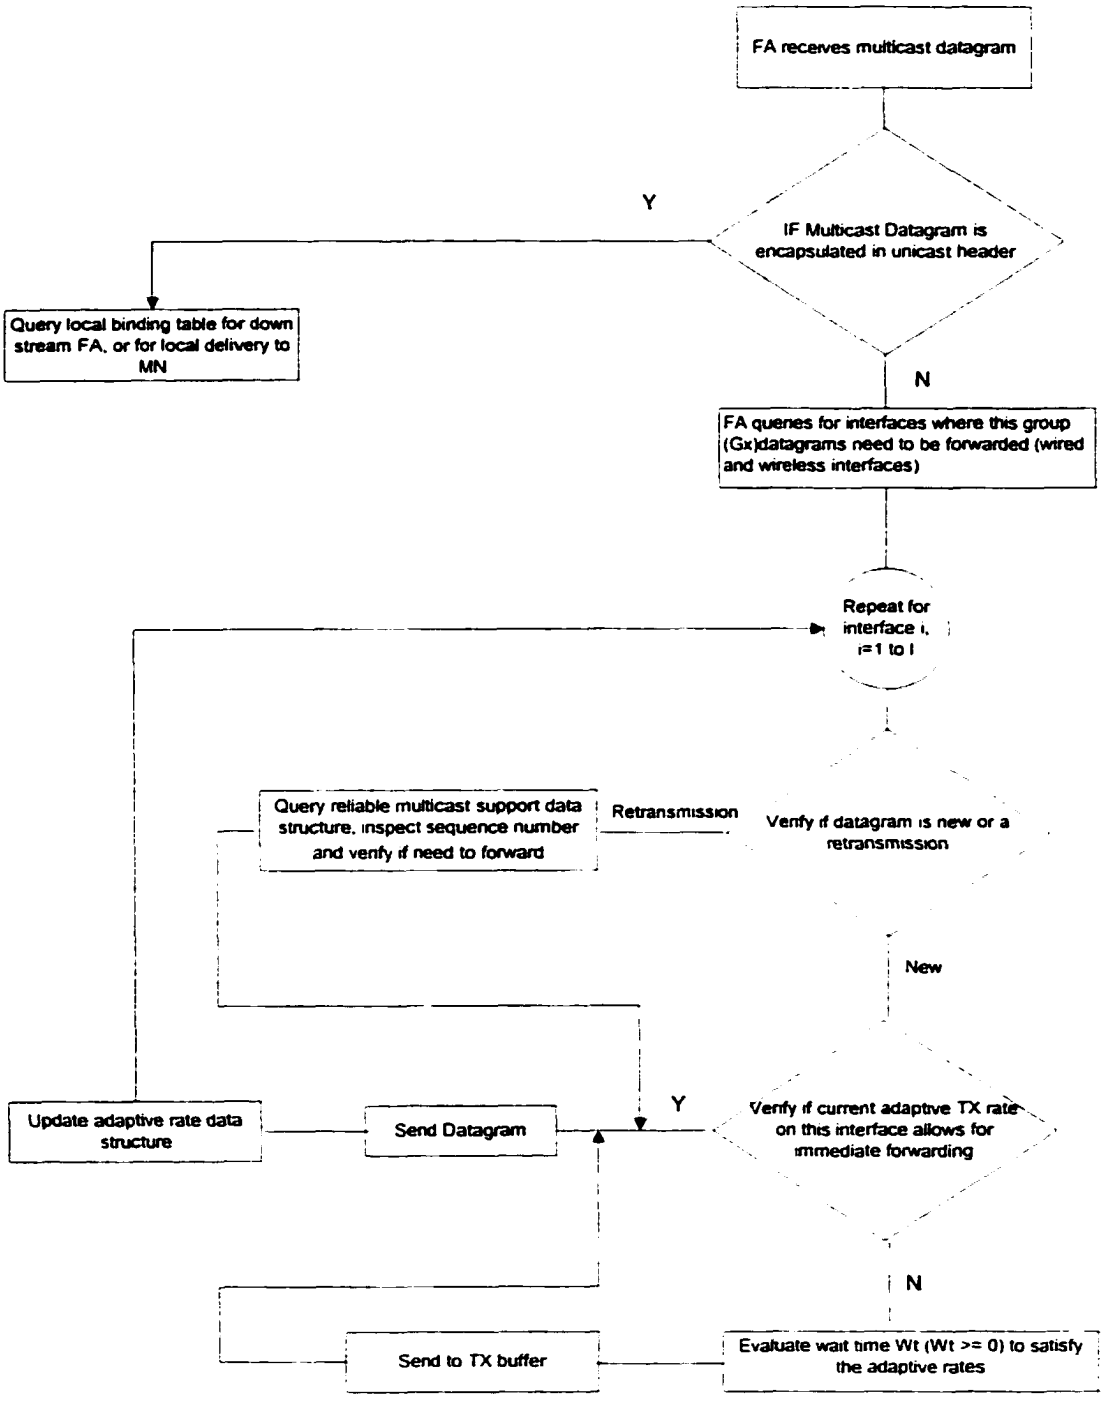

Figure 5.3 Conceptual process flow for the processing of a received multicast datagram

5.3.3 Retransmission Requests

When a FA receives a multicast datagram, it will check the sequence number and will forward the datagram to the local MNs and/or FAs on the lower level. On detecting a missing datagram, the FA will need to send a retransmission request with information regarding the requested sequence number.

In a network with stationary hosts, links congestion and transmission error events represent the primary source of retransmission requests. In a mobile network, datagrams loss during the handover process represents an additional source for those requests. In this work, we propose two types of retransmission requests: *the aggregated* and *the expedited retransmission requests*.

The aggregated retransmission request capitalizes on both the multicast support platform and the inherent characteristics of the regional registration, such that minimum number of requests are propagated. Each FA will keep a bit map for missing sequence numbers, which is updated by ACKs and new retransmission requests from MNs and downstream FAs. Each of those ACKs contains an indication of the requested and last received datagrams sequence numbers, and will act as a retransmission request for the missing datagram(s). This ACK message, with the embedded retransmission request, is sent when a specific number of datagrams is received from the upstream FA. A timer, with an expiration value that depends on the current receiving rate, will trigger the generation of this ACK when lower receiving rate is observed. In those situations, where a large number of requests concerning the same group are expected to be received in a

short period of time from different FAs, the aggregated approach shows strong scalability features.

In the case where the mobility pattern of different mobile nodes is not expected to be correlated, the aggregated retransmission scheme may not present the best approach to recover those datagrams lost because of mobility. Considering the fact that a MN sends a regional registration request upon moving into the new service area, it is possible to embed the retransmission request in the registration message with minimum extra overhead. To take advantage of this opportunity the registration request with the expedited retransmission extension is generated by a MN, when moving to a new service area while engaged in a reliable multicast session. This request differs from the aggregated request in not being delayed by the ACK timer action, reducing the retransmission delay. Also this request in not considered in the congestion indication generation, with the benefit of avoiding interpreting loss due to mobility as an event that needs to be tolerated by lowering the forwarding rate.

If a MN detects packet loss immediately after moving to a new FA, the MN will enable the expedited flag in the retransmission request. The FA responding to this request will retransmit the datagram and encapsulate it using a unicast header, with the destination field set to the address of the new current FA. If this flag is disabled, indicating an aggregated request, the FA responding to the request will forward the retransmitted datagrams to the adjacent FA that forwarded the request. The expedited request scheme will reduce the delay associated with the retrieval of datagrams lost due to mobility events.

5.3.4 Aggregated Acknowledgment

More than one approach can be considered to set the time when a copy of a datagram is to be removed from the cache and when to send a corresponding ACK. If the LMSP was considered responsible for retransmission to all members in its service domain, the LMSP will need to wait for each FA and MN in this domain to acknowledge receiving the datagram before being able to recover the corresponding cache space. This approach is not efficient, in particular in the case of an extended hierarchy. On the other hand, if each FA acknowledges datagrams without waiting for the corresponding ACKs from downstream FAs, this may lead to a situation where retransmission requests cannot be satisfied.

In our proposed protocol, the acknowledgment scheme is designed with two objectives. First, an ACK is generated when at least two FAs on a lineage have received a copy of the same datagram. Care needs to be considered not to affect the forwarding rate by slowing the generation of ACKs. The second objective is to tolerate the situation when a FA fails immediately after receiving and acknowledging datagram to the upstream FA, and before forwarding it successfully to the downstream receivers.

To satisfy those two objectives, we are extending the ACK message to provide two functions. First, the ACK serves as an indication for the successfully received and the requested sequence numbers, using the ACK window bitmap. In addition, the ACK message will include a cache management bit map. The ACK window bitmap is used to adjust the forwarding rate and as a retransmission request indication. The cache

management bitmap permits the receiver to remove datagrams, already acknowledged by downstream receivers, from its local cache. In the upward trip, the ACKs from different members of the immediate receiver set are aggregated at each FA.

Figure 5.4 illustrates the approach we considered in generating and aggregating ACKs. FA1 forwards G1 multicast datagrams to the downstream FA2. FA2, referring to its reliable multicast database, finds that the members in the immediate receivers set have already acknowledged sequence numbers 1,2 and 3. FA2 is instructed by the ACK window mechanism to generate an upstream ACK. The bitmap associated with the ACK is constructed such that FA1 can update its local database with the information that FA2 has received sequence numbers 1 to 6. The cache management bit map instructs FA1 that only datagrams with sequence numbers 1 to 3 can be removed from the cache.

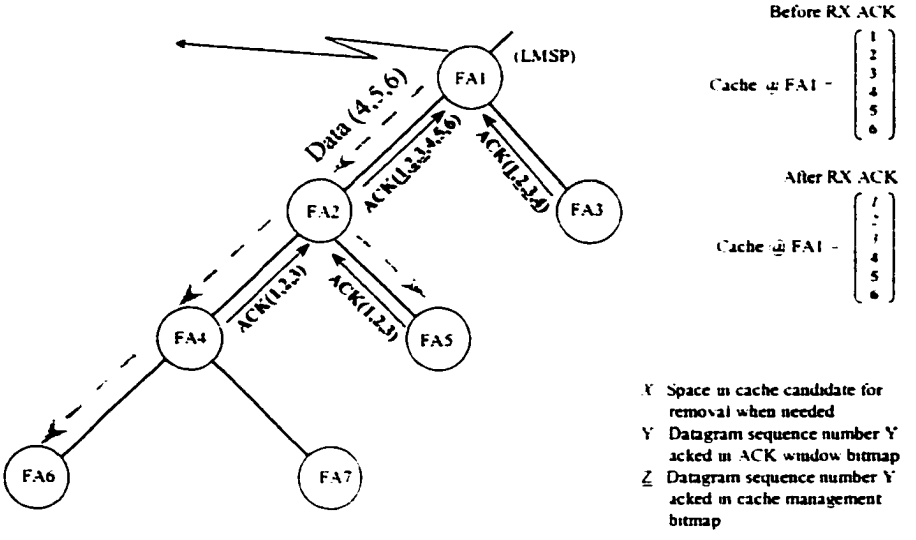

Figure 5.4 Acknowledgement for reliable multicast data

5.3.5 Immediate and Delayed Data Retransmission

On receiving a retransmission request, the FA will decide whether the request is generated by a local MN. The receiving FA does not expect to receive, on the same interface, multiple retransmission requests concerning the same datagram from multiple FAs due to the request aggregation process. Accordingly, if the request is transmitted from a downstream FA and can be satisfied, the receiving FA will immediately retransmit the requested datagram. If the request is generated from a MN, the FA will implement a policy to refrain from responding immediately to the request. Instead, the FA will wait for an amount of time, limited by an upper bound, expecting to receive other requests from other MNs concerning the same datagram.

The FA will have a record of the round trip delay (RTD) expected for each MN member of the immediate receivers list. Before retransmitting the datagram, the FA will wait for an amount of time that is equal to the difference between the maximum RTD value in the table and the RTD value associated with the MN requesting retransmission. This waiting time is limited by an upper bound to accommodate the case when the RTD values associated with the MNs vary widely.

5.3.6 Dynamic Logical Hierarchy

To provide an end-to-end reliable multicast service, the LMSP will need to join the underlying reliable delivery session. We will refer to the node, to which the LMSP is

sending retransmission requests, as the external reliable multicast provider. In a mobility support environment, a critical design aspect of the reliable delivery service is to isolate or reduce the effect of typically slow mobile receivers on the external multicast service provider. If this provider detects that an LMSP is not capable of handling the current transmission rate, a situation usually indicated by higher rate of retransmission requests or congestion indications, it may react by denying the reliable service to this LMSP.

If the LMSP has a member in its domain that is not capable of receiving the multicast stream with the current transmission rate, this situation will have one of two possible outcomes. If the LMSP continues in providing the reliable service to this element, the LMSP will start sending multiple retransmission requests, and eventually will be denied reliable service from the external provider. This situation needs to be avoided since it will deny the service associated with this particular group, for all FAs and MNs serviced by this LMSP. Another option for the LMSP is to deny the local service to the misbehaving element. Although this seems to be a practical solution, there may be a situation where the MN(s) causing the problem has a QoS contract that includes a reliable multicast service. It is beneficial to try to accommodate this element as a way of providing a premium service.

In this work, the dynamic logical hierarchy feature is proposed as a possible solution to address the two issues mentioned above, and to allow the platform to show better efficiency against temporary network conditions. When the LMSP detects that the available cache threshold has been exceeded, it will refer to its data structure to identify the particular group and the particular member of the immediate receivers set causing the

problem. If the LMSP determines that this member is accepting traffic within its service profile, the LMSP will start a logical restructure for the affected area of the domain.

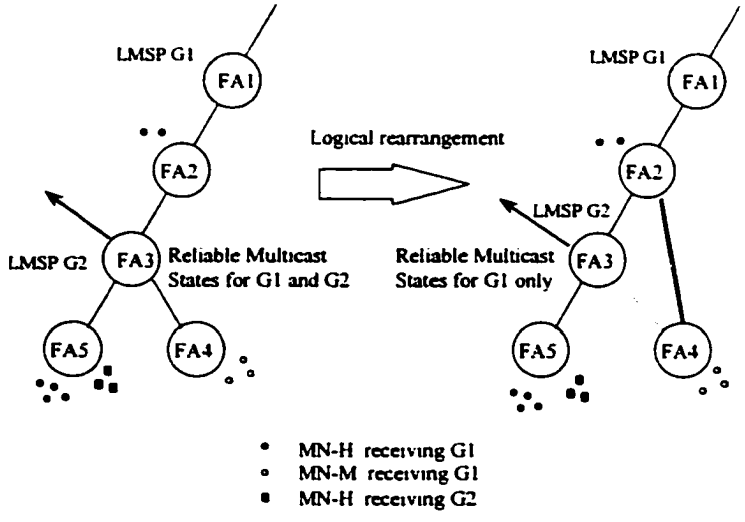

Figure 5.5 Dynamic logical hierarchy

The operation of this scheme may be best described by an example. In figure 5.5, FA1 and FA3 are is the LMSPs for groups G1 and G2 respectively. MN-H is a MN with a contract for a high class of reliable multicast service, while MN-M is a MN with a medium class of service. FA3 is not serving any local MNs. FA5 is receiving traffic for both G1 and G2, while FA4 is receiving traffic for G1. When the traffic associated with group G1 starts to exhibit bursty characteristics that cannot be accommodated by the MNs served by FA4, FA3 will start to receive congestion indications. FA3 recognizes that elements on the downstream direction of FA4 are affecting the system and that MNs with higher class of service will be denied the service eventually. Accordingly, FA3 will send a `change_logic_HIR` message, instructing FA4 that his new reliable multicast

send a `change_logic_HIR` message, instructing FA4 that his new reliable multicast provider for G1 has changed from FA3 to FA2. This message will trigger FA4 to register with FA2 as a reliable multicast receiver for G1.

In this new arrangement, FA3 will not keep any reliable delivery states associated with G1 for FA4, but may only forward the multicast traffic. This will allow FA3 to recover cache space occupied by datagrams stalled in the cache because of FA4. This recovered space will also provide FA3 with more capacity, and accordingly will be able to reduce the number of retransmission requests to be transmitted to the external multicast provider for G1, G2 and other supported groups.

5.3.7 Handling Mobility of MNs

When a mobile node moves from one area to another, it may lose those datagrams forwarded to its previous FA during the handoff process and during the time needed for setting the entries for multicast subscription. Accordingly, multicast datagrams delivered to the previous FA, instead of the current FA, need to be compensated for. Also, entries associated with the reliable multicast support for this MN, which correspond to its previous location, need to be removed.

An approach to retrieve a missing datagram is to find the closest FA that still has a copy of the requested sequence number. This task can be addressed by introducing the Reliable Multicast Forwarding Request (RMFR) extension. This message is embedded by a MN in its registration request along with its membership report, generated after moving

to a new service area. A regional FA intercepting this message can either forward it without modification, or forward it after turning the multicast forward flag off or stop forwarding after processing it according to the availability of requested datagrams.

When a MN leaves its previous service area, it will not receive a number of datagrams forwarded to its previous FA. The previous FA can store those datagrams such that the MN on the move can retrieve them at a later time upon arriving to the new FA. When a MN registers with its regional FA, a value for datagrams credit will be negotiated in addition to the traditional fields in the membership report. This credit represents the number of datagrams that the FA is willing to store, such that a MN moving away from this FA can retrieve later. When a MN moves to a new FA, it will estimate the number of datagrams lost during the handoff and compare it to its credit with the previous FA. If the credit accommodates for those lost datagrams then the MN may elect to ask the previous FA for retransmission. The MN will mark the expedited retransmission request, such that if the request cannot be satisfied by a FA on the path between the current FA and the common FA, it will be forwarded to the previous FA. If this retransmission request still could not be satisfied, the request will be directed to the underlying reliable protocol.

The details of this scheme can be illustrated with an example as shown in figure 5.6. and considering the following initial states:

- MN1 is located in the area of FA5 and receives the traffic associated with the multicast group G1.

- FA2 is the current LMSP, that provides the reliable multicast service to FA3 for G1, but has already removed datagram sequence 15 from cache.
- FA6 and FA4 do not receive G1 traffic.

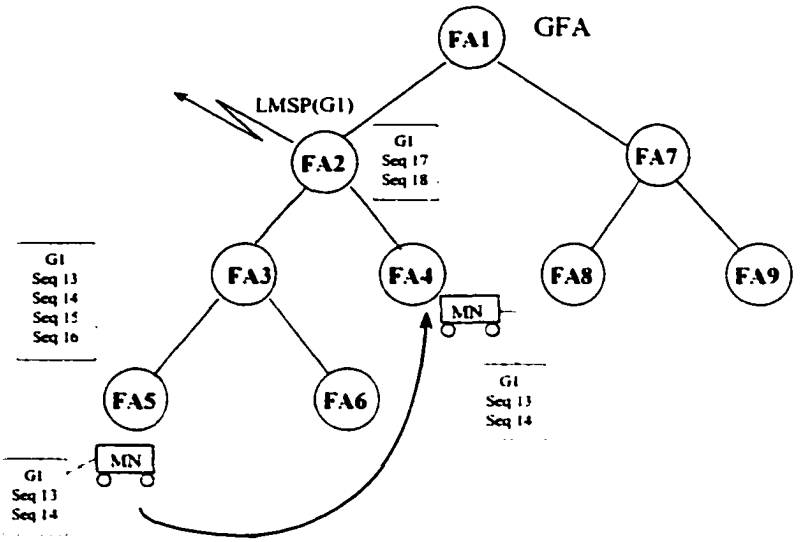

Figure 5.6 Accommodating mobility within the reliable delivery support

When the MN moves to FA4, it will send regional registration request to FA4. This registration request will embed the expedited retransmission request, asking for retransmission starting from datagram sequence number 15. The request will include the RMFR pointing to FA5 as the previous FA. Since we assumed that local cache management process on FA2 has removed sequence 15 and 16 of G1 multicast datagrams, only datagrams starting from sequence 17 will be forwarded towards the MN at FA4. To retrieve the remaining datagrams, the RMFR will be forwarded downstream

towards FA5. FA3 will intercept this message, locate the required sequence numbers, and send the corresponding datagrams using unicast to the current FA (FA4). At this point, the RMFR will not be forwarded any further.

The case of interdomain mobility is handled in a similar manner. When the MN moves to the new domain, it will generate a registration request towards the GFA. This request will be forwarded to the HA to inform about the new root FA. If the requested sequence number cannot be found in the new domain, the RMFR will be forwarded to the GFA associated with the previous domain. The most upward FA in the previous domain, that has the requested sequence numbers available, should retransmit the corresponding datagrams. Since moving to a new domain is associated with the MN sending a registration request destined to the GFA, using the expedited mode of retransmission requests proofs very efficient in reducing the associated signaling.

