

Performance of IP Micro-Mobility Management Schemes using Host Based Routing

K. Daniel Wong, Hung-Yu Wei, Ashutosh Dutta, Kenneth Young

Telcordia Technologies Inc., 331 Newman Springs Road, 3X-363, Red Bank, NJ 07701, USA

dwong@research.telcordia.com, <http://www.argreenhouse.com/bios/dwong>

Henning Schulzrinne

Columbia University, New York, NY 10027, USA

Abstract

Global IP mobility solutions using protocols like Mobile IP (MIP) and SIP are not optimized to handle micro-mobility management, where low-latency handoffs are essential, to avoid inefficiencies and performance degradation. Host based routing (HBR) schemes, such as HAWAII, Cellular IP and MMP are one of two main classes of schemes for IP micro-mobility management, the other being hierarchical Mobile IP-derived schemes. We look at performance issues of HBR schemes, both qualitatively and quantitatively. Various simulation results and prototype system measurements demonstrate the superiority of HBR schemes over both MIP and hierarchical MIP-derived micro-mobility schemes in terms of fewer packets dropped per handoff for UDP traffic and better TCP throughput under a variety of scenarios.

Key words

Mobile communications, wireless IP, HBR, Mobile IP, SIP, micro-mobility management.

1. Introduction and Background

The Internet Protocol (IP) occupies an increasingly dominant position in computer networking. As its usage base expands, so does the list of requirements, and much research and development is being done to enable IP to meet various needs. For example, the next major version of IP, IPv6, is being developed to address concerns such as the size of the address space. Another important example is the work being done on meeting the particular needs of mobile users, the intersection of IP and wireless access. These developments will enhance the attractiveness of using IP in an array of network scenarios, such as commercial 3rd generation (3G) mobile networks, fixed broadband wireless access networks, enterprise and campus intranets, and tactical networks. Many new and enhanced applications and services would be possible. All-IP wireless networks are being designed in standards bodies like the Internet Engineering Task Force (IETF), 3G Partnership Project (3GPP) and 3G Partnership Project 2 (3GPP2). These all-IP wireless networks will allow roaming subscribers to access integrated data, voice and multimedia services of the Internet via their wireless IP terminals and appliances. One vision is that an end-to-end wireless/wireline IP platform comprising 3G wireless access networks and a wireline IP backbone will support real-time and non-real-time multimedia services in the future.

The cornerstone of the work in addressing the needs of mobile users is the Mobile IP [1] (MIP) framework and its derivatives, variations, and auxiliary protocols. An overview

of MIP is therefore an appropriate starting point, and this is provided in Section 1.1. Application-layer mobility management is an alternative to network-layer mobility management (as provided by MIP for example), and it is discussed in Section 1.2. Both MIP and SIP-based application-layer mobility schemes are more suitable for macro-mobility management than micro-mobility management. Micro-mobility management schemes are introduced in Section 1.3, although the introduction of schemes for micro-mobility management based on host-based-routing (HBR) is deferred to Section 2 for more detailed coverage. Following that discussion of HBR schemes, the performance-related issues are explored in Section 3. Section 4 contains selected performance results from our simulations and prototype test-bed that illustrate the performance of HBR schemes, especially in comparison with MIP. This is followed by discussions and conclusions in Section 5.

| | | | |
|--------|---|---------|---|
| 3G | 3 rd Generation Mobile Systems | HBR | Host-Based Routing |
| 3GPP | 3G Partnership Project | IDMP | Intra-Domain Mobility Management Protocol |
| 3GPP2 | 3G Partnership Project 2 | IETF | Internet Engineering Task Force |
| BS | Base Station | MA | Mobility Agent |
| CH | Correspondent Host | MH | Mobile Host |
| CIP | Cellular IP | MIP | Mobile IP |
| COA | Care-Of-Address | MIP-RO | MIP with Route Optimization |
| DHCP | Dynamic Host Configuration Protocol | MIP-RR | MIP with Regional Registration |
| DNS | Domain Name System | MMP | Micro-mobility Management Protocol |
| DSDV | Destination-Sequenced Distance-Vector | NAI | Network Access Identifier |
| FA | Foreign Agent | RFA | Regional FA |
| GFA | Gateway FA | RTP | Real-time Transfer Protocol |
| HA | Home Agent | SIP | Session Initiation Protocol |
| HAWAII | Handoff-Aware Wireless Access Internet Infrastructure | TeleMIP | Telecommunications-enhanced MIP |

Table 1: Glossary of Acronyms

1.1. Mobile IP Overview

IP-based networking is designed such that each host is identified by a unique IP address¹. Standard IP routing assumes that IP addresses are distributed hierarchically. For example, a host with a certain subnet prefix is assumed to be located at the subnet referenced by that prefix, the home network. This dual use of IP addresses is fine when hosts are not mobile, as each host can be assigned its unique IP address according to the hierarchical structure needed for IP routing. However, it creates a problem when hosts need to be mobile. If a host moves to a foreign network, packets for it will still be routed to its home network. Furthermore, a host may obtain a temporary address in the foreign network for routing purposes, but there is no association between its temporary and permanent addresses. The Mobile IP framework has been developed to support IP mobility through a network layer solution.

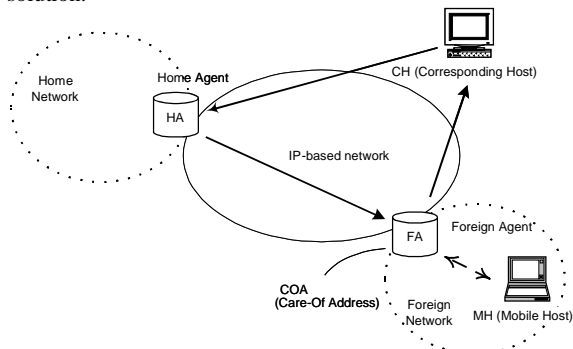


Figure 1: Mobile IP with Foreign Agent

In MIP, each Mobile Host (MH) is still identified by its permanent IP address. However, for routing purposes, when an MH is roaming it obtains a temporary care-of-address (COA), which is a foreign network address that identifies the location of the MH. The MH registers this COA with a mobility agent in its home network known as its Home Agent (HA). The HA then stores the COA of the MH in a binding cache. Nodes communicating with the MH send packets addressed to its permanent address. These packets are routed to the MH's home network, where its HA intercepts them and tunnels them (encapsulated) to its COA. The MH registers its latest COA with its HA whenever its COA is changes, which occurs when the MH moves to another foreign network. It should also refreshes the registration with its HA periodically.

MIP can operate in two modes, namely with foreign agents or with co-located COAs, illustrated respectively in Figure 1 and Figure 2. In the mode with foreign agents, the visited network has a Foreign Agent (FA). The FA broadcasts its IP addresses that can be used as COAs. The MH picks a valid IP address of the FA as its COA and registers this with its HA (in this mode, the registration goes through the FA rather than directly to and from the HA). When packets arrive for the MH at the FA tunneled from the

¹ There are exceptions to this statement, e.g. multi-homed hosts and the use of Network Address Translation (NAT). Multi-homed hosts face the same mobility problems as non-multi-homed hosts, individually for each of their interfaces (and corresponding home networks). The relationships between NAT and IP mobility are beyond the scope of this paper.

HA, they are un-encapsulated and forwarded to the MH through its layer 2 address previously registered with the FA. On the other hand, in the mode with co-located COAs, the MH would obtain a temporary IP address at the foreign network using a protocol such as DHCP (Dynamic Host Configuration Protocol). The MH would use this temporary IP address as its COA and registers this with HA.

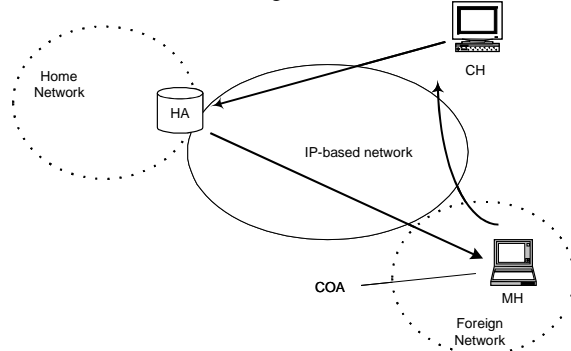


Figure 2: Mobile IP with co-located COA

When the MH sends packets to the CH, however, it does not need to route them via its HA. It sends packets directly to the CH's IP address. Hence, as seen in both Figure 1 and Figure 2, the routing path resulting from using MIP is "triangular".

The strengths of MIP include its transparency to the CHs (who do not need to know that the MH is mobile), its transparency to higher layers in the protocol stack, and the fact that the MH keeps its IP address for identity purposes, allowing it to continue to function as a server (e.g. email server) without the need for troublesome patches whenever it moves, e.g. changes Domain Name System (DNS) entries for the MH whenever it moves. The weaknesses of MIP include triangular routes, single point of failure (at the HA), potentially high latency handoffs (when the MIP registration takes a long time because of long latency in the communication path between FA and HA) and potentially high signaling load (if there are many idle MHs moving rapidly between foreign subnets).

1.1.1. MIP with Route Optimization

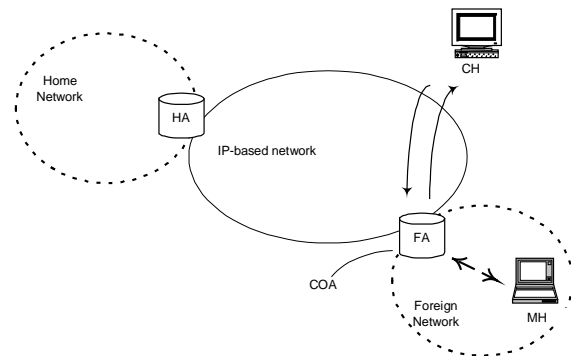


Figure 3: Mobile IP with Route Optimization

To deal with the problem of triangular routing, MIP with Route Optimization [2] (MIP-RO) has been proposed. In order to use MIP-RO, a CH must understand binding updates and be able to tunnel packets to a COA, while the MH must

send binding updates to the CH to update it on the MH's location. The binding update informs the CH of the COA of the MH and hence the CH can tunnel packets to the COA without going through the HA. If there is no binding cache entry in a CH for a given MH, packets still need to go through the HA as is the case in basic Mobile IP. Several new messages, including "binding warning", "binding update", "binding request", and "binding acknowledge", are used to maintain the correct COA binding. While MIP-RO deals with the triangular routing problem, it does not address the issue of micro-mobility management.

