

HAWAII: A Domain-based Approach for Supporting Mobility in Wide-area Wireless networks

R. Ramjee, T. La Porta, S. Thuel, K. Varadhan
Bell Labs, Lucent Technologies
ramjee,tlp,thuel,kvaradhan@bell-labs.com

S.Y. Wang
Dept of EECS, Harvard
shieyuan@eecs.harvard.edu

Abstract

Mobile-IP is the current standard for supporting macro-mobility of mobile hosts. However, in the case of micro-mobility support, there are several competing proposals. In this paper, we present the design, implementation, and performance evaluation of HAWAII: a domain-based approach for supporting mobility. HAWAII uses specialized path setup schemes which install host-based forwarding entries in specific routers to support intra-domain micro-mobility. These path setup schemes deliver excellent performance by reducing mobility related disruption to user applications. Also, mobile hosts retain their network address while moving within the domain, simplifying QoS support. Furthermore, reliability is achieved through maintaining soft-state forwarding entries for the mobile hosts and leveraging fault detection mechanisms built in existing intra-domain routing protocols. HAWAII defaults to using Mobile-IP for macro-mobility, thus providing a comprehensive solution for mobility support in wide-area wireless networks.

1. Introduction

Mobile-IP is the current standard for supporting macro-mobility in IP networks [4]. Mobile-IP defines two entities to provide mobility support: a *home agent* (HA) and a *foreign agent* (FA). The HA is statically assigned to the mobile host based on the permanent home IP address of the mobile host. The FA is assigned to the mobile host based on its current location. The FA has associated with it an IP address called the *care-of address*. Packets destined for a mobile host are intercepted by the HA and tunneled to the FA at the care-of address. The FA then decapsulates the packets and forwards them directly to the mobile host. Thus, Mobile-IP provides a good framework for allowing users to roam outside their home networks.

When Mobile-IP is used for micro-mobility support, it results in disruption to user traffic during handoff, and high control overhead due to frequent notifications to the HA [1].

Recently, Mobile-IP Route Optimization (RO) with support for micro-mobility was proposed [5]. In this proposal, packets are forwarded from the old FA to the new FA to reduce disruption during handoff. Still, the mobile device's care-of-address changes each time the user moves between neighboring base stations, resulting in undesirable notifications to the HA and the correspondent hosts on every handoff. Furthermore, in the case of a Quality of Service (QoS) enabled mobile host, acquiring a new care-of address on every handoff would trigger the establishment of new QoS reservations from the HA to the FA even though most of the path remains unchanged.

Thus, while Mobile-IP should be the basis for mobility management in wide-area wireless data networks, it has several limitations when applied to wide-area wireless networks with high mobility users that may require QoS. Our goal is to extend Mobile IP to address these limitations. Our proposal for supporting mobility hinges on the assumption that most user mobility is local to a domain, in particular, an administrative domain of the network. Therefore, we consider optimizations in routing and forwarding for more efficient support of intra-domain mobility.

In this paper, we present the design, implementation, and performance evaluation of our Handoff-Aware Wireless Access Internet Infrastructure (HAWAII). In HAWAII, mobile hosts retain their network address while moving within a domain. The HA and any corresponding hosts are unaware of the host's mobility within this domain. Routes to the mobile host are established by *specialized path setup schemes that update the forwarding tables with host-based entries in selected routers* in that domain. HAWAII defaults to using traditional Mobile-IP schemes for inter-domain mobility. We demonstrate that the HAWAII approach results in quantitative gains (such as less disruption to user traffic during handoff and fewer updates to the home agent) as well as qualitative gains (ease of QoS support and robustness) over the Mobile-IP schemes. The price that HAWAII pays for these advantages is the limit of number of host-based forwarding entries that can be supported in the access routers. We show that this concern can be addressed by appropriate

sizing of the domain and by carefully choosing the routers that are updated when a mobile host is handed off.

The remainder of the paper is organized as follows. We begin by enumerating the design goals of our protocol in Section 2, and contrast the related work in the context of these goals in Section 3. We then present an overview of our solution, HAWAII, in Section 4. In Section 5, we introduce several path setup schemes for supporting mobility within a domain. In Sections 6 and 7, we compare the performance of these path setup schemes with the Mobile-IP and the Mobile-IP RO schemes through measurements obtained from simulation and implementation. In Section 8, we describe how our design simplifies providing QoS in the wired portion of the network. In Section 9, we illustrate the impact of our design on reliability. In Section 10, we present our conclusions.

2. Design Goals

We have five design goals in HAWAII:

- *Limit disruption to user traffic.*
- *Enable efficient use of access network resources.* This includes avoiding inefficient routing and tunneling where possible.
- *Enhance scalability* by reducing updates to the home agent (enabling it to support a large number of mobile users) and avoiding addition of state in backbone routers.
- *Provide intrinsic support for QoS* in the mobility management solution. This includes allowing per flow QoS and limiting the number of reservations that must be re-established when hosts move.
- *Enhance reliability.* We require HAWAII to be no less fault tolerant than existing Mobile-IP proposals, and we explore additional mechanisms to improve the robustness of mobility support.

While there has been a large body of prior work in this area, previous solutions only address a subset of the above goals, often at the expense of negatively impacting others. We believe that HAWAII is the first comprehensive solution that jointly addresses these goals.

We next survey the related work in this area, identifying the goals addressed by each particular solution. For convenience, we refer to our design goals as *disruption*, *efficiency*, *scalability*, *QoS*, and *reliability*, respectively. Note that we are trying to limit *disruption* while enhancing the measure of the other goals.

