

HAWAII: A Domain-Based Approach for Supporting Mobility in Wide-Area Wireless Networks

Ramachandran Ramjee, *Member, IEEE*, Kannan Varadhan, Luca Salgarelli, Sandra R. Thuel, Shie-Yuan Wang, and Thomas La Porta, *Fellow, IEEE*

Abstract—Mobile IP is the current standard for supporting macromobility of mobile hosts. However, in the case of micromobility 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 micromobility. 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 quality-of-service (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 macromobility, thus providing a comprehensive solution for mobility support in wide-area wireless networks.

Index Terms—Handoff, micromobility, Mobile IP, wireless.

I. INTRODUCTION

MOBILE IP is the current standard for supporting macromobility in IP networks [1]. 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.

However, the Mobile IP paradigm needs to be enhanced to cope with micromobility, i.e., movement across multiple subnetworks within a single network or domain. As specified, Mobile IP can result in disruption to user traffic during handoff; it also has high control overhead due to frequent notifications to the

HA [2]. A recent proposal for route optimization (RO) in Mobile IP [3] allows an FA to forward packets to a new FA following a handoff, thereby reducing traffic disruptions due to handoffs. Still, the mobile device's care-of address changes each time the user moves between neighboring base stations, resulting in notifications to the HA and the correspondent hosts on every handoff; this can be undesirable when there is high user mobility. 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, Mobile IP has some limitations when applied to wide-area wireless networks with high mobility users that may require QoS. Our approach for addressing these limitations hinges on the assumption that most user mobility is local to a domain, in particular, an administrative domain of the network. Therefore, we extend Mobile IP through 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. Whereas it is reasonable to add host-based route entries in wireless access networks, it is unscalable to add such routes in backbone networks; for mobility across backbone networks, or inter-domain mobility, HAWAII defaults to using traditional Mobile IP schemes. This combination of HAWAII for micromobility within a domain and Mobile IP for macromobility across domains provides for scalable and robust mobility across all levels.

Later in this paper, we will 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. While it could be argued that the overhead of host-based forwarding entries in the access routers is excessive, 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 II,

Manuscript received April 11, 2000; revised July 9, 2001; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor J.-Y. LeBoudec.

R. Ramjee, L. Salgarelli, S. R. Thuel, and T. La Porta are with Bell Labs, Lucent Technologies, Holmdel, NJ 07733 USA (e-mail: ramjee@bell-labs.com; salga@bell-labs.com; thuel@bell-labs.com; tlp@bell-labs.com).

K. Varadhan is with Procket Networks, Inc., Milpitas, CA 95035 USA (e-mail: kannan.varadhan@procket.com).

S.-Y. Wang is with the Department of Computer Science and Information Engineering, National Chiao Tung University, Hsinchu, Taiwan 30050, R.O.C. (e-mail: shieyuan@csie.nctu.edu.tw).

Publisher Item Identifier S 1063-6692(02)05234-2.

and contrast the related work in the context of these goals in Section III. We then present an overview of our solution, HAWAII, in Section IV. In Section V, we introduce several path setup schemes for supporting mobility within a domain. In Sections VI and VII, 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 VIII, we describe how our design simplifies providing QoS in the wired portion of the network. In Section IX, we illustrate the impact of our design on reliability. Before concluding in Section XI, we briefly discuss the interactions between HAWAII and Mobile IP in Section X.

II. 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.
- *Simplify 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.

III. 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 [4], [5]. The extension approach also forms the basis of the Mobile IP RO proposal [3]. However, in the case of mobility solutions proposed for connection-oriented ATM networks [5], 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 remapped locally during handoffs. In the case of connectionless IP networks, a change in the 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 [6]. 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. One approach that addresses the *scalability* issue of multicasting, while supporting host mobility, is through the use of distributed core multicast (DCM) [7]. DCM uses a set of distributed core routers (DCR) placed strategically at the edge of the backbone. These DCRs send user multicast data through, for example, point-to-point tunnels to other DCRs, thus avoiding the need for multicast state in the backbone routers. DCM can be used to provide host mobility as follows: when a mobile host enters a domain, it is assigned a multicast address as a care-of address for the domain. The mobile host maintains this multicast address as long as it is in the domain, similar to HAWAII, where the unicast address remains unchanged in the domain. Inside the domain, per-group multicast state is kept, similar to per-host state in HAWAII. A correspondent host would send packets to the multicast address of the mobile host. A DCR in the correspondent host's domain would then tunnel the packet to the DCR that is serving the mobile host. Multicast routing entries within the domain would ensure the delivery of packets to the mobile host. In order to avoid *disruption*, DCM uses advance multicast joins. However, this approach does not have the flexibility of choosing a path setup scheme adapted for a specific wireless link from the multiple path setup schemes as in HAWAII.

A common technique to enhance *scalability* is to introduce a hierarchy. In traditional Mobile IP-based architectures, there is no notion of a domain and the mobile node is directly attached either to the home domain root router (called the home agent) or the foreign domain root router (called the foreign agent). Thus, every handoff causes a change of the globally routable IP address for the mobile, resulting in *scalability*, *efficiency*, and *QoS* support difficulties. To address the *scalability* issue, one proposal is to build a hierarchy of foreign agents [2]. 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 [8] is similar in spirit to HAWAII; it uses specialized domain routers with host-based entries for local mobility and the use of Mobile IP for inter-domain mobility. Thus, updates can be localized, enhancing the *scalability* of update mechanisms and limiting *disruption*. Unlike HAWAII, however, the routers in this proposal automatically detect that a mobile user has been handed off by snooping actual data packets. In addition, Cellular IP relies on the gateway to act as an FA that decapsulates the packets before delivering them to the user. The presence of the Gateway FA