1.1.2. MIPv6

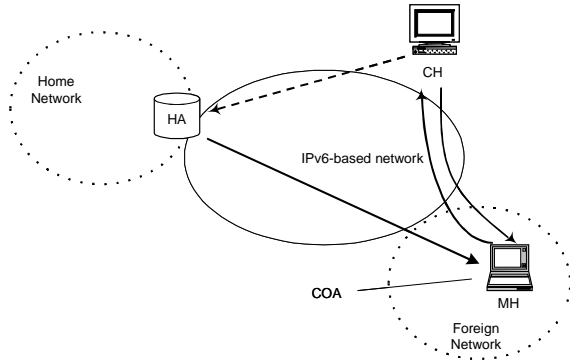


Figure 4: MIPv6

MIP is designed to work with IPv4. MIPv6 is the corresponding framework for IPv6. Since address auto-configuration is a standard part of IPv6, the MH will always be able to obtain a COA routable to the foreign network. Furthermore, there are more than enough address in the IPv6 space that the network designer doesn't need to consider whether to conserve addresses by using an FA address as a common COA for roaming MHs (one of the advantages of using an FA in MIP). Hence there is only a co-located COA mode in MIPv6, and no FAs. To better support route optimization, MIPv6 takes advantage of IPv6 destination options to provide binding updates and binding acknowledgments (replies to binding updates) directly to CHs as well as to the HA.

Three advantages of MIPv6 are apparent: (a) route optimization is facilitated, without needing to be concerned about whether the CHs can understand binding updates, as with MIP-RO; (b) explicit binding updates or MIP registration messages become unnecessary, as the destination options are naturally piggy-backed on IP data packets; and (c) packets from CH to MH need not be encapsulated but are sent directly to the MH with its COA in the source route. The 3rd advantage just mentioned also is due to the way IPv6 makes source routing possible.

It should be noted that the HA is still needed, since the MH need not send binding updates to all CHs. So packets may still be tunneled from the HA to the MH, coming from CHs that do not know the COA of the MH. That is why in the figure, there is a dotted line from CH to HA, for the case that the CH is unaware of the current COA of the MH. While MIPv6 enjoys certain improvements over MIP, it still does not adequately address the micro-mobility problems of MIP.

1.2. The Application-Layer Macro-Mobility Management Alternative

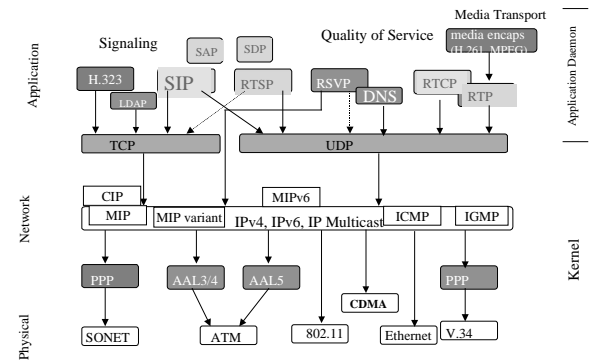


Figure 5: A possible future IP multimedia stack

One of the strengths of a network layer IP mobility solution like MIP is that it is transparent to, and serves, all the application above it. On the other hand, if the mobility solution were to be implemented at a higher layer, e.g. separately by each application, it might be argued that this would be inefficient. However, this may not apply if a widely used application layer protocol were to be able to handle mobility.

Indeed, Session Initiation Protocol [3] (SIP) is rapidly gaining widespread acceptance (e.g. in IETF, 3GPP) as the signaling protocol of choice for Internet (and wireless Internet) multimedia and telephony services. It fits into a possible future IP multimedia stack as shown in Figure 5. SIP allows two or more participants to establish a session consisting of multiple media streams, e.g. audio, video or any other data communications. SIP components, i.e. User Agents, Servers (Proxy and Redirect) and Registrars, provide an application layer mobility management solution while interacting with other network protocols such as DNS and DHCP. While SIP supports personal mobility (see Section 1.2.1 on personal, service and session mobility) as part of its signaling mechanism, its feature set can also be extended to provide adequate means of terminal, service and session mobility. Handoff, registration, configuration, dynamic address binding, location management are key requirements for a SIP based mobility management scheme [4].

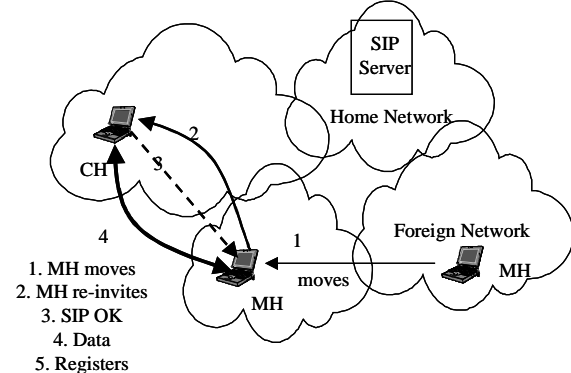


Figure 6: SIP Terminal Mobility Illustrated

In principle, mobility management in the wireless Internet may involve terminal, session, service and/or personal mobility. MIP and its derivatives, variations and auxiliary

schemes are basically network layer solutions that provide continuous media support when nodes move around, dealing with the terminal mobility problem. However MIP and related schemes by themselves do not provide means of device independent personal, session or service mobility. For delay sensitive real-time application, a MIP-based solution suffers from several limitations such as triangle routing, triangle registration, encapsulation overhead and need for a HA in the home network. MIP-RO helps alleviate the triangular routing problem but it also tunnels the binding update through the HA and it requires changes in the operating system of the end hosts. MIPv6 has a lot of similarities with SIP-based terminal mobility in terms of updating the IP address on the CH directly, but it still needs to carry a 16-byte Home Address destination option.

Multimedia traffic can be categorized as real-time or non-real-time, based on delay and loss characteristics. Different transport mechanisms may be used to carry each type of traffic. Most of the real-time traffic should be carried over RTP/UDP whereas non-real-time traffic has traditionally been carried over TCP. SIP-based terminal mobility provides a means of subnet and domain hand-off while a session is in progress. The SIP-based scheme provides a different approach to achieving terminal mobility by means of application layer signaling unlike the traditional MIP approach. This scheme does not rely on the mechanism of the underlying network components in the core of the network, but rather proxy servers instituted by any third party service providers can provide mobility support.

When the MS moves from one subnet to another within the same administrative domain, SIP would support subnet hand-off during the session as described below:

- The MH obtains a new temporary IP address through a protocol like DHCP
- The MH re-invites the CH to its new temporary address. The identifier of the outbound proxy in the visited network should be inserted in the Record-Route field of this SIP INVITE messages.
- In case of domain hand-off a complete registration takes place.

A complete handoff procedure for SIP session would consist of SIP signaling between the corresponding entities and actual media delivery. Delay associated with handoff would consist of several factors such as delay due to layer two detection, IP address acquisition by the mobile, activating the SIP signaling with the new address parameters and actual delivery of media.

If the MH and CH are situated wide apart, then it may take some time for the Re-Invite to reach the CH. Reference [5] proposes some methods similar to many micro-mobility approaches (see Section 1.3) where an RTP translator can be affiliated with a SIP proxy server that would intercept the traffic and would send the media to the current location of the mobile host. Thus RTP translator reduces the end-to-end handoff delay (due to traversal of the INVITE request) to a one-way delay between the MH and the SIP proxy. In cases when both the communicating hosts move during a session each side would have to issue INVITE requests through their respective home proxy servers, where the MHs register their new location address after the movement.

While the RTP translator concept may reduce the micro-mobility problem somewhat, SIP does not in itself provide an

optimized, targeted solution to the micro-mobility problem. Like MIP, it is optimized for macro-mobility. Based on our brief examination of MIP-based and SIP-based macro-mobility management, it can be deduced that a highly desirable property for a micro-mobility scheme is flexibility to work with a variety of macro-mobility schemes, and not just MIP-based macro-mobility.

1.2.1. SIP Support for Other Types of Mobility

In addition to terminal mobility, SIP also supports other mobility concepts, namely personal mobility, service mobility and session mobility. It could be argued that for subscribers interested in these other mobility concepts, SIP offers a more unified macro-mobility management scheme than MIP and its variants and derivatives, which are more limited.

Personal mobility is the ability of users to originate and receive calls and access the subscribed network services on any terminal in any location in a transparent manner, and the ability of the network to identify end users as they move across administrative domains. This is achieved by personal mobility feature inherent in SIP. The URI scheme and registration mechanism are some of the main components used in providing personal mobility. A roaming subscriber is accessible independent of the device the subscriber uses. Service mobility refers to the subscriber's ability to maintain ongoing sessions and obtain services in a transparent manner regardless of the subscriber's point of attachment. Service mobility includes the ability of the service home provider to either maintain control of the services it provides to the user in the visited network or transfer their control to the visited network. Session mobility allows a user to maintain a media session even while changing terminals such as transferring a session that began on a mobile device to a desktop PC after entering an office.

1.3. Micro-Mobility Management

The requirement for MIP registration to be performed every time an MH moves between subnets may cause high handoff latency that could significantly affect data throughput performance of the MH. Various solutions have been proposed to solve this problem. The proposals generally implicitly or explicitly use a concept of micro-mobility regions where Micro-mobility regions comprise numerous subnets, and registrations with the HA are not necessary for movement of the MH within these regions. Registration with the HA would still be necessary for movement of the MH between micro-mobility regions. Typically, MIP would handle the macro-mobility (mobility between micro-mobility regions), while a micro-mobility management scheme would handle micro-mobility (mobility within micro-mobility regions). Micro-mobility management schemes are designed to reduce the high handoff latency of MIP by handling mobility within micro-mobility regions with low-latency local signaling.

1.3.1. Hierarchical Mobility Agent schemes

Micro-mobility solutions that use a hierarchy of mobility agents include MIP with Regional Registration [6] (MIP-RR), and TeleMIP/Intra-Domain Mobility management Protocol [7] (IDMP).

MIP-RR perhaps involves the fewest modifications to MIP. In a foreign network, the two level mobility hierarchy contains the upper-layer GFA (Gateway Foreign Agent) and several lower-layer RFAs (Regional Foreign Agent). All MHs under the GFA share the same COA. When a MH moves to another FA under the same GFA, it only needs to register with the new RFA and with the GFA. This is because its HA already knows how to route packets addressed to the MH to that GFA. It does not need to register with its HA unless it moves under a new GFA.