The mobility of MNs can indirectly affect the reliability of datagram delivery, by affecting the LMSP election process. The handover process, from the previous to the new LMSP, has to be carefully designed to avoid negatively affecting the performance. The handover process starts by an LMSP, intercepting membership reports, decides that the number of MNs does not satisfy the requirement for keeping the LMSP responsibility for a particular group. Accordingly, the FA will forward the summarized report upstream. When the report is processed by a FA that can assume the LMSP responsibility, this FA will join the requested multicast group. The previous LMSP will give up the LMSP functionality only when it starts to receive the multicast traffic forwarded from the new LMSP.

5.3.8 Tolerating FAs Failures

Due to the nature of forwarding traffic on the hierarchy, the failure of a FA has a severe effect on the multicast delivery to the MNs. The failure of a single FA will deny the multicast service to a large number of FAs and MNs. In this work, we will apply a scheme to support the fault tolerance feature in the proposed reliable multicast platform.

The current information, regarding the regional FAs interconnectivity in the region served by a GFA, is maintained by the Hierarchy Registry that is a logical function co-located with the GFA. A FA is responsible for monitoring the health of FAs on the immediate lower level of the hierarchy. The FA will send a Hello message on regular intervals towards those FAs. A FA will respond with a Hello_ACK for those messages. On receiving those ACK messages, the FA will calculate the round trip delay to the corresponding FAs.

Service recovery upon the failure of FAs

The failure of a leaf FA will cause service interruption only to local MNs. Service can be recovered if the MNs in the affected area can still register with another FA, similar to the case where service areas are overlapping. Reliable service can be restored with a procedure similar to that of handling MNs mobility described earlier.

In addition to service outage to the local MNs, the failure of an intermediate FA event will prevent service to the downstream FAs. After detecting the failure of the FA, the parent FA (FA_H) will query the Hierarchy Registry for the affected FAs list (FA_List)

that contains FAs in the downstream level relative to the faulty FA. On receiving this list, FA_H will send the Change_Hir message to each member, instructing the FAs to consider FA_H as their new parent FA instead of the faulty FA. In a message acknowledging the receipt of this request, each FA will include a list of the current MNs serviced locally, allowing FA_H to update its forwarding table. The ACK will be generated immediately by the FAs in the FA_List, without waiting for the Change_Hir to be intercepted by the leaf FAs. In addition, the FAs will provide FA_H with the reliable multicast states that include the group membership and the last received sequence numbers. It is expected that the failure event will prevent the ACKs, corresponding to packets received just before the failure and sent by members of FA_List, from reaching FA_H. The information regarding the last received sequence numbers, embedded in the Hir_Change_Ack, will compensate for those lost ACKs and will avoid having the FA_H to forward already received datagrams.

The recovery from a failure of an LMSP FA is more complicated than that of a non-LMSP FA failure event. In the LMSP failure case, an affected FA has no guarantee that any of the FA in the higher level is supporting the required multicast group, leading to longer recovery time. The membership report will travel upstream till intercepted by a FA that is able to satisfy the LMSP requirements. In this case, the FA_List member will join the domain serviced by an upper LMSP, or will become an LMSP itself. If such LMSP is not supporting the requested multicast group, it will need to use the underlying multicast protocol to join the multicast tree.

Figure 5.7 illustrates the control messages exchanged to support the failure recovery process. The states associated with unicast, multicast and reliable multicast need to be delivered to the upstream FA. Service is recovered as soon as the upstream FA receives the information. Information propagated to the MNs is needed to advertise the lineage leading to the GFA, and directed to newly arriving MNs.

Retransmission for packets lost after the failure of a FA

The case when the FA fails immediately after receiving datagrams and sending the corresponding ACK to the upper FA, but before forwarding those datagrams down stream, needs to be considered. This situation is tolerated by the cache management bitmap scheme illustrated earlier in figure 5.3. A FA, receiving an ACK message, will need to inspect the cache management bitmap vector and accordingly marks datagrams as candidate for removal. If FA1 is forwarding multicast datagrams to FA2, then FA2 will not instruct FA1 to remove the datagrams from cache immediately after receiving the datagrams. FA2 will wait for FA4 to ACK, and then it will forward a cache bitmap to FA1. Only after receiving the ACK with the embedded bitmap, FA1 can clear the corresponding space in the cache.

This arrangement allows for 1:1 like data redundancy, such that there are always copies of the same datagrams on two different FAs. One FA is allowed to remove its copy when a third FA, in the downstream direction on the hierarchy, ACKs the receiving of those datagrams. This feature is particularly useful in an environment with relatively higher FAs failure probability, and in an environment where the FAs are mobile.

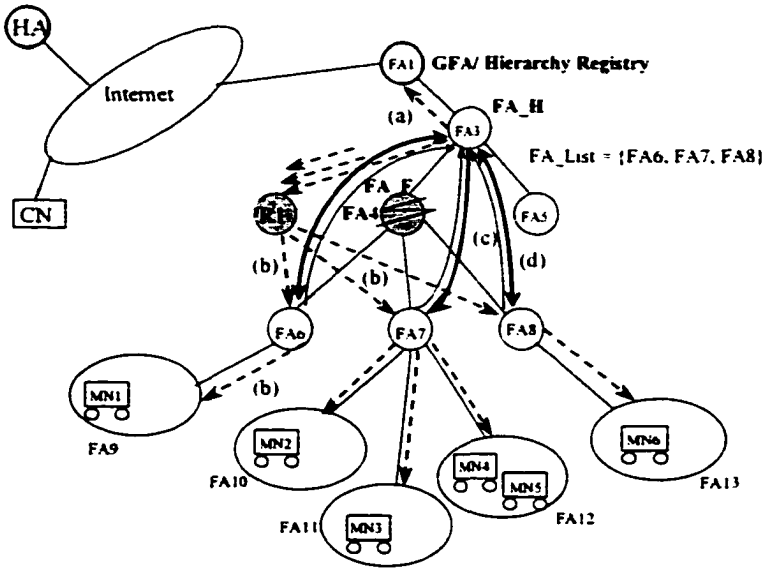

(a) Query for FA_List
(b) Change_Hir message propagated downstream
(c) Change-Hir_Ack message to FA_H
(d) Date can be forwarded from FA_H to FAs in FA_List

Figure 5.7 Tolerating failure of FA

5.3.9 Reliable Multicast and Quality of Service Support

In addition to delay and jitter, packet loss is an important parameter that is often considered by application requiring quality of service (QoS). It is very possible that the mobility environment will include MNs receiving multicast traffic for the same group with different QoS profiles. Reliable delivery may be provided as a premium service that may not be supported in all profiles. A premium QoS service profile will possibly specify

lower delivery delay for multicast traffic, for original and retransmitted datagrams, in addition to low jitter and low packet loss. Accordingly, the presence of MNs with such premium profiles may suggest assigning the LMSP responsibility to those FAs in close proximity to such MNs. In addition, larger datagram storage capacity credit on previous FAs can be made available to MNs if their profiles specify such preference. In this approach, the FA receiving the reliable multicast request will authenticate and authorize the request with the element providing the resource management function in the platform supporting the QoS.

5.4 Performance Evaluation and Simulation Results

To illustrate the behavior of the proposed scheme, we considered discrete events simulation to evaluate aspects of the reliable multicast support performance. The objective of this simulation is to highlight some of the characteristics of the proposed schemes under different conditions.

Hierarchical network model: The Network environment used for simulation has a hierarchy that has two GFA, two HAs and seven levels of intermediate FAs. The simulation environment includes 4 HAs and 100 MNs moving among 25 FAs. Each MN can request to receive multicast traffic from up to 2 groups out of 10 available groups, unless otherwise mentioned. The FAs and MNs send an ACK for each received 16 packets. If missing datagrams are detected in two ACK windows in sequence, a congestion indication is sent upstream. Receiving a congestion indication will cause the

FA to reduce its current transmission rate by 20%, while two ACKs with no retransmission requests will increase the transmission rate by 10 %. When the available free cache falls below 5% of the total size, the FA will start dropping datagrams.

The delay over links between two FAs in the hierarchy is set to 10 msec, and the delay over wireless links is set to 25 msec. The delay between a HA and the root FA ranges from 50 to 80 msec. Different values for the delay should not affect the correctness of the approach. The only restriction is that the delay from the MN to the HA should be larger than that from the MN to the root FA. This assumption matches the expected topology and is the primary motivation behind introducing the local registration approach. The multicast traffic is generated from four CNs to the different multicast groups, with average of 250 datagrams/sec. Simulation time of 150-0 minutes was considered, with uniformly random distributed mobility rate between 1 and 5 moves/minute. The packet loss probability on the wireless link ranges between 0.1 and 0.01, while on the wired link it is 0.001.

Effect of introducing the expedited retransmission request

In this experiment, we will investigate the delay and the overhead associated with recovering missed datagrams during mobility events. The first performance aspect considered is the signaling overhead associated with generating the expedited retransmission requests. We evaluate the cost associated with the retransmission requests as follows.

We compare between the case of using aggregated request and that of the expedited request. When generating an expedited request, it is embedded in the local registration request as an RMFR, such that no additional message is needed to carry the request. This is true if the request can be satisfied before or when the common FA receives the registration request. If the request is not satisfied by this time, the RMFR will be forwarded as a separate message. A feature of this message is that it is not subjected to the aggregation process, thus it is forwarded faster. The possible drawback of this approach may appear when a large number of MNs move to the same service area within a short time, while receiving the same multicast group traffic. Depending on the location of the element responding to the request, this may lead to large number of RMFR on some links of the local regional domain. This effect is greatly reduced by the fact that requests from MNs, generated within a short time from each other, are aggregated locally by the current FA. Also the possibility of this situation to happen is reduced with relatively large number of MNs, where the probability of finding a FA that can retransmit those packets is higher.

Using aggregated requests has the advantage of minimizing the signaling overhead, where FA retransmission requests are aggregated from FAs downstream and from local MNs, and only one request is transmitted upstream. The possible drawback of this approach is the inherent delay, associated with forwarding the aggregated request upstream.

We define the load of retransmission request as the summation of the total number of retransmission requests traveled over each individual link. The load is affected by

different factors associated with topology and network dynamics. For example, less number of hops traveled by a request will lead to a reduced cost. Also, aggregating those requests will produce less number of messages leading to similar effect. In our proposed protocol, we are using expedited retransmission requests to recover datagrams lost because of mobility. To justify this selection, we evaluate the associated load when using the expedited request, and we compare the result to that of the case of using aggregated requests. Figure 5.8 illustrates the results of this experiment, where retransmission requests considered here are only those due to mobility events. The ratio of the retransmission load in both cases is evaluated. In the case of the expedited requests, only those requests where local registration requests have been processed by a downstream FA are counted. The load was evaluated in this way to account for the fact that an expedited request is embedded into the local registration message, and only travels as separate message after being processed by the common FA.

With small number of MNs, it is less likely that MNs will be able to form clusters, or that a MN using the aggregated request will find a FA on the new lineage to respond to its retransmission request. On the other hand, the local registration will carry the expedited request at least till the common FA. Accordingly, the expedited mode will outperform the aggregated approach in this range. With increasing number of MNs, the aggregated mode will start to show better performance due to more requests available for aggregation. With sufficient number of MNs, it is very possible that a FA capable of satisfying the request can be found near the new FA, causing both approaches to perform

better. The expedited scheme will still generate fewer messages, capitalizing on the local registration messages.

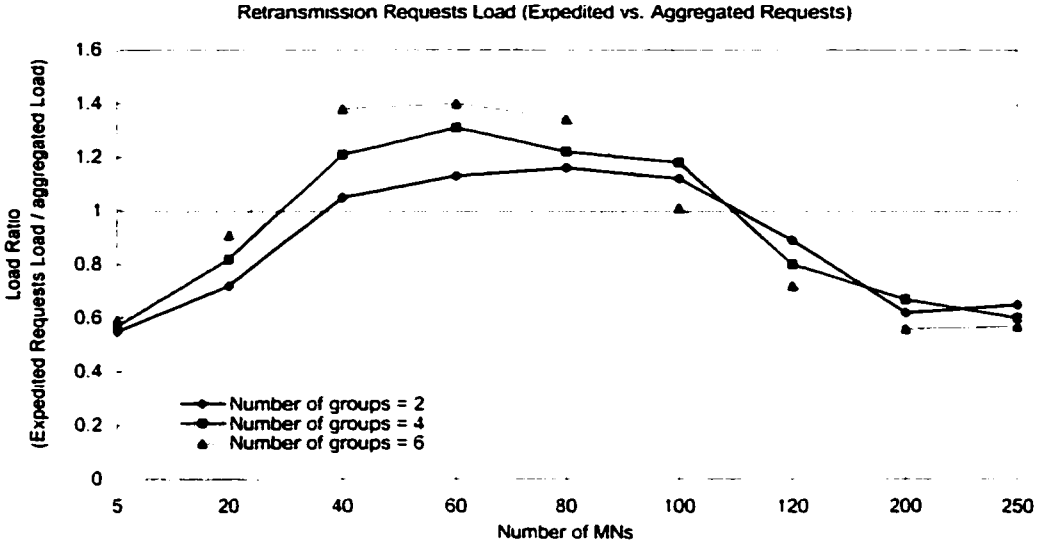

Figure 5.8 Retransmission requests load with number of MNs

The number of groups available for the MNs to join affects both retransmission request schemes. With number of MNs between 40 and 90, the larger number of groups assists the aggregated requests scheme in performing better. With larger number of MNs, and accordingly larger probability that FAs responding to the retransmission requests will be within few hops from the current FA, the effect of the number of groups becomes less apparent. The results show that overhead associated with the signaling in the expedited requests is not significant, and does not show scalability problems.

The second performance aspect that we will consider is the delay in delivering the retransmitted datagrams to the MN. The delay is evaluated from the time when a retransmission request is issued by a MN till the time when the MN receives the requested datagram. Figure 5.9 illustrates the average delay values, in delivering the retransmitted packets after mobility events, with different values for MNs in the system. The results show that the expedited approach provides less delay. One reason is that the aggregated scheme benefits appear when considerable number of requests are sent around the same time from different FAs. This is not a characteristic of retransmission requests generated upon roaming.

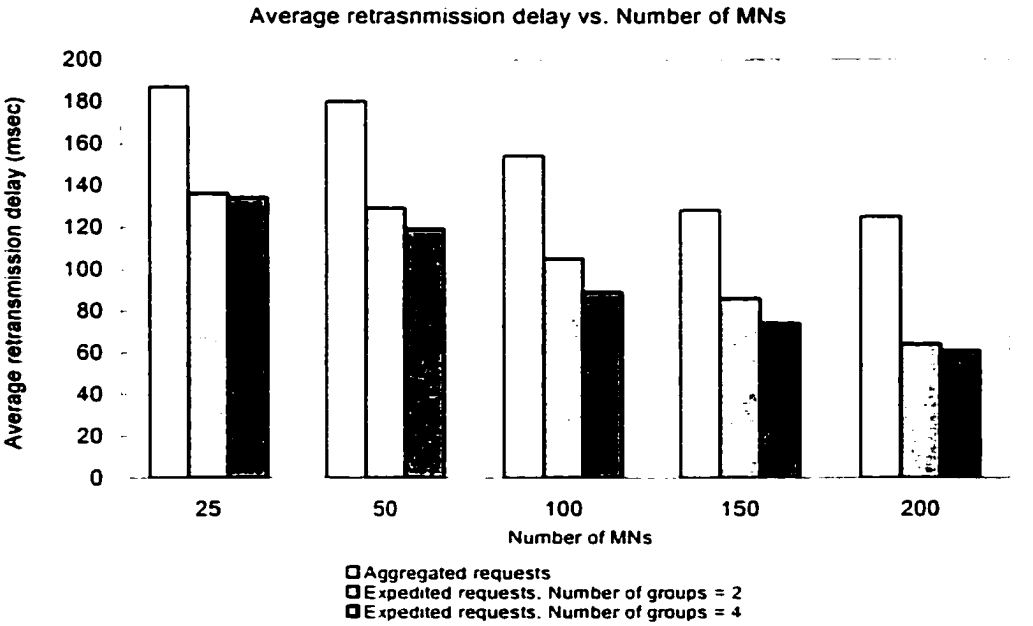

Figure 5.9 Retransmission delay with number of MNs

Retransmission delivery cost

Figure 5.10, illustrates a scalability aspect of the proposed system. We evaluate the cost of delivering the retransmitted datagrams. The cost is defined as the summation of number of hops traveled by each retransmitted datagram copy. Accordingly, the less the number of duplicates and the distance traveled by each copy, the less is the cost. In this experiment, the cost is associated with datagrams retransmitted because of links congestion and mobility.

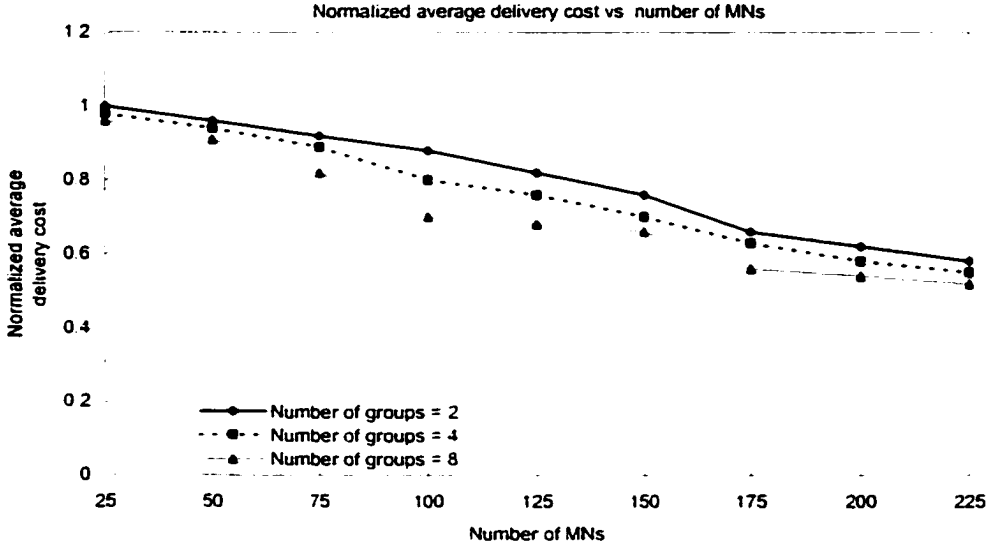

Figure 5.10 Scaling characteristics with number of MNs

The cost is normalized using a reference cost value that is equal to the average cost of retransmission with 25 MNs, and two multicast groups. The larger number of MNs will lead to more aggregated and expedited retransmission requests. The aggregation effect will reduce the average cost when the number of MNs increases. Also the distance

traveled by both types of requests will be reduced with larger number of MNs. Increasing the number of groups assists in locating a FA which is closer to the current FA and can respond to the retransmission request.

At first, it may appear that the expedited scheme can cause explosion in the number of retransmissions received, if a large number of MNs move simultaneously into the same FA and start requesting the same sequence numbers. This is not the case since the system effectively reduces the number of requests, using aggregation on the CFA, such that multiple requests concerning the same group will be combined.

A datagram retransmitted, in response to an aggregated request, is duplicated at FAs where recipients on multiple interfaces have requested the same sequence number. This is in contrast to the unicast retransmission in the expedited request case. The aggregated scheme is well suited to respond to retransmission requests upon congestion or losses on interfaces and transmission links, where multiple requests are expected to be transmitted in a short period of time.

Effect on external reliable multicast provider

An efficient reliable multicast scheme will try to avoid reducing the transmission rate of the multicast source, by trying to locate the requested datagrams locally. In the experiment shown in figure 5.11, we look at the number of retransmission requests observed by the external reliable multicast source. We monitor the ratio of the number of requests satisfied locally to the total number of requests. We consider different numbers of MNs and different mobility patterns.