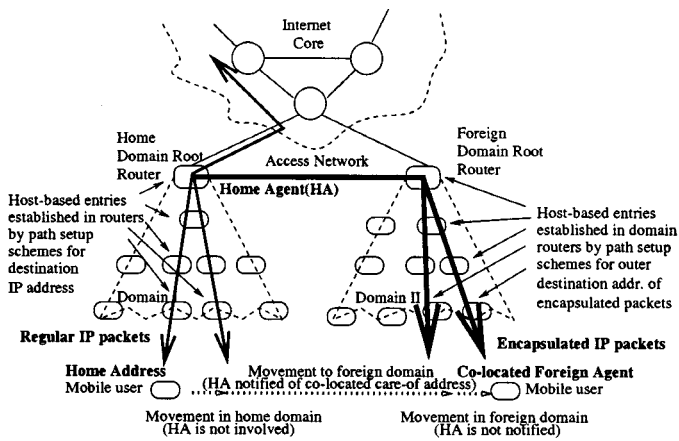


Fig. 1. Hierarchy using domains.

can complicate *QoS* management since the backbone nodes cannot distinguish easily between the packets sent to different mobile hosts, but are tunneled to the same Gateway FA. The Gateway FA also potentially impacts *reliability*.

A comparison of the performance of Cellular IP and HAWAII can be found in [9]. The authors find that the performance of HAWAII and Cellular IP are similar for networks with tree topologies but when the network topology is not a tree, different crossover routers are chosen during handoff in the two approaches, resulting in small differences in performance. In Cellular IP, optimal downlink paths are ensured at the cost of propagating update messages higher up in the hierarchy (leading to potential bottlenecks), while, in HAWAII, routing updates are kept close to the access points resulting in localized update processing, but at the cost of possibly nonoptimal routing. Overall, their conclusion is that the choice of a micromobility protocol should be dictated more by deployment considerations such as whether to do implicit or explicit signaling, etc., than the small differences in terms of handoff performance.

We next present an overview of the HAWAII protocol, highlighting how the various design choices help toward achieving these goals.

IV. 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 Fig. 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 using host-based routes in routers in the domain to the mobile host. The establishment of these paths in a single domain are discussed in detail in Section V.

When the mobile host first enters into a foreign domain, we revert to traditional Mobile IP mechanisms and the mobile host

is assigned a co-located care-of address¹ using DHCP. Packets are tunneled to the care-of address by a home agent in its home domain. If the foreign domain is also based on HAWAII, then for subsequent movements within the foreign domain, the mobile host retains its care-of address unchanged, and connectivity is maintained using dynamically established paths of HAWAII.

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

A mobile host that first powers up and attaches to a domain sends a *path setup powerup 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 toward 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.²

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 selected 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 V, 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, send periodic *aggregate hop-by-hop refresh messages* toward 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 V, ensure that data disruption during handoff is limited. The disruption caused by various schemes for audio and video traffic is quantified in Section VI.

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 VI.

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.

¹In this paper, we assume the co-located care-of model of Mobile IP even though HAWAII can work with the network-based foreign agent model of Mobile IP as well.

²In 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 toward the domain root router. Any intermediate router that then has the route to the destination host will forward the packet downstream toward that host. In the worst case, that intermediate router could be the domain root router that then forwards the packet to the mobile host.

One potential concern is the number of mobile hosts that can be attached to, and supported by, a single domain as the domain root router can become a processing bottleneck. In Section VII, we present a numerical example showing how a single domain in HAWAII can include over 100 base stations in a typical wide-area wireless network. In order to enhance the scalability further, multiple domain root routers can be used in a single domain (with mobile hosts sending updates toward different root routers) in a manner similar to using multiple DCRs in DCM [7].

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 VIII. One drawback of using the co-located care-of address option is the need for two IP 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 IX, 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.

The reader is referred to [10] for a more detailed description of the protocol.

V. HAWAII PATH SETUP SCHEMES

We first describe the path setup messages that are initiated after powerup, assuming an address has been obtained by the mobile host. We then 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. 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. We use a tree-based topology in our examples for clarity. Note that our schemes will work in any general topology. In particular, Section IX illustrates how recovery is accomplished in non-tree-based topologies.

A. Path Setup Message After Powerup

Fig. 2 illustrates the sequence of path update messages during powerup. The forwarding table entries are shown adjacent to the routers. These entries are prepended with a message number indicating what message was responsible for establishing the entry (a message number of zero indicates a pre-existing entry). The letters denote the different interfaces. The mobile host sends an update message to its current base station to set up path state. The mobile host sets the destination field of the message to the domain root router.

When the current base station receives the path setup message, as illustrated in Fig. 2, it adds a forwarding entry for the

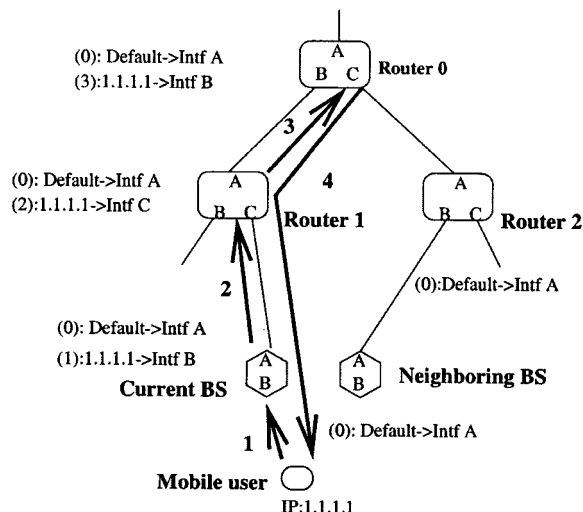


Fig. 2. Path setup message after powerup.

mobile host setting the outgoing interface to the interface on which it receives the message (the wireless interface in this case). It then forwards the path setup message to the next hop router, Router 1, along its default route to the domain root router. Router 1 performs similar processing and forwards the message to Router 0 (domain root router in this example). Router 0 adds an entry for the mobile host, and since it is the intended destination for the update message, sends an acknowledgment back to the mobile host (shown as message 4 in Fig. 2). At this time, packets destined for the mobile host arrive at the domain root router based on the subnet portion of the mobile host’s IP address. The packets are routed within the domain to the mobile host, using the host-based forwarding entries just set up.