3. Related Work

The vast majority of prior work has focused on limiting the *disruption* to user traffic during handoff. One common

approach for reducing *disruption*, proposed originally for ATM-based networks, is extending connections from the previous base station [2, 6]. The extension approach also forms the basis of the Mobile-IP RO proposal [5]. However, in the case of mobility solutions proposed for connection-oriented ATM networks [6], the goals of *scalability* and *QoS* can be easily achieved since each connection is identified by a pair of triplets at each switch (port/Virtual Path Identifier(VPI)/Virtual Connection Identifier (VCI)); these triplets can be re-mapped locally during handoffs. In the case of connection-less IP networks, a change in mobile host's IP care-of address during handoff (as in the Mobile-IP RO proposal) requires updates to the home agent, that then introduces *scalability* concerns; it also impacts *QoS* support, requiring the establishment of new QoS mappings end-to-end even though mobility is typically localized.

Another common approach for reducing *disruption* is through the use of multicasting [9]. However, join latency and group management issues in multicasting-based solutions could result in loss of *efficiency* due to wasted bandwidth. These considerations also impact *scalability* in the backbone routers where every mobile host's multicast address needs to be managed. A workaround to address *scalability* is to tunnel packets to a domain FA after which multicasting is used within the domain. However, since packets for several mobile hosts are now tunneled to the domain FA, being able to isolate those flows that require per-flow *QoS* is difficult. Furthermore, the failure of the domain FA can impact *reliability*, and must be addressed.

A common technique to enhance *scalability* is to introduce a hierarchy. One such proposal is to build a hierarchy of foreign agents [1]. This approach is effective in managing mobility locally using multiple foreign agents: this limits *disruption* of traffic during handoffs, and enhances *scalability* by limiting updates to the home agent. However, this scheme needs to address *reliability* considerations to recover from the failure of these additional FAs, possibly through new fault recovery mechanisms. Further, since data packets traverse multiple tunnels, providing *QoS* support and maintaining data transfer *efficiency* are difficult as well.

The recent Cellular IP proposal [11] uses specialized domain routers with host-based entries for local mobility and the use of Mobile-IP for inter-domain mobility. These routers automatically detect that a mobile user has been handed off by snooping actual data packets; thus updates can be localized, enhancing the *scalability* of update mechanisms and limiting *disruption*. However, Cellular IP relies on the gateway to act as a FA that decapsulates the packets before delivering them to the user. This has two drawbacks: failure of the Gateway FA can impact *reliability*; also, the presence of the Gateway FA can complicate *QoS* management.

HAWAII explores a different approach to address the goals outlined earlier. We next present an overview of the protocol, highlighting how the various design choices in HAWAII help towards achieving these goals.

4. Protocol Overview

A common approach for providing transparent mobility to correspondent hosts is to divide the network into hierarchies. HAWAII uses a similar strategy, segregating the network into a hierarchy of domains, loosely modeled on the autonomous system hierarchy used in the Internet. The network architecture is illustrated in Figure 1. The gateway into each domain is called the *domain root router*. Each host is assumed to have an IP address and a home domain. While moving in its home domain, the mobile host retains its IP address. Packets destined to the mobile host reach the domain root router based on the subnet address of the domain and are then forwarded over special dynamically established paths to the mobile host.

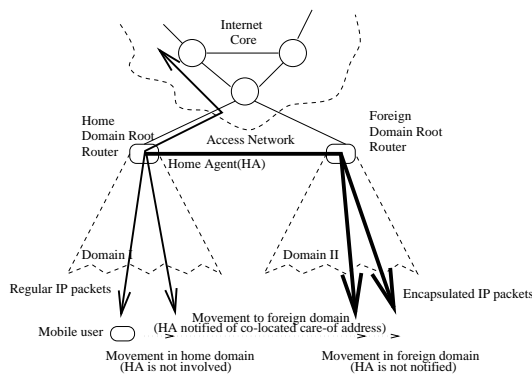


Figure 1. Hierarchy using domains

When the mobile host moves into a foreign domain, we revert to traditional Mobile-IP mechanisms. If the foreign domain is also based on HAWAII, then the mobile host is assigned a co-located care-of address from its foreign domain. Packets are tunneled to the care-of address by a home agent in its home domain. When moving within the foreign domain, the mobile host retains its care-of address unchanged, and connectivity is maintained using dynamically established paths.

The protocol contains three types of messages for path setup: power-up, update and refresh.

A mobile host that first powers up and attaches to a domain sends a *path setup power-up message*. This has the effect of establishing host specific routes for that mobile host in the domain root router and any intermediate routers on the path towards the mobile host. Thus, the connectivity from that domain root router to the mobile hosts connected through it forms a virtual tree overlay. Note that

other routers in the domain have no specific knowledge of this mobile host’s IP address.¹

While the mobile host moves within a domain, maintaining end-to-end connectivity to the mobile host requires special techniques for managing user mobility. HAWAII uses *path setup update messages* to establish and update host-based routing entries for the mobile hosts in selective routers in the domain so that packets arriving at the domain root router can reach the mobile host with limited *disruption*. The choice of when, how, and which routers are updated constitutes a particular path setup scheme. In Section 5, we describe four such path setup schemes.