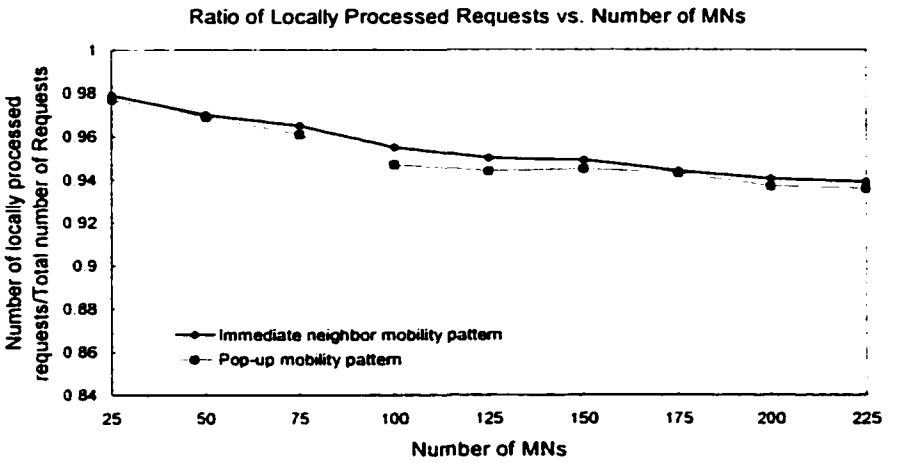

Figure 5.11 Effect on the external multicast source

In the first mobility pattern (immediate neighbor), the MN is allowed to move to an immediate neighbor FA. In the pop-up mobility pattern, the MN can move to any new FA. The first observation is that, in both mobility patterns, the number of requests processed locally in the regional domain will be reduced slightly with increasing number of MNs. With smaller number of MNs, less stress is being applied on the cache. Accordingly, larger number of datagrams can be stored locally on the previous FA and can be made available to be retrieved by the MN upon arrival to the new FA. Although increasing the number of MNs imposes larger load on the FAs caches, it also increases the possibility of finding a responding node within the domain. Overall, the platform shows a robust performance, and the majority of requests will be satisfied within the regional domain. Since the pop-up mobility pattern involves longer time to move from

one cell to another, more datagrams need to be recovered. This situation imposes more stress on the cache and leads to slightly less number of requests processed locally.

Effect of mobility rate on retransmission delay

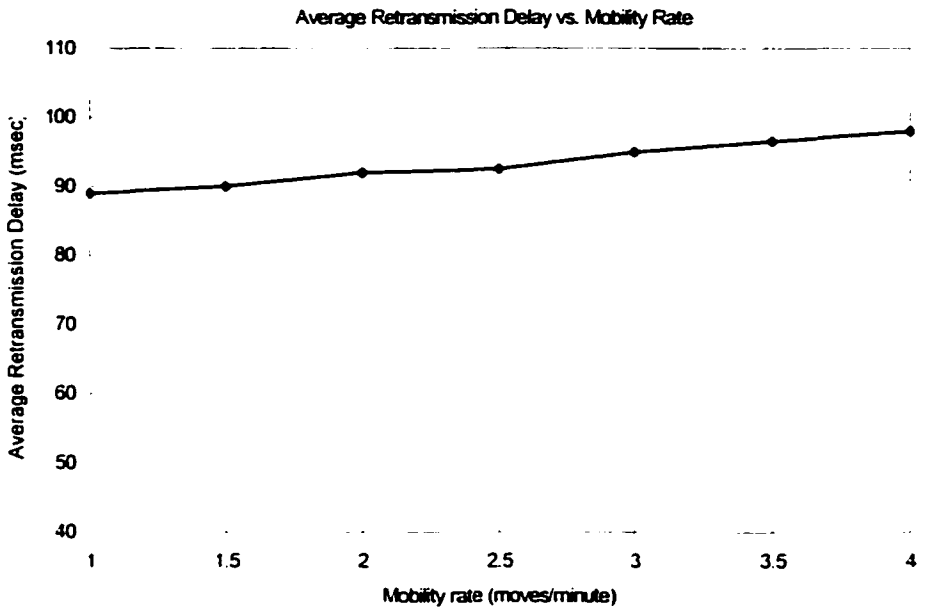

Figure 5.12 Average retransmission delay

Figure 5.12 illustrates the effect of increasing the mobility rate of MNs on the average retransmission delay. The increase in the number of mobility events imposes larger load on the cache where more datagrams need to be stored, leading to larger average retransmission delay. This behavior is not severe and can be tolerated for the following reasons. To produce a major negative effect on the cache because of mobility, a large number of MNs need to move simultaneously, and also should be subscribed in a large number of multicast groups. Another limiting factor is that high mobility translates to

more LMSPs subscribing to multicast groups, and those will be able to respond to the retransmission requests more effectively.

Effect of loss rate

As illustrated in figure 5.13 , we study the effect of different error rates on the average time needed to complete a successful session reception. The error rate on some wireless links (50% of the total number of links in the domain) is varied between 0.01 and 0.1, while kept as 0.01 for the remaining number of links. The link is simulated such that the average good state duration is 100 times that of the average bad state. A multicast session average delay is evaluated for different network configuration. Sessions of 100,000 datagrams each are considered in this simulation. As expected, if one forwarding rate is used over all interfaces of the FA, higher error rates will cause consistent delay increase for all MNs. On the other hand, when the adaptive forwarding rate scheme on the different interfaces is used, the system will tend to reduce the forwarding rate only on interfaces receiving larger number of retransmission requests. This arrangement allows the system to tolerate the temporary error conditions, without unnecessary delay to other receivers. If the mobility pattern of the MNs exhibits clustering characteristics, less average delay can be achieved. In this clustering pattern, the MNs move such that 75% of the total number of MNs in their new location have neighbor FAs that each is serving at least one MN. Clustering will enhance the performance since retransmission requests are aggregated in the current FA. Accordingly, less number of requests will be propagated and lower load on the cache will be observed.

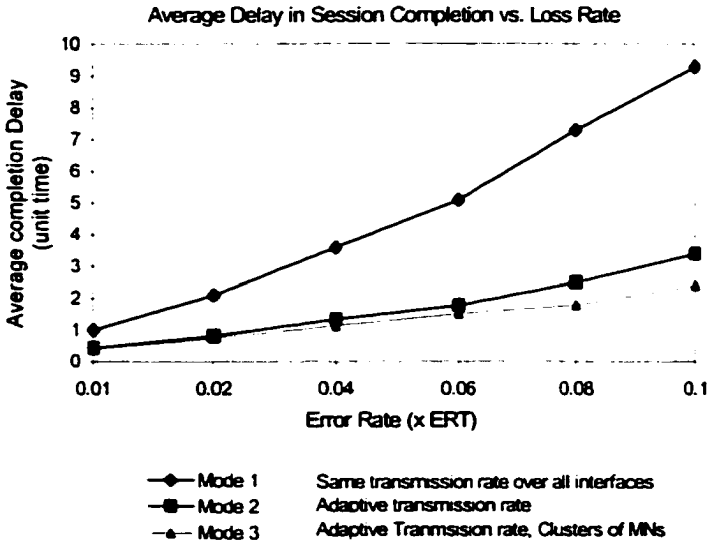

Figure 5.13 Effect of loss rate on session completion time

Support for multiple forwarding rates

In this experiment, bursty traffic is generated from a multicast source, and delivered to the MNs. The generated rate from the multicast application alternates between two levels, TXRH and TXRL (where TXRL is 35% of TXRH). Two categories of MNs exist, one with low rate receiving capacity RXL (50% of TXRH) and the other with a higher rate RXH (100% of TRXH). When no adaptive scheme is used, the FA transmission rate is limited by the slowest of the receivers. In the mode of operation, where the FA can transmit with different rates on the different interfaces, FAs will reduce transmission rates only on those interfaces where congestion indications are received. This arrangement allows FAs and MNs with higher receiving rates to make use of their available resources.

The FA will continue to support multiple rates as long as local resources, limited by the cache sizes, will allow. Considering the fact that FAs on higher levels of the hierarchy carry more traffic and service more groups than those FAs on the lower levels, an engineering guideline may specify larger cache sizes for those FAs. As shown in the diagram, when considering the multi transmission rates mechanism, faster receivers are able to receive higher rates and capture more of the real time characteristics of the original transmitted traffic. If the burst size is too large, cache limitation will force the FA to reduce its rate over all its interfaces as shown in the last third portion of the time line in figure 5.14.

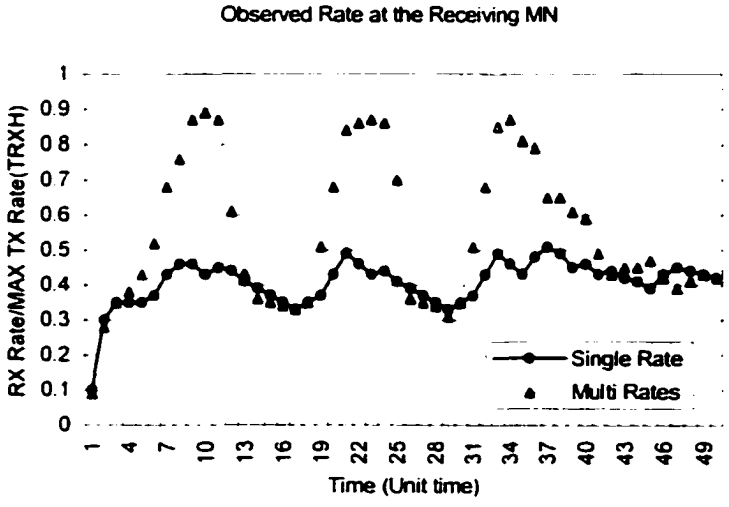

Figure 5.14 Support for MNs with different receiving rates

6. Tolerating the Failures of Mobility Support Agents

The support of the Mobile-IP protocol relies on services provided by the home agents and the foreign agents. The failure of any of those agents will impact the service provided to the mobile nodes. Mobile-IP systems, utilizing the local registration approach, are more sensitive to foreign agents failure such that the failure of a single agent may disrupt the service provided to mobile nodes in other FAs service areas.

The issue of tolerating the failure of elements supporting mobility is critical for reliable services, and thus it has been addressed by researchers from different points of view. In [45], information stored at a mobile station is replicated at several secondary support stations. Challenges faced by the checkpointing approach to support distributed applications, such as the MN disconnected mode, are described in [46]. A protocol that tolerates both base station failures and the corruption of time stamp is illustrated in [47]. In [48] and [49], approaches based on checkpointing and search are demonstrated to support failure restoration.

Different from those mentioned protocols, our proposed protocol for failure recovery [50] was designed to capitalize on the regional registration characteristics. For example, the GFA usually has partial location information regarding the MNs currently residing in its domain. Such information can be used to provide a simple and quick recovery scheme for HA failure. Also, partial location information available at the intermediate FA can be used to recover from FAs failure.

6.1 Fault Tolerance Challenge in the Local Registration Environment

The effect of failure event of mobility support element can be described as follows:

- Foreign Agent failure

In the case of Non-Hierarchical Mobile-IP, the failure of a FA will cause a loss of network connectivity that is limited in scope to the MNs currently serviced by this particular faulty FA. The failure scope is wider in the case of HLRM-IP where the failure of a single FA will prevent the packets flowing through the hierarchical system from reaching this FA and any other FA in a lower level of the hierarchy. It is clear that a FA failure exhibits more negative effect in the HLRM-IP case.

- Home Agent failure

For a Non-Hierarchical Mobile-IP, the HA needs to be continuously updated with the MNs' exact current location information, which will be reflected as larger activities associated with the local MN-FA binding table. On the other hand, consider a HA which provides the HA service for a number of MNs using HLRM-IP. The binding table on the HA will experience less activity. This behavior suggests that HA redundancy can be implemented with less demanding requirement in the case of HLRM-IP.

In the current specification of Mobile-IP it is clear that the MN relies on the HA for connectivity and for cache maintenance when using the route optimization extension. The HA is expected to be responsible for the hosting of multiple MNs, representing a single point of failure for those mobile nodes. The HARP protocol presented in [51] allows two

or more HAs to cooperate and share registration information associated with the MNs. Each HA element in the redundancy set is configured with information about the other peers. In this work, we will propose a simple mechanism to create HA redundancy that is less demanding than the HARP protocol, taking advantage of the inherent characteristics of the local registration.

6.2 Proposed Solution

In this section, we will propose solution for failure recovery in local registration Mobile-IP systems. Home agents and foreign agents failures are considered in this work. First we describe an approach to recover form HA failure, then we illustrate two schemes to tolerate the failure of FAs.

6.2.1 HA Fault Tolerance

In local registration systems, the regional movement of a MN is made transparent to the HA by binding the location information at the HA to the gateway FA servicing this region. Accordingly, the HA will have entries in its MN-FA binding table pointing to gateway FAs. If the mobility of each of the MNs is confined to one hierarchy (or one gateway FA), which is a reasonable assumption at least for a limited period of time, then the HA binding table will not experience significant maintenance activities. This is a

direct result of the fact that the effect of the MNs local mobility has been isolated from the HA.

According to the above discussion, we propose providing HA redundancy in such system using a secondary redundant HA (HA_S) that can be located on a partitioned or a non-partitioned home subnet along with the primary HA (HA_P). The inherent features of the HLRM-IP suggest the use of less demanding HA redundant system. As a matter of fact, our proposed HA redundant system does not require the secondary system to be dedicated all the time.

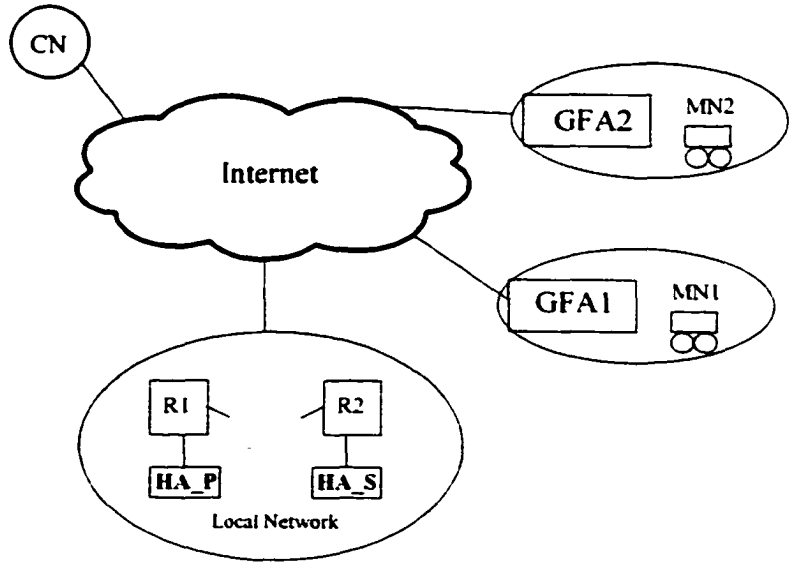

Figure 6.1 HA redundancy in the regional registration environment

As shown in figures 4.1 and 4.2, redundancy can be provided to the dedicated primary HA by considering a non-dedicated secondary HA. HA_P is a standard HA, while the HA_S is a host capable of implementing the HA functionality when needed and have a

mechanism to monitor the health of the primary HA. HA_S is configured with the list of MNs supported by the HA_P and with a list of the supported gateway FAs, which is expected to be a short list. The detection of the HA_P failure by the HA_S will cause an update to the internal routing mechanism such that traffic destined for the HA address will be directed to the HA_S instead of HA_P.

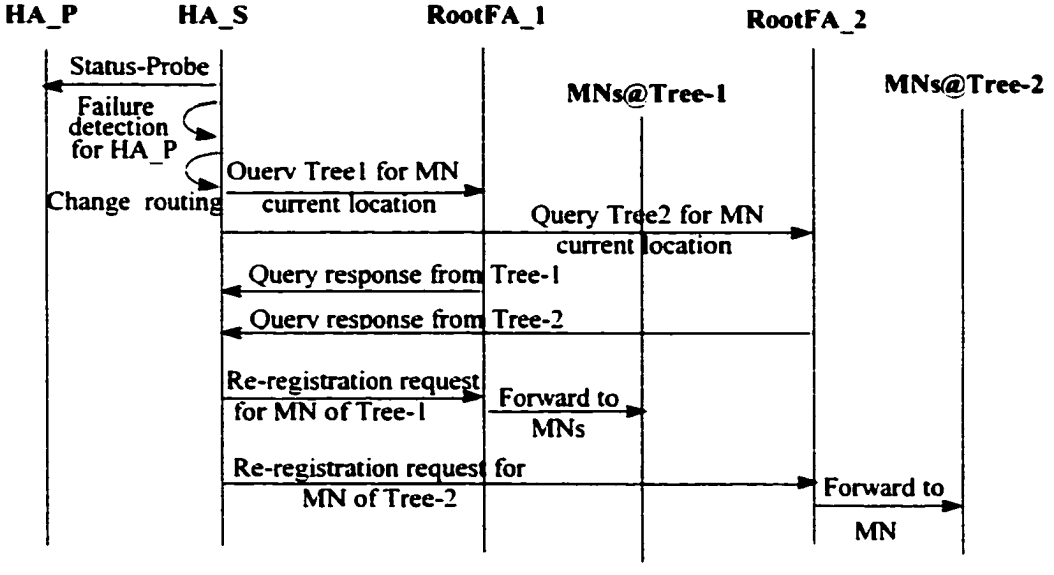

Figure 6.2 Messages associated with supporting HA redundancy

Under normal operation conditions, the HA_P will be responsible for the HA functionality. Upon the HA_P failure detection, the HA_S will send a request to the gateway FAs asking for the current location of the supported MNs. On receiving the gateway FAs responses, the HA will be able to build its binding table and the HA_S can restore communications with its MNs and ask them to re-register. The intention behind

this mechanism is not to provide a hot standby system but is to provide a minimally demanding redundant system that takes advantage of the HLRM-IP features. The list of the supported GFAs can be synchronized between the two HAs capitalizing on the regular health check messages. On receiving the gateway FAs responses, HA_S will be able to build its binding table and will restore communications with its MNs instructing them to re-register. The routine running on HA_S to control the responsibility between the HA_S and the HA_P can be illustrated as follows:

```
While (send status_probe to HA_P) returns HA_P is OK
      Query HA_P for GFA_Active_List
If  (send status_probe to HA_P) returns HA_P is down
      Look up last version of the GFA_Active_List
      For each GFA_i ∈ {GFA_Active_List}
            Query GFA_i  for List_of_current_MNs
            Update Entries in HA_S
            Change routing entries such that HA_S intercepts the new
            registration requests
            HA_S assume HA_P responsibility
If  (send status_probe to HA_P) returns HA_P is OK
            Upload current data structures to HA_P
            Wait for HA_P solicitation for HA responsibility resumption
```

When the HA_P comes back on line, the HA_S will detect this new event. The Ha_S will start uploading the current binding information to the HA_P. After information transfer,

the internal routing is set such that new registration requests are forwarded to the HA_P that is now back on-line.

6.2.2 FA Fault Tolerance: Revert to HA Registration Mode

In the case of a FA failure, all MNs serviced directly by this FA or by a FA located in a lower level in the hierarchy and on a path on which the faulty FA is an intermediate point will suffer from loss of service. In this approach, all MNs affected by the FA failure will be notified and are instructed to send Home Registration to the HA. This message will be refereed to as the Non-local Registration Request Solicitation.

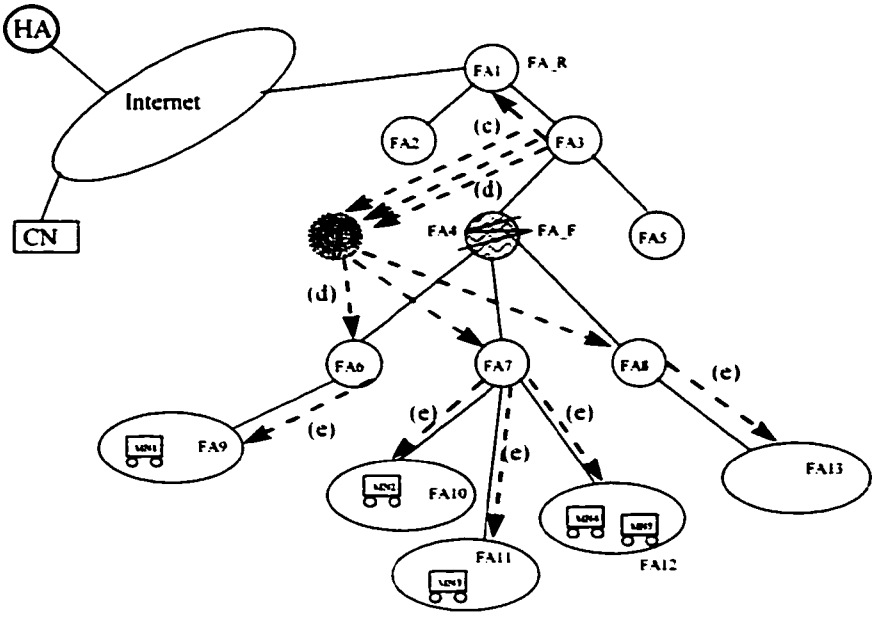

Figure 6.3 FA failure recovery using the Revert to HA registration mode

Considering a system as the one shown in Fig 6.3, the mechanism can be illustrated as follows:

a. Consider the faulty FA (FA_F) and the FA in the higher hierarchical level (FA_H). FA_H will detect the failure of FA_F. (FA_F=FA4, FA_H=FA3)
b. FA_H will construct a list (MN_List) of those MNs affected by this failure. (MN_List=MN1, MN2, MN3, MN4 and MN5)
c. FA_H will contact the Hierarchy Registry to get a list (FA_List) of the FAs directly connected to the faulty FA. FA_List (FA_List=FA6, FA7 and FA8)
d. For each FA in FA_List, the FA_H will send the MN_List.
e. Upon receiving the MN_List, each FA will perform one of two possible actions. If none of the entry in the MN_List exists in the visitor list of the FA, no further processing is needed and the MN_List is disregarded. If one or more of the MNs included in the MN_List is found in the visitor list of the FA, the list will be forwarded to the next lower FAs in the hierarchy (as indicated in the visitor list).
f. On receiving the Non-local Registration Request Solicitation, each of the affected MNs will send a non regional Registration Request to its HA. When this request is received by the HA, the HA will update its table to register the serving FA as the current FA in place of the gateway FA.
g. The MN will need to send a Registration Request with the Previous Foreign Agent Notification option. This will have the effect of removing the binding information for the affected MN on the gateway FA. When route optimization is implemented the CN will have a binding information pointing to the gateway FA as the current FA serving the MN. On receiving packets from the CNs, the root FA will generate a Binding Warning to the HA.
h. It is expected that the faulty FA will come back in service after the time needed for repair and for failure recovery. The FA_H may regularly examine the status of the FA_F. When FA_F recovery is detected, the FA_H may send a message of the type "Local Registration Solicitation" directed towards the affected MNs announcing that those MNs can start using local registration. The issuance of this message can be delayed using a timer if it is required to keep the MNs for extra time in the Non HLRM-IP mode.