Note that other routers in the domain have no knowledge of the mobile host’s IP address. If these routers receive packets for the mobile host, say, from another mobile host, they forward the packets on their default route to the domain root router. The domain root router will forward the packets to the mobile host.

If the mobile host is in its home domain, the powerup procedure is complete. If the mobile host is in a foreign domain, it will register its IP address with its home agent. Once this is complete, the aggregate refresh messages from the base station contain the mobile host IP address.

For the remaining subsections, let us define the crossover 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.³ A router is able to identify itself as a crossover router during a given handoff as follows. If the interfaces corresponding to the host route for the mobile host before and after handoff are different from the interface for the default route (in other words, the interface continues to be “downstream” after handoff), then the router is a crossover router for this mobile host.

The four path setup schemes considered in this paper can be classified into two types, *forwarding* and *nonforwarding*, based on the way packets are delivered to the mobile host during a

³The schemes considered in this paper can be modified to work with other definitions of the crossover router as well.

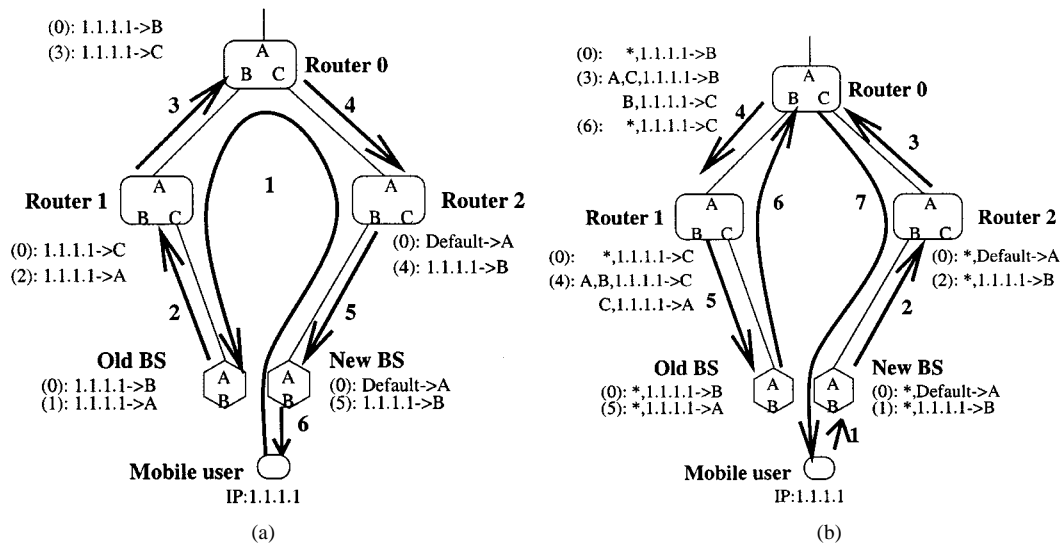


Fig. 3. Forwarding schemes. (a) MSF. (b) SSF.

handoff. In the forwarding schemes, packets are forwarded from the old base station to the new, whereas in the nonforwarding schemes, they are diverted at the crossover router. Thus, forwarding schemes are independent of the wireless link and rely on the wired network to buffer packets and forward them to the new base station for seamless delivery, and the nonforwarding schemes take advantage of properties of certain wireless links where both old and new base stations can maintain connectivity to the mobile host for seamless delivery during handoff. We first describe two forwarding schemes followed by two nonforwarding path setup schemes.

B. 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 crossover router.*

The idea of forwarding packets during handoff is not new [2]–[5]. In the case of ATM networks [4], [5], 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 ARP mechanisms if the user stays within the same broadcast network [2] or through tunneling [3]. 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 Fig. 3(a). The forwarding table entries are shown adjacent to the routers. 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 crossover 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 toward the new base station. Eventually, Message 5 reaches the new base station which changes its forwarding entry and sends an acknowledgment of the path setup message to the mobile host, shown as Message 6.

Note that this order of updating the routers can lead to the creation of multiple streams of misordered 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 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 misordered streams and routing loops exist for extremely short periods of time. The duration of these anomalies is a function of the protocol timers, and can be tightly controlled by having fairly small timeout values before forwarding entries are deleted. The main benefit of this scheme is that it is simple and results in no loss.

As an alternative, the 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*




Fig. 4. Nonforwarding schemes. (a) UNF. (b) MNF.

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). In Fig. 3(b), Messages 1–5 establish these entries resulting in packets arriving at the old base station and being forwarded to the new base station as a single stream. The old base station subsequently sends Message 6 to Router 0 for diverting the stream at the crossover router. After processing Message 6, Router 0 sends an acknowledgment of the path setup message to the mobile host, shown as Message 7, and 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 until the diversion is performed at the crossover router (until Message 6), it is somewhat complex to implement. In Section VI, 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, which involve routers that are only one or two hops away.

C. Nonforwarding 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 crossover router to the new base station, resulting in no forwarding of packets from the old base station.*

There are two variants of the nonforwarding scheme, motivated by two types of wireless networks. The unicast nonforwarding (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 nonforwarding (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, nonforwarding schemes have been studied in the context of ATM networks [5] and in the multicasting-based approaches [6]. 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 crossover router multicasts data packets to at most two of its interfaces during handoff.

The UNF scheme is illustrated in Fig. 4(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 crossover 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 which then changes its forwarding entry and sends an acknowledgment, Message 6, back to the mobile host.

The MNF scheme is very similar to the UNF scheme. The main difference is that the crossover router, Router 0, multicasts data packets for a short duration. In Fig. 4(b), Router 0 dualcasts 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.