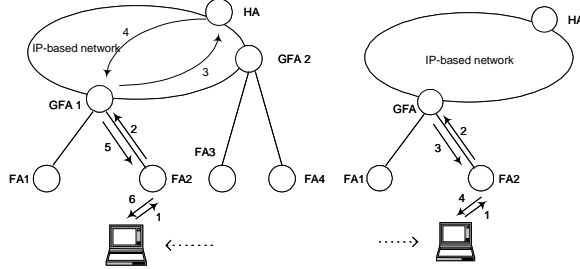


Figure 7: MIP Regional Registration; (a) movement between regions; (b) movement within a region

Suppose an MH moves between subnets under a GFA with which it is already registered. As shown in Figure 7(b), the MH initiates its registration with FA2. Then the registration request is sent to GFA1. Since MN is already registered with GFA1, GFA1 does not initiate a home registration to HA, but just sends the registration reply to the MH through FA2. Since the HA does not need to be contacted in this scenario, MIP-RR reduces the handoff latency caused by registration with Home Agent.

If the MH changes its GFA, it needs to register with its HA. As shown in Figure 7(a), the MH moves from FA3 to FA2, and its GFA is no longer GFA2. The MH sends a registration request to its new RFA, which is FA2, and then GFA1. Because GFA1 is a new GFA, it has to register with the HA. The HA sends the registration reply all the way through GFA1 and FA2 to the MH.

If the MH moves frequently within the same GFA domain, it does not need to perform the time-consuming registration procedure with its HA. Therefore, the average handoff latency is reduced.

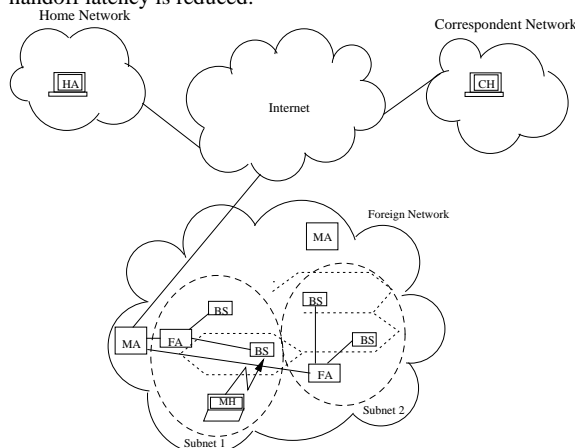


Figure 8: TeleMIP's IDMP

Like MIP-RR, TeleMIP is a scalable and hierarchical IP-based architecture that provides lower handoff latency and signaling overhead than MIP. TeleMIP uses MIP as a macro-mobility management protocol and can inter-work with SIP based mobility management for real-time traffic. Intra-Domain Mobility management Protocol (IDMP) is the micro-mobility management protocol used in the TeleMIP scheme. IDMP offers intra-domain mobility by using multiple COAs that are taken care of by Subnet Agents (SA) and the Mobility Agent (MA) at the subnet and domain level respectively. SAs are FAs or DHCP servers at the subnet level that provide an MH with a locally scoped address, and are analogous to MIP-RR RFAs. The locally scoped address provided by an SA identifies the MH's location within the domain. An IDMP MA is similar to a MIP-RR GFA and acts as a domain-wide point for packet redirection.

The serving MA provides an MH with a global COA that stays constant as the MH moves within the domain. MAs are distributed within the domain. Unlike MIP-RR, multiple MAs can be provisioned for load balancing and redundancy within the domain. All packets from the global Internet are tunneled to the MA (or one of the MAs, in the case of multiple MAs), which acts as a single point of enforcement and accounting. The serving MA forwards packets to the MH using regular IP routing, with the local COA (co-located or FA) as the destination. It does this by un-encapsulating the packet and then performing a second encapsulation of the IP packet, with the local COA as the destination address.

On subsequent movement within the domain, the MH only obtains a new local COA. At that point there is no need to update the HA or CHs. However, the MH does update its MA with its new local COA. By limiting intra-domain location updates to the MA, it reduces the latency associated with intra-domain mobility. This two-level mobility management scheme allows the use of private addressing and (possibly non-IP address scheme) within the provider's own domain. In addition IDMP also provides the added advantage (over MIP-RR) of dynamic load balancing, paging and fast handoff scheme using multicast within a domain.

2. HBR Overview

A class of micro-mobility management schemes is those employing host based routing (HBR), including CIP [8], HAWAII [9] and MMP [11]. HBR schemes for micro-mobility could be considered a class of auxiliary schemes that deal with the handoff latency problem of MIP. However, they have grown beyond being just a class of auxiliary schemes, e.g. MMP is designed to be usable with SIP mobility for real-time traffic, and a MIP variant for non-real-time traffic. This flexibility is an advantage over micro-mobility schemes based on a hierarchical MIP structure, e.g. MIP-RR or TeleMIP. A second major advantage of HBR schemes is that they offer the lowest latency networking re-routing solution for micro-mobility management. Heuristically, this is because hierarchical MIP-derived schemes like MIP-RR and IDMP only reduce the latency problem inherent in MIP registration. The GFAs or MAs still need to be over a reasonably large area to be scalable and cost-effective. Location updates still need to reach them in these schemes. On the other hand, with HBR schemes, take an optimal path to the closest node that should handle the location/route update, namely the crossover node, as well be explained shortly. Later, the performance

results in Section 4 will also support these assertions with numerical results.

The distinctive characteristics of HBR schemes for micro-mobility are that (a) host-based routing is used within the micro-mobility regions; (b) very low-latency handoffs are possible since the update message only needs to propagate to the crossover node for both the location management and routing update to be completed; and (c) one or more special nodes (known as gateways or root nodes) are used as the demarcation point between each micro-mobility region and the rest of the Internet.

With host-based routing schemes, forwarding behavior is specified separately for each host. For example, nodes may route packets according to tables or caches indexed by unique host identifiers (e.g. their IP address). MMP, HAWAII and CIP are examples of HBR where forwarding caches are used, indexed by host IP addresses. HBR differs from group-based routing schemes, where forwarding behavior is specified for groups of hosts. For example, nodes may route packets according to tables or caches indexed by group identifiers (e.g. IP address prefix and netmask), i.e. packets with different destination addresses, but where the destination addresses match the prefix and netmask, will be routed in the same way.

A critical advantage of HBR schemes for micro-mobility management is that location management and routing can be integrated. In a framework like MIP, location management is handled by MIP registrations, while routing is overlaid on the existing IP-based network routing. The location management requires possibly long-latency registrations to a potentially distant HA whenever subnet boundaries are crossed, while the use of overlay routing over standard IP routing creates problems like triangular routes and encapsulation overhead. Using HBR, on the other hand, gives us the power to update routes simultaneously with location management, precisely because the routes are host specific and do not affect routes to/from other MHs when changed. For lowest latency handoffs, the intuition would be to update the routing information for an MH at the closest intersection with the old route (also known as the crossover node), whenever it performed a handoff. And this is indeed what happens with the HBR schemes for micro-mobility management.

Rather than present three separate descriptions of the three HBR schemes (MMP, HAWAII, CIP), we shall introduce a generic, bare-bones HBR scheme in Section 2.1, and then discuss differences between MMP, CIP and HAWAII in Section 2.2. This approach brings out the essence of the HBR schemes before explaining the minor differences between them. A comparison with ad-hoc routing protocols that use HBR-like routing follows in Section 2.3.

2.1. A Generic HBR Solution for IP Micro-Mobility

We present here a generic HBR scheme to illustrate the essential workings of HBR in facilitating fast low-latency handoffs for IP micro-mobility. Actual protocols (CIP, MMP, HAWAII) differ from this generic scheme in some aspects, especially where enhancements are provided by the actual protocols. For example, actual protocols use various enhancements to provide more seamless handoffs. There may be multiple non-overlapping micro-mobility domains in a given network, but here we only consider how movement within a single micro-mobility domain is handled. Movement

between micro-mobility domains may be handled by a macro-mobility protocol like MIP or SIP, with appropriate modifications.

In the generic HBR scheme, each HBR micro-mobility domain has one and only one root router that serves as the interface between the micro-mobility domain and the rest of the IP-based network. The rest of the infrastructure nodes in the HBR micro-mobility domain are arranged in a strict inverse tree structure beneath the root router. In other words, every node is a child of one and only one other node (possibly the root router), and may be a parent of zero, one or more other nodes. The inverse tree structure can be seen in Figure 9, which shows an abstraction of a generic HBR scheme for IP micro-mobility management. Some of the nodes at the bottom of the hierarchy are base stations, with wireless interfaces, and whose coverage areas are called cells. For a well-designed HBR domain, the cells would adequately cover the great majority of the geographical region with which the domain is associated. Each infrastructure node other than the root router has one upstream interface and zero, one or more downstream interfaces. The upstream interface is the interface towards the root router, while the downstream interface(s) is/are towards the base stations and MHs. The way that nodes know which are the upstream, and which are the downstream, interfaces is unspecified in generic HBR. However, one way this may happen is by listening to beacon messages sent periodically down by the root router (the interface through which the beacon arrives is recorded as the upstream interface and is used as the next-hop for routing of any packet to the root router; the rest are considered downstream interfaces). Another way might be if the HBR scheme is implemented as an overlay over standard IP routers, as with HAWAII (see next section). Alternatively, they may be pre-configured.

When a MH first enters an HBR micro-mobility domain, the network would use access control mechanisms, such as authentication of the MH. Some form of macro-mobility signaling would be initiated, so that macro-mobility can be handled by the relevant protocol. The initiation of this signaling is typically by the MH or the root router. This initial signaling is unspecified as far as generic HBR is concerned, and will become clearer in Section 2.2. The MH may use its home IP address, or it may obtain a temporary address for use in the HBR domain. The temporary address, if applicable, is often called a COA, although it is not necessarily exactly like a MIP COA. It may be obtained through a variety of unspecified means, and may be a globally routable address of the root router, or be co-located. In any case, the MH uses one IP address as long as it is in the same HBR domain, even when its moves from cell to cell.

Upstream routing of packets (from MH to gateway) is simple – once packets from a particular MH are admitted into the micro-mobility domain infrastructure, each node (base stations included, root router excluded) merely forwards upstream packets to the root router through their upstream interface. Therefore, the root router eventually gets upstream packets. It then routes them normally through the IP-based network, or it sends them back through its micro-mobility domain infrastructure, depending on the destination address.

Downstream routing of packets depends on routing caches maintained in each node. By definition, for HBR, there are separate cache entries for each MH. The cache entries are set whenever upstream data passes through a node.

The node will read the source IP address to identify the MH, and then bind it in the cache to the interface address (plus MAC address, if necessary) of the incoming packet, with the assumption that that is the right interface (and MAC address) to use for downstream packets destined to that MH. In order to facilitate simpler handoffs (as will be explained shortly), the cache entries have a soft state and need to be periodically refreshed. Since it cannot be assumed that upstream data is always being transmitted, periodic control packets called route updates are transmitted by each MH to refresh the cache entries after a period of no upstream transmissions from that MH.