The number of control messages, needed to restore the service, affects the cost associated with failure recovery. Also, the delay coupled with the delivery of retransmitted datagrams affects the efficiency of the recovery scheme. The recovery cost upon a single failure event can be expressed as

$$C_t = C_H + C_Q + C_I + C_{IMN} + C_{HA\,Reg} \tag{6.1}$$

Where:

C_H is the cost of sending one Hello message

C_Q is the cost of one query session between a FA and HIR Registry

C_I is the cost of the FA_H contacting the FA members in the FA_List

C_{IMN} is the cost of propagating the Non-local Registration Request Solicitation to the MNs

C_{HAReg} is the cost of sending registration requests and replies between the MNs and the corresponding HAs.

The recovery cost, associated with a period of time T seconds, can be expressed as follows:

$$C_T = \sum_{All\_FA\_H} T.R_H.C_H \quad + \quad \sum_{FA_i} \int_{t=0}^{T} P_{f(FA_i,t)}.(C_Q + C_I + C_{IMN} + C_{HA\,Reg})dt \tag{6.2}$$

FA considered in 6.2 are those supporting the local registration mode of operation at the time of the failure. R_H represents the rate by which FA_H will generate the downstream Hello messages. The other cost components can be expressed as follows, where $D_{x\text{-}y}$ is the propagation delay between elements x and element y.

$$C_H = 2.(D_{FA\_H-FA\_F}) \tag{6.3}$$

Expression in (6.3) does not account for the time period when the FA_F is not responding to Hello messages. Under normal operation condition is number should be small.

$$C_Q = 2.(D_{FA\_F-FA\_HIR\_Registry}) \tag{6.4}$$

$$C_I = \sum_{FA\_List\_members(FA_\tau)} D_{FA\_H-FA_\tau} \tag{6.5}$$

$$C_{IMN} = \sum_{FA\_List\_members(FA_\tau)} \sum_{MN\_Interfaces} D_{FAx-MN} \tag{6.6}$$

D_{FAx-MN} corresponds to the delay needed to deliver the Non-local Registration Request Solicitation to the MNs. The aggregated delay takes into consideration the delay encounters by the individual solicitation copies generated at branching points.

$$C_{HA\,Reg} = 2.\ (\sum_{MN\_Revertingback\_to\_HA\_Reg} D_{MN-HA(MN)}\ )\ D_{MN-HA(MN)}\ ) \tag{6.7}$$

6.2.3 FA Fault Tolerance: The Self-Healing Mode

This approach may be used when reverting back to Non HLRM-IP is not preferred, which may be the case if the relatively higher delay associated with non-local registration can not be tolerated. The basic idea of this solution is to heal the breakage in the hierarchy tree caused by the faulty FA. This can be accomplished by bypassing this faulty FA such that the FA in the hierarchical level just above the faulty FA will remove the faulty FA from his copy of the hierarchy, and consider the FA in the level just below the faulty FA as its tunnel end. The same steps will be repeated for each FA connected to the faulty FA.

In an environment as the battlefield, where the probability of FAs failure is relatively high and location dependent, it is expected that other surrounding FA in the hierarchy

will be subjected to similar attack. The higher FA may point to further lower FA and not necessarily a FA in the level just below the faulty FA. Steps involved in supporting this mechanism are as follows considering figure 6.4:

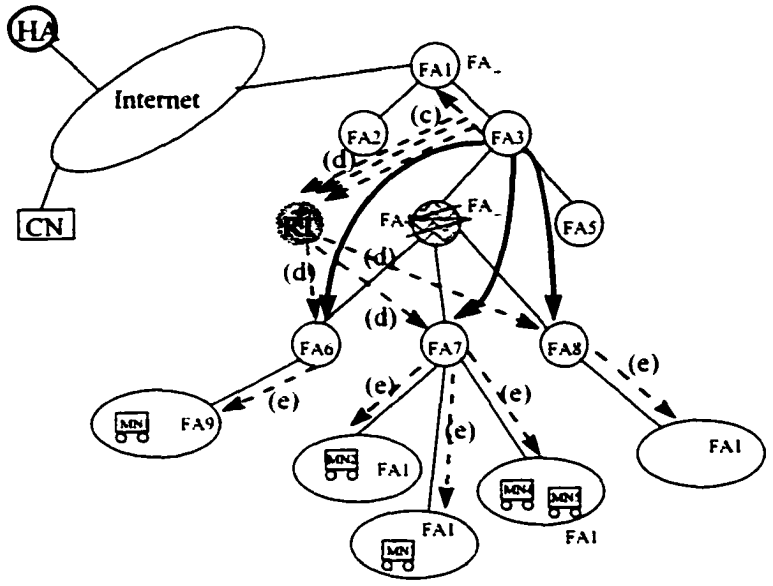

Figure 6.4 FA failure recovery using the Self-Healing mode

a) The FA_H detects the failure of the FA below (FA_H=FA3, FA_F=FA4)
b) FA_H will construct the MN_List of those MNs affected by this failure. (MN_List=MN1, MN2, MN3 and MN4)
c) FA_H will contact the Hierarchy Registry to get a list (FA_List) of the FAs directly connected to the faulty FA (FA_List=FA6, FA7 and FA8)
d) FA_H will send the message `FA_Change_HIR` to members of the FA_List. This message carry the list of MNs that FA_H has binding for. On receiving this message, each member of the FA_List will change its hierarchy information such that FA_H will become the FA in place of the faulty FA. In addition, a `FA_Change_Hir_Confirm` message will be sent from each member of the FA_List containing information about which MNs it has binding for such that the FA_H can update its binding and forward packets to the correct FA.
e) FA_Ls propagate the message downwards. FAs will need to know the new hierarchy to be used when announcing their local registration support.
f) The same information is sent to the Hierarchy Registry.

In some situation, or if desired, only one FA from the FA_List can be used to heal the hierarchy. For example, if connectivity exists between FA6 and FA7 and between FA6 and FA8 then FA7 and FA8 can be attached to the hierarchy through FA6.

The recovery cost can be expressed as follows:

$$C_t = C_H + C_Q + C_I + C_{IFA\_Leaf} \tag{6.8}$$

$$C_T = \sum_{All\_FA\_H} T.R_H.C_H + \sum_{FA_i} \int_{t=0}^{T} P_{f(FAi,t)}.(C_Q + C_I + C_{IFA\_Leaf})dt \tag{6.9}$$

C_Q and C_I are calculated as in (6.4) and (6.5), while C_{IFA_Leaf} can be expressed as:

$$C_{IFA\_Leaf} = \sum_{FA\_List\_members(FA_t)} \sum_{MN\_Interfaces} D_{FAx-FA\_leaf} \tag{6.10}$$

An observation regarding the expression in (6.8) is associated with the C_{IFA_Leaf} cost component. The `FA_Change_HIR message` represent a signaling load, but is not affecting the service recovery speed. The service can be restored as soon as the message is intercepted by the FAs in the FA_List. Propagating the message downstream all the way to the leaf FAs is included for the topology correctness of the base local registration Mobile-IP.

6.2.4 Comparison and Performance Aspects for the Fault Tolerance System

In the previous subsections we have demonstrated two approaches to tolerate the FA failure in Hierarchical Local Registration M-IP systems. In the following we will

highlight the features associated with each approach. We will refer to the Revert to Non HLRM-IP mode as approach 'A' and the Self-Healing as approach 'B':

Simulation Environment

The simulation environment is composed of 20 MNs moving in a regional domain. The hierarchy has eight levels, where the FAs supporting M-IP at the top two levels were equipped with hot redundant FAs to minimize the impact, in the event of a failure, on the elements of the lower levels. The delay over a link between two FA on the hierarchy is set to one msec. The delay between any FA and a routing element is set to 2 msec, and the delay over wireless link is set to 2.5 msec. The delay between the HA and the root FA is set to 4 msec. The delay figures used here represent an aggregation for all delay components such as propagation and processing. Different values for the delay values should not affect the correctness of the approach. The only restriction is that the delay from the MN to the HA should be larger than that from the MN to the root FA, which is the expected topology and is the primary motivation behind the local registration. A unidirectional uniform traffic is generated from two CNs to the different MNs. The emulated failure pattern corresponds to multiple FA failures in sequence on different levels of the hierarchy. The failure rate is Gaussian distributed. For the Revert mechanism, the system is allowed to remain in this mode for 0.5 minutes after the last failure before returning to the local registration mode. Simulation time of 1500 minutes was considered, with randomly distributed mobility rate between 1 and 5 moves/minute. The case of the Base M-IP is considered in this simulation. The decision to consider

either mode is based on number of MNs supported by the faulty FA. If this number is not greater than a threshold, of a value equals to two in this simulation, the revert mode is considered. In addition, the FAs were made to gain some knowledge about the expected failure pattern such that the failure on two or more FAs on the same vertical lineage will be considered a triggering event to use the revert mode.

1. Number of packets dropped under high probability of failure of multiple FAs on adjacent levels

It is expected that the FAs failure pattern will be location dependent in a battlefield environment. For example, if a FA housed in a mobility support station was hit in an attack, it is expected that FAs on higher and lower adjacent levels will be exposed to similar attacks. Approaches A and B behalf differently under such conditions.

Approach A:

When in the Non HLRM-IP mode, MNs will not be affected by the failure of FAs on adjacent hierarchical levels as long as another route exists to forward packets from the HA to the current FA. When such multiple failures are expected, it is more efficient for the MNs to stay in Non HLRM-IP mode. System can return to the HLRM mode when the probability of multiple failures is low, or after the expiration of a configured timer. This can be implemented using a timer function on individual MNs and FAs. On the other hand, the service is only recovered when the registration request is processed by the HA. This cause relatively larger number of packets to be dropped.

Approach B:

This approach is sensitive to multiple level failures. Each failure event will cause the control messages associated with the recovery mechanism to be generated and processed such that multiple failures will be associated with more control messages and more service loss time intervals. Figure 6.5 presents such observation.

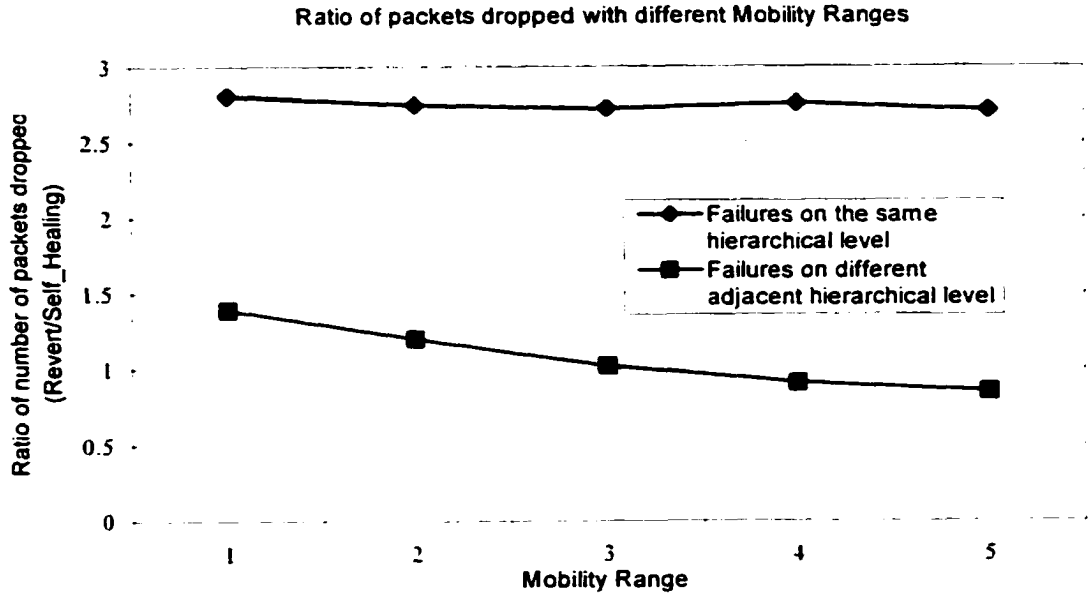

Figure 6.5 Ratio of packets dropped over different mobility ranges with different failure patterns

2. Number of control messages

Approach A will generally produce more control messages, due to forcing affected MNs to use home registration. The revert mode requires individual MNs to issue home agent registration request. The larger the number of affected MNs, the larger the number of control messages needed. Figure 6.6 illustrates this behavior. Only when very small

number of MNs exist in the system, then the revert mode will outperform the self-haling scheme.

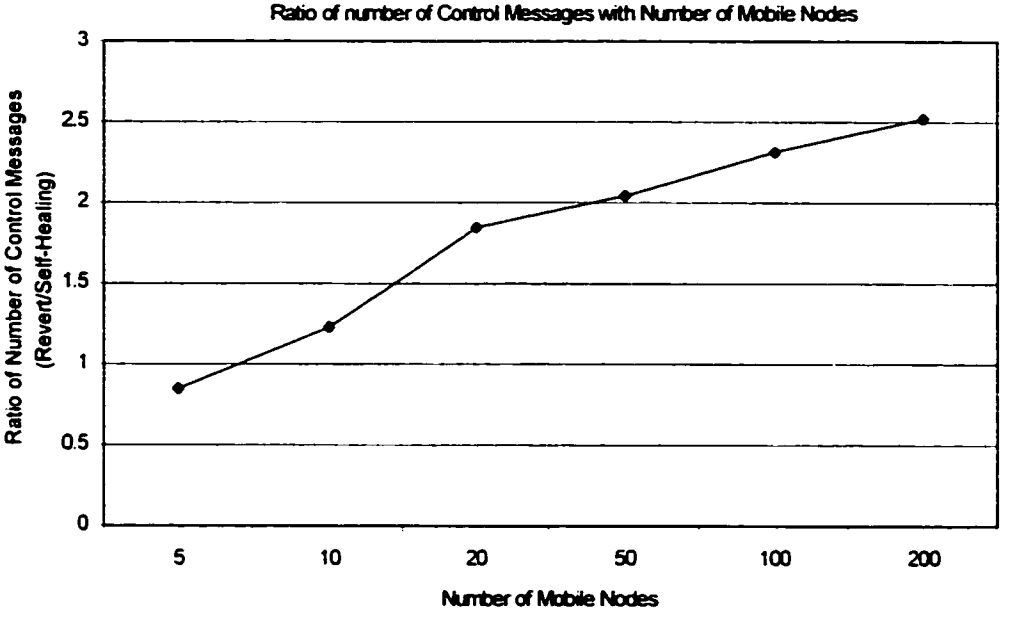

Figure 6.6 Ratio of number of control messages over different number of MNs

3. Number of packets dropped with different topologies

Approach A is not sensitive to changes in topologies. On the other hand, the self-healing approach provides a performance that differs with the considered topology. The topology in HIR-1 is less sensitive to the failure of FAs than that of HIR-2. The failure of one FA in figure 6.8 (HIR-2) can deny the service for a large number of FAs. Such differences in topology may suggest the applicability of one recovery scheme over another. For example, a large number of FAs on a lower level directly connected to an upstream FA, may suggest using the revert approach instead of the self healing to avoid the need to reattach back the considerable large number of FAs on the lower level to the hierarchy.

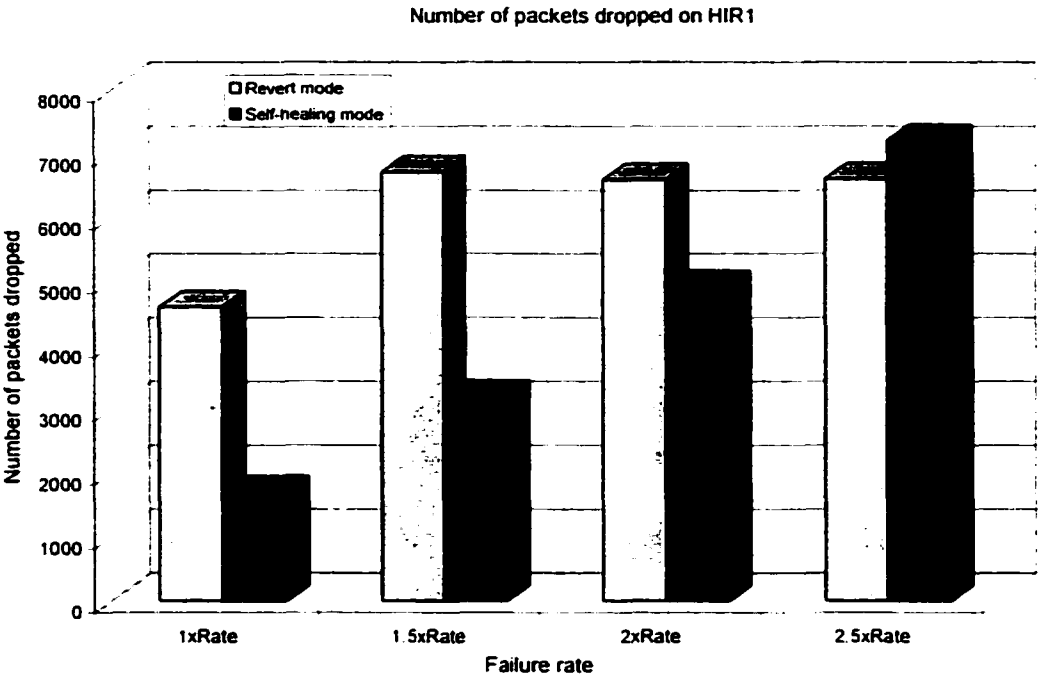

Figure 6.7 Number of packets dropped with different rates on hierarchy 1.