We characterize the HAWAII path state maintained in the routers as “soft-state”. This increases the robustness of the protocol to router and link failures. The mobile host infrequently sends periodic *path refresh messages* to the base station to which it is attached to maintain the host based entries, failing which they will be removed by the base station. The base station and the intermediate routers, in turn, sends periodic *aggregate hop-by-hop refresh messages* towards the domain root router. As we shall see in the following two sections, path setup messages are sent to only selected routers in the domain, resulting in very little overhead associated with maintaining soft-state.

We conclude this section with a few observations about HAWAII in the context of the design goals stated earlier.

Disruption: specialized path setup schemes, described in Section 5, ensure that data disruption during handoff is limited. The disruption caused by various schemes for audio and video traffic is quantified in Section 6.

Efficiency: when the mobile host is in its home domain, data transfer efficiency is maintained since the home agent is not involved; thus, IP packets are delivered to the mobile host without any tunneling. The impact of tunneling on web and FTP traffic is discussed in Section 6.

Scalability: the home agents and correspondent hosts are unaware of intra-domain mobility. This enhances the scalability of home agents in supporting a large number of mobile users. One potential concern is the state required to maintain multiple host specific entries in the routers. The concern is specifically for the number of mobile hosts that can be attached to, and supported by, a single domain. In Section 7, we present a numerical example showing how a single domain in HAWAII can include over a hundred base stations in a typical wide-area wireless network.

QoS: the design choices of using co-located care-of addresses and maintaining the mobile host address unchanged within a domain simplifies per flow qos support, and is discussed in further detail in Section 8. One drawback of using the co-located care-of address option is the need for two IP

¹In the case of mobile to mobile communication, packets arriving at a router that has no specific host-based entry are routed using a default route to the domain root router that then forwards the packet to the mobile host.

addresses for each mobile host that is away from its home domain. One possible optimization is to adapt the “dialup” model used by ISPs to wireless networks and assign the home address via DHCP.

Reliability: as we shall see in Section 9, HAWAII does not define a new protocol for failure detection. In fact, HAWAII relies on standard intra-domain routing protocols such as RIP or OSPF to detect router and link failures. When a failure is detected, HAWAII simply triggers soft-state refresh messages to restore connectivity thereby achieving reliability amidst link and router failures. The robustness of HAWAII is also increased because single points of failure such as home agents are eliminated while a host is in its home domain.

Please refer to [7] for a more detailed description of the protocol.

5. HAWAII Path Setup Schemes

The HAWAII handoff procedures are only activated when the mobile host’s next hop IP node is changed during the handoff. Thus, for discussion, we assume base stations have IP routing functionality in the remainder of the paper.

Let us now consider a mobile host that is powered up in a domain. As described in the previous section, this implies that host-based forwarding entries have been setup at the domain root router and any intermediate routers to the mobile host. The path from the domain root router to the different mobile hosts forms a virtual tree topology. We now describe the operations of four path setup schemes used to re-establish path state when the mobile host moves from one base station to another; we use a tree-based topology for clarity. Note that our schemes will work in any general topology. In particular, Section 9 illustrates how recovery is accomplished in non-tree-based topologies.

For the remaining subsections, let us define the cross-over router as the router closest to the mobile host that is at the intersection of two paths, one between the domain root router and the old base station, and the second between the old base station and the new base station.² The four path setup schemes considered in this paper can be classified into two types based on the way packets are delivered to the mobile host during a handoff: in the first type, packets are forwarded from the old base station to the new, whereas in the second, they are diverted at the crossover router.

5.1. Forwarding Schemes: MSF and SSF

In these path setup schemes, *packets are first forwarded from the old base station to the new base station before they are diverted at the cross-over router.*

²The schemes considered in this paper can be modified to work with other definitions of the cross-over router as well.




Figure 2. Forwarding Schemes

The idea of forwarding packets during handoff is not new [1, 2, 6, 5]. In the case of ATM networks [2, 6], each switch has a unique mapping of (input interface, VPI, VCI) to (output interface, VPI, VCI) for each connection. Thus forwarding can be accomplished by creating a new set of such mappings from the old to the new base station. In IP networks, since there is no such per-connection mapping that includes the incoming and outgoing interface, forwarding has so far been accomplished by either proxy mechanisms if the user stays within the same broadcast network [1] or through tunneling [5]. Since we would like to maintain the user’s IP address unchanged for easier QoS support between handoffs across wide-area base stations not connected to the same broadcast network and also avoid tunneling as far as possible to maintain data transfer *efficiency* we adopt a new approach in HAWAII to implement data forwarding. We propose two variants of forwarding schemes in HAWAII, one that works with standard IP routing tables to update the host-based entries and, another scheme in which we extend the IP routing table to accommodate interface-based information thereby adapting the ATM per-connection entries to the IP per-host entries. These schemes, Multiple Stream Forwarding (MSF) and Single Stream Forwarding (SSF), are described below.