VI. DISRUPTION

In this section, we use simulation to compare the disruption performance of the four HAWAII and two Mobile IP schemes.

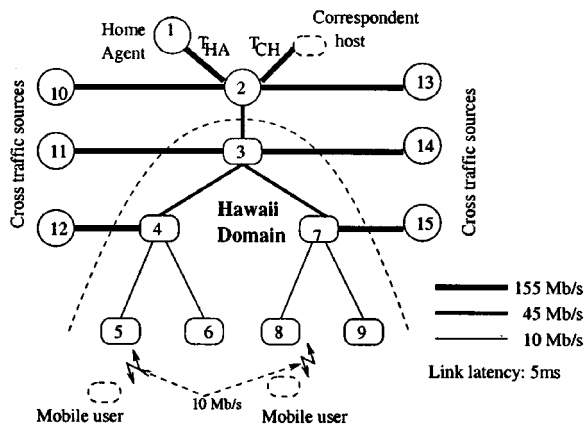


Fig. 5. Simulation topology.

These were simulated using the HARVARD simulator [11]. 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 need to be notified before packets go directly to the new base station, while in HAWAII, local updates result 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 semilocal (Mobile IP RO) and nonlocal schemes (Mobile IP) for supporting IP mobility.

The simulation topology is shown in Fig. 5.

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.⁴ We begin by presenting the results for UDP-based audio and video applications during intra-domain handoffs. We then go on to describe our results for TCP traffic of mixed duration flows, including a mix of FTP- and Web-based traffic under identical handoff conditions.

A. UDP—Audio and Video

In the case of audio experiments, the correspondent host transmits 160-byte UDP packets every 20 ms (64-kb/s audio) to the mobile host. On every handoff of the mobile host, we collect statistics on the incoming UDP packets in the downlink direction.⁵

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

⁴For nonoverlapping cells, both HAWAII and Mobile IP schemes would need to be augmented with buffering capabilities to avoid user level disruption.

⁵In the uplink direction, the data path from the mobile host to the correspondent host is identical in all the schemes, resulting in similar performance.

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}.

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

Fig. 6(a) plots the disruption caused to an audio application during handoff when the crossover 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 of getting dropped, resulting in a 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 five packets per handoff are lost in the network. This is because, in our configuration, the registration update from the mobile host takes about 100 ms (link delay of 50 ms and queueing delay of 50 ms) to reach the HA. In this interval, about five packets are sent to the old base station and are lost.⁶

Let us now compare the remaining five handoff schemes. Consider a playout delay value of 100 ms in Fig. 6(a). In this case, the Mobile IP RO scheme results in a total loss of about three packets per handoff, while the HAWAII schemes result in the total loss of less than one packet per handoff. This is because the HAWAII schemes switch over very quickly to the new route, while in the 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 nonforwarding 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

⁶If we assume that the MH can listen to both base stations until the HA is informed, then there could be no loss.




Fig. 6. Packet loss during two-hop handoff. (a) Audio. (b) Video.

in packet loss in the network (and duplicates) of about 0.25 packets per handoff.⁷ Note that for higher values of playout delay, the RO scheme may match the total loss values of the HAWAII schemes. For example, in Fig. 6(a), with 150 ms of playout delay, the Mobile IP RO scheme results in comparable total loss as the HAWAII UNF scheme with 100 ms of 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 requires 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 Fig. 6(b), are similar to the results for audio except for slightly lower total losses due that fact that (4-kB) UDP packets are sent every 33 ms rather than the 20-ms 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 because they operate locally. For Mobile IP (Mobile IP RO) schemes, as shown in Fig. 7, when T_{HA} decreases, the performance approaches that of the HAWAII nonforwarding (forwarding) schemes. The impact of T_{CH} on the Mobile IP RO scheme is similar, except for an increase in the end-to-end delay.

⁷The reason for packet loss in the MNF scheme is subtle; a packet that is delayed arriving at the base station before the handoff may not reach the mobile host if the host has since completed the handoff.

Summarizing the UDP results, the localized HAWAII schemes result in less 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 semilocal 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.

B. TCP—Web and File Transfers

We first consider the effect of the mobility schemes on Web browsing. A typical Web page contains several components, each of which, when using a protocol such as HTTP 1.0, requires a separate TCP connection to be established. Because each TCP connection is short-lived, disruptions due to handoff occur very infrequently, and are therefore not a great concern. However, the tunneling of TCP packets may have a side effect of increasing the latency of page downloads.

When a Web page is being transferred, typically the first TCP data packet uses the full MTU. Then, when a tunneling header is added by an HA, the MTU size is violated. If the Don't Fragment flag is set, as is typically the case, an ICMP error is sent from the HA to the corresponding host, resulting in an extra




Fig. 7. Impact of T_{HA} on Mobile IP schemes. (a) Mobile IP. (b) Mobile IP RO.

round-trip delay. This procedure may be repeated for each new TCP connection,⁸ resulting in a cumulative delay in the Mobile IP schemes when downloading a single Web page. Recall that in HAWAII, tunneling is rare since it is used only when users move out of their domain. Thus, this phenomenon has minimal impact on performance in HAWAII.

We next consider the impact of handoff on file transfer applications where FTP is used to download a very large file to a mobile host. We use the same simulation topology shown in Fig. 5. In this case, 20 mobile users are performing file downloads while moving amongst the four base stations. Each user is handed off, on the average, every 10 s. T_{CH} is 50 ms and T_{HA} is 5 ms.

The average aggregate throughput of all these users for the various schemes is shown in Fig. 8(a). We find that the HAWAII schemes deliver sizable 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. Among the HAWAII schemes, the UNF and SSF schemes deliver the best performance for mobile devices which can listen to two or one base stations, respectively. The MSF and the MNF schemes approach the performance of the UNF and SSF schemes when the link buffer sizes (at the routers) are

⁸If the Web server caches the path MTU value and multiple invocations are served by the same server, then there would be no penalty after the first invocation.

large. Note that the magnitude of the improvement of HAWAII schemes over Mobile IP schemes depends on the rate of user handoffs. In Fig. 8(b), as the handoff frequency in the domain is decreased, the schemes deliver similar aggregate throughput.