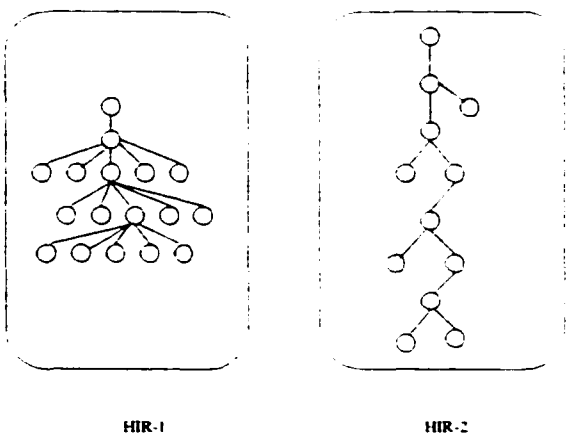

Figure 6.8 Topologies for HIR-1 and HIR-2

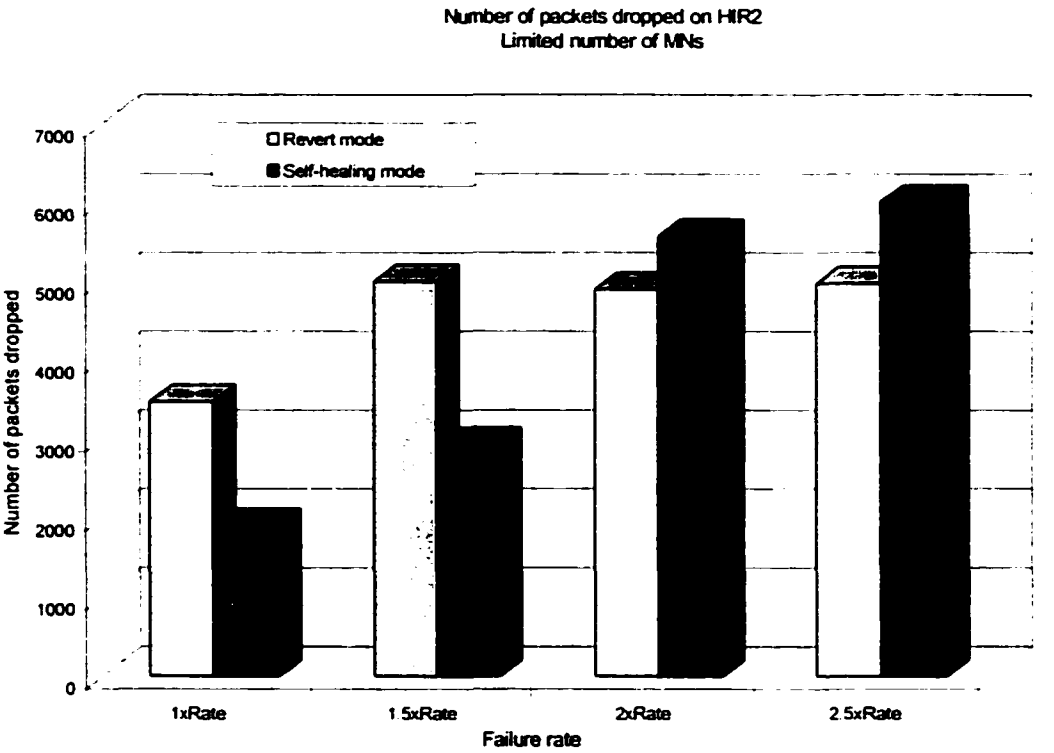

Figure 6.9 Number of packets dropped with different rates on hierarchy 2

The number of packets dropped is shown in figures 6.7 and 6.9 for the two topologies in figure 6.8. In Figure 6.9, we can see that the self-healing scheme provides better performance than the revert mode for most of the range of the failure rate. When the rate failure is relatively large, then more packets will be dropped (in both HIR-1 and HIR –2).

4. Transparency of the recovery mechanism to the MNs

Approach A: Not transparent since MNs will need to receive and process the Non-local Registration Request Solicitation.

Approach B: Transparent to MNs since the `FA_Change_Hir` messages are destined to the FAs and not the MNs.

The Adaptive system:

From the above discussion, it is clear that both recovery schemes have their characteristics and perform differently, according to the network topology and the associated dynamics. For a network to optimize the recovery performance, it may be necessary to be able to implement both schemes and to adapt the mode of operation accordingly. The adaptive mode was implemented such that when failures happen on the two levels next to the very top levels supporting the hot redundancy, the self-healing mode will be used to avoid forcing the MNs downstream to use the revert mode. On the remaining lower level, the decision to consider either mode is based on number of MNs supported by the faulty FA. If this number is not greater than a threshold, then the Revert mode is considered. Figure 6.10 shows the benefits gained from applying an adaptive approach. With lower failure rates, the adaptive mode follows closely the performance of the self-healing mode. Since higher failure rate will cause the self-healing mode to produce a larger amount of signaling messages, the adaptive mode will converge to a performance similar to that of the revert mode.

A better performance can be gained if additional information can be made available to the adaptive system. An example is information regarding the expected failure rate and the expected number of MNs. Such information, if made available to the system in advance, can be used to direct the system towards the appropriate mode of operation. It is

important also to keep this information updated. Information available in the regional registration data structure is already available to provide this parameter. This is true since a FA has information regarding the MNs population in the downstream direction. A simple aggregation function can be used to keep track of this number considering the number of current registration instances.

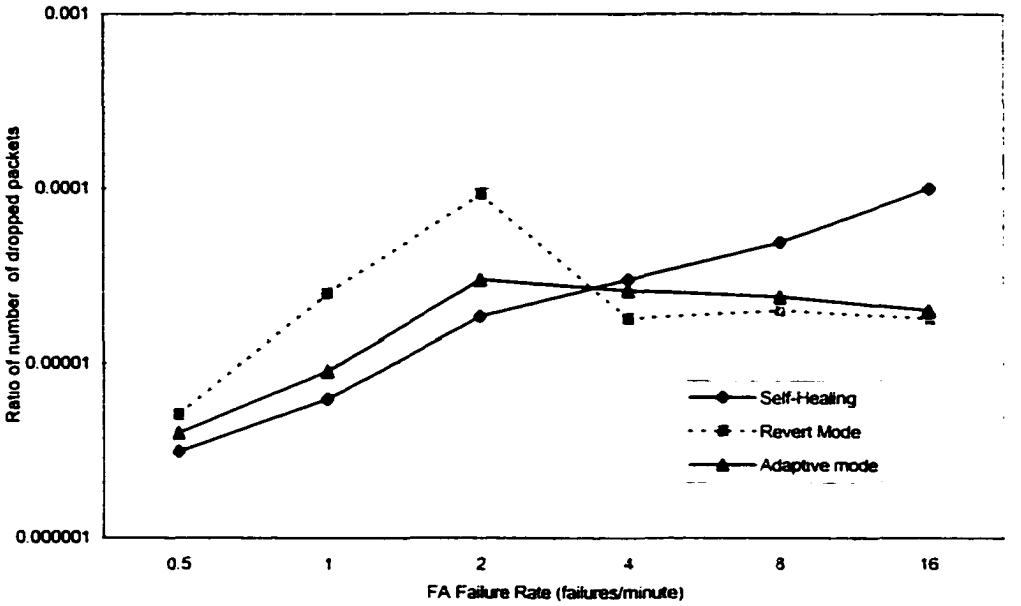

Figure 6.10 Ratio of number of dropped packets considering the three different approaches

7. QoS Differentiation in the Regional Registration Model

The Differentiated Services (DiffServ) approach [52] promises a scalable solution for the QoS support over the Internet. Extending QoS support to mobile users is a challenging task, considering the dynamic nature of mobility. Flexible, robust and efficient schemes are needed to overcome the additional challenge associated with supporting additional features such as multicast and reliable multicast support.

In this work, we analyze the problem of supporting DiffServ in the regional registration Mobile-IP environment. Different Mobile-IP configurations such as route optimization [7] and bidirectional-tunneling [8] are considered. We address the additional complexity associated with supporting a multicast service, such as Differentiated Service Code Point (DSCP) [52] mapping and resource management signaling. We propose a protocol that allows the efficient application of the DiffServ approach, in this environment, with support for both unicast and multicast traffic. In our proposed scheme, existing elements and control messages supporting the Mobile-IP are utilized to support the QoS service, with minimum additional overhead. We show how this integrated platform can be extended to support added features such as reliable multicast.

To support QoS over IP networks, two architectures have been suggested. The Integrated Services architecture (IntServ) [53] uses a signaling mechanism to allocate per-flow resources, between the datagram source and the corresponding destination, through the network elements. Due do the maintenance overhead associated with

maintaining the per-flow states, the IntServ model does not provide attractive scalability features and thus limiting its implementation to medium size networks. The DiffServ [52] model is positioned as a simple and scalable solution that pushes the complexity of per flow processing to the edges of the IP networks. In the core of the DiffServ network, the differentiated service is provided per traffic aggregate.

The DiffServ architecture was designed with no particular support for mobility neither for multicast. In a mobile environment, resources associated with mobile nodes need to be allocated and released according to the mobility of the hosts. In the Mobile-IP platform, a home agent and multiple foreign agents will support the forwarding of the traffic between the mobile node and the corresponding node. Accordingly, multiple DiffServ domains will need to be considered in the provisioning process. In addition to the base Mobile-IP case, the route optimization and the bi-directional tunnel schemes need to be supported in the proposed platform.

Supporting multicast groups with constantly changing memberships and different classes of service present additional challenges to the QoS support scheme. If the DSCP (Differentiated Service Code Point) value is copied from the original datagram to the corresponding generated copies at a multicast branching point, several problems can affect the performance of the system [54].

In this work we address the issues related to the support of DiffServ in the regional registration Mobile-IP environment. In particular, we focus on issues associated with the control messages needed for resource management in the DiffServ domains traversed by datagrams. Those domains include the local regional domain, the home agent domain and

other intermediate domains. In addition, we describe an approach to accommodate mobility while forwarding both unicast and multicast datagrams while providing the contracted QoS. We propose a protocol that capitalizes on the inherent features of the local registration scheme to integrate the QoS support with the existing schemes to support location information binding and multicast group management.

7.1 QoS and Mobility Support

Proposed network architectures to support different levels of services to IP traffic in an environment that support mobility, can be classified into two main categories. The first category relies on the IntServ model [52], where specific resources are allocated on a per-flow basis using the RSVP signaling protocol. An example of such category is the MRSVP [55], which is an extension of RSVP to handle mobile users. The mobility independent mode in MRSVP provides two QoS levels of guarantee, dependent and independent of the mobility of the hosts, using both active and passive reservation. Another example of the reservation approach is the INSIGNIA [56], which is an in-band signaling protocol that supports QoS in mobile ad hoc networks. INSIGNIA relies on encoding command, using the IP option field, to support fast reservation and restoration.

In [57], the ITSUMO, a QoS architecture framework based on differentiated services is proposed. The ITSUMO model is characterized by a central server (QoS global server) which has global information of the administration domain, and local nodes presenting the ingress node of the DiffServ domain. The QoS global Server, based on the mobility

pattern and local resources information provided by the local nodes, allocate resources based on the required level of guarantee. Issues associated with tunneling and DiffServ are discussed in [59], while QoS support requirements are illustrated in [60]. A framework for QoS and mobility is presented in [61]. An analysis for the DiffServ approach in the mobile environment is illustrated in [62]. A scheme for receiver control in DiffServ is presented in [73].

Considering that the regional registration environment provides both attractive scalability features and reduced registration overhead, it is positioned as a possible deployment candidate. In this work, the DiffServ support is designed with the intention to take advantage of the inherent features and the existing control messages of the regional registration environment, providing minimum signaling with different Mobile-IP configuration. A scheme using temporary allocation is used to complement the resource manager based approach, providing faster response to QoS requests generated after mobility events. Our proposed scheme accommodates the encapsulation decapsulation process needed for datagram forwarding in the regional registration domain. The Support for DiffServ, for both unicast and multicast traffic in an integrated platform, is another feature that differentiates this proposed platform from the other proposals mentioned above.

7.2 Differentiated Services

The DiffServ (DS) architecture is based on a simple model where traffic entering a network is classified and possibly conditioned at the boundaries of the network. DiffServ

treatment of datagrams relies on the forwarding behaviors of packets, Per Hop Behaviors (PHBs), where each behavior aggregate is identified by a single DS codepoint. Based on network policies, different types of traffic can be marked for different forwarding priorities, and resources can be assigned according to the marking and the applicable policies.

The architecture in [52] describes the functions and elements supporting the QoS model in a stationary environment. In the following, we will point to those elements and how mobility affects the corresponding functions.

Differentiated domain: A contiguous set of DiffServ nodes with a well-defined boundary. In the local registration environment, a DiffServ domain will correspond to the domain served by a GFA.

DS Boundary node: A DS node that connects one DS domain to a node in another DS domain. The GFA can act as the boundary node for unicast traffic destined for MNs, while GFA or intermediate FAs may act as boundary nodes to support multicast traffic depending on the multicast operation mode.

Interior and edge nodes: When a MN moves to a FA, this FA becomes an edge node for traffic originated from or destined to this MN. When forwarding traffic received from the upstream or the downstream directions, the FA acts as an interior node.

Packet classification: For traffic originating from a MN, the current FA can act as the classifier. For ingress traffic, the GFA may classify and remark unicast traffic even if it has been classified earlier in the source DS domain.

Shaping/dropping: Those functions are needed to control non-conforming traffic, and can be implemented by the FA in the local registration domain.

Resource Manager (RM): The resource manager provides a function similar to the Bandwidth Broker (BB) [58]. The RM main functions are to set up the routers at the edge of its domain, and to support the control of resource management across boundaries to adjacent regions. Considering the regional registration model, and the mobile nodes moving among all FAs in the hierarchy, all FAs may appear as leaf (edge) routers.

Service Level Agreement (SLA): A service contract between a customer and a service provider that specifies the forwarding service a customer should receive. To support the different possible configurations of Mobile-IP, the agreement will include other information associated with the mobility support. For example, if it is required to have QoS service in both directions between the MN and the other host involved in the session. The agreement also may indicate if resource allocation on only part of the path between the source and destination is acceptable.

7.3 Requirements for Efficient DiffServ Support in the Mobile-IP Environment

An environment that supports mobility imposes additional requirements on the differentiated service approach. The DiffServ was designed with no support for the mobile environment. To provide such support, the following points have to be

considered, while additional requirements which are specific to multicast support will be mentioned in the corresponding section.

- The DiffServ platform shall support hosts mobility, such that the same level of QoS can be provided in the different service areas visited by the MN.
- The platform shall support dynamic service level agreement (SLA) renegotiations such that it can react to the dynamics associated with mobility. The available applications and the corresponding requested QoS are expected to vary with the MN's location. In addition, the MN may prefer to use one wireless access technology over another depending on the visited network, triggering the request of different QoS.
- The system should accommodate for packet encapsulation in its classification process, such that packets are marked and processed according to the intended DSCP.
- The platform shall support MN initiated QoS request, for both upstream and downstream directions relative to the MN.
- The DiffServ platform should capitalize on the characteristics associated with the regional registration environment, providing minimum interruption in QoS support at the time of handover.

7.4 DiffServ and Unicast Traffic

To support the resource management function associated with the DiffServ model, we introduce the Domain Resource Manager (DRM) function. Since FAs in a particular regional domain are served by the same GFA, a reasonable assumption is that the FAs in a local domain will be associated with the same DiffServ domain. The DRM is responsible of intercepting and processing QoS requests from the MNs and the FAs, and keeps track of the current available resources in the domain served by the GFA. The fault monitoring function and the HIR Registry described in [50], along with the DRM, can be co-hosted on the same network element. The DRM function is similar to that of the Bandwidth Broker [58]. In this work, we will consider the case where the DRM function is co-located with the GFA. The following cases describe the operation of the proposed protocol, to support end-to-end DiffServ, under different configurations for the Mobile-IP and different service specifications. The base Mobile-IP, the route optimization extension and the bi-directional tunneling are supported. In addition, the MN can specify QoS treatment for one direction only or for both incoming and outgoing unicast traffic.

When the MN first moves to the domain serviced by GFA, it will send local registration request to the current FA. When a MN receives a packet from a CN at the start of a session that requires QoS, it will send a receiver-initiated QoS request message. This message contains additional information such as the MN ID, the HA address and the intended session duration. The QoS request can be embedded in future local registration requests, instead of dedicating a separate message for this purpose. This QoS message will cause the DRM, associated with the domain where the current FA is located (DRM-CFA), to request a copy of the SLA from the MN's HA. The acquired

specifications will remain stored in the DRM-CFA , such that no further communications with the HA is needed for future QoS requests from this particular MN.

The following cases illustrate the resource allocation activities, considering different service configurations. For example, a Mobile-IP platform may specify using tunneling only for incoming traffic relative to the MN. Another configuration, designed to accommodate firewall security rules [71], may require bi-directional tunneling. In some situation, QoS support may be required only in one direction between the CN and the MN. Figure 7.1 illustrates the different domains associated with traffic forwarding in the regional registration environment.

Case 1: Tunneling for traffic to MN, regular routing for traffic from MN, QoS is required in both directions.

Resource allocation upstream (relative to MN):

When the DRM-CFA receives the QoS request from a MN, it will start the process of allocating resources along the path from the CN to the MN. This process includes the following sub-processes as follows:

A. Resource allocation from the GFA to the MN

The DRM-CFA will verify the availability of resources from the GFA to CFA. The RM-CFA keeps a database that contains information regarding the physical link capacities, and the current bandwidth assignment on the different links of the regional hierarchy. Upon receiving and verifying resource availability, the RM-CFA will update the

corresponding entry in the database considering the newly allocated resources. As an example, consider that current allocation for the FA1-FA3 and FA3-FA5 links (shown in figure 7.1) was 1.6 Mbps and 0.9 Mbps respectively. When FA5 receives a QoS request from a newly arrived MN requesting class 1 treatment for its 0.1 Mbps ingress traffic, the allocation needs to be updated according to this new request. RM-CFA will increase the allocation to 1.7 and 1 Mbps, provided that the new request will not exceed the total assigned allocation for this class.

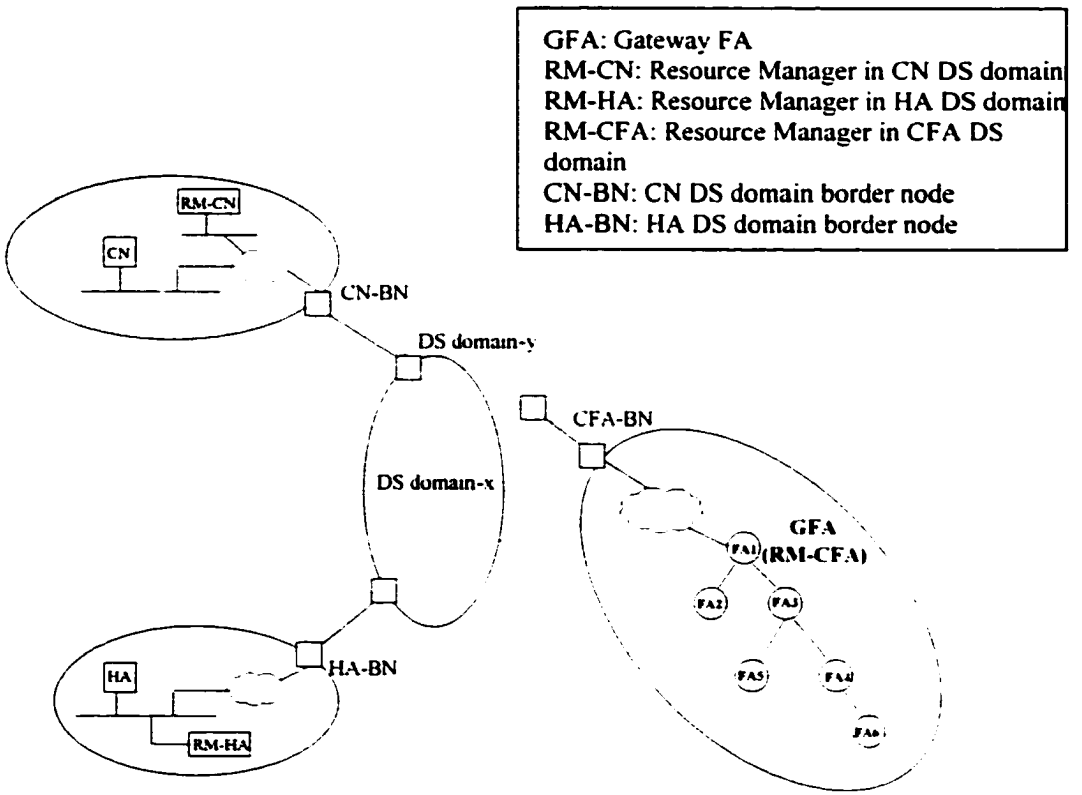

Figure 7.1 DS domains in the regional registration environment

B. Resource allocation from the HA to the GFA, and from the CN to the HA

At the same time when the DRM-CFA is verifying the local resources in its domain, it will forward QoS requests to the HA and the CN. On receiving the QoS request, the HA will forward it to its RM-HA, while the CN will forward it to the RM-CN. The RM-HA will try to allocate resources from the HA to the GFA, while RM-CN will verify resources from the CN to the HA network. The RM-CN and RM-HA may need to contact other RMs to complete the allocation to the respective destinations. The success or failure of the allocation on those two segments will be compiled by the DRM-CFA.

On receiving the responses for the QoS, the DRM-CFA will select an action according to the MN's Service Level Agreement (SLA). An SLA may specify that only end to end resource allocation is accepted, while another MN may have an SLA that permits partial allocation. If the DRM-CFA decides to continue with the allocation process, it will then set up a profile meter on the border router in the local domain. The RM-CN and RM-HA will start the allocation process also in their respective domains. When needed, the DRMs will use an underlying inter-domain signaling protocol, to contact the RMs in the adjacent domains to arrange for cross boundaries resource allocation.