The MSF scheme is illustrated in Figure 2(a). The forwarding table entries are shown adjacent to the routers. These entries are prepended with a message number indicating which message was responsible for establishing the entry (a message number of zero indicates a pre-existing entry). The letters denote the different interfaces. The path setup message is first sent by the mobile host to the old base station. Message 1 contains the new base station’s address. The old base station performs a routing table lookup for the new base station and determines the interface, interface A, and next hop router, Router 1. The base station then adds a forwarding entry for the mobile host’s IP address with the outgoing interface set to interface A. It then forwards Message 2 to Router 1. Router 1 performs similar actions and forwards the message to Router 0. Router 0, the

cross-over router in this case, adds forwarding entries that result in new packets being diverted to the mobile host at the new base station. It then forwards the message towards the new base station. Eventually Message 5 reaches the new base station that changes its forwarding entry and sends an acknowledgement of the path setup message to the mobile host, shown as Message 6.³

Note that this order of updating the routers can lead to the creation of multiple streams of mis-ordered packets arriving at the mobile host. For example, during transient periods newer packets forwarded by Router 0 may arrive at the mobile host before older packets forwarded by Router 1 which might in turn arrive before even more older packets forwarded by the old base station. The creation of multiple streams during handoff could adversely impact both audio and TCP applications. Also, this scheme can result in creation of transient routing loops (for example, after old base station has changed its entry to forward packets but before Router 1 processes Message 2). However, note that the mis-ordered streams and routing loops exist for extremely short periods of time.⁴ The main benefit of this scheme is that it is simple and results in no loss.

As an alternative, the Single Stream Forwarding (SSF) scheme updates the forwarding entries in a method that is similar to the Mobile-IP RO scheme in which packets are forwarded from the old base station to the new base station in a single stream. In order to achieve this without the use of tunneling, we use a technique we term *interface-based forwarding*. This requires more descriptive routing table entries. A routing table typically has an entry of the form (IP address \rightarrow outgoing interface). In this scheme, the router must be able to route based on an additional field, the incoming interface of the packet. The resulting routing entry is of the form (incoming interface(s), IP address \rightarrow outgoing interface). Thus, in Figure 2(b), after Message 5 is processed, packets arrive at the old base station and are forwarded to the new base station as a single stream. After processing Message 6, Router 0 diverts new data packets directly to the new base station. This redirection is similar to what would happen in the Mobile-IP RO scheme except that the redirection in this case happens quickly (after Message 6) without the corresponding host or the HA being aware of the handoff. While this scheme is also lossless and maintains a single stream of forwarded packets, it is somewhat complex to implement. In Section 6, we show that the added complexity of interface-based forwarding in SSF improves performance over the simpler MSF approach but the improvement is not significant enough for typical handoffs, that involve routers that are one or two hops away.

³Pseudo-code executed at the routers for the different schemes can be found in [7].

⁴This can be tightly controlled by having fairly small timeout values before forwarding entries are deleted.

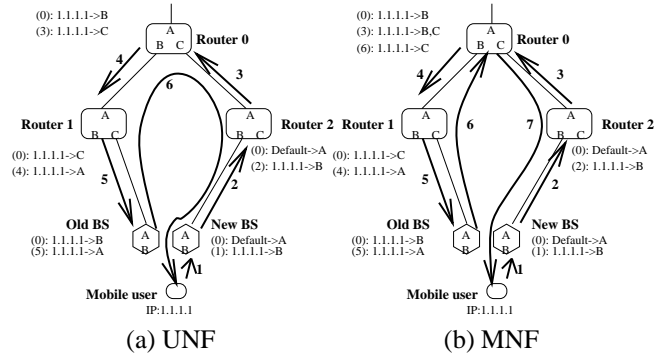


Figure 3. Non-Forwarding Schemes

5.2. Non-Forwarding Schemes: UNF and MNF

In these path setup schemes, as the path setup message travels from the new base station to the old base station, *data packets are diverted at the cross-over router to the new base station, resulting in no forwarding of packets from the old base station.*

There are two variants of the Non-Forwarding scheme, motivated by two types of wireless networks. The Unicast Non-Forwarding (UNF) scheme is optimized for networks where the mobile host is able to listen/transmit to two or more base stations simultaneously for a short duration, as in the case of a WaveLAN or Code Division Multiple Access (CDMA) network. The Multicast Non-Forwarding (MNF) scheme is optimized for networks where the mobile host is able to listen/transmit to only one base station as in the case of a Time Division Multiple Access (TDMA) network.

Again, Non-Forwarding schemes have been studied in the context of ATM networks [6] and in the multicasting-based approaches [9]. HAWAII differs from these approaches in that our schemes perform the redirection based on host-based entries rather than relying on general purpose multicast routing protocols. Therefore our handoff latencies are less than the join latencies of multicast-based approaches. Furthermore, the MNF scheme, where we do use multicasting, is a custom-designed “dual-casting” scheme, in which the cross-over router multicasts data packets to at most two of its interfaces during handoff.

The UNF scheme is illustrated in Figure 3(a). In this case, when the new base station receives the path setup message, it adds a forwarding entry for the mobile host’s IP address with the outgoing interface set to the interface on which it received this message. It then performs a routing table lookup for the old base station and determines the next hop router, Router 2. The new base station then forwards Message 2 to Router 2. This router performs similar actions and forwards Message 3 to Router 0. At Router 0, the cross-over router in this case, forwarding entries are added such that new packets are diverted directly to the mobile host at

the new base station. Eventually Message 5 reaches the old base station that then changes its forwarding entry and sends an acknowledgement, Message 6, back to the mobile host.

The MNF scheme is very similar to the UNF scheme. The main difference is that the cross-over router, Router 0, multicasts data packets for a short duration. In Figure 3(b), Router 0 dual-casts data packets from interface A to both the new and old base stations after it receives Message 3 and until it receives Message 6. This helps limit packet loss in networks in which the mobile host can only listen to a single base station.

6. Disruption

In this section, we use simulation to compare the disruption performance of the four HAWAII and two Mobile-IP schemes. These were simulated using the HARVARD simulator [12]. The transfer of a packet in the simulated network is achieved through execution of real TCP/UDP/IP code in the kernel, resulting in high-fidelity simulation results.