VII. 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 nonhierarchical 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 toward its destination. Soft-state refresh messages are sent independently by each of the nodes every T_R seconds. Typically, processing the refresh message simply involves updating the expiry timer in the HAWAII daemon and can be performed very efficiently.

Table I lists the processing time of HAWAII and Mobile IP update and refresh messages measured on a Pentium II 333-MHz CPU running the FreeBSD 2.2.7 operating system. The reason for the relatively large processing time at the




Fig. 8. Aggregate TCP bandwidth. (a) Impact of link buffer size. (b) Impact of user speed.

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 proxy ARP 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 II. 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 (20 per router). We now consider two different approaches: 1) the Mobile IP approach, where FAs are present at each base station and are served by an HA, and 2) the HAWAII approach where the HA is at the domain root router (DRR).

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 $\rho A_D = 38220$. This is also the same as the number of tunneling entries in the case of the nonhierarchical Mobile IP approach at the HA.⁹ Note that 40 K entries are well within the capability of modern routers.

⁹While one could potentially have as many HAs as base stations in the Mobile IP approach, two practical reasons would preclude it: 1) HAs need to be very reliable, as they are single points of failure, and 2) administration and management of multiple HAs and user profiles can become complicated.

TABLE I
CPU PROCESSING TIMES

Item	Message type	time (μs)
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 II
EXAMPLE CONFIGURATION VALUES

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

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.

From the derivation in the Appendix, the processing load at the HA in the Mobile IP approach, CPU_M , is given by the formula

$$\text{CPU}_M = P_{\text{MU}} \frac{\rho \nu L_B B_D}{\pi} + P_{\text{MR}} \frac{\rho L_B^2 B_D}{16T_M} \quad (1)$$

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

Similarly, the processing load at the domain root router in HAWAII, CPU_H , is

$$\begin{aligned} \text{CPU}_H = & P_{\text{MU}} \frac{\rho \nu L_B \sqrt{B_D}}{\pi} + P_{\text{MR}} \frac{\gamma \rho L_B^2 B_D}{16T_M} \\ & + P_{\text{HU}} \frac{\rho \nu L_B \sqrt{B_D R_D}}{\pi} + P_{\text{HR}} \frac{[\rho L_B^2 B_D / 16Y]}{T_R} \end{aligned} \quad (2)$$

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 (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 (2)] varies with *the square-root of the number of base stations* $O(\sqrt{B_D})$. Recall that the processing of Mobile IP updates is more expensive than HAWAII updates ($P_{\text{MU}} > P_{\text{HU}}$; see Table I). Thus, term 1 in (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 (1) and terms 2 and 4 in (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 (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 (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.

The numerical results for the configuration shown in Table II are summarized in Table III. 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 nonhierarchical 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 an HA is still almost three times the number of messages received by a domain root router in HAWAII.

TABLE III
RESULTS

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 Regn.	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

VIII. QUALITY-OF-SERVICE SUPPORT

Methods for providing QoS support for wired hosts include per-flow reservation approaches such as RSVP [12]. 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 [13] may limit the reservation restoration latency, they still have a high overhead because of reservations along multiple paths.

The interaction between HAWAII and RSVP is shown in Fig. 9. The case when the mobile host is a receiver is shown in Fig. 9(a). The state in the solid box represents the HAWAII forwarding state, while the state in the dotted box represents the RSVP state, comprising a destination address (DEST), a previous hop (PHOP), and a next hop (NHOP). After Router 0 processes a HAWAII path setup update, its RSVP daemon receives a path change notification (PCN) (Message 1) using the routing interface for RSVP [14]. In standard RSVP, the router must now wait a time interval before generating the RSVP PATH message to allow the route to stabilize; this time interval is set to 2 s by default. In HAWAII, the RSVP PATH message (Message 2) can be triggered immediately on receiving a PCN since the route to the mobile host is stable at that point. This allows for a faster reconfiguration due to mobility. The PATH message follows the new routing path (Messages 2 and 3), installing PATH state on all the routers toward the new base station. When this PATH message reaches the mobile host, a QoS agent on the host generates an RSVP RESV message upstream that follows the reverse forwarding path (Messages 5, 6, and 7). Router 0 stops forwarding the RESV messages upstream since there is no change in the reservation state to be forwarded. Thus, reservations are restored locally in a timely manner. The case when the mobile host is a sender is fairly




Fig. 9. Interaction of RSVP and HAWAII. (a) Mobile host as receiver. (b) Mobile host as sender.

simple and is illustrated in Fig. 9(b). A RSVP PATH message is sent by the mobile host after handoff as soon as the HAWAII path setup is complete, resulting in reservations along the new path.

Note that 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. Thus, intra-domain handoffs in HAWAII are handled efficiently; since they are localized, they result in fast reservation restorations for the mobile user. In the case of inter-domain handoffs, HAWAII defaults to Mobile IP for mobility management; therefore, reservation restorations would follow along the procedures elaborated by the Mobile IP working group.