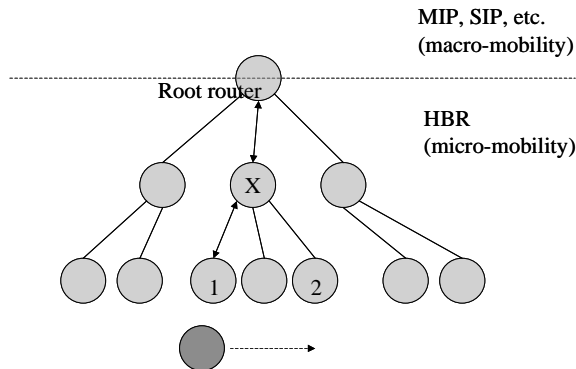


Figure 9: Abstraction of generic HBR scheme for IP micro-mobility management

To aid in our discussion of how the generic HBR scheme would handle handoffs, Figure 9 shows an MH (the unattached node at the bottom of the figure) moving from the cell covered by the base station labeled 1 (BS1) to the cell covered by the base station labeled 2 (BS2). The node marked “X” is the crossover node for this handoff. The crossover node is defined as the lowest node that is in the upstream of the paths from both base stations to the root router. In some cases, it may be the root router itself, but often the crossover node would be below the root router. The MH would initiate handoff by sending a route update through the new base station, which is BS2 in this case. As this route update propagates up to the root router, for each node below the crossover node, the routing cache may not have any binding for the MH, and the appropriate entry is added. At the crossover node, the critical switch occurs as the binding for the MH is updated to point towards the interface heading towards BS2. As the route update continues up to the root router from the crossover node, the routing cache entry at each node is updated as normal, as though no handoff had occurred.

In its purest form, HBR does not require any signaling with the old base station, BS1, nor any of the nodes between BS1 and the crossover node (some signaling occurs at through these nodes in HAWAII, though, in one of the enhancements to provide seamless handoffs). Since the routing caches contain soft-state entries, the entries will naturally time-out and be removed without the need for additional signaling. This is one of the reasons why soft-state routing cache entries are used. Another reason is that the MH does not need to perform any kind of de-registration when it leaves the HBR domain, which it would need to do to update the routing caches if they had hard-state entries. Furthermore, even if the MH were willing to do that, if for any reason the

MH loses connectivity or crashes, the hard-state entries would not be removed.

The generic HBR solution could also have some kind of paging scheme, to provide paging gains over MIP (see Section 3.1.1). One way this could be done would be by using paging caches, similar to but parallel to the routing caches. The paging caches would have longer expiry times, so that they need not be refreshed as often as routing caches. When an MH is idle, it only needs to send paging updates occasionally.

2.2. Comparison Between HBR Schemes

In this section, we discuss specific differences between HBR schemes like HAWAII, CIP and MMP. Although each of these protocols has much in common with the generic HBR scheme just described, their design goals are not all the same. The differences in the design objectives explain some of the differences between the schemes. For HAWAII, the design goals are [9]: to limit disruption to user traffic, to enable efficient use of access network resources, to enhance scalability by reducing updates to the HA, to provide intrinsic support for QoS and to enhance reliability. Another implicit goal is to provide ways to achieve seamless handoffs in a range of network scenarios. For CIP, the design principles are [8]: to use universal building blocks as nodes, to be a plug-and-play solution, to be scalable, to minimize the burden on the MH, and to support passive connectivity. For MMP, the design principles are: to limit disruption to user traffic, to be robust, reliable and survivable, to enable efficient use of low-bandwidth access network resources, to be a plug-and-play solution, to be scalable, to minimize the burden on the MH, to support passive connectivity and to facilitate QoS support.

The MH in HAWAII is a MIP client. The HAWAII domain looks like an FA to the MH, and its root router looks like an FA to the HA of the MH. On the other hand, the MH in CIP needs to use CIP signaling, e.g. route updates and the root router takes care of the MIP signaling. As for MMP, the MH needs to use MMP signaling, and the root router acts on behalf of the MH to perform signaling for the macro mobility scheme, whether it be MIP, SIP or something else. The advantage of having the MH be an ordinary MIP client is that no changes need to be made to MHs that already have MIP client software. However, this choice is not made in CIP, since that would go against the “universal building block” principle because then the BSs would have extra IP-level functionality in addition to CIP node functionality, i.e. they would have to send route and paging updates to the root router on behalf of the MH. Furthermore this would reduce the ability to put together a plug-and-play wireless IP access network. This second reason applies to MMP as well. In addition, however, MMP MHs cannot be MIP MHs, since MMP is designed to work with a variety of macro-mobility schemes, not just MIP.

Even though both CIP and HAWAII are designed to work with MIP as the macro-mobility management scheme, each works with a different mode of MIP. In HAWAII, the MH acts as a MIP MH in co-located care-of-address (COA) mode, while in CIP, the gateway acts as an enhanced FA that takes care of MIP registration on behalf of the MH. Another major choice is whether the HBR scheme is implemented as an overlay over standard IP routers running in intra-domain

routing protocol (e.g. RIP, OSPF), as with HAWAII, or whether the HBR routing is the only routing that the HBR nodes are capable of, as with CIP. The overlay approach allows more sophisticated seamless handoff techniques, as it allows HBR nodes to communicate with one another to forward buffered packets, etc. It also may be able to rely on some of the reliability mechanisms of the underlying intra-domain routing protocol, and may be easier to implement using existing equipment. However, it is not as lightweight and scalable over a range of wireless access environments as the non-overlay approach, since full-fledged IP routers are used. The CIP designers claim that their non-overlay, simpler approach can be implemented with low-cost “layer two switches”

One of the design goals of MMP is to be robust, reliable and survivable. Because of the efficiency in signaling and distribution of the routing information in the generic HBR scheme, only the nodes along the path between the serving BS and the root router know how to route to the MH. This could be result in disruption in communications if a node or link fails, or if the root router fails. The problem is shared by CIP, and to some extent, by HAWAII. With MMP, however, there is the option to use multiple root routers, and there is also the option for some or all the nodes to have more than one parent node. This provides for robustness and survivability, at the cost of a little more complexity and signaling traffic within the HBR domain.

Another difference between the schemes is in how handoffs are treated. The basic HBR handoff scheme is equivalent to the hard handoff scheme of MMP and CIP. Although HBR schemes can achieve fast, low-latency handoffs by using only their very fast local updates, much faster than MIP, packets could still be dropped. Therefore, various seamless handoff schemes have been proposed to improve performance even further. In HAWAII, there are 4 schemes for setting up the new path when handoffs occur. With the Multiple Stream Forwarding (MSF) and Single Stream Forwarding (SSF) schemes, the old BS receives the initial handoff message and signals with the new BS to set up a path with the new BS to forward packets there that have been buffered at the old BS. Using MSF could result in multiple out-of-order streams arriving at the BS through the new BS, as some earlier packets being forwarded from the old BS to the new BS may arrive later than newer packets from the crossover node to the new BS for very short periods of time. SSF is more sophisticated, and results in a single stream of packets being forwarded to the new BS. The crossover node in this case needs to be informed by the old BS when it has cleared its buffers of packets for the MH, and only then would the crossover node switch packets for the MH over to the new BS. With the other two handoff schemes in HAWAII, Unicast Non-Forwarding (UNF) and Multicast Non-Forwarding (MNF), it is the new BS that receives the initial handoff message. UNF, in the network, is like the basic hard handoff except that the old BS is informed through signaling from the new BS and it sends an acknowledgment back to the MH. The difference from simple hard handoff is that the MH is assumed to be able to communicate with both BSs during the handoff period, to reduce packet losses associated with hard handoff. As for MNF, it is like UNF except it uses a special dual-cast scheme (from crossover node to both BSs) for a short period of time from the time the crossover node receives the handoff message so the MH

doesn't have to talk to two BSs simultaneously. On the other hand, the only seamless handoff scheme with CIP, semi-soft handoff, is somewhat like MNF. The difference is that the MH switches back to the old BS after sending the initial handoff message, to reduce packet loss while the message is traveling to the crossover node.

There are also other differences between the actual protocols, such as paging schemes, but for more details, the reader is referred to the source documents, e.g. [12].

2.3. Comparison with Ad-hoc Mobility Schemes

Another class of routing problems where various non-hierarchical, HBR-like schemes have been proposed for handling mobility is that of ad-hoc routing in Mobile Ad-hoc Networks (MANET) [13]. In ad-hoc networks, the nodes are not arranged in a fixed infrastructure, typically because they are mobile and are constantly changing positions with respect to one another. The distinction between the fixed infrastructure and mobile hosts may vanish, and every node may well be a mobile router. The lack of a fixed infrastructure makes the route discovery problem more difficult than in the case of the HBR schemes for micro-mobility management. We are not aware of any publications that compare HBR domains with MANETs, and this section is an attempt to do that.

A variety of hierarchical routing protocols have been proposed, using concepts of ad-hoc clusters of nodes based on factors like proximity and geographical location. However, such protocols may depend on assistance from the Geo-Positioning System (GPS) [14] or depend on the existence of a core of “backbone” nodes (e.g. Core-Extraction Distributed Ad-Hoc Routing, CEDAR, [15]) or use some heuristics for selecting cluster heads. Other ad-hoc routing protocols are non hierarchical but flat, which is closer to the routing within HBR micro-mobility domains. Among the flat ad-hoc routing protocols, some like Ad Hoc On-Demand Distance Vector routing (AODV) are reactive. These differ from our HBR schemes in that the routes are computed in an on-demand manner, as needed.

The flat, proactive ad-hoc routing protocols such as Destination-Sequenced Distance-Vector routing (DSDV) may be closer to the HBR schemes. In DSDV, all nodes maintain a routing table that contains separate entries for all the possible destinations, which are periodically refreshed. Two differences between DSDV and HBR for micro-mobility are noticeable. While in DSDV, all nodes maintain a routing table for all the destinations, in HBR only the infrastructure nodes maintain these tables, and then only for MHs being served by one of their children or descendent nodes. The other difference is that the refresh problem is an order of magnitude more challenging in DSDV, since the nodes all are moving. Infrequent transmissions of full dumps are needed, where full dumps contain all routing information, and periodic transmissions of incremental packets are used to relay information on changes occurring since the last full dump. On the other hand, for HBR, route updates are always incremental and specific to mobile hosts, very quickly providing up-to-date information on MH location after movement occurs. Even the periodic refreshes are meant to be for optimizing network resource utilization more than to handle significant network topology changes.