While one would expect the HAWAII schemes which operate locally to outperform the basic Mobile-IP scheme, the performance differences between the HAWAII schemes and the Mobile-IP RO scheme is less clear. Recall that in the Mobile-IP RO scheme packets are forwarded from the old base station to the new base station just like the forwarding path setup schemes in HAWAII; the only difference lies in the fact that in Mobile-IP RO, the HA and the correspondent host needs to be notified before packets go directly to the new base station while in HAWAII, local updates results in packets being quickly redirected to the new base station. To our knowledge, this is the first quantitative comparison of truly local update handoff schemes (such as HAWAII schemes) with semi-local (Mobile-IP RO) and non-local schemes (Mobile-IP) for supporting IP mobility.

The simulation topology is shown in Figure 4.

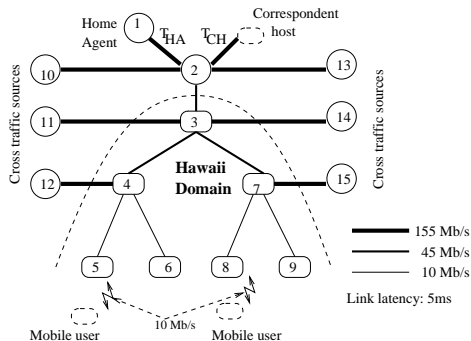


Figure 4. Simulation topology

Since we are mainly interested in wide-area wireless networks where cell coverage usually overlaps, we assume that

the mobile host is able to gracefully handoff from one base station to another.⁵ Due to space constraints, we present detailed results only for UDP-based audio and video applications during intra-domain handoffs. In the case of long-duration TCP flows, we find that the HAWAII schemes deliver sizeable improvements over the basic Mobile-IP scheme of around 15% and a small improvement over the RO scheme, which varies between 0-6% in aggregate TCP bandwidth [8].

In the case of audio experiments, the correspondent host transmits 160 byte UDP packets every 20ms (64Kb/s audio) to the mobile host. On every handoff of the mobile host, we collect statistics on the incoming UDP packets in the down-link direction⁶ such as delay, loss, etc., for all the handoff schemes.

We consider a scenario in which there are several cross-traffic sessions in the network topology, competing with the aforementioned UDP session. This would be the case, for example, when we have a shared wired/wireless access network. We introduce bursty web traffic from nodes {11, 14} to other users under base stations {5, 6, 8, 9}. We also introduce greedy FTP traffic from nodes {12, 15} to {11, 14} and {10, 13} to {11, 14}.

In order to compare the disruption caused during a handoff by the various schemes quantitatively, consider the operation of an interactive audio application. The application typically uses a playout delay to overcome network jitter. The packet playout time at the receiver is set to packet-send-time + playout delay. If the packet arrives after its playout time, the packet is dropped. Note that this *packet drop* is different from *packet loss* that might occur in the network during a handoff. We are interested in the *total packet loss* which includes both packets dropped due to late arrival as well as packets lost in the network.

In Figure 5, we plot the total of dropped and lost packets per handoff (averaged over 100 or more handoffs) versus playout delay for all the 6 handoff schemes. In this simulation, the value of link delay to correspondent host (T_{CH}) is 5ms and link delay to the home agent (T_{HA}) is 50ms. Therefore the propagation delay from correspondent host to mobile host is 125ms for the basic Mobile-IP scheme and 25ms for the other schemes.

Figure 5(a) plots the disruption caused to an audio application during handoff when the cross-over router is two hops away from the base station (the disruption for one hop handoffs, not shown, is similar in shape but with lower loss values). As the playout delay is increased along the x-axis, late arriving packets get buffered at the mobile host instead

⁵For non-overlapping cells, both HAWAII and Mobile-IP schemes would need to be augmented with buffering capabilities to avoid user level disruption.

⁶In the uplink direction, the data path from the mobile host to the correspondent host is identical in all the schemes, resulting in similar performance.




Figure 5. Packet loss during 2-hop handoff

of getting dropped, resulting in smaller number of dropped packets. However, the number of packets lost in the network during handoff is unaffected by the playout delay.

In the case of basic Mobile-IP scheme, about 5 packets/handoff are lost in the network. This is because, in our configuration, the registration update from the mobile host takes about 100ms (link delay of 50ms and queueing delay of 50 ms) to reach the HA. In this interval, about 5 packets are sent to the old base station and are lost.⁷

Let us now compare the remaining five handoff schemes. Consider a playout delay value of 100ms in Figure 5(a). In this case, the Mobile-IP RO scheme results in a total loss of about 3 packets/handoff while the HAWAII schemes result in the total loss of less than 1 packet/handoff. This is because the HAWAII schemes switch over very quickly to the new route, while in Mobile-IP RO scheme, the HA and then the correspondent host must be notified before packets use the new route. Among the HAWAII schemes, UNF and MNF perform the best for the case of mobile hosts with capabilities to receive from multiple and single base stations respectively. The forwarding schemes, SSF and MSF, have slightly lower performance than the non-forwarding schemes. Between SSF and MSF, SSF slightly outperforms MSF; the difference is due to the creation of multiple flows in the MSF scheme that results in older packets getting delayed beyond their playout time.

For higher values of playout delay, the forwarding schemes outperform the MNF scheme. The forwarding schemes result in no packet loss in the network whereas the MNF scheme results in packet loss in the network (and duplicates) of about 0.25 packets/handoff.⁸ Note that for

⁷If we assume that the MH can listen to both base stations until the HA is informed, then there could be no loss.