IX. 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

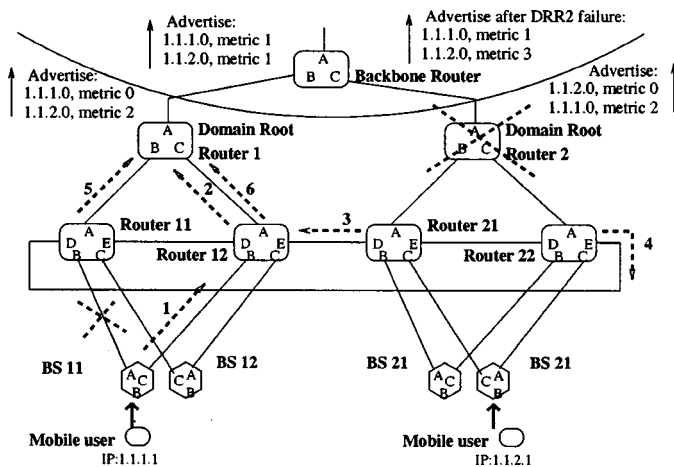


Fig. 10. Link/domain root router failures.

and router failures other than the domain root router, and 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 the link connecting BS 11 and Router 11 in Fig. 10. 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 Fig. 10. When DRR 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 and 4) to be sent toward Routers 12 and 11, respectively, which would then trigger immediate refreshes (Messages 5 and 6) to DRR 1. Meanwhile, the backbone router would also detect

the failure of DRR 2 and start forwarding packets destined for 1.1.2.0 to DRR 1. Thus, connectivity to hosts in 1.1.2.0 would be restored.

To summarize, 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 anomalies, is the subject of future study.

X. MOBILE IP INTERACTIONS

Many a dragon hides in the complexities of multiprotocol interaction. HAWAII and Mobile IP protocols operate in parallel at the domain root router and they interact with each other when the mobile host performs an inter-domain handoff.

Furthermore, in order to make HAWAII transparent to the mobile host, it is possible for the mobile host to communicate with the network by using only Mobile IP and sending Mobile IP registrations during each handoff. Transparent to the host, HAWAII path setup messages need to be generated within the domain. This type of interaction between HAWAII and Mobile IP would occur at the base stations during intra-domain handoff.

In this section, we consider each of these cases.

A. Issues During Inter-Domain Mobility

Recall that, in HAWAII, the domain root router acts as the HA for hosts that are in a foreign domain. Therefore, HAWAII processing is required if the host is within the domain, and Mobile IP processing when it is roaming in a foreign domain. The protocols should coordinate with each other to maintain forwarding entries for the mobile host, so that interleaved arrival of protocol messages do not leave the forwarding tables in an inconsistent state.

After the domain root router has processed the Mobile IP registration message, indicating that the host has moved from its home to a foreign domain, HAWAII must ignore refresh messages for this host from downstream routers; these are stale refresh messages and will eventually be timed out. Similarly, when the host has moved back from the foreign domain to its home domain, the domain root router should process the HAWAII update; however, Mobile IP should process any subsequent deregistration messages from the host only to remove its internal state, without affecting the forwarding entries.

Another issue is that Mobile IP requires the home agent to add proxy ARP entries [15] for those mobile hosts that are in the foreign domain, and remove them when it receives an explicit deregistration message. The HA on the domain root router need not set such ARP entries, since data packets will reach the domain root router based on the address allocation architecture of HAWAII. Moreover, such ARP entries could interfere with connectivity when the host revisits its home domain; if the Mobile IP deregistration message is lost, the ARP entry causes packets to be encapsulated and forwarded to the previous foreign domain by the HA. Thus, no packets can be sent to the mobile host until the Mobile IP state is removed following a timeout. Hence, the home agent on the domain root router must disable

its proxy ARP processing, as it is unnecessary when the host is in a foreign domain, and will interfere with HAWAII processing when the host revisits its home domain.

B. Issues During Intra-Domain Mobility

Allowing mobile hosts to communicate with the base stations using Mobile IP messages, instead of specialized HAWAII messages, provides a number of advantages. There is an existing base of Mobile IP host implementations, and these hosts need not support another new protocol. It allows for the incremental deployment of HAWAII in various access networks. Security considerations are simplified, in that we can adopt the same security models as defined in the Mobile IP RO approach [3].

In this approach, the mobile host runs the standard Mobile IP protocol with NAI, route optimization and challenge/response extensions. To reduce the frequency of updates to the HA and avoid high latency and disruption during handoff, we split the processing and generation of Mobile IP registration messages into two parts: between the mobile host and the base station and between the base station and the HA. Note that this separation is needed for any approach that desires to reduce updates to the HA. For example, similar separation at the foreign agent is proposed in the Mobile IP Regionalized Tunnel Management approach as well [16].

A similar issue exists when the host is roaming in a foreign domain that is HAWAII-enhanced. Mobile hosts in a HAWAII domain use a co-located care-of address (CCOA); such hosts, according to Mobile IP, are required to always contact their HA directly. Again, this is in conflict with reducing the frequency of updates to the HA. We advocate that the mobile hosts register with a base station even while using the CCOA option. The base station helps reduce the frequency of updates to the HA by processing the registrations locally and also ensures smooth handoffs by forwarding packets if necessary. This approach also allows networks to enforce security and authentication measures in their domain. Thus, data packets are sent directly from the HA to the mobile host, while registrations are processed in two stages: at the base station and the HA.

XI. CONCLUSION

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. QoS 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 selected 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 micromobility support and seamlessly works with Mobile IP in order to support wide-area user mobility.

An interesting aspect of the protocol presented in this paper is the coupling between HAWAII and Mobile IP at the domain root router. Operational or administrative policy concerns might dictate the need for less coupling. It is interesting to conjecture the effects of such a system, possibly with the use of a distinct Mobile IP HA. This could provide greater flexibility in deployment as well as other advantages, such as scalability in terms of memory and processing power at the domain root router. However, there will be the added complexity of interactions between these protocols. It remains to future work to investigate the coupling between HAWAII and Mobile IP, and to more systematically evaluate the tradeoffs between these two protocols.