Despite these advantages of HBR, ad-hoc schemes must still be used in cases where the mobility situation demands it. However, whenever HBR can be used, it should, e.g. in less mobile or semi-mobile networks (e.g. a tactical network where some “infrastructure” nodes move infrequently and remain on the same hierarchy even after moving), since it also has additional advantages over ad-hoc schemes. The root router provides a natural transition point between the micro-mobility region and the macro network. Moreover, MANETs run into a scalability problem with as few as 50-100 nodes, because of all the updating and exchange of route information that goes on, whereas HBR domains can handle thousands of nodes [10].

An interesting and open area of research is where the crossover between an HBR network and a MANET might occur. For situations where there is a relatively stable infrastructure, HBR makes sense, whereas ad-hoc routing protocols would need to be used in more mobile, fluid networks. A key question is how to qualify and quantify what is meant by a relatively stable infrastructure with core nodes, in which the network can take advantage of the core nodes to reduce signaling overhead, etc. Ad-hoc routing protocols like CEDAR that assume that core, high-bandwidth backbone nodes can be found, are closer to HBR in a sense. One could imagine a self-organizing protocol where nodes perform some self-discovery of the network and switch into an HBR mode or an ad-hoc routing mode depending on certain conditions. It could periodically check if the conditions have changed, and switch modes if necessary.

In summary, the ad-hoc routing protocols may at first appear similar to the HBR schemes for micro-mobility management because host-based routes are used. However, the routes may be hierarchical and even if they are flat, they may be on-demand, unlike the HBR schemes. The flat, proactive ad-hoc routing protocols may be the closest to HBR schemes, and it is here that some fundamental differences are most evident. The differences are because of the different mobility assumptions. The ad hoc networking problem is a more difficult problem, and ad hoc routing solutions should be used where ad hoc networking assumptions do apply. However, when networks are not as mobile, or are semi-mobile, HBR solutions should be considered because they are more efficient (in terms of the signaling for updating routes) and optimized for those scenarios. They also provide a natural way to get out of the mobile part of the network, and are more scalable than MANETs.

3. Performance Issues

Performance is examined qualitatively in Section 3.1. Previously reported quantitative results are discussed in Section 3.2, as a prelude to discussing our performance results in Section 4.

3.1. A Qualitative Perspective

The major goal of HBR is providing fast, low-latency handoffs. This may result in fewer packets dropped during handoffs, leading to less disruption of UDP traffic and better TCP throughput performance, etc.

While obtaining significant reductions in handoff latency compared to MIP is a major accomplishment of micro-mobility management schemes, other advantages have been claimed. These include:

- (a) reduction in signaling overhead through paging concepts;
- (b) reduction in packet header overhead in the low-bandwidth radio access network through not using encapsulation;
- (c) easier to integrate with QoS provisioning
- (d) better use of scarce IP address resources by using fewer IP addresses than MIP or hierarchical MIP derivatives

These advantages will be discussed, qualitatively, in sections 3.1.1 to 3.1.4. Other issues, such as scalability, will be discussed in sections 3.1.5 to 3.1.6.

3.1.1. Paging Gains

The idea behind paging gains is that MIP does not differentiate between active and idle MHs. It requires that MHs go through the registration process whenever MHs move between subnets, regardless of the activity level of the MH. There are two problems with this. Firstly, the signaling overhead is high, even if the MH is idle (communications-wise) but moving around rapidly. Secondly, each of the registrations with movement between subnets would consume power from the MH's battery. The first of these problems is largely dealt with in an HBR domain even without paging, because macro-mobility signaling would not need to be invoked upon every handoff occurring. However, the second problem can be dealt with using an idea (paging) borrowed from traditional wireless cellular networks.

Networks such as Global System for Mobile communications (GSM, [16]), have long differentiated between active and idle states of a MH. When an MH is idle, it registers less often with the network, the tradeoff being that the network knows the MH's position with less precision, and needs to page the MH to reach it when it needs to communicate with it. The less often the MH registers, the less precisely the network would know its position and the larger the paging area. The HBR schemes implement variants of the paging concept. However, this is not a fundamental flaw of MIP, nor a fundamental advantage of HBR schemes. Moreover, a paging extension to MIP has been proposed [17].

3.1.2. Reduction in Packet Header Overhead

MIP adds at least 8 to 12 bytes per packet for minimal encapsulation, and more for alternative encapsulation schemes. Under circumstances, e.g. with small packets, this can make a significant difference compared to a scheme without encapsulation overhead [18]. Since micro-mobility regions are often in wireless access networks where bandwidth efficiency may be at a premium, it is an advantage of HBR schemes that there is no encapsulation overhead. This advantage is not only over MIP, but also over other non-HBR micro-mobility schemes that use encapsulation, such as TeleMIP/DMP and MIP-RR.

3.1.3. QoS

In the Integrated Services (IntServ) model for providing QoS, resource reservation protocols like RSVP are used to reserve network resources. The resource reservation, however, assumes that the endpoints have unchanging IP addresses. When an endpoint changes IP address (e.g. an MH obtains a new COA), the old reservations cannot be used, and new reservations need to be made. The problems with using MIP

with RSVP, then, are that with each handoff to another subnet, there is an interruption in the availability of resources for QoS, and unnecessary signaling is occurring. With HBR schemes, MHs keep the same IP address within an HBR domain, providing a more stable endpoint for RSVP.

3.1.4. Use of IP Addresses

There is a shortage of IP addresses in the IPv4 address space. Using MIP for micro-mobility would require a pool of IP addresses to be set aside for use as COAs in every subnet. Since MHs with HBR can keep one IP address as they move within an HBR domain, a pool of COAs can be set aside for the entire HBR domain, if necessary, resulting in a more efficient use of IP addresses. Interestingly, TeleMIP/IDMP also has a way to conserve IP addresses, i.e. by using a private address space within the micro-mobility region.

3.1.5. Scalability

HBR has a potential scalability problem in that the forwarding cache grows linearly with the number of hosts. To deal with the scalability problem, one solution is to use a group-based routing scheme. The Internet is an example of a network with group-based routing. Furthermore, the groups are hierarchical, with smaller groups as subsets of larger groups, providing a very efficient and flexible way to specify routing behavior. At the minimum, a routing table may contain a default route that specifies how to route all packets.

The Internet uses a hierarchical routing scheme because host-based routing does not scale. However, an alternative way to deal with the scalability issue is to restrict the number of hosts involved, for example to just the hosts roaming within a certain region. Constraints like radio link capacity will put practical limits on the number of hosts involved. Then HBR is usable. For the micro-mobility schemes that use HBR, it works because it is confined to a definite region, with a gateway between that region and the rest of the Internet. In that definite region, there are practical limits on how many mobile hosts there can be simultaneously.

3.1.6. Communications Between Two MHs in the Same HBR Domain

For cases where the CH is not in the same HBR domain as the MH, the routing within the HBR domain is optimal in both directions. Uplink packets go straight to the root router as they should, and downlink packets go straight to the correct base station, as they should. However, what happens when the CH is another MH in the same HBR domain? The way that CIP, HAWAII and MMP have been specified at this time, the route would be through the root router and down to the other base station. Even if the two MHs had a crossover node that was the direct parent of their respective base stations, or potentially if the two MHs were in the same cell, packets would still go to the root router. The reason is that the HBR nodes along the path simply forward any uplink packets to the root router regardless of final destination.

Should this routing “inefficiency” be removed? A straightforward solution might be to check the destination address of every uplink packet with the contents of the routing cache to see if it should be forwarded down rather than to the root router. However, the problem is that the uplink forwarding would become less efficient. It is unclear

if the tradeoff is worthwhile, given that the percentage of traffic from one MH to another in the same domain might be very low. Even if that percentage is not that low, the “inefficient” route through the root router is not a serious inefficiency because the added latency would not be very high, and would be naturally bounded by the size of the HBR domain. Nevertheless, a solution has been proposed that introduces the concepts of optimizing CIP node, proxy route-update packet and optimizing teardown packet [19].

3.2. Quantifying Performance

Attempts have also been made to quantify the performance of HBR schemes. Reference [2] describes an experimental prototype and measurements made thereupon, which show that TCP throughput decreases as the handoff rate increases. Measurements also show how semi-soft handoff (a.k.a. advanced binding handoff) performs better than hard handoffs, experiencing less of a TCP throughput decrease as the handoff rate increases. It also contains analysis on fine-tuning the choice of route update and paging update intervals, and some scalability analysis.

Reference [9] provides performance results using a novel network simulator developed at Harvard. It compares HAWAII with Mobile IP in terms of average number of dropped packets (UDP case) per handoff. TCP throughput of HAWAII is compared with that of MIP as handoff frequency varies from 0.5 to 4 times per second. It also contains various other results on processing load at various nodes.

While the first order performance improvements are based on much reduced handoff latency (over MIP) using HBR, second order performance improvements may be available through various additional optimizations. For example, attempts at seamless (or almost) handoffs, including semi-soft, MNF, UNF,

Our simulation results extend, as well as complement, the performance results of [2] and [9].

4. Performance Results

In this section, performance results (mostly from computer simulations using the ns-2 simulator, but also including some analytical and laboratory prototype results) are discussed. Section 4.1 introduces the simulation environment. Various results are discussed in Section 4.2, which is divided into two main types of simulations, namely UDP simulations and TCP simulations. Finally, Section 4.3 describes some measurement results from our laboratory prototyping.

4.1. Simulation Environment

The base simulation set-up for our simulations of HBR schemes for micro-mobility management is illustrated in Figure 10. A simple wireless model is used that assumes perfect overlapping coverage, no propagation delay, and no transmission errors. Furthermore, handoffs are smooth and instantaneous at layer 2 and below. The link latency for the plain (straight-line) links in the micro-mobility domain are 2 ms, whereas the link latency of the dash-dot links in the Internet are 10 ms. The link bandwidths are 375 kbit/s within the micro-mobility region, and 1.544 Mbit/s in the “Internet”. Routing and paging cache entries need to be refreshed, and the route update and paging update intervals we use are 3 s

and 60 s, respectively. The size of each update packet was 100 bytes. Lightly-loaded network conditions were simulated, with only one MH and one CH. TCP Tahoe [20] is used as the transport protocol, with a flat application data rate of 200 kbit/s. The application is assumed to not be delay sensitive, i.e. it is equally acceptable for the instantaneous throughput to fluctuate a lot as it is for the instantaneous throughput to be relatively constant, provided that the average throughput is the same. The handoff rate is once every 5 seconds on average, with exponentially distributed inter-handoff intervals, and the handoffs are back and forth between the two base stations labeled BS 1 and BS 2. Each simulation was run at least as long as needed for 200 handoffs to occur.