⁸The reason for packet loss in the MNF scheme is subtle; a packet that

higher values of playout delay, the RO scheme may match the total loss values of the HAWAII schemes. For example, in Figure 5(a), with 150ms of playout delay, the Mobile-IP RO scheme results in comparable total loss as the HAWAII UNF scheme with 100ms playout delay. In the case of a stored audio application, where maintaining a small playout delay is not critical, the Mobile-IP RO scheme will deliver similar performance as the HAWAII schemes. However, in an interactive audio application which require small playout delays, mobile hosts using the Mobile-IP RO will need a larger playout delay than that needed in HAWAII. This affects the quality of the interactive application.

The results for video traffic, shown in Figure 5(b), are similar to the results for audio except for slightly lower total losses due that fact that (4KB) UDP packets are sent every 33ms rather than the 20ms interval for audio packets. We also examined the effect of T_{HA} , the link delay to the HA, on performance. The HAWAII schemes are unaffected since they operate locally. For Mobile-IP (Mobile IP RO) schemes, when T_{HA} decreases, the performance approaches that of the HAWAII non-forwarding (forwarding) schemes. Details can be found in [8].

Summarizing the UDP results, the localized HAWAII schemes result in smaller disruption to audio/video traffic compared to the Mobile-IP schemes. In particular, HAWAII has fewer dropped packets (or lower values for the average playout delays) compared to the semi-local Mobile-IP RO scheme. Among the HAWAII schemes, UNF performs best for mobile hosts that can listen to two base stations simultaneously, while MNF performs best for mobile hosts that can listen to only one base station at any given time. SSF and MSF are lossless and deliver good performance but require slightly larger values of playout delay.

7. Scalability

In this section, we first briefly describe our implementation and present performance numbers for processing different types of messages in our testbed. We then present a numerical example to illustrate the scalability advantages of using HAWAII over a non-hierarchical approach based on Mobile-IP.

We have implemented a HAWAII daemon that is currently integrated with “routed”, the routing daemon. This daemon processes the path setup update and refresh messages. The processing of an update message is fairly simple: on receiving the message, modify the forwarding entry for the mobile host in the kernel and forward the update message towards its destination. Soft-state refresh messages are sent independently by each of the nodes every T_R seconds.

is delayed arriving at the base station before the handoff may not reach the mobile host if the host has since completed the handoff.

Typically, processing the refresh message just involves updating the expiry timer in the HAWAII daemon and can be performed very efficiently.

Item	Message type	time (μ secs)
P_{HU}	HAWAII update (or power-up)	156
P_{HR}	HAWAII refresh (25 entries)	166
P_{MU}	Mobile-IP update	1590
P_{MR}	Mobile-IP renewal	120

Table 1. CPU processing times

Table 1 lists the processing time of HAWAII and Mobile-IP update and refresh messages measured on a Pentium II 333MHz CPU running the FreeBSD 2.2.7 operating system. The reason for the relatively large processing time at the home agent for an update registration, P_{MU} , as compared to a HAWAII update, P_{HU} , is because the home agent has to perform several actions when processing a Mobile-IP registration: authenticate the message, enable proxyarp for the mobile host, remove the old entry from the home list, and add the new care-of address for the mobile host.

We now illustrate the advantages of managing mobility locally through a numerical example. Consider a domain with configuration parameters as shown in Table 2. The domain is in the form of a tree with three levels: at the highest level there is a single domain router; at the second level there are seven intermediate routers; at the third and lowest level, there are 140 base stations (twenty per router). We now consider two different approaches: 1) Mobile-IP approach where FAs are present at each base station and are served by a HA and 2) the HAWAII approach where the HA is at the domain root router.

First note that the coverage area of this domain is quite large: $A_D = B_D L_B^2 / 16 = 980 \text{Km}^2$. The number of forwarding entries at the domain root router in HAWAII, which is the same as the number of active users in the domain, is $\rho A_D = 38,220$. This is also same as the number of tunneling entries in the case of the non-hierarchical Mobile-IP approach at the HA.⁹ Note that 40K entries are well within the capability of modern routers. Furthermore, a majority of these entries are completely specified entries of hosts from a particular domain/subnet. In this case, perfect hashing is possible resulting in $O(1)$ memory access for IP route lookup. Thus, route lookup for data forwarding can be done efficiently at the domain routers.

We now compute the CPU utilization for the two Mobile-IP and HAWAII approaches.

⁹While one could potentially have as many HA's as base stations in the Mobile-IP approach, two practical reasons would preclude it: a) HA's need to be very reliable as they are single points of failure, and b) Management of multiple HA's and user profiles will be complicated.

Item	Type	Value
B_D	Base stations per DRR	140
R_D	2nd level routers per DRR	7
ρ	Active user density	$39/\text{Km}^2$
v	User speed	112 Km/hr
L_B	Perimeter of base station	10.6 Km
T_R	Refresh timer for HAWAII	30 secs
Y	No. of entries in aggregate refresh	25
T_M	Mobile-IP binding lifetime	300 secs
γ	Percent users outside domain	0.1

Table 2. Example configuration values

The processing load at the HA in the Mobile-IP approach, CPU_M , is shown in [8] to be

$$CPU_M = P_{MU} \frac{\rho v L_B B_D}{\pi} + P_{MR} \frac{\rho L_B^2 B_D}{16 T_M} \quad (1)$$

where the first term in equation (1) is due to Mobile-IP registration updates during handoff and the second term is due to Mobile-IP registration renewals or refreshes.