APPENDIX

DERIVATION OF PROCESSING LOAD DUE TO UPDATES AND REFRESHES

In this Appendix, we derive the update and refresh processing loads for systems based on Mobile IP and HAWAII. We assume that the coverage area of a base station is a square with a perimeter of L_B and area of $L_B^2/16$. We also assume that there are B_D base stations in the domain structured as a tree, resulting in a domain coverage area of $L_B^2 B_D/16$. If ρ is the density of active users, the number of users in the domain is

$$N = \rho L_B^2 B_D/16.$$

First, consider a system based on Mobile IP. Let M_U denote the rate of mobility related updates at the HA from the B_D base stations. Assuming the direction of user movement is uniformly distributed over $[0, 2\pi]$ and using a fluid flow mobility model [17], the rate of mobile hosts crossing a boundary of perimeter l at a speed v is

$$R(l) = \rho v l / \pi.$$

Since user handoffs between any two base stations in the domain generates an update registration at the HA, $M_U = R(L_B) B_D$ or

$$M_U = \rho v L_B B_D / \pi.$$

Let the rate of registration renewals (or refreshes) at the HA be M_R . We assume here that the renewal period is not reset even if the host sends an intermediate registration. During each renewal period, T_M , every user in the domain sends out one renewal request. Thus, $M_R = N/T_M$, or

$$M_R = \frac{\rho L_B^2 B_D}{16 T_M}.$$

The processing load at the HA, CPU_M , is simply

$$\begin{aligned} \text{CPU}_M &= P_{MU} M_U + P_{MR} M_R \\ &= P_{MU} \frac{\rho v L_B B_D}{\pi} + P_{MR} \frac{\rho L_B^2 B_D}{16 T_M} \end{aligned}$$

where P_{MU} and P_{MR} are CPU processing times for Mobile IP registration updates and renewals, respectively.

Now consider a system based on HAWAII. The domain root router is the most heavily loaded router in this system, as it has to process both path setup messages as well as Mobile IP messages.

Let the rate of Mobile IP updates at the domain root router be H_{MU} . Updates at the HA occur only when users cross domain boundaries. The rate of domain boundary crossing is $R(l)$, where l , the perimeter of the domain coverage area, is $4\sqrt{L_B^2 B_D/16} = L_B \sqrt{B_D}$. Thus, $H_{MU} = R(L_B \sqrt{B_D})$, or

$$H_{MU} = \frac{\rho v L_B \sqrt{B_D}}{\pi}.$$

In HAWAII, Mobile IP registration renewals are sent by only those users who are away from their home domain. Let γ be the fraction of users away from their domain. Then, the rate of registration renewals in HAWAII, H_{MR} , is γM_R or

$$H_{MR} = \frac{\gamma \rho L_B^2 B_D}{16 T_M}.$$

Let the rate of path setup updates at the domain root router be H_{PU} . These updates are generated whenever a user is handed off between base stations attached to two different second level routers. Thus, $H_{PU} = R_D R(l)$ where $l = 4\sqrt{L_B^2 B_D/16 R_D}$ is the perimeter of a second-level router coverage area and R_D is the number of second-level routers in the domain. Substituting gives

$$H_{PU} = \frac{\rho v L_B \sqrt{B_D R_D}}{\pi}.$$

Finally, let the number of path setup refreshes at the domain root router be H_{PR} . Let a single refresh message be an aggregate for Y mobile hosts. During the refresh period, T_R , the path state of every user in the domain is refreshed. Thus, $H_{PR} = \lceil N/Y \rceil / T_R$, or

$$H_{PR} = \frac{\lceil \rho L_B^2 B_D / 16 Y \rceil}{T_R}.$$

The processing load at the domain root router CPU_H is

$$\begin{aligned} \text{CPU}_H &= P_{MU} H_{MU} + P_{MR} H_{MR} + P_{HU} H_{PU} + P_{HR} H_{PR} \\ &= P_{MU} \frac{\rho v L_B \sqrt{B_D}}{\pi} + P_{MR} \frac{\gamma \rho L_B^2 B_D}{16 T_M} \\ &\quad + P_{HU} \frac{\rho v L_B \sqrt{B_D R_D}}{\pi} + P_{HR} \frac{\lceil \rho L_B^2 B_D / 16 Y \rceil}{T_R} \end{aligned}$$

where P_{HU} and P_{HR} are CPU processing times for HAWAII path setup updates and refreshes, respectively. \square

REFERENCES

- [1] C. E. Perkins, "IP mobility support," RFC 2002, Oct. 1996.