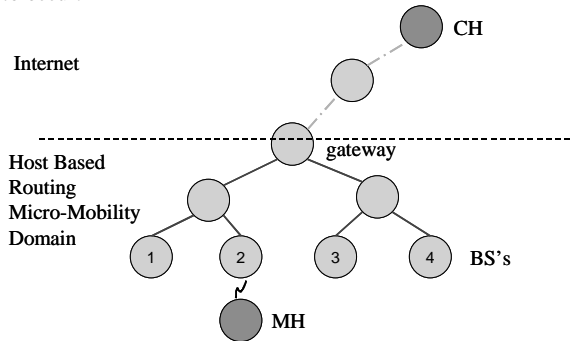


Figure 10: Base simulation set-up for HBR micro-mobility simulations.

Simulations were run comparing performance in terms of (a) average number of packets dropped per handoff for UDP traffic; and (b) the TCP throughput. In different simulations, the effects of varying link latencies, handoff frequencies, application data rate, link bandwidths and other parameters were investigated. While TCP Tahoe [20] was the default TCP used, TCP Reno [20] was also simulated for comparison. Since there have been various seamless handoff schemes proposed, it has been decided that for this paper only the basic HBR hard handoff, and one representative seamless handoff scheme, the semi-soft handoff scheme, be simulated. The reason for including basic HBR hard handoff is to bring out the performance gains of HBR schemes resulting simply from the low-latency route updates even without any auxiliary schemes for seamless handoffs. The reason for including one representative seamless handoff scheme is merely to confirm and illustrate that such schemes do indeed help further improve performance. In this, the choice of semi-soft handoffs is somewhat arbitrary, partly because its performance is expected to be somewhat moderate compared with the other seamless handoff schemes.

The performance of HBR schemes has been compared with the performance of MIP where the base simulation set-up for MIP simulations is shown in Figure 11. The wireless model, handoff model, and other parameters are almost identical to those for the HBR simulations, and the topology is similar to that in Figure 10, in order to allow for meaningful comparisons. Each base station now becomes a different subnet, and has either an FA or an HA. The MIP registration request and reply messages are each 100 bytes long. Minimal encapsulation [21] is simulated, adding only 12 bytes to the un-encapsulated packets. MIP periodic registrations are not simulated in this model, because the

intervals typically are long, in the order of many minutes, and so would hardly impact the simulation results. As for the dashed link, the link latency on that link was varied (from the original 2 ms to 10 ms to 100 ms), to simulate the real possibility that the HA is further away. It should be noted that the case that the dashed link has only 2 ms of latency should be useful to indicate the performance of either (a) MIP where the HA is very close to the FA; or (b) micro-mobility management by MIP-RR or TeleMIP/IDMP (without seamless handoff enhancements). In the case of MIP-RR or TeleMIP/IDMP, the HA in Figure 11 would be analogous to the GFA or the MA, in that it is very close to the FA (which would be the RFA or SA for MIP-RR or IDMP, respectively).

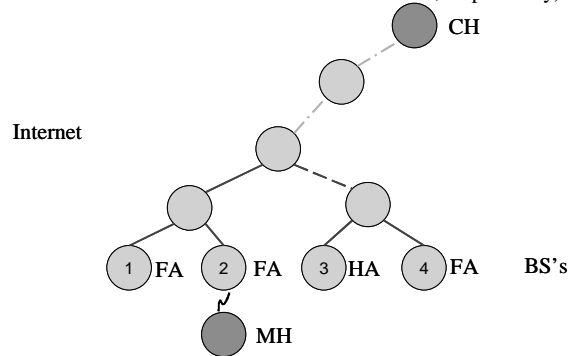


Figure 11: Base simulation set-up for MIP micro-mobility simulations

4.2. Simulation Results

4.2.1. Constant-Bit-Rate UDP Traffic

Constant-bit-rate UDP traffic was applied from CH to MH to investigate HBR performance compared with MIP, in terms of number of packets dropped per handoff. The simulation environment was as described in Section 4.1. It could be expected that the results would be little impacted by varying the handoff rates, for reasonable handoff rates (not too large). Indeed, simulations verified this assumption. Recall that in the base case for both HBR and MIP simulations, described in Section 4.1, the application data rate is 200 kbit/s. For the base UDP simulations, this rate is accomplished by sending packets 1000 bytes every 40 ms. This can be described as a “heavy traffic” scenario, since the 375 kbit/s links are more than 50% loaded. A “light traffic” scenario will also be investigated next, where the application data rate is 20 kbit/s.

Some results for the simulations in the heavy traffic scenario are shown in Figure 12. As can be expected, increasing the link bandwidth would decrease the number of packets dropped per handoff, because the packets spend less time in transit. Similarly, increasing the size of the update packets would increase the number of packets dropped per handoff, as the update packets will take longer to arrive at their destination. These two effects were investigated by simulating micro-mobility domains with 10 times larger bandwidths, as well as cases of 10 times larger update packets. The plots marked “3.75Mbps” have micro-mobility domains with link bandwidths of 3.75 Mbit/s, while the plots marked “1kbyte” use large update packets 1 kbyte long. Curves whose labels are prefixed by “MIP” are those in which

MIP was run, and the rest use HBR by default (this statement applies to the rest of the performance results as well, not just this figure). The x-axis shows the link latency, which is the transmission latency of the straight-line links in the simulation set-up. The y-axis shows the packets dropped per handoff.

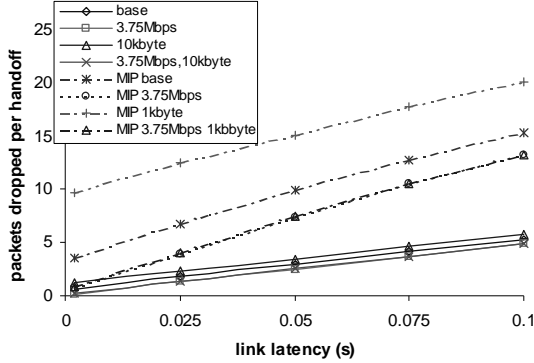


Figure 12: Packets dropped per handoff in a heavy-traffic scenario

Since this is a “heavy traffic” scenario, the packets dropped per handoff can quite significant. However, it is the relative performance of HBR and MIP that is of interest. With MIP, the performance is worse (more packets dropped per handoff), even for this best-case MIP scenario where the FA and HA are close together. Also noteworthy is the spread between the performance of the different cases when the link bandwidths are modified and/or the update packet size is modified. With HBR, the spread is very slight, from the best case (large link bandwidths and regular size update packets) to the worst case (regular size bandwidths and large update packets). With MIP, the spread is much more pronounced, showing that it is much more sensitive to the settings of such variables.

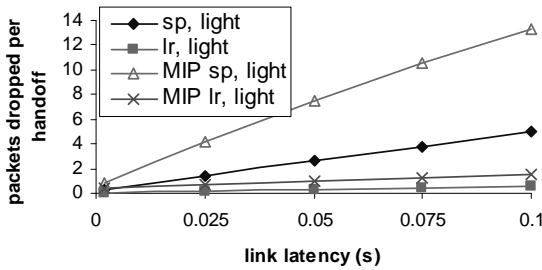


Figure 13: Packets dropped per handoff in a light-traffic scenario.

Turning our attention to a “light traffic” scenario, we compare two ways of reducing the traffic load. Firstly, reducing the size of the UDP packets 10-fold but keeping the rate of the packets at one every 40 ms can reduce traffic load from 200 kbit/s to 20 kbit/s. Secondly, keeping the same packet size (1000 bytes), but reducing the rate of the packets by 10-fold, can also reduce traffic load to 20 kbit/s. As would be expected, the number of packets dropped per handoff in the first case would be significantly higher than in the second case. Indeed, this is the case, as illustrated in Figure 13. In this figure, “light” refers to the light traffic scenario, and “sp”

refers to “short packets” whereas “lr” refers to “low rate” (the two ways to reduce the traffic load). It can be seen that for both HBR and MIP, “sp” has more dropped packets per handoff. However, as in the heavy traffic scenario, the variation is greater for MIP, in addition to the actual numbers being worse.




Figure 14: Comparing simulations of random handoffs with analytical model

Having had a flavor of some simulation results, it is appropriate to address the issue of how generally the results can be interpreted. One question that arises is whether the results would be limited only to the unlikely case that an MH just moves back and forth between two BSs, rather than more realistic movement, e.g. randomly moving between all four BSs in the base simulation set-up. This is a valid question, but we claim that the “to-and-fro” movement results are more generally useful because they can be extended to more general movement patterns according to the following methodology.

- Suppose the HBR domain is an n -tier domain, so the crossover node could be 1 level above the BSs, or up to n levels above the BSs.
- For each level from $i=1$ to n , to-and-fro handoffs are simulated between any two BSs whose crossover node is i levels above the BSs. Let λ_i be the number of packets dropped per handoff.
- Given the HBR domain (or sub-region within it) and mobility pattern, compute $E[h_i]$, the expected number of i th-tier handoffs, for each $i=1$ to n , and define $h = \sum_{i=1}^n E[h_i]$
- The overall expected number of packets dropped per handoff is then computed as the weighted average

$$\lambda = \sum_{i=1}^n \frac{E[h_i]}{h} \lambda_i \quad (1)$$

For example, for our base simulation set-up, $n=2$. For a handoff distribution that is uniform over the other 3 BSs, the crossover node would be 2 levels up for 2 of them and 1 level up for the other BS. Hence the result λ_2 should be weighted by $2/3$ and λ_1 by $1/3$. It would be expected that the resulting value would be similar to what could be obtained by actually simulating handoffs under such conditions of random motion. Figure 14 shows the results. The curve labeled “1-up” is for to-and-fro handoffs when the crossover node is 1 level above the BSs. The curve labeled “2-up” is similarly for to-and-fro handoffs when the crossover node is 2 levels above the BSs. The curve labeled “analytical” is where the performance in the more general random case is computed analytically according to Equation (1). The curve labeled “random” is for

the same case, but with results from actual simulations. It can be seen that “analytical” and “random” are almost the same curve, demonstrating that the analytical methodology works. One other curve can be seen in the figure, labeled “analytical gradient”. This was obtained for the case where the “2-up” simulation results were not used, but a “2-up” scenario was emulated by simulating a “1-up” scenario with double the link latency (up to 0.1s instead of 0.05s). The reason this underestimates the actual number of packets dropped per handoff for the random case is that it excludes the time for store-and-forward, and processing, at the intermediate nodes (just one such node in this case). Since this time is relatively constant, the offset of the resulting estimate from the real values is also roughly constant. Nevertheless, it at least provides the gradient of the correct curve, and so is labeled “analytical gradient”.