The processing load at the domain root router in HAWAII, CPU_H , is shown in [8] to be

$$CPU_H = P_{MU} \frac{\rho v L_B \sqrt{B_D}}{\pi} + P_{MR} \frac{\gamma \rho L_B^2 B_D}{16 T_M} + P_{HU} \frac{\rho v L_B \sqrt{B_D R_D}}{\pi} + P_{HR} \frac{[\rho L_B^2 B_D / 16 Y]}{T_R} \quad (2)$$

where the first two terms represent the Mobile-IP registration updates (term 1) and renewals (term 2) at the HA in the domain root router, and the last two terms represent the HAWAII path setup updates (term 3) and refreshes (term 4), and typically $R_D \ll B_D$.

First consider the impact of mobility related updates. Observe that the processing load due to mobility related updates in the Mobile-IP approach (term 1 in equation 1) varies *linearly with the number of base stations in the domain*, $O(B_D)$, while the processing load due to mobility related updates in HAWAII (terms 1 and 3 in equation 2) varies with the *square-root of number of base stations*, $O(\sqrt{B_D})$. Furthermore, the processing of Mobile-IP updates is more expensive than HAWAII updates ($P_{MU} > P_{HU}$, see Table 1). Thus, term 1 in equations 1 and 2 is the dominant term. Since the dominant term is reduced by a factor of $\sqrt{B_D}$ in HAWAII, the processing load due to updates in the HAWAII approach is significantly lower than in the Mobile-IP approach.

Now consider the impact of refresh messages. In both approaches, the processing load due to refresh messages (term 2 in equation 1 and terms 2 and 4 in equation 2) varies linearly with B_D . However, note that these terms are averaged down by the refresh interval (T_M and T_R), thus reducing the overhead impact. Furthermore, in the case of

the HAWAII approach, the processing load due to Mobile-IP renewals (term 2 in equation 2) is further reduced from the corresponding term in the Mobile-IP approach by a factor of $\gamma \ll 1$, representing the fraction of users who are away from their home domain. The rate of path setup refreshes (term 4 in equation 2) in HAWAII is reduced by a factor of Y because of aggregation. Thus, the processing load of refresh messages in HAWAII is also lower than in the Mobile-IP approach.

Message	HAWAII @ DRR		Mobile-IP @ HA	
	Freq/s	CPU %	Freq/s	CPU %
H Update	127.8	1.92	0	0
H Refresh	51.3	0.85	0	0
MIP Reg.	48.4	7.6	574	91.2
MIP Renewal	12.74	0.15	127.4	1.5
Total	240.2	10.5	701.4	92.7

Table 3. Results

The numerical results for the configuration shown in Table 2 are summarized in Table 3. In this case, *HAWAII's approach to managing mobility locally results in almost ten times lower processing overhead at the most heavily loaded router as compared to using a non-hierarchical approach based on Mobile-IP.* Even if the processing time for a Mobile-IP registration (P_{MR}) is optimized to a much lower value, the total number of control messages received by a HA is still almost three times the number of messages received by a domain root router in HAWAII.

8. Quality of Service support

Methods for providing QoS support for wired hosts include per-flow reservation approaches such as RSVP [13]. Rather than designing new QoS mechanisms for mobile hosts, we contend that HAWAII's localized mobility management enables an efficient adaptation of the wireline QoS mechanisms to wireless access networks.

Per-flow QoS reservation in the network requires identifying the address of both end-points of the flow. If either end-point changes its address, possibly because of mobility, then fresh end-to-end reservations have to be established. Protocols such as RSVP assume that hosts have fixed addresses; they use the destination address of the end node, i.e. the mobile host's care-of-address, to identify a session. Therefore, when the mobile host's care-of-address changes as it moves, one has to redo the resource reservation along the entire path from the correspondent host (or HA) to the mobile host. This must be performed even though most of the path is probably unchanged, as handoff is a local phenomenon. This results in increased reservation restoration

latency and unnecessary control traffic. While solutions such as flow extension via RSVP tunnels [10] may limit the reservation restoration latency, they still have a high overhead because of reservations along multiple paths.

In the case of HAWAII, support for QoS is simplified since a mobile host's address remains unchanged as long as the user remains within a domain. The details of the interaction between HAWAII and RSVP may be found in [7]. The straightforward integration of RSVP and HAWAII is due to the fact that RSVP was designed to blindly follow the routing path established and maintained by an independent routing entity. The HAWAII path setup messages for a mobile host handoff are no different from any other routing changes to which RSVP was designed to respond. Since HAWAII routing changes are localized, they result in fast reservation restorations for the mobile user

9. Reliability

In this section, we examine the impact of failure of each one of the HAWAII components. Failure of home agents is a concern for any approach that is based on Mobile-IP. In HAWAII as well as Mobile-IP, this failure could be tackled through the configuration and advertisement of backup home agents. Note that this could result in no connectivity to the mobile host for the renewal period. However, recall that in HAWAII, in the common case of a mobile host not leaving its home domain, there is no HA involved. This greatly reduces HAWAII's vulnerability to HA failure as compared to the Mobile-IP schemes.

We next examine failures of links/routers inside the HAWAII domain. In these cases, HAWAII relies on standard intra-domain routing protocols such as RIP or OSPF to detect router and link failures. When a failure is detected, HAWAII triggers soft-state refresh messages to restore connectivity. Let us examine this in more detail. These failures can be divided into two cases: link and router failures other than the domain root router, and domain root router failures.