- [2] R. Caceres and V. N. Padmanabhan, "Fast and scalable handoffs for wireless internetworks," in *Proc. ACM MOBICOM*, Nov. 1996, pp. 56–65.
- [3] C. E. Perkins and D. B. Johnson. (2000, Feb.) Route optimization in mobile IP. Internet draft, work in progress. [Online]. Available: <http://www.iprg.nokia.com/~charliep/txt/optim/optm.txt>
- [4] K. Y. Eng *et al.*, "BAHAMA: A broadband *ad-hoc* wireless ATM local-area network," *J. Wireless Networks*, vol. 1, no. 2, pp. 161–174, 1995.
- [5] R. Ramjee, T. La Porta, J. Kurose, and D. Towsley, "Performance evaluation of connection rerouting schemes for ATM-based wireless networks," *IEEE/ACM Trans. Networking*, vol. 6, pp. 249–261, June 1998.
- [6] S. Seshan, H. Balakrishnan, and R. H. Katz, "Handoffs in cellular wireless networks: The Daedalus implementation and experience," *Int. J. Wireless Pers. Commun.*, vol. 4, no. 2, pp. 141–162, Mar. 1997.
- [7] L. Blazevic and J.-Y. Le Boudec, "Distributed core multicast (dcm): A multicast routing protocol for many groups with few receivers," in *Comput. Commun. Rev.*, vol. 29, Oct. 1999, pp. 6–21.
- [8] A. T. Campbell, J. Gomez, and A. Valko, "An overview of cellular IP," in *IEEE Wireless Communications and Networking Conf. (WCNC)*, Sept. 1999, pp. 606–611.
- [9] A. T. Campbell, J. Gomez, S. Kim, Z. Turanyi, C.-Y. Wan, and A. Valko, "A comparison of IP micromobility protocols," *IEEE Wireless Commun. Mag.*, vol. 9, no. 1, Feb. 2002.
- [10] R. Ramjee, T. La Porta, L. Salgarelli, S. Thuel, K. Varadhan, and L. Li, "IP-based access network infrastructure for next generation wireless data networks," *IEEE Pers. Commun.*, vol. 7, pp. 34–41, Aug. 2000.
- [11] S. Y. Wang and H. Kung, "A simple methodology for constructing an extensible and high-fidelity TCP/IP simulator," in *Proc. IEEE INFOCOM'99*, 1999, pp. 1134–1143.
- [12] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: A new resource ReSerVation protocol," *IEEE Network Mag.*, vol. 7, pp. 8–18, Sept. 1993.
- [13] A. Terzis, L. Zhang, and E. Hahne, "Reservations for aggregate traffic: Experiences from an RSVP tunnels implementation," in *Proc. Int. Workshop QoS (IWQoS)*, 1998, pp. 23–25.
- [14] D. Zappala and J. Kann. (1998, July) RSRR: A routing interface for RSVP. Internet Draft, work in progress. [Online]. Available: <http://www.ietf.org/proceedings/99jul/I-D/draft-ietf-rsvp-routing-02.txt>
- [15] S. Carl-Mitchell and J. S. Quarterman, "Using ARP to implement transparent subnet gateways," RFC 1027, Oct. 1987.
- [16] E. Gustafsson, A. Jonnson, and C. Perkins. (2002, Mar.) Mobile IPv4 Regional Registration. Internet Draft, work in progress. [Online]. Available: <http://search.ietf.org/internet-drafts/draft-ietf-mobileip-reg-tunnel-06.txt>
- [17] S. Mohan and R. Jain, "Two user location strategies for personal communications services," *IEEE Pers. Commun.*, vol. 1, no. 1, pp. 42–50, 1994.



Ramachandran Ramjee (M'97) received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Madras, India, and the M.S. and Ph.D. degrees in computer science from the University of Massachusetts, Amherst.

He has been a Member of Technical Staff at Bell Labs, Lucent Technologies, Holmdel, NJ, since 1996. His research interests are signaling, mobility management, and quality-of-service issues in wireless and high-speed networks. He has published

over 20 technical papers and holds seven patents.

Dr. Ramjee serves as an associate editor of the IEEE TRANSACTIONS ON MOBILE COMPUTING and as a technical editor of the *IEEE Personal Communications Magazine*.

Kannan Varadhan received the Ph.D. degree from the University of Southern California, Los Angeles, in 1998.

He is currently with Procket Networks, Milpitas, CA. His research interests are in robustness issues in protocol design and routing protocol behavior in the Internet, and its observable behavior at the edges of the network and consequent effects on end-to-end protocols.



Luca Salgarelli received the Laurea (Dr. Eng. degree) in electronic engineering from the Milan Polytechnic University, Milan, Italy, in 1995, and the M.Phil. degree in computer science from CEFRIEL, Milan, also in 1995.

He was a Researcher in the Networking Department of CEFRIEL beginning in 1995, where he worked in the areas of broadband IP networks and QoS provisioning. Since 1998, he has been with Lucent Technologies, first in the Wireless R&D Department in the U.K., and then in the Mobile

Networking Department of Bell Labs Research, Holmdel, NJ. His research covers design, development, and evaluation of systems and protocols for data networks, in particular when they involve mobility, QoS provisioning, and security.



Sandra R. Thuel received the B.S. degree in computer engineering from the University of Puerto Rico, Mayaguez, in 1986, and the M.S. and Ph.D. degrees in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1988 and 1993, respectively.

Since 1993, she has been a Member of Technical Staff at Bell Labs, Lucent Technologies (formerly AT&T), Holmdel, NJ. Her research interests include mobility management, wireless voice and data networking, quality of service, scheduling, and

resource management.



Shie-Yuan Wang received the Ph.D. degree from Harvard University, Cambridge, MA, in 1999, the M.S. degree from National Taiwan University, Taipei, Taiwan, R.O.C., in 1992, and the B.S. degree from National Taiwan Normal University in 1990, all in computer science.

He is currently an Assistant Professor in the Department of Computer Science and Information Engineering, National Chiao Tung University (NCTU), Taiwan. His research interests include computer networks, network simulations, Internet

technologies, and operating systems. He developed the Harvard network simulator in 1999, and is now leading his team at NCTU in developing the NCTUNS network simulator.



Thomas La Porta (S'85–M'87–SM'99–F'02) received the B.S.E.E. and M.S.E.E. degrees from The Cooper Union, New York, NY, and the Ph.D. degree in electrical engineering from Columbia University, New York.

He joined AT&T Bell Laboratories in 1986. He is currently Director of the Mobile Networking Research Department of Bell Laboratories, Lucent Technologies, Holmdel, NJ, where has worked on various projects in wireless and mobile networking for the past several years. He is also the Director of

the Advanced Mobile Networking Department in the Wireless Business Unit of Lucent. He has published over 50 technical papers and holds 23 patents. He has been an Adjunct Member of Faculty at Columbia University since 1997 where he has taught courses on mobile networking and protocol design. His research interests include mobility management algorithms, signaling protocols and architectures for wireless and broadband networks, and protocol design.

Dr. La Porta is a Bell Labs Fellow, and received the Bell Labs Distinguished Technical Staff Award in 1996. He received an Eta Kappa Nu Outstanding Young Electrical Engineer Award in 1996. He is the Editor-in-Chief of the IEEE TRANSACTIONS ON MOBILE COMPUTING, and an Associate Editor for the *ACM/Baltzer Journal of Mobile Networking and Applications* and *KICS Journal of Communications and Networks*. He served as Editor-in-Chief of the *IEEE Personal Communications Magazine* for three years and is currently a Senior Advisor.