C. Resource allocation in the MN to CH direction

In this mode of operation, the data traffic from the CFA will follow the optimal route to the CN. The DRM-CFA will refer to its local database to verify that sufficient resources exist from the CFA to the corresponding border router on the path to the CN. If resources exist in all intermediate domains, then the corresponding profile meter will be placed on

the CFA. Using similar steps as described earlier, in the opposite direction allocation, the corresponding profile meter will be placed on the appropriate border routers across the domains leading to the CN network. The operation of this procedure, as applied to case 1 above, can be illustrated in the diagram in figure 7.2.

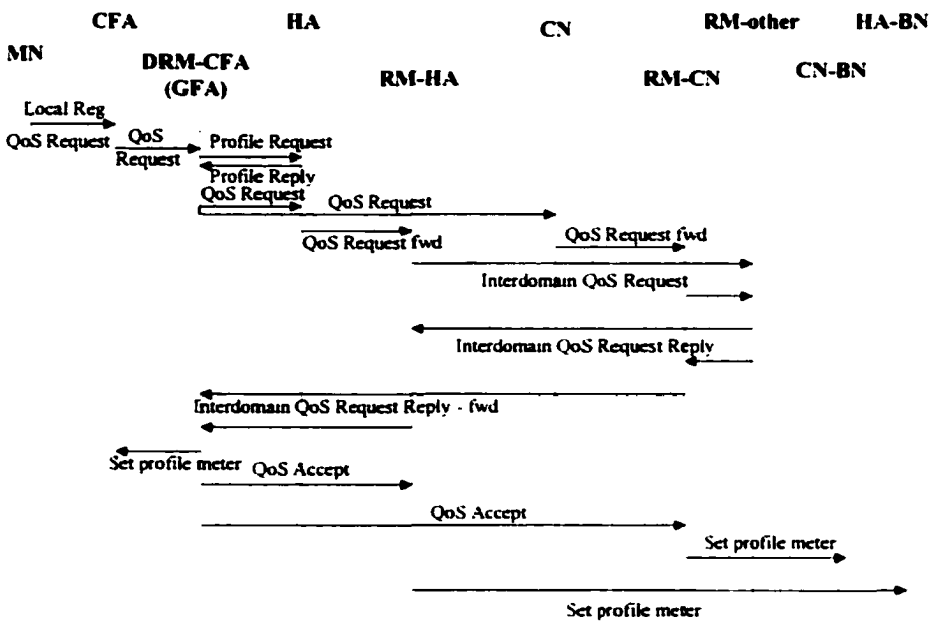

Figure 7.2 Control messages associated with case 1

Case 2: Unicast, tunneling for traffic to MN, tunneling for traffic from MN, QoS is required in both directions.

This case can be handled similarly to the procedures in case 1, except when allocating resources from the MN to the CN. In this case, the DRM-CFA will verify resources from

the CFA to the border router on the path to the HA. Also, the DRM-HA will verify resources to the CN network.

Case 3: Route optimization, QoS is required in both directions

An enhanced CN, implementing the route optimization extension, will tunnel datagram directly to the current FA. In this case, the DRM-CFA will forward the QoS request to the CN. Accordingly, the request will trigger the RM-CN to allocate the resources from the CN to the GFA network. Traffic from MN can be handled as in case 1.

7.5 Packet Marking and Tunneling Accommodation

As described in the DiffServ specifications [52], a datagram may be marked by the datagram source, or by any of the intermediate node, or by the egress node in the source domain. As illustrated in figure 7.1, a datagram will cross different DS domains in its trip between the CN and the MN.

When a HA or an enhanced CN is tunneling datagrams to the GFA, datagrams are encapsulated with a new header. Several methods of encapsulation are available for use by the HA and the FA, and also by CN when using route optimization. IP-in-IP encapsulation [15], Minimal encapsulation [16], and GRE [17] are examples for available schemes.

When using tunneling, information regarding the original datagram header will not be available for inspection in the outside header. For example, a HA serving three MNs that

are residing in the same CFA, will encapsulate datagrams with the same external header. If the HA-BN is to mark the traffic for different QoS treatment, it will not be able to differentiate between the individual MNs sessions by only inspecting the external header. Instead, it will need to inspect the internal header. This can lead to unacceptable effect on the performance. The same observation is applicable on traffic tunneled by the CN when route optimization is used. In the following discussion, we will assume that the CN and the MN (and or the current FA) have the capability to mark the generated traffic.

The DiffServ architecture defines the marker, as the node setting the DS codepoint in a packet based on a defined rule. The encapsulation process involves adding an extra header at the entry point of the tunnel, while this header is removed at the tunnel exit point during the decapsulation process. The DSCP value in the encapsulated/decapsulated datagram should be treated carefully to assure appropriate traffic treatment. Accordingly, the particular Mobile-IP configuration plays a part in deciding the preferred node to do marking. To study the effect of encapsulation on the DSCP processing, we will consider datagrams forwarding on the individual domains mentioned in the previous cases.

In case 1 mentioned earlier, the datagram forwarding path from the CN to the CFA can be partitioned into three segments. No tunneling takes place over the CN-HA segment. Accordingly, the CN can assign the appropriate DSCP value for this traffic. This traffic will be intercepted by the HA-BN, where it is metered and marked accordingly. On the segment from the HA to the GFA, the HA tunnels the datagrams received from the CN to the GFA. Since the tunneling process involves adding a new header to the received datagram, the HA may as well map the DSCP value in the inner header1 to the same or to

a new value in the DSCP field of the external header. The last segment, from the GFA to the CFA, involves a decapsulation/reencapsulation process at each level of the hierarchy. It is not necessary to implement actual decapsulation, and simple header swapping may be used instead. In either case, the FA needs to ensure that datagrams forwarded downstream have the appropriate DSCP values.

When route optimization is used, as in case 3, the CN will encapsulate its datagrams destined to the MN. The internal header will have the MN as the destination, while the destination IP in the external header will be that of the GFA. The CN will set the appropriate DSCP in the external header.

7.6 Handling Mobility

In a typical DiffServ model, the MN's traffic will not be marked and processed according to its SLA till the CFA receives a control message from the DRM-CFA. This message will setup the appropriate profile meters and markers in the domain.

When a MN moves to a new FA in the same local registration domain, it is highly probable that the common ancestor is one or few hops away. Accordingly, if the MN is interested in keeping the same QoS in its new location, resources need to be allocated only in the path between the common and the new FAs. The MN's traffic is not guaranteed to be serviced as contracted, until the control messages from the DRM-CFA are received by the border routers and the new CFA. Holding the MN from experiencing the contracted QoS, while waiting for the DRM-CFA acceptance message, may not be

efficient in this situation. To provide a better service to the MNs, we propose the following scheme.

When sending a local registration request to the common FA, the MN will append it with the QoS request extension. When the common FA intercepts a QoS request extension, it will forward it to the DRM-CFA for regular processing. At the same time, the common FA does not wait for the DRM-CFA's response. Instead, it spawns another process that aims at providing a temporary QoS allocation that matches as possible the MN's SLA. We will refer to the process of allocating temporary resources, to provide DiffServ service to the mobile till resources are allocated by the DRM-CFA, as the temporary allocation. Temporary allocation has the objective of allocating required resources, temporarily, on the segment of the hierarchy that is common between the previous and new lineages.

Temporary allocation support does not necessitate new messages, instead it relies completely on the local registration messages. When the MN arrives to the new FA and generates the local registration request with the embedded QoS report, the request message will be extended to carry resource availability indication. On its trip upward, each FA will verify its local resources, and if resources exist, the extended registration request will be marked accordingly. When the common FA receives the local registration request, it will be processed normally except that the registration reply will be used also for resource allocation downstream. The operation of this scheme can be illustrated using the following example, considering figure 7.3.

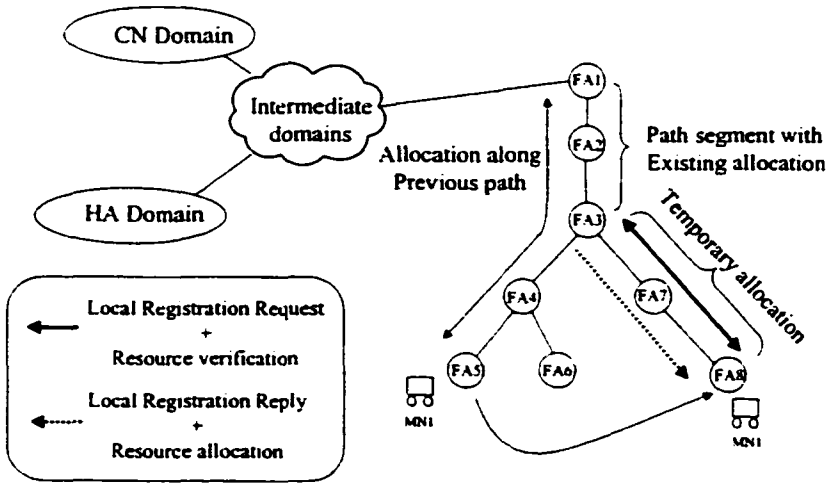

Figure 7.3 Resource allocation upon mobility

Assume that MN1 is residing at FA5. MN1 then decides to communicate with a CN1 in the CN DS domain. MN1 specifies 20 Kbps of traffic in a bi-directional tunnel mode, with service class 1. The QoS request will propagate eventually to the DRM-CFA and corresponding resources will be allocated, including those in the regional domain. When MN1 moves to FA8, it will send a local registration request destined to FA3. Intercepting this request, Intermediate FAs (FA8, FA7 and FA3) will verify if resources are available to satisfy this request. For example, if FA7 egress interface facing FA3 is configured with static maximum allocation of 2 Mbps for class 1, then the 20 Kbps additional requests can be satisfied if current allocation is 1.5Mbps. On receiving the registration request and the embedded information, FA3 will decide to accept or deny the request. If accepted, the resources will be allocated only on a temporary basis till strong allocation message is received from the DRM-CFA. Considering that new strong allocation can be processed

and accepted for other traffic during the lifetime of other temporary allocations, traffic accepted based on temporary allocation is marked differently from traffic with strong resource allocation. This arrangement is needed to ensure that service provided to traffic associated with strong allocation is not degraded because of a competition from the traffic with the temporary allocation.

Differentiated treatment between the temporary and strong allocation types of traffic can be accomplished using a queuing mechanism similar to class based queuing (CBQ) and a congestion avoidance mechanism such as Random Early Discard (RED). Figure 7.4 illustrates that approach. Initially, we will assume a system that used DSCP values from 4 to 7 to represent traffic with strong QoS allocation. Traffic associated with temporary allocation is marked with DSCP from 0 to 3. This will allow 1:1 correspondence between the two types of traffic for up to 4 streams. On the FA, traffic is marked and directed to queue 1 where a RED scheme is applied. The RED objective is to prioritize DSCP value of 4 over the value of 0, if congestion is occurring on the egress interface. This prioritization is accomplished by assigning a higher drop priority to the datagrams with the DSCP value of 0.

To use this feature more efficiently, this temporary allocation can be made in advance before the MN moves to the new location. If information is available regarding the mobility pattern of the MNs and the future location, then temporary allocation can be directed towards the new FA candidates.

Low latency handoffs represents one aspect of service quality associated with the mobile environment. A MN may select to subscribe to a premium service where low

latency handoff with QoS support is provided. To support this service, traffic destined to a MN is forwarded simultaneously to more than one FA. The duplicate packets are marked with the contracted QoS class. Another way to provide different levels of service, is to provide different forwarding scopes for duplicate datagrams. A mobile user who is interested in QoS lossless handover, may subscribe to a service level that provide advance registration to a wider scope of neighboring FAs.

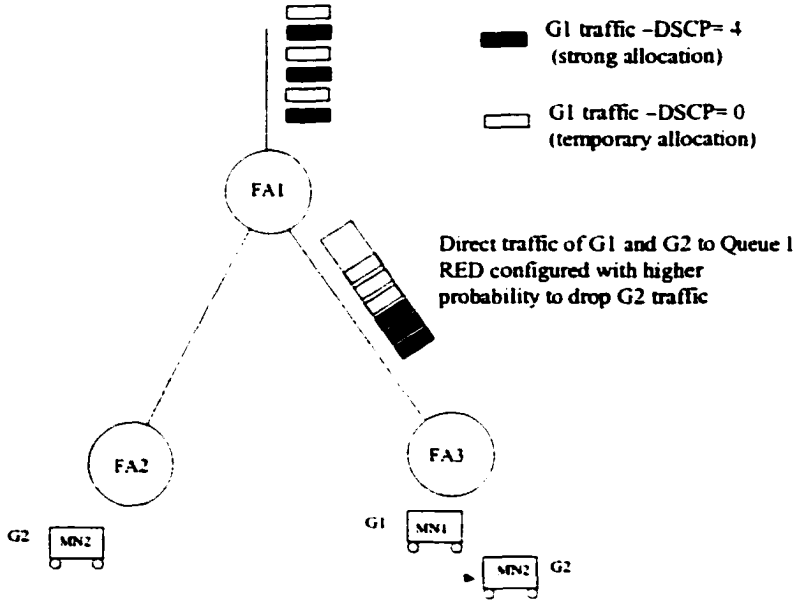

Figure 7.4 Temporary allocation and mobility support

7.7 DiffServ and Multicast Traffic

7.7.1 Mobile-IP and Multicast

Our proposed approach, based on the FA subscription approach, avoids the tunnel convergence behavior and provides route optimality that is better than that of the HA

subscription approach. In the context of providing different quality of service, the proposed protocol can be extended to allow MNs to select the multicast service provider that satisfies their QoS profiles.

In our proposed scheme, the Gateway FA subscription, the GFA will act as multicast service provider to the MNs residing in its domain. This approach aims at reducing the overhead associated with the group join/prune process. The main functional characteristics of the GFA subscription scheme can be described as follows. Figure 7.5 illustrates the elements associated with the multicast support.

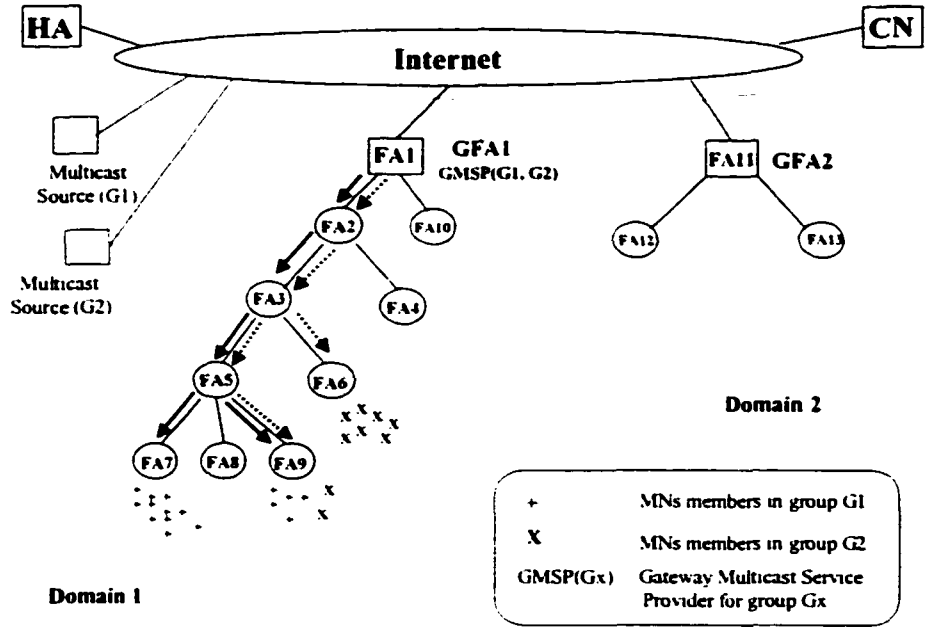

Figure 7.5 Multicast support in a Mobile-IP local registration system

Group subscription: A MN will send its membership report to its current FA. The reports from different MNs will be summarized by the FA, and forwarded upstream. The

same process is repeated at each intermediate FA, till the summarized report is received by the GFA. When processing this membership report, each intermediate FA will update its local membership table. An entry in the table specifies that a downstream FA(s), associated with a particular interface, is receiving the traffic of one or multiple multicast groups. An entry is created also to support MNs local to this FA, and is associated with the wireless interface. On receiving the summarized report, the GFA will join the requested groups utilizing the underlying multicast routing protocol.

Datagram delivery: After joining the multicast group, the GFA will start receiving the multicast datagrams. The GFA will consult its table and will forward the datagrams downward towards the interested FAs. Eventually, the datagrams will not be forwarded anymore when received by the last interested FA in the domain.

Handling mobility: Upon arriving to the new FA, the MN will send a regional registration request. The MN will embed its membership report in this same request. The membership report is extended with a field that contains the identification of the previous FA. If the new FA is receiving the traffic associated with group(s) requested by the newly arrived MN, then it will suppress this membership report but the information regarding the previous FA is summarized and forwarded with the request. If any of the requested groups is not currently supported by the current FA, the FA will need to include the corresponding groups addresses in its summarized report.

To avoid unnecessary delivery of datagrams to the previous FA, entries on FAs between the common and the previous FAs need to be cleared. Two mechanisms work together to remove unneeded entries. First, the FA regularly queries its local MNs for

changes in the memberships. The result is compared with last transmitted report, and will be forwarded upward only if there is a change to report. A new report will be forwarded also if a membership from the downstream direction warrants an updated new summarized report. This mechanism will eventually provide the GFA with updated membership reports. If the MNs mobility and the group memberships are highly dynamic, then the GFA may decide to utilize a complementary scheme. On receiving the local registration request, the common FA will extract the information regarding previous FAs with group membership changes. The common FA will use this information, to send a group membership report query to the previous FA. In addition, the GFA may periodically send a domain wide solicitation for membership reports. In this case, the request will be propagated to the leaf FA, and a summarized report is created and forwarded upstream towards the GFA.

7.7.2 Issues with DiffServ and Multicast Support

The multicast environment represents a considerable challenge to the DiffServ environment. On the multicast traffic delivery tree, each datagram is copied and forwarded over the links attached to a branching point. Without additional accommodation, datagrams will be treated according to the DSCP assigned by the originating domain. This situation limits the flexibility of providing end-to-end QoS. In this subsection, we will describe issues affecting the DiffServ model in a multicast environment.

Accommodating different classes of service for the same group and for multiple groups:

In a mobile environment, with different wireless access rates, it is very possible that different MNs in the same domain will request to receive multicast data traffic for a particular group with different QoS classes. In our proposed system, each FA will include the highest class of service requested from local MNs and from the FAs downstream. This arrangement allows MNs with high class of service to receive traffic with this contracted class. At the appropriate branching point, the DSCP will be mapped to a value that corresponds to the lower service classes. Figure 7.6 illustrates this concept. The approach will be more efficient, if the mobility pattern of the MNs in a domain was such that the MNs are clustered into one or few lineages. A lineage of a FA is defined as the path (or the set of intermediate FAs) from the GFA to the FA.

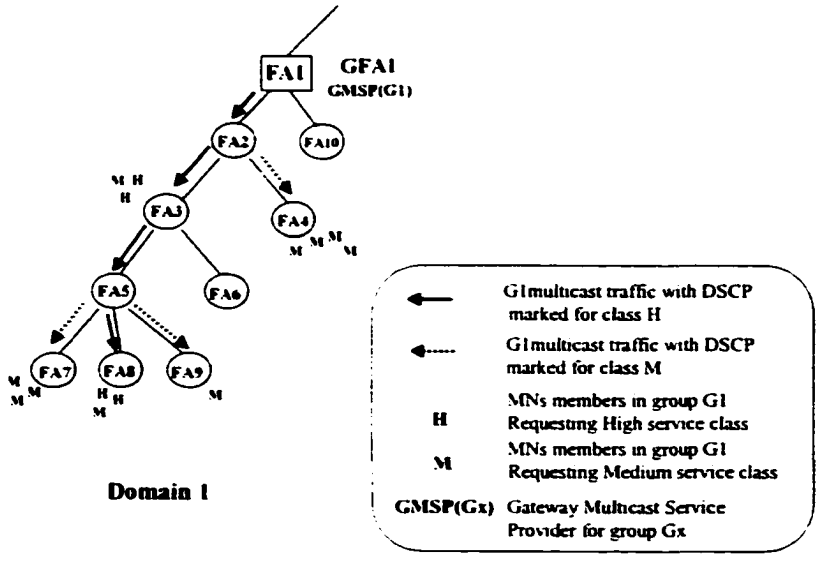

Figure 7.6 Supporting DiffServ for multicast traffic

A MN can subscribe simultaneously to different groups, receiving traffic associated with applications requiring different QoS. Considering that our approach for DiffServ request is receiver initiated, different QoS for different groups can be supported by the MNs. As mentioned earlier, the lower QoS requests will be satisfied as long as it does not conflicts with higher class requests on the same lineage.