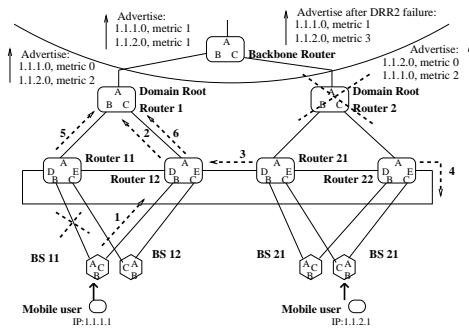


Figure 6. Link/Domain Root Router failures

Link and router failures other than the domain root

router are overcome without the involvement of any external routers. For example, consider the failure of link connecting BS 11 and Router 11 in Figure 6. This would trigger a change in the default route in BS 11 by a routing daemon. The change in default route would result in a soft-state refresh being sent to Router 12 (message 1). Router 12 would also trigger an immediate soft-state refresh to Domain Root Router 1 (message 2) and end-to-end connectivity would be re-established.

Recovery from the failure of the domain root router is also illustrated in Figure 6. When Domain Root Router 2 fails, it results in the update of default routes by the routing daemon on Routers 21 and 22. The change in default route triggers soft-state refreshes (messages 3, 4) to be sent towards Routers 12 and 11 respectively, which would then trigger immediate refreshes (messages 5, 6) to the Domain Root Router 1. Meanwhile, the backbone router would also detect the failure of Domain Root Router 2 and start forwarding packets destined for 1.1.2.0 to Domain Root Router 1. Thus, connectivity to hosts in 1.1.2.0 would be restored.

Summarizing, two design aspects of HAWAII that help achieve high reliability are the use of soft-state refreshes and, in some cases, the elimination of the HA. The robustness of HAWAII under subtle failures, such as route instability caused by misbehaving routing daemons, is the subject of future study.

10. Conclusions

In this paper, we presented the design, implementation, and performance evaluation of HAWAII: a domain-based approach for supporting mobility in wide-area wireless networks. The five design goals of HAWAII were *scalability*, *efficient routing*, *limited disruption*, *QoS support*, and *reliability*. We showed through simulation and implementation measurements how the HAWAII path setup schemes perform better than the Mobile-IP and Mobile-IP RO schemes in terms of reduced disruption to audio/video traffic, better TCP throughput, and reduced update traffic generated due to user movements. Quality of Service support is simplified through the design choices of using co-located care-of addresses and maintaining the mobile host address unchanged within the domain. This helps ensure that each mobile host is uniquely identifiable for classification purposes and does not affect reservations in external domains due to local mobility. Furthermore, reliability is achieved through maintaining soft-state forwarding entries for the mobile hosts and leveraging fault detection mechanisms built in existing intra-domain routing protocols.

These advantages are achieved at the expense of propagating host specific routes in selective routers within the domain. However, by judiciously limiting the number of host entries through appropriate sizing of the domain, and

limiting updates by managing mobility locally, we illustrated how large domains can be supported without the involvement of Mobile-IP. Thus, we conclude that HAWAII is a comprehensive solution for micro-mobility support and seamlessly works with Mobile-IP in order to support wide-area user mobility.

We are currently examining security issues with respect to HAWAII. One straightforward solution is to adopt the same security model as defined in the Mobile-IP RO approach [5]. In this case, the mobile host would simply send regular Mobile-IP registrations during handoff into the network. Transparently to the mobile host, HAWAII path setup messages would be used inside the domain.

References

- [1] R. Caceres and V.N. Padmanabhan, "Fast and Scalable Handoffs for Wireless Internetworks," in ACM MOBI-COM, November 1996.
- [2] K.Y.Eng, et al, "BAHAMA: A Broadband Ad-Hoc Wireless ATM Local Area Network," ACM/Baltzer Journal on Wireless Networks, Vol 1., No. 2, 1995.
- [3] S. Mohan and R. Jain, "Two User Location Strategies for Personal Communications Services," IEEE Personal Communications, Vol 1., No. 1, 1994.
- [4] C.E. Perkins, "IP Mobility Support," Request for Comments 2002, October 1996.
- [5] C.E. Perkins, and D.B. Johnson, "Route Optimization in Mobile IP," Internet Draft, November 1997.
- [6] R. Ramjee, T. La Porta, J. Kurose, and D. Towsley, "Performance Evaluation of Connection Rerouting Schemes for ATM-based Wireless Networks," in IEEE/ACM Transactions on Networking, Vol 6., No. 2, June 1998.
- [7] R. Ramjee et al, "IP Micro-Mobility support through HAWAII," Internet Draft, March 1999.
- [8] R. Ramjee et al, "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>.
- [9] S. Seshan, H. Balakrishnan, and R.H. Katz, "Handoffs in cellular wireless networks: The Daedalus implementation and experience," International Journal on Wireless Communication Systems, 1996.
- [10] A. Terzis, L. Zhang, and E. Hahne, "Reservations for aggregate traffic: experiences from an RSVP Tunnels Implementation," International Workshop on QoS, IWQoS, 1998.
- [11] A. Valko, A. Campbell, and J. Gomez, "Cellular IP," Internet Draft, November 1998.
- [12] S.Y. Wang and H. Kung, "A Simple Methodology for Constructing an Extensible and High-Fidelity TCP/IP Simulator," in Proceedings of Infocom 1999.
- [13] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: A New Resource ReSerVation Protocol," IEEE Network Magazine, September 1993.