4.2.2. TCP Traffic

TCP throughput vs. link latency is shown in Figure 15. “Base HBR” and “base MIP” are the results for the base simulations already described. “HBR semi-soft” is for the case where the semi-soft handoff scheme is used to further reduce handoff latencies. The figure shows the TCP throughput as link latency (of the straight-line links in Figure 10 and Figure 11) varies. HBR with semi-soft handoffs tolerates the most link latency, and provides the highest throughput for any given link latency. Base MIP shows a sharp deterioration in TCP throughput, that gets worse as the HA moves further away (not shown in this figure), which is more realistic.

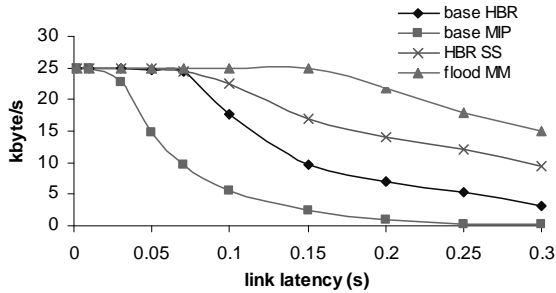


Figure 15: TCP throughput versus link latency

The remaining curve in Figure 15 is labeled “flood MM”. In this case, the root router floods the whole HBR domain with downstream packets, and upstream packets can arrive from any BS to the root router. This is not an HBR scheme, but is included to act as a performance bound. It is expected that no packets would be lost, since every BS is receiving the same stream, except for packets actually in the middle of being transmitted over the air when a handoff occurs. Notice that the throughput of flood MM starts to decrease at a link latency of just over 0.15s. We conjecture that the reason is because that is where it runs into the so-called bandwidth-delay product bound. The TCP window size in the simulations is 20 kbytes, so the TCP “pipe” from sender to receiver can only hold that much data. Let link latency (in the HBR domain) be x seconds, and recall that the link latency of the “Internet” dash-dot links is 0.01s. There are two of each type of link between the CH and MH. Therefore, the bandwidth-delay product from sender to receiver is (in kbits)

$$\beta = 2 \times 0.01 \times 1544 + 2x \times 375 \quad (2)$$

To be within the window size constraint, $\beta < 160$, and so $x < 0.1722$. However, it would be expected that the throughput would start dropping for link latencies even a little less than 0.1722s because the sender needs to allow time for acknowledgments from the receiver (ideally, therefore, it would need the bandwidth-delay product to be slightly less than the window size to allow room for it to keep sending without interruptions while waiting for the ACKs). Therefore, the performance of “flood MM” is reasonable, and shows the bounds in performance for this scenario. It can be seen that HBR SS gets the closest to this bound. Furthermore, it can be expected that “flood MM” will not perform as well when there are multiple MHs, as flooding will affect other MHs the most, whereas HBR SS would still provide good results.

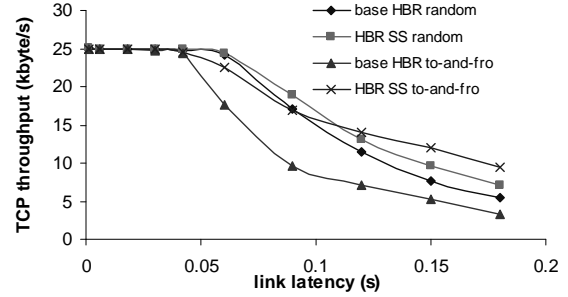


Figure 16: Comparing to-and-fro with random handoffs

As in Section 4.2.1 for UDP, we also compare the results of the to-and-fro handoffs with that of actually simulating random handoffs, where the MH hands off from any BS with a uniform distribution to any of the other three. In Figure 16, the results are plotted for both HBR (with hard handoffs) and “HBR SS” (with semi-soft handoffs). For “base HBR random” and “HBR SS random”, the link latency is as given on the x-axis. For “base HBR to-and-fro” and “HBR SS to-and-fro”, the actual link latencies used for the points on the curve are $5/3$ the values on the x-axis.

This is because for the random handoff cases, the crossover node is expected to be 2 levels above the BSs $2/3$ of the time, and 1 level above the BSs $1/3$ of the time, and through a similar reasoning process as for the UDP simulations, the $5/3$ weighting factor for the to-and-fro simulations can be derived. Unlike the case of the UDP simulations, it is not expected that the curves will line up so well. The reason is that the number of packets dropped per handoff in the UDP simulations is linearly related to the handoff latencies, whereas the relationship for the TCP case is non-linear. The results confirm this. Therefore, for TCP simulations, to-and-fro simulations may be useful to get an approximate understanding of performance for specific handoff situations like random handoffs, but actual simulations of the actual handoff situations are necessary to get specific performance results for specific scenarios.




Figure 17: Packet retransmission behavior

Although packets dropped per handoff is a useful performance indicator for UDP traffic, packet retransmission behavior is arguably more useful for TCP traffic. This is because throughput is more closely related to the packet transmission and retransmission behavior. Slight differences in number of packets dropped may result in timeouts occurring in one situation and not another, resulting in rather significant differences in packet retransmission behavior, and hence in throughput performance. The packet retransmission behavior for the different schemes is illustrated in Figure 17. For each of the 3 schemes (base HBR, base MIP and HBR SS), both the retransmit ratio and the number of retransmissions per handoff are plotted. The retransmit ratio refers to the ratio of retransmitted packets to total transmitted packets, and it is plotted as vertical bars with the values on the left y-axis. The number of retransmissions per handoff, on the other hand, are plotted as regular curves, with the values on the right y-axis. The results for “flood MM” were also computed, but not plotted in the figure, because 0 retransmissions occur throughout. Even the throughput decline for higher latencies is due to the bandwidth-delay product being constrained by the window size, not packet losses.

Looking at the retransmissions per handoff, it can be seen that all three curves follow the same basic pattern of increasing first, and then decreasing. The reason is that as the link latency increases, TCP is able to keep up with the increasing number of dropped packets by increasing the retransmissions and varying the instantaneous throughput so that can exceed 25 kbyte/s, allowing the average throughput to be still about 25 kbyte/s. However, there is a point where TCP cannot “catch up” because of too many dropped packets, and the (average) throughput drops as a result. Since the average rate of packets is decreasing in this phase, the number of dropped packets per handoff, and retransmitted packets per handoff, also declines. It is interesting to note that MIP has the worse performance in that the peak retransmission per handoff is the highest, and it occurs with the smallest link latency before entering the declining throughput phase. Similarly, HBR SS performs the best in having the lowest peak occurring with the largest link latency. As for the retransmit ratios, these tend to increase as link latency increases. In the worst case, with MIP, somewhere between a link latency of 0.2s and 0.25s, over ¼ of the packets arriving at the MH are retransmitted packets!

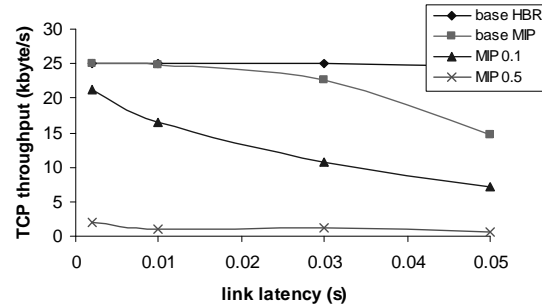


Figure 18: Performance degradation as latency between HA and FA increases

In order to see how the TCP throughput is affected by the network latency between the HA and FA in MIP, the latency on the dashed link in Figure 11 is varied from the base value of 2 ms to 100 ms to 500 ms. The resulting TCP throughput is shown in Figure 18, where the dashed link has a latency of 100 ms and 500 ms for “MIP 0.1” and “MIP 0.5”, respectively. These are not unreasonable values for MIP, where HA and FA could be very far apart. The resulting degradation on performance is evident. The performance of “base HBR” is also included in the figure for reference. As expected, in this region where the link latency is 0.05s or less, HBR can achieve 25 kbyte/s throughput, unlike MIP with the various latencies.




Figure 19: Effect of varying handoff rate

In the base simulations, the handoff rate is one handoff every 5 seconds on the average. To investigate the impact of different handoff rates, simulations were run (for HBR, HBR SS and MIP) with 3 average handoff intervals, 1 s, 5 s, and 10 s. The results are plotted in Figure 19. The solid lines show the performance of HBR. The dashed lines show the performance of HBR SS. The dotted lines show the performance of MIP. For each of these cases, the lines marked with crosses show the performance with handoff interval of 10 s, the lines with circles show the performance with handoff interval of 5 s, and the lines with triangles show the performance with handoff interval of 1 s. By looking at the 3 curves with crosses, it can be seen that the best performance is with the longer handoff intervals, whereas by looking at the 3 curves with triangles, it can be seen that the worst performance is with the shorter handoff intervals. This is expected because the more frequent the handoffs, the more frequent the dropping of packets during handoff, resulting in more frequent transition of TCP into “fast retransmit” and “slow start”. It should also be noticed from the figure that for

any given handoff rate, HBR SS performs best, followed by HBR, and followed by MIP.

Simulations were also run with TCP Reno. TCP Reno improves on the performance of TCP Tahoe when single packets are dropped, because of how it deals with congestion. In both cases, lost packets would result in duplicate ACKs being sent back to the transmitter, resulting in a retransmission of the lost packet (the “fast retransmit” algorithm). However, this is followed by “slow start” with TCP Tahoe, which can have a big impact on the TCP throughput. With TCP Reno instead, the “fast recovery” algorithm is used, going back to “congestion avoidance” instead of “slow start”. Despite performing better with single dropped packets, it has been previously found that TCP Reno performs poorly in general when multiple consecutive packets are dropped [22]. We verified that this is the case for TCP Reno over HBR and MIP as well.

Various other simulations were also run. These include cases where the basic simulation topology is modified, through vertical expansion (adding more layers of hierarchy) and horizontal expansion (adding more leaf nodes to each parent node). Also, cases of uneven link latencies have been explored. The results are similar to those with our basic simulation environment.

4.3. Results from Laboratory Prototype

The laboratory prototype that we used is shown in Figure 20. The purpose of the prototype was to confirm the simulation results with actual measurements performed on a prototype implementation. The nodes labeled “root node”, “HBR BS1” and “HBR BS2” run the Linux (kernel 2.2.14) operating system, as do the MH and CH. The router is a standard Cisco router. The HBR prototype is based on the CIP version 1.1 software from Columbia University [23], while the MIP prototype is based on the Sun Laboratories MIP prototype [24].