Packet replication using the same codepoint:

In addition to the effect mentioned earlier of copying the same DSCP in the replicated multicast packets, another side effect is possible degradation of service provided to other classes of traffic. This degradation is resulting from the fact that multicast group members may receive traffic with a class of service, with no corresponding allocated resources. This issue was referred to as the neglected reservation subtree problem, and a proposed corresponding solution was described in [54].

In our scheme, the MN when first subscribing to the multicast group, it has to specify the QoS level required. This QoS request, embedded in a multicast group report, will be propagated upstream till processed by a FA that supports the requested group. If the requested QoS class is higher than what is currently received, the FA will generate a QoS upgrade request message in the upstream direction. If the FA receiving the QoS upgrade request is capable of satisfying the request, then it will adjust the DSCP mapping corresponding to the service level requested considering temporary allocation. The FA, which is capable of providing the required class of service, will forward the request to the DRM-CFA.

7.7.3 Basic Operation

When a MN, which is residing in a foreign network, decides that it needs to receive a multicast service for a particular group with a specific class of service, it will send a QoS request to the current FA. The request is forwarded upstream till processed by a FA that receives the requested multicast group traffic. We will refer to this FA as the Closest Local Provider (CLP_FA). This FA will forward the request to the DRM-CFA, which verifies resources from the CLP_FA to the CFA. If resources are available then appropriate profile meters and markers are installed on the CLP_FA. The previous scenario is applicable if the CLP_FA is receiving the multicast traffic with a class of service that is equal to or higher than the service class requested by the MN.

If the CLP_FA is receiving traffic marked for a class of service lower than the level requested by the MN, the CLP_FA will not forward the QoS request to the DRM-CFA. Instead, it will send a QoS upgrade request upstream. This upgrade request will propagate till processed by a FA supporting the requested class of service. At that time, this FA will forward the QoS request to the RM_GFA to verify resources. In most cases, the DRM-CFA and the RM-GFA are the same entity, except when using hierarchical resource management architecture. If a FA that is capable of satisfying the QoS request can not be found, then the request will be intercepted eventually by the GFA.

To be able to support the highest requested class of service in the domain, the GFA will need to request same or higher service level from the DRM-CFA. To support DiffServ in the environment described in the previous section, resources need to be

allocated from the point where the GFA joins the group multicast tree (MTJP) to the current FA where the MN is residing. Resources need to be allocated over segments, from the MTJP domain to the GFA domain and from the LMSP to the MN. The multicast group join request, generated by the GFA, will be intercepted by the border node in the domain where the multicast joining point is located. The request, along with the required QoS is forwarded to the RM-MTJP, for resources verification and allocation.

7.7.4 DiffServ and Enhanced Services for Multicast

In the following subsections, we will investigate three additional extensions to the proposed protocol. Those extensions allow more flexibility in deploying the service and provide enhanced performance. Since applications and access modes can differ widely in their QoS requirements, the design of those extensions does not necessitate applying a particular feature domain wide. Those features can be applied on demand and to the extent required by the subscriber and permitted by the service provider.

7.7.4.1 The LMSP Approach for Delay Sensitive Multicast Applications

One observation on the scheme described to support the multicast service, is that the GFA joins the requested multicast groups in behalf of the MNs residing in its local domains. Accordingly, the multicast traffic distributed in this regional domain will be marked according to the highest service class requested from the downstream direction.

Also the delay encountered in the regional domain, will be proportional to the position of the current FA relative to the GFA. To address those issues we use the LMSP (Local Multicast Service Provider) functionality described earlier.

The operation of this scheme may be illustrated with the following example. Figure 7.7 illustrates the support for DiffServ in the multicast environment. First, we will define the LMSP as an intermediate FA that can use the underlying multicast protocol to join the requested group. LMSPs can be selected in advance, or can be elected dynamically. A possible criterion to select LMSPs is based on MNs distribution and group memberships. To support this scheme, the summarized membership report propagating upward will carry the count of MNs interested in receiving the traffic of a particular group. Each intermediate FA will monitor this count, and will assume the LMSP responsibility when this observed count exceeds a threshold.

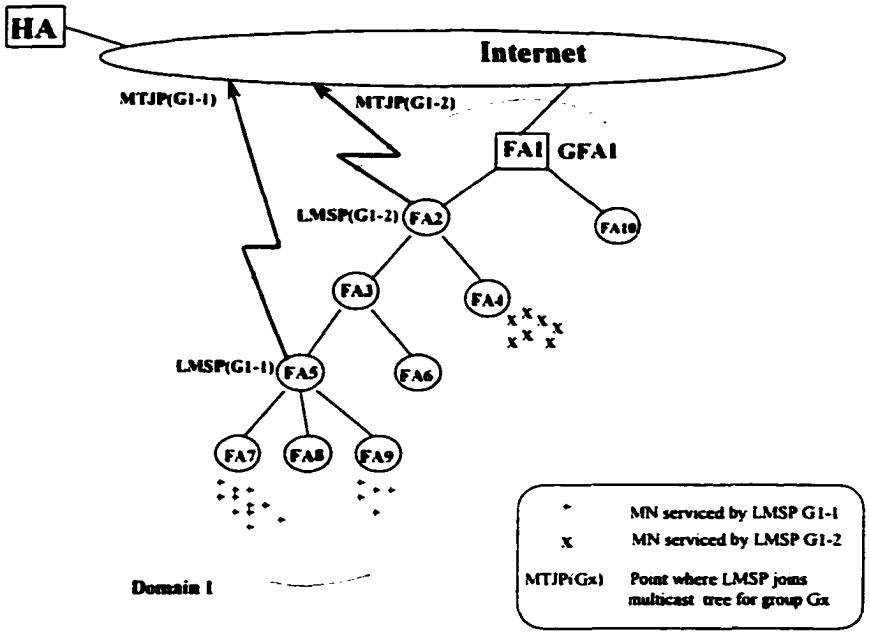

Figure 7.7 Multiple multicast providers and QoS

Assume that a number of MNs moved to domain 1, and are running application that requires the least possible delay. Accordingly, they request a higher class of service in their QoS request. In the GFA subscription mentioned earlier, those requests would be aggregated till intercepted by the GFA. In this LMSP scheme, QoS requests with delay sensitive indication are treated differently. The request is allowed to propagate upstream for a specific number of hops. If the request is not intercepted by a LMSP by that time, the FA where this number of hops threshold is crossed will assume the LMSP responsibility. This arrangement allows the MNs to be served by a FA that is closer to its current FA, and permit the FA acting as an LMSP to join the multicast tree using optimum routing. This option allows more flexibility in the case when resources need to be allocated across domains. Figure 7.7 illustrates a scenario where a multicast group is supported by two LMSPs.

7.7.4.2 Reliable Multicast as a Service Differentiation Aspect

A MN may be running an application that can not tolerate packet loss. If such application is based on multicast service then a reliable multicast session is requested by this MN. The capability to support reliable multicast, and the degree by which the network is willing to compensate for lost packets, differ with the particular location where the MN is residing. In a mobility support environment, a critical design aspect of the reliable delivery service is to isolate or to reduce the effect of typically slow mobile receivers on the external multicast service provider. If this provider detects that the

LMSP is not capable of handling the current transmission rate, indicated by higher rate of retransmission requests or congestion indications, it may react by denying the reliable service to this LMSP.

If the LMSP has a member in its domain that is not capable of receiving the multicast stream with the current transmission rate, the LMSP may deny the local service to the misbehaving element. Although this seems to be a practical solution, there may be a situation where the MN(s) causing the problem has a QoS contract that includes a reliable multicast service. It is beneficial to try to accommodate this element as a way of providing a premium service.

Dynamic logical hierarchy is a scheme that provides flexibility in supporting different quality levels of the reliable multicast service. When a FA detects that the available cache dedicated for the reliable service support is depleting, it will start a corrective action. The system will try to change the logical hierarchy of the delivery tree, such that MNs with higher class of service are serviced accordingly. Also, MNs with lower receiving capacity are serviced through a different provider that can accommodate this lower rate. To start this process, the FA identifies the particular group and the particular member of the immediate receivers set that represent the problem. If the FA determines that this member is accepting traffic within its service profile, the LMSP will start a logical restructure for the affected area of the domain.

In figure 7.8, FA3 is the LMSP for groups G1 and G2 serving all FAs downstream. MN-H is a MN with a contract for a high class of reliable multicast service (low loss rate), while MN-M is a MN with a medium class of service and low receiving rate. MN-

Ms eventually will cause FA3 to deplete the available cache. In order to continue to service MN-H receiving G2 with the requested high class, FA3 will stop acting as the LMSP for G1. Instead, FA2 will consider this responsibility. Also FA3 will not keep reliable service states for FA4. FA2 will appear as the new reliable multicast provider for G1 servicing FA4, and will keep the corresponding states.

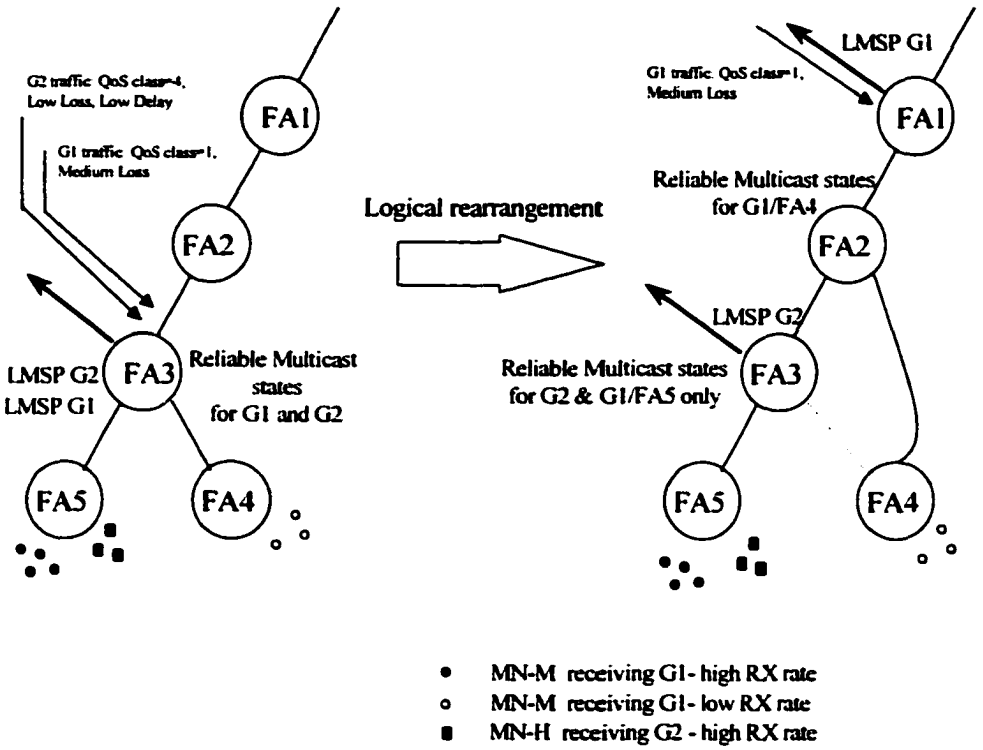

Figure 7.8 Dynamic logical hierarchy and supporting differentiated services

In this new arrangement, G1 traffic received by FA5 will experience more delay than in the first case. This is acceptable, in this case, since no particular loss specification was requested. If FA5 requests a low delay service in the future, FA3 may need to resume the

LMSP responsibility for G1, without extending the reliable service to the misbehaving FA4.

7.7.4.3 Advanced Registration and Reservation

A MN running applications that is both loss and QoS sensitive, will need to utilize advanced registration and reservation. In advanced registration, the MNs register with the multiple locations that may be visited in the next mobility event. A MN with higher class of service may have a profile that specifies the number of candidate FAs to consider in the advanced registration. Another MN with lower class of service may have less number of candidates or none at all. The first needed process to perform advance registration is to determine the members in the next FA candidate set, which can be inferred by the network from the mobility pattern and the topology. The next step to complete the advance registration is to deliver the local registration requests to the candidate FAs.

If the MNs have the knowledge of the IDs and the hierarchy lineage of the members in the next FA candidate set, then the requests can be delivered as follows. Comparing the lineage of the CFA to that of each of the candidates, the MN will send one or more compressed regional registration report depending on the network topology. As an example, considering figure 7.5 and a MN residing with FA7 and moving towards either FA6 or FA4. The MN will send an advance registration request that will be processed by both FA2 and FA3.

If the MN has knowledge only of the IDs of the members in the next FA candidate set, the MN will need to send individual registration message to each candidate FA.

8. Conclusion

The next generation Internet will provide high-quality and high bandwidth connectivity, providing support for both traditional and new services. Future Internet users will demand the ability to use the full range of Internet applications, regardless of the access mode. Mobile users will expect to be able to use applications and to experience levels of service similar to those available on their stationary networks.

The regional registration Mobile-IP is a promising scheme that is characterized by lower signaling overhead. The associated hierarchical environment causes the regional Mobile-IP to be sensitive to the failure of FAs. In addition, the inherent characteristics of the regional registration approach may impede the support of services like multicast, reliable multicast and differentiated QoS.

The mobility environment is a challenge for protocol designed for stationary nodes. Existing protocols supporting services like multicast and QoS were not designed with mobility support in mind. Accordingly, different accommodations are needed to support those services. Research and implementations efforts have to be directed towards tackling those issues, allowing the regional registration environment to support the demanding applications that users will continue to use on the move.

In this work, we analyzed the challenges associated with the regional registration environment. We then proposed an integrated platform that supports essential services in the regional domain. Multicast, reliable multicast and differentiated QoS services are

supported in this platform. Capitalizing on existing control messages and data structures, the support for those services was provided with minimum additional overhead.

Scalability is an important aspect in any solution supporting mobility. The solution should provide an acceptable performance with different number of mobile nodes, with different mobility patterns and with different transmission error rates. Support for services like reliable multicast and DiffServ should tolerate an environment where nodes may have different receiving capacity.

In our work, we evaluated the performance of our proposed schemes in a dynamic environment. In many aspects, our proposed integrated platform exhibits strong scalability features. Popular applications similar to audio and video conferencing, distributed databases and multimedia contents rely on underlying services similar to multicast and reliable multicast delivery. Our proposed integrated platform provides support for those services for mobile hosts, in addition to enabling the DiffServ model operation in the regional domain. Such capabilities position this platform as a strong candidate for wireless network supporting the Mobile-IP.

Appendix A. Adaptive Route Optimization

Early research work addressed enhancement regarding routing in mobile environment [64, 65]. The need to include adaptive action to accommodate the dynamics associated with mobility has been addressed in [14, 66, 67].

The route optimization extension [7] describes the binding update and binding warning messages. Those messages are generated for each tunneled datagram received by the HA or the previous FA, indicating that a CN does not have the updated location information. If the generation of those messages are not limited, the network links associated with the HA, PFA and the CN can be congested with a flood of control messages. The need for a scheme to control the generation of those messages is clear [74]. We propose an adaptive scheme that considers the network status and provides an efficient procedure to manage the generation of those control messages.

The Adaptive System limits when and which CN should receive the Binding Update messages and when to send the Binding Warning and Binding Update messages. The basic idea behind the Adaptive System is flexible enough to accommodate one or multiple criteria to consider when constructing the group of CNs that the Adaptive System identifies as valued hosts and justify the cost paid to maintain their caches up-to-date. In this system each MN keeps track of the current transmission rate associated with the CN currently involved with a session with this MN. For a CN exchanging traffic with a rate exceeding a threshold, the MN will identify this CN as a valued host. Figure A.1

illustrates the elements and messages involved in the adaptive scheme. A detailed description of the adaptive system can be found in [68].

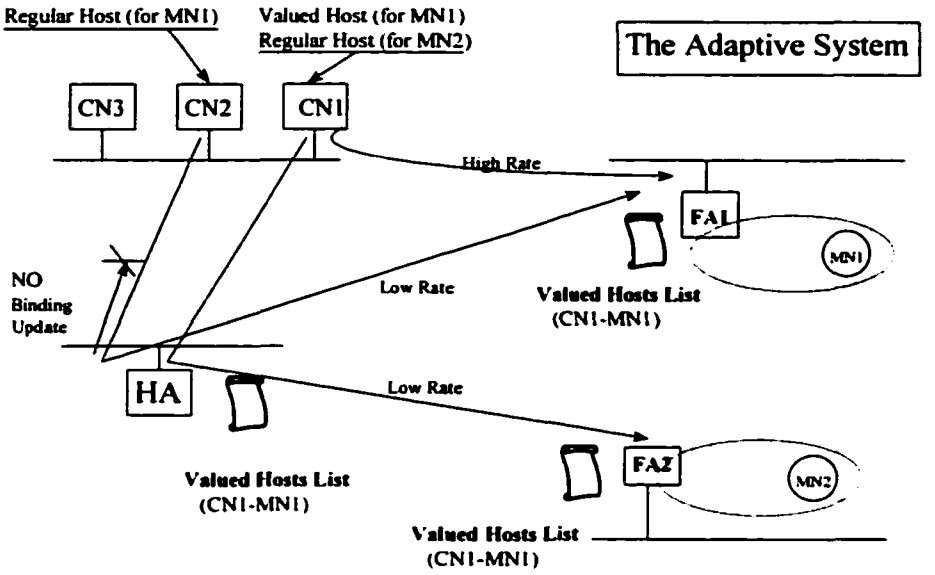

Figure A.1 Adaptive control messages limiting

The Adaptive Rate Limiting

If the HA Binding Update Rate Limiting is implemented using a constant rate threshold, then the HA will prevent the generated Binding Update rate destined to a CN concerning a specific MN from exceeding this threshold. Using the same fixed threshold all the time for all communication sessions is not the most efficient way to deal with this issue, since different sessions are characterized with a different HA-CN delay, and also with different rates transmitted from different CNs. A more efficient approach is to consider each session as an individual one and adjust the Home Agent Binding Update

Rate Limiting threshold according to the session characteristics. In our proposed scheme, the measured round trip delay between the HA and the CN is used to dynamically adjust this threshold, such that the time duration between any two successive Binding Updates to the same node regarding a specific mobility binding should not be less than the measured delay.

To implement the Adaptive Foreign Agent Binding Warning Rate Limiting, the aggregation of the delay values from FA to HA, from HA to CN and CN to FA should be considered when selecting a threshold. We have to notice that a reasonable accuracy in the delay measurement will suffice for this purpose. The overhead associated with implementing the Adaptive System can be categorized into processing overhead and control messages overhead. According to the Route Optimization specification [7], the Foreign Agent keeps information of associated with previous MN guests to be used when implementing both the Previous Foreign Agent Notification and the Foreign Agent Binding Warning Rate Limiting. Additional information of the same nature can be integrated easily to the already existing database or data structures to support the adaptive action such that the management overhead for the adaptive system can be justified and easily integrated into the Route Optimization approach.

Simulation environment for the Adaptive System

The mobility support system was constructed using discrete event simulation [69] to follow the guidelines of the Base Mobile-IP [1] and the Route Optimization draft [7].

The mobility support elements used for the simulation are three CNs, one HA, fourteen FAs and six MNs. The network used for simulation consists of 18 cells, where each cell has a subnet housing only one stationary element either a HA, a FA or a CN. Any of the six Mobile Nodes can roam between cells serviced by the HA or by any FA. The MNs belong to the same Home network, thus serviced by the same Home Agent. Each subnet is connected to at least three of its neighboring subnets. The delay over a link between two neighbor subnets is set to one msec and that over the wireless link to 2.5 msec. Packet processing delay of 0.3 msec was considered for the mobility support elements. If (n) cells surround the current cell of the MN, the MN can move to any of the neighbor cells with a probability of (1/n). A bi-directional traffic is flowing between each CN-MN pair. Each of the Correspondent Nodes generating traffic with a variable rate destined to different Mobile Nodes. The packet generation rate of the CNs is selected to alternate in a value above and below a certain threshold. This threshold plays a role in the adaptive mechanism of the system, since it is used to identify valued hosts entitled for updated location information. The value of the traffic generated by the CN is uniformly distributed between 20 and 100 packets per second for 80% of the simulation time, and no traffic is generated for the remaining 20% of the 60 minutes simulation time. The previous model was applied on the traffic associated with 50% of the MNs, while the other half is associated with traffic characterized by a uniformly distributed rate between 100 and 200 packets per second for 40% of the time. The rate of issuing Binding Update messages and Binding Warning messages is not allowed to exceed a specific threshold that was set to 150 messages per second. It is assumed here that the New Foreign Agent

will accept packets tunneled from the Previous Foreign Agent and intended to a Mobile Node before receiving the registration reply from the Home Agent.

Simulation results and analysis

The objective of the simulation was to verify the performance of the Adaptive System compared to the other two non-adaptive systems (The Base Mobile IP and The Mobile IP with the Route Optimization) over a range of different mobility rates. The comparison between different schemes will be based on the following aspects: the number of control messages, the number of dropped packets for traffic generated from CNs and received by MNs. The graphs below describe those performance aspects and compare between the different schemes.

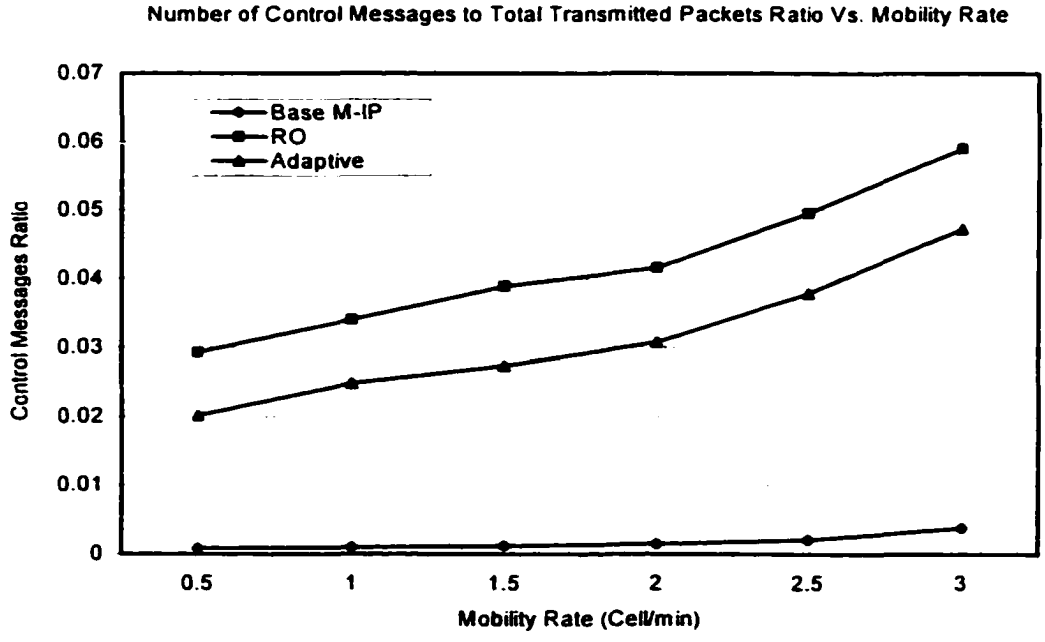

Figure A.2 Effect on number of control messages

Figure A.2 shows the ratio of the number of control messages to the total number of transmitted packets over different mobility rates. Although it is obvious that the Base Mobile IP provides the least number of location management messages, we notice that the Adaptive System provided a noticeable reduction of the number of location management messages compared to the case of the Mobile IP with the Route Optimization. This reduction is accomplished by limiting the transmission of the Binding Update and the Binding Warning Messages. This reduction is not associated with a costly penalty since those filtered control messages are intended to provide some of the correspondent nodes with the most updated location information. Those CNs that are not receiving the update information, are currently generating traffic rates that are not high enough to justify the high cost of providing them with the premium location information.

Figure A.3 shows the ratio of the dropped packets to the total number of transmitted packets. Since that the Adaptive System follows the Route Optimization guidelines in using the Previous Foreign Agent Notification technique to reduce the number of dropped packets, the Adaptive System provides a very comparable performance to that of the Route Optimization. This is expected since the penalty of limiting the number of binding update messages in the Adaptive System may cause a slight increase in the average packet delay, but it will not increase the number of packets dropped.

From the previous results, it is clear that the adaptive scheme provides an efficient and needed control features to limit the amount of signaling associated with the route optimization binding.

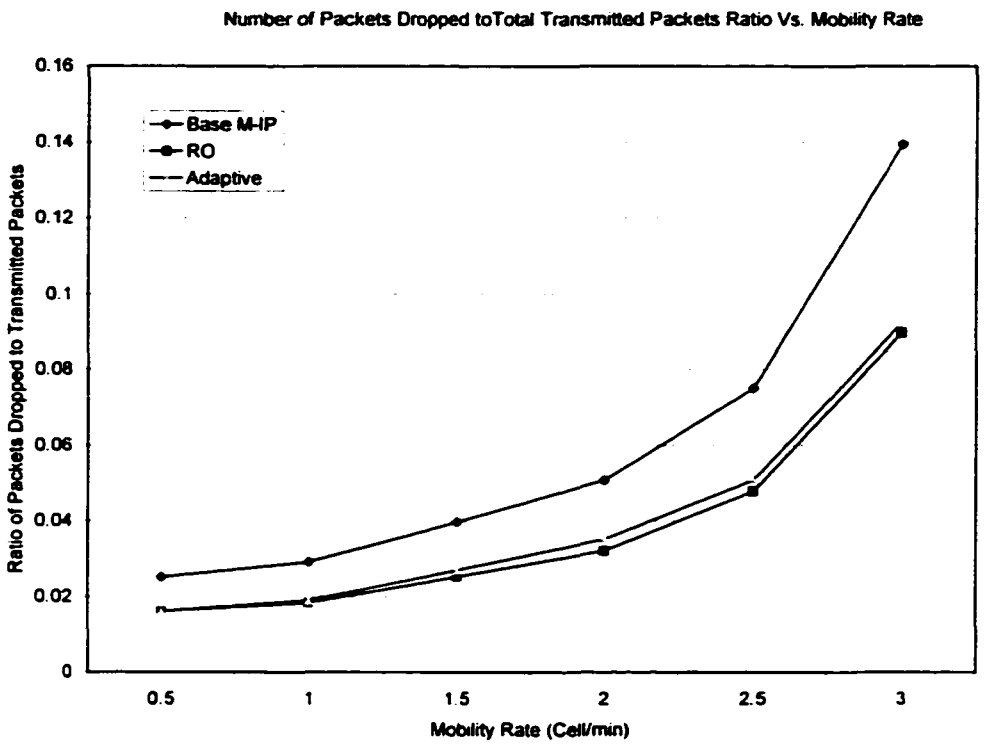

Figure A.3 Effect of mobility rate on number of dropped packets

Appendix B. Abbreviations and Acronyms

BB: Bandwidth Broker

CBQ: Class Based Queuing

CLP_FA: Closest Local Provider FA

CN: Correspondent Node

DiffServ: Differentiated Services

DRM: Domain Resource Manager

DRM-CFA: Domain Resource Manager for the Current FA domain

DSCP: Differentiated Service Code Point

DVMRP: Distance Vector Multicast Routing Protocol

FA: Foreign Agent

GFA: Gateway Foreign Agent

GMSP: Global Multicast Service Provider

GRE: Generic Routing Encapsulation

HA: Home Agent

HA_P: Primary HA

HA_S: Secondary HA

HA-BN: Home Agent Border Node

HARP: Home Agent Redundancy Protocol

HIR-REG: Hierarchy Registry

HLRM-IP: Hierarchical Local Registration Mobile-IP

ICMP: Internet Control Message Protocol

IGMP: Internet Group Management Protocol

IntServ: Integrated Services

IP: Internet Protocol

LMSP: Local Multicast Service Provider

M-IP: Mobile IP

MN: Mobile Node

MTJP: Multicast Tree Joining Point

PFA: Previous Foreign Agent

PHB: Per Hop Behavior

QoS: Quality of Service

RED: Random Early Discard

RM: Resource Manager

RMFR: Reliable Multicast Forwarding Request

SLA: Service Level Agreement

TCP: Transmission Control Protocol

TOS: Type of Service

TTL: Time to live

UDP: User Datagram Protocol

References

[1] Charles Perkins, "IP Mobility Support", RFC-2002, Mobile IP Working Group, October 1996.

[2] Eva Gustafsson, Annika Jonsson and Charles Perkins, "Mobile IP Regional Registration" (work in progress), Internet Draft, draft-ietf-mobileip-reg-tunnel-03.txt, July 2000.

[3] S. Deering, C. Patridge and D. Waitzman, "Distance Vector Multicast Routing Protocol", Internet RFC 1075, Nov 1988.

[4] H. Eriksson, "MBONE: The Multicast Backbone", Commun. ACM, vol 37, no. 8, Aug 1994, 54-60.

[5] J. Wroclawski, "The use of RSVP with IETF integrated services", IETF RFC 2210, 1997.

[6] Y. Bernet, J. Binder, S. Blake, M. Carlson, E. Davies, B. Ohlman, D. Werma, Z. Wang and W. Weiss, "A framework for differentiated Services", draft-ietf-diffserv-framework-02.txt, February 1999.

[7] C. Perkins and D. Johnson, "Route Optimization in Mobile IP", draft-ietf-mobileip-optim-09.txt (work in progress), February 2000.

[8] Gabriel Montenegro, "Reverse Tunneling for MobileIP", IETF Draft, 'draft-ietf-mobileip-tunnel-reverse-06.txt', April 13, 1998.

[9] R. Bless and K. Wehrle, "IP Multicast in Differentiated Services Networks", <draft-bless-diffserv-multicast-01.txt>, work in progress, November 2000.

[10] J. Postel, "Internet Protocol", RFC 791, September 1981.

[11] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, January 1996.

[12] D. Johnson and C. Perkins, Mobility Support in Iv6", Internet Draft (work in progress), November 2000 (draft-ietf-mobileip-ipv6-13.txt)

[13] Vipul Gupta and S. Glass, "Firewall Traversal for MobileIP: Guidelines for Firewalls and MobileIP entities", IETF Draft, 'draft-ietf-mobileip-firewall-trav-00.txt', March 17, 1997.

[14] S. Chesire and M. Baker, "Internet Mobility 4x4.", in proceedings of the ACM SIGCOMM 96 Conference. August 1996.

[15] Charles Perkins, "IP encapsulation within IP", RFC 2003, 1996.

[16] Charles Perkins, "Minimal Encapsulation within IP", RFC 2004, 1996.

[17] H. Stan, T. Li, D. Farinacci and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 1701, 1994

[18] C. Perkins, "Mobile IP Design Principles and Practices," Wireless Communications Series. Reading, MA: Addison Wesley Longman, 1997.

[19] K. El Malki and H. Soliman, "Fast Handoffs in Mobile IPv4", Internet Draft (work in progress), September 2000 (draft-elmalki-mobileip-fast-handoffs-03.txt).

[20] K. El Malki and H. Soliman, "Hierarchical Mobile IPv4/v6 and Fast Handoffs". Internet Draft (work in progress), March 2000, (draft-elmalki-soliman-hmipv4v6-00.txt).

[21] S.F. Fo and K.C. Chu, "Regional Aware Foreign Agent (RAFA) for Fast Local Handoffs <draft-chuafoo-mobileip-rafa-00.txt>.

[22] P. McCann, T. Hiller, J. Wang, A. Casati, C. Perkins and P. Calhoun, "Transparent Hierarchical Mobility Agents". Internet Draft (work in progress), March 1999, (draft-mccann-thema-00.txt).

[23] J. Moy, "Multicast Extensions to OSPF", RFC 1584, March 1994.

[24] A. Ballardie, P. Francis and J. Crowcroft, "Core Based Trees (CBT): an architecture for scalable inter-domain multicast routing", in: Proc. 1993 ACM SIGCOMM, San Francisco, CA (September 1993) pp.85-95.

[25] D. Estrin et al., "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specifications," RFC 2117, June 1997.

[26] W. Fenner, Internet Group Management Protocol, version 2, RFC 2236, (November 1997).

[27] V. Chikarmane, C. Williamson, R. Bunt and W. Mackrell, "Multicast support for mobile hosts using Mobile IP: Design issues and proposed architecture", Mobile Networks and Applications 3 (1998) 365-379

[28] C. Williamson, T. Harrison, W. Mackrell and R. Bunt, "Performance evaluation of the MoM mobile multicast protocol", Mobile Networks and Applications 3 (1998) 189-201.

[29] G. Xylomenos and G. Polyzos, IP multicast for mobile hosts, IEEE Communications 35(1) (1997) 54-58.

[30] C. Liu, M. Lee and T. Saadawi, "Mobility Support for the Core-Manager Based Scalable Multicast Routing", in: Proc. MILCOM 1998, Bedford, MA (October 1998).

[31] C.L. Tan and S. Pink, "MobiCast: A Multicast Scheme for Wireless Networks", Mobile Networks and Applications, Dec 2000.

[32] S. Deering, D. Estrin, D. Farinacci and V. Jacobson, "An architecture for wide-area multicast routing", Proc. ACM SIGCOMM Conference, London, UK (August 1994) pp. 126-135.

[33] H. Omar, T. Saadawi and M. Lee, "Multicast Support for Mobile-IP with the Hierarchical Local Registration Approach", Proc. WOWMOM 2000, Boston, MA (August 2000).

[34] B. Levine, D. Lavo and J. Garcia-Luna-Aceves, "The Case for Reliable Concurrent Multicasting Using Shared Ack Trees", in Proc. Fourth AM International Conference on Multimedia, Boston, MA (November 1996)

[35] B. Levine, J. Garcia-Luna-Aceves, "A comparison of Known Classes of Reliable Multicast Protocols", in Proc. IEEE International Conference on Network Protocols, October 1996.

[36] D. Chiu, S. Hurst, M. Kadansky and J. Wesley, "TRAM: A Tree-based Reliable Multicast Protocol", Sun microsystems, Technical Report TR-98-66, July 1998.

[37] B. Rajagopalan, "Reliability and Scaling issues in Multicast Communications", in Proc. of ACM SIGCOMM'92, pp. 188-198, August 1992.

[38] A. Mankin, A. Romanow, S. Bradner, V. Paxson, "IETF Criteria for Evaluating Reliable Multicast Transport and Application Protocols", RFC 2357, Network Working Group, June 1998.

[39] K. Brown and S. Singh, "RelM: Reliable Multicast for Mobile Networks," Journal of Computer Communications, 1998, pp. 1379-1400

[40] S. Gupta and P. Srimani, "An Adaptive Protocol for Reliable Multicast in Mobile Multi-hop Radio Networks", in Proc. of the Second IEEE Workshop on Mobile Computer Systems and Applications.

[41] I. Nikolaidis and J.J. Harms, " A logical Ring Reliable Multicast Protocol for Mobile Nodes," in Proc. of IEEE ICNP 1999, pp. 106-113

[42] M. Yajnik, J. Kurose and D. Towsley, "Packet Loss correlation in the Mbone Multicast Network", in Proc. of IEEE Global Internet Conference, London, November 1996.

[43] H. Omar, T. Saadawi and M. Lee, "Multicast with Reliable Multicast Support in the Regional Mobile-IP Environment", Proc. ISCC 2001, Tunisia, July 2001.

[44] John C. Lin and S. Paul, "RMTP: A Reliable Multicast Transport Protocol", in proceedings of IEEE INFOCOM'96, March 1996, pp. 1414-1424.

[45] Sridhar Alagar, Ramki Rajagopalan, S. Venkatesan, "Tolerating Mobile Support Station Failures", Proc of IEEE Conf. on Fault Tolerant Systems, pp. 225-231, Dec. 1995.

[46] A. Acharya, B. Badrinath and T. Imielinski, "Checkpointing Distributed Applications on Mobile Computers", Proc of the 3rd International Conference on Parallel and Distributed Information Systems, pages 73--80, September 1994.

[47] S. Rangarajan, K. Ratnam and A. Dahbura, "A Fault-Tolerant Protocol for Location Directory Maintenance in Mobile Networks", Proc of the 25th International Symposium on Fault-Tolerant Computing, June 1995, Pasadena, CA.

[48] S. Biaz and N. H. Vaidya, "Tolerating Location Register Failures in Mobile Environments", Texas A&M University, Technical Report 97-015.

[49] Y. Lin, "Failure restoration of mobility databases for personal communication networks", Wireless Networks (1995), pp. 365-372.

[50] H. Omar, T. Saadawi and M. Lee, "Support for Fault Tolerance in Local Registration Mobile-IP Systems", Proc. MILCOM 1999, Atlantic City, New Jersey (October 1999).

[51] HARP – "Home Agent Redundancy Protocol" <draft-chambless-mobileip-harp-00.txt>- B. Chambless Portlant State University, J. Binkley Oregon Graduate Institute October 27,1997.

[52] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weiss, "An architecture for Differentiated Services", IETF RFC 2475, 1998.

[53] R. Braden, L. Zhang and S. Berson, "Resource ReSerVation Protocol (RSVP)-Version 1 Functional Specification. (work in progress) draft-ietf-rsvp-spec-14.txt

[54] R. Bless and K. Wehrle, "IP Multicast in Differentiated Services Networks", <draft-bless-diffserv-multicast-01.txt>, work in progress, November 2000.

[55] Talukdar. A.K., Badrinath. B. R. and Acharya. A, "On Accommodating Mobile Hosts in a Integrated Services Packet Network", In the Proceedings of the IEEE INFOCOM'97, April 1997.

[56] S-B. Lee, G-S Ahn, X. Zhang and A. Campbell, "INSIGNA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks, Journal of Parallel and Distributed Computing (Academic Press), April 2000.

[57] J. Chen, A. McAuley, A. Caro, S. Baba and P. Ramanathan, "QoS Architecture Based on Differentiated Services for Next Generation Wireless IP Networks", <draft-itsumo-wireless-diffserv-00.txt>, work in progress, July 2000.

[58] K. Nichols, V. Jacobson and L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet", RFC 2638, July 1999.

[59] D. Black, "Differentiated Services and Tunnels", Interent Engineering Task force, RFC 2982, October 2000.

[60] H. Chaskar, "Requirements of a QoS Solution for Mobile-IP" (work in progress), draft-chaskar-mobileip-qos-requirements-00.txt, June 2001.

[61] V. Rexhepi, G. Karagiannis and G. Jeijenk, "A Framework for QoS & Mobility in the Internet Next Generation", Proceedings EUNICE 2000, University of Twente, Enschede, the Netherlands, September 2000.

[62] T. Braun, C. Castellucia and G. Stattenberger, "An Analysis of the DiffServ Approach in Mobile Environments", In the Prpceedings of IWQiM'99 .

[63] H. Omar, T. Saadawi and M. Lee, "An Integrated Platform for Reliable Multicast Support in the Regional Mobile-IP Environment", to appear in ACM Mobile Computer and Communication Review.

[64] G. Cho and L. Marshall, "An Efficient Location and Routing Scheme for Mobile Computing Environments", IEEE Journal on Selected Areas in Communications, Vol (13), No (5), Pages: 868-879, June 1995.

[65] Robit Dube, Ibrahim Korpeoglu and Satish Tripathi, "Reduced Router-Crossing in a Mobile Intranet", University of Maryland, Technical Report, CS-TR-3735, UMUACS-TR-97-1, Jan '1997.

[66] Subhashini Rajagopalan and B. Badrinath, "An Adaptive Location Management Strategy for MobileIP", in the Proceedings of the first ACM Mobicom '95, November 1995.

[67] Ruixi Yuan, "An Adaptive Routing Scheme for Wireless Mobile Computing", NEC Systems Lab, September 1993.

[68] H. Omar, H. Elsherif, T. Saadawi and M. Lee " An Adaptive IP Mobility System for Enhanced Performance", Proc. 4th workshop on Multiaccess, Mobility and Teletraffic for Wireless Communication, Washington, D.C., October 1998.

[69] ModSim, "Reference Manual and User's Manual", CACI Products Co.

[70] R. Droms, "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[71] P. Ferguson, "Ingress Filtering in the Internet", (work in progress) <draft-ferguson-ingress-filtering-01.txt> November 1996.

[72] Postel, J. B. "User Datagram Protocol", RFC 768, August 1980.

[73] B. Ohlman and P Koskelainen, "Receiver control in Differentiated services", <draft-ohlman-receiver-ctrl-diff-01.txt>, March 1999.

[74] H. Omar, T. Saadawi and M. Lee, "Supporting Reduced Location Management Overhead and Fault Tolerance in Mobile-IP Systems", Proc. The Fourth IEEE Symposium on Computers and Communications, July 1999.

[75] H. Omar, T. Saadawi, M. Lee, "Multicast Scheme with Reliable Delivery Support and Fault Tolerance for Mobile Hosts Using a Regional Distribution Approach," PP. 4-49 - 4-56, in Book; Advanced Telecommunications and Information

Distribution Research Program (ATIRP), Final Report 1996 - 2001, ISBN: 0-9711916-0-3.

Index

E

F

G

H

I

L

T