Figure 20: Laboratory Prototype Configuration

If the HBR measurements are performed with one set of platforms and the MIP measurements with another, then some irrelevant differences might creep in. For example, irrelevant differences might include differences in processor speeds, cross-traffic, etc., that would reduce the accuracy of the comparisons. In order to reduce irrelevant differences, both the HBR scheme and MIP were therefore run on the same platforms and with the same hardware configuration and network connections. Switching back and forth between HBR and MIP is a matter of typing a few commands to change a few interface configurations and routing table entries, and turn IP forwarding off (for HBR; the routing/paging caches are used instead) or on (for MIP).

Since the same hardware was used, both setups (HBR and MIP) are shown on the same diagram (Figure 20). Where the functionality differs, the MIP functionality is shown in square brackets beneath the HBR functionality (e.g. HBR BS1 becomes the BS with HA in the MIP case).

TCP throughput between MH and CH was measured for both HBR and MIP, using `tcp`. Some of these results are shown in Figure 21. The values are each averages over several measurements made at the `tcp` receiving process, for CH to MH communications. The throughput of MH to CH traffic has also been measured and shows similar behavior. Throughput with MIP for over 6 handoffs per minute has not been included because the results become unstable in that region.




Figure 21: TCP throughput for data transfer from the CH to the MH

Measurements of TCP throughput were also made for traffic from the MH to the CH. The throughput degradations using MIP with an increasing number of handoffs per minute are similar in this case as the previous. However, HBR performs better in this upstream direction. This is because even before the crossover node is aware of the handoffs, data packets following the handoff message are already taking the right path up to the gateway (TCP acknowledgments may be lost during this time, though, accounting for the slight degradation in throughput as the handoff rate increases). Typically, the studies on HBR schemes focus on the downstream to the MH, because that is more critical for many applications (e.g. MH obtaining streaming video from a network server). However, it should be noted that HBR schemes can perform even better in the upstream (this assertion is predicated on the assumption that data packets can immediately follow a route update packet. If security measures are in place that do not allow that, then upstream performance will be affected).

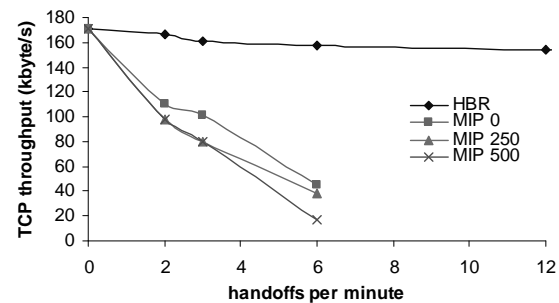


Figure 22: TCP throughput for data transfer from MH to CH

5. Discussion, Summary & Conclusions

This paper deals with HBR schemes, and in particular on how they are designed for reducing IP-level handoff latency to possibly orders of magnitude less than what may be experienced with pure macro-mobility schemes based on MIP and its variants, or SIP. This minimizes IP-level handoff disruptions, resulting in significantly fewer dropped packets and higher throughput. This paper has provided an overview of the macro-mobility schemes and their shortcomings, especially where the shortcomings relate to handoff latency problems and the fact that they are not optimized to handle micro-mobility. An overview of two classes of micro-mobility schemes has been provided. The first class is the hierarchical MIP-derived schemes like MIP-RR and TeleMIP/IDMP. The second class is that of the HBR schemes, and it has been discussed in more detail. First, a generic HBR scheme has been described, that contains the essential features of the HBR schemes. Second, differences between actual HBR schemes like MMP, CIP and HAWAII have been explained. Some of the differences arise from differences in the design objectives, and the range of implementations of HBR schemes should provide network architects with sufficient flexibility to choose a micro-mobility solution that best fits a particular network. Third, comparisons were made between HBR domains and MANET. It was concluded that for certain types of very mobile networks, MANET ad-hoc routing protocols must be used, but where a certain degree of fixed infrastructure exists that can support the HBR micro-mobility schemes, HBR would be preferred, because of reasons such as scalability and more efficient distribution of routing information.

Next, performance issues of HBR micro-mobility schemes were discussed. In addition to providing the lowest latency handoffs of all the protocols discussed, the HBR schemes also reduce signaling overhead through paging concepts, improve bandwidth efficiency in the low-bandwidth radio access network by not using encapsulation in the radio access network, facilitate QoS reservations by using an unchanging IP address while moving within the HBR domain, and use IP address resources more efficiently than other schemes like MIP with co-located COA. Results of simulations and laboratory prototype measurements were also reported. The goal of the performance studies reported in this paper is to demonstrate the improvements of using HBR over MIP or SIP for micro-mobility management. It has also been explained how the MIP results could be applied to hierarchical MIP derived schemes like MIP-RR, demonstrating that HBR performs better because of the lowest latency handoffs it provides. One implication of the simulation results is that using TCP over an IP micro-mobility management scheme magnifies the differences in handoff latencies between the schemes, because of the workings of the congestion control mechanisms like slow start.

Close to the top of the list of further simulations to do are simulations where there are multiple MHs in an HBR domain, where the traffic to/from other MHs would impact the performance of each MH. One expectation is that "flooding" micro-mobility management might not perform as well as it did in the results in this paper (which is the best case scenario for using flooding), because it would have the most impact on other MHs, so it would not be practical.

Further study is needed on transport protocols other than TCP, such as Real-Time Protocol (RTP [25]) and Stream Control Transmission Protocol (SCTP [26]). Several modifications of TCP itself have been proposed for the problem of TCP interpreting errors and delays on wireless links as congestion [27]. The performance of such schemes over HBR micro-mobility schemes might be worth investigating.

This paper has focused on the routing aspects of HBR schemes. However, it is very important to consider the implications of using HBR schemes together with schemes for Quality-of-Service (QoS) and security. Additionally, network management issues related to the HBR schemes should be investigated.

HBR schemes may be more challenging to implement commercially than some other micro-mobility schemes, but can be implemented slowly, in steps. Additional practical considerations are beyond the scope of this paper. However, it is hoped that the good performance of HBR schemes for IP micro-mobility management would serve as an incentive for further investigation in standards bodies like the IETF, and for engineers to work out the implementation issues.

6. References

- [1] C. Perkins, "IPv4 Mobility Support", RFC 2002, IETF, October 1996.
- [2] C. Perkins, D.B. Johnson, "Route Optimization in Mobile IP", Internet Draft, IETF, November 2000. Work in Progress.
- [3] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg, "SIP: Session Initiation Protocol," RFC 2543, IETF, March 1999.
- [4] A. Dutta, F. Vakil, J.-C. Chen, M. Tauil, S. Baba, N. Nakajima, and H. Schulzrinne, "Application Layer Mobility Management Scheme for Wireless Internet", *3G Wireless International Conference*, San Francisco, May 2001.
- [5] H. Schulzrinne, E. Wedlund, "Application Layer Mobility Support using SIP", *ACM Mobile Computing and Communications Review*, vol. 4, no. 3, pp. 47-57, July 2000.
- [6] E. Gustafsson, A. Jonsson, C. Perkins "Mobile IP Regional Registration", Internet Draft, IETF, March 2001. Work in Progress.
- [7] S. Das, A. Misra, P. Agrawal, and S.K. Das, "TeleMIP: Telecommunications Enhanced Mobile IP Architecture for Fast Intra-domain Mobility," *IEEE Personal Communications Magazine*, Vol 7, pp. 50-58, Aug. 2000
- [8] A. Valko, *Design and Analysis of Cellular Mobile Data Networks*, Ph. D. Dissertation, Technical University of Budapest, 1999.
- [9] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan, and S. Wang, "HAWAII: A Domain-based Approach for Supporting Mobility in Wide-area Wireless Networks", available from <http://www.bell-labs.com/user/ramjee/papers/hawaii.ps.gz>.

- [10] A. Campbell, J. Gomez, S. Kim, A. Valko and C.-Y. Wan, "Design, Implementation, and Evaluation of Cellular IP", *IEEE Personal Communications*, pp. 42-49, August 2000.
- [11] A. Dutta, R. Jain, K. D. Wong, J. Burns, K. Young and H. Schulzrinne, "Multilayered Mobility Management for Survivable Network", to appear, *Proceedings of Milcom*, October 2001.
- [12] R. Ramjee, T. La Porta and L. Li, "Paging support for IP mobility using HAWAII", Internet Draft, IETF, June 1999.
- [13] E. M. Royer and C.-K. Toh, "A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks", *IEEE Personal Communications*, pp. 46-55, April 1999.
- [14] M. Joa-Ng and I. T. Lu, "A Peer-to-Peer Zone-based Two-level Link State Routing for Mobile Ad-hoc Wireless Network", *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 8, pp. 1415-1425, August 1999.
- [15] <http://timely.crhc.uiuc.edu/Projects/cedar/cedar.html>
- [16] M. Mouly and M. Pautet, *The GSM System for Mobile Communications*, Cell. and Sys. Correspondence, France, 1992.
- [17] X. Zhang, J. Castellanos, A. Campbell, K. Sawada and M. Barry, "P-MIP: Minimal Paging Extensions for Mobile IP", Internet Draft, IETF, July 2000. Work in progress.
- [18] M. Joa-Ng and K. D. Wong, "IP Mobility Management for the ACN Platform", *Proceedings of Milcom*, pages 465-469, October 2000.
- [19] Z. Shelby, P. Mahonen, D. Gatzounas, A. Inzerilli, V. Typpo, "Cellular IP Route Optimization", Internet Draft, IETF, June 2001. Work in progress.
- [20] W. Stevens, *TCP/IP Illustrated, Volume 1: The Protocols*, Addison-Wesley, Reading, Massachusetts, 1994.
- [21] C. Perkins, "Minimal Encapsulation within IP", RFC 2004, IETF, October 1996.
- [22] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno and SACK TCP", *Computer Communications Review*, vol.26(3), pp. 5-21, July 1996.
- [23] <http://comet.ctr.columbia.edu/cellularip/>
- [24] <http://playground.sun.com/pub/mobile-ip/sunlabs/>
- [25] H. Schulzrinne, S. Casner, R. Frederick and V. Jacobsen, "RTP: A Transport Protocol for Real-Time Applications", RFC 1889, IETF, February 1996.
- [26] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang and V. Paxson, "Stream Control Transmission Protocol", RFC 2960, IETF, October 2000.
- [27] H. Balakrishnan, V. Padmanabhan, S. Seshan, R. Katz, "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links", *IEEE/ACM Transactions on Networking*, December 1997.