

Vertical Handoffs in Wireless Overlay Networks

Mark Stemm

stemm@CS.Berkeley.EDU

*Computer Science Division, Department of Electrical Engineering and Computer Science,
University of California at Berkeley, Berkeley, CA 94720-1776.*

Abstract

We present extensions to a traditional cellular [Ses95] handoff system to handle the simultaneous operation of multiple wireless network interfaces. This new system allows mobile users to roam in a “Wireless Overlay Network” structure consisting of room-size, building-size, and wide-area data networks. In this structure, the user can connect to the wired network through multiple wireless subnets, and offers the best possible connectivity given the user’s geographic location and local wireless connectivity. We present the basic handoff system and show that the handoff latency is bounded by the amount of time that the mobile host takes to discover that it has moved in or out of a new wireless overlay. To efficiently support applications that can not tolerate these disruptions, we present optimizations to this basic scheme that assume no knowledge about specific channel characteristics. For handoffs between room-size and building-size overlays, these optimizations lead to a handoff latency of approximately 170ms with a 1.5% overhead in terms of network resources. For handoffs between building-size and wide-area data networks, the handoff latency is approximately 600ms with a similarly low overhead. For these wide-area data networks, more specialized optimizations that attempt to take advantage of a user’s network traffic patterns have little advantage due to the low-bandwidth, high-latency nature of these networks.

1. Introduction

Wireless networking is becoming an increasingly important and popular way to provide global information access to users on the move. Current technologies vary widely in their bandwidths, latencies, frequencies, and media access methods. Despite this heterogeneity, most existing wireless network technologies can be divided into two categories: those that provide a low-bandwidth service over a wide geographic area and those that provide a high bandwidth service over a narrow geographic area. Unfortunately, neither technology in and of itself makes possible the best available network at all times. Wireless local area networks only provide limited coverage, and a mobile host equipped only with a wide-area data interface does not have the opportunity to take advantage of existing high-bandwidth infrastructure such as in-building RF networks or wired networks. No single network technology simultaneously provides a low-latency, high-bandwidth, wide-area connection to a large number of users simultaneously.

Our solution to the global information access problem is to use a combination of these wireless networks to provide the best possible coverage and bandwidth over a variety of geographic ranges. A mobile device with multiple wireless network interfaces has many ways of accessing the wired infrastructure through alternative wireless subnets. This allows it to overcome the problems of accessing information in the best possible manner for its current environment. This combination of wireless network interfaces, spanning in-room, in-building, campus, metropolitan, and regional cell sizes, fits into a hierarchy of network interfaces which we call a *wireless overlay network* structure.

We have implemented a *vertical handoff* scheme that allows a mobile user to roam between multiple wireless networks in a manner that is completely transparent to applications and disrupts connectivity as little as possible. For example, in the course of a single day, a typical user may move from her office, where her PDA is connected via an in-room Infrared network, to elsewhere in the building, where it is connected via a building-wide RF network. The same user may then move outside, where her connectivity is via a wide-area data network, and then inside another building which is connected via a different building-wide RF network.

Such a system provides seamless coverage; the typical handoff latency between networks is (a few) hundred milliseconds with minimal overheads in terms of bandwidth and power consumption.

The organization of the rest of this paper is as follows: in Section 2, we describe in more detail the concept of wireless overlay networks and the challenges to be met in our handoff scheme. In Section 3, we describe in detail our approach for vertical handoff. In Section 4, we present our definition of the metrics

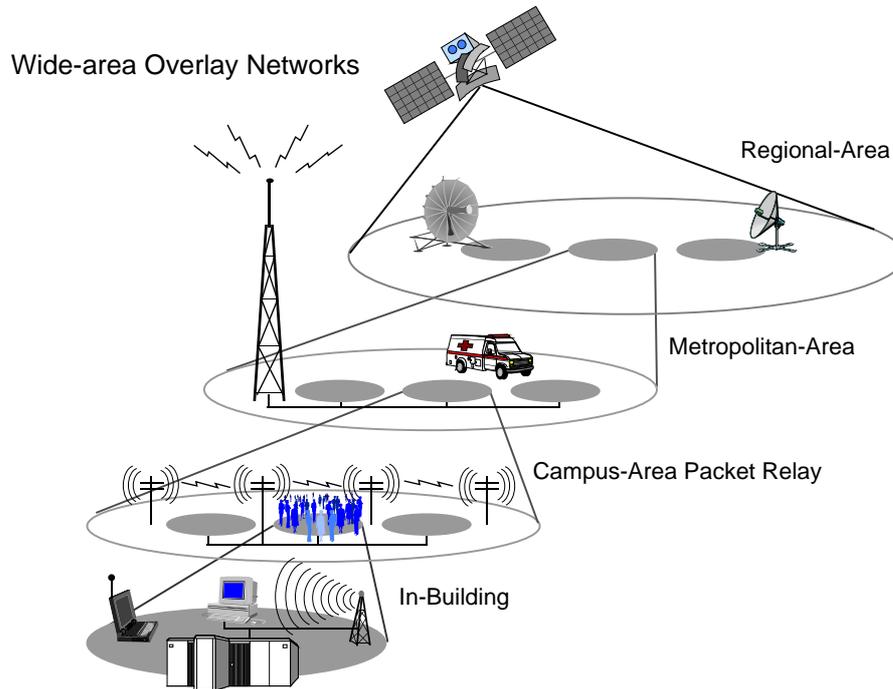


Figure 1. The Wireless Overlay Network Structure

used to quantitatively show the performance and cost of our system. In Section 5, we describe our wireless testbed and present initial performance results for the base handoff system and show that the dominant contributor to latency is the “rendezvous time”, the amount of time before the mobile discovers that it has moved into a new environment. In Section 6, we present several optimizations that can be used to decrease the handoff latency for those applications that are sensitive to disruption. In Section 7, we discuss related work dealing with low-latency handoff and the use of multiple network interfaces, and in Section 8, we conclude and describe some ongoing and future projects in our system.

2. Details about Vertical Handoffs

In this section, we describe in detail the wireless overlay network concept and some of the challenges to be met in our system.

2.1 The Wireless Overlay Network Structure

Figure 1 shows an example of a wireless overlay network. In our system, there are three levels of wireless overlays. The first level is a collection of disjoint room-size high bandwidth networks. This level provides the highest bandwidth per unit area: 1 megabit or more per room. The second level consists of building-size high bandwidth networks. This level provides approximately the same bandwidth as the room-size networks covering a larger area (for example, a single floor of a building). The final level is a wide-area data network. This provides a much lower bandwidth connection (tens of kilobits) over a much wider

geographic area.

2.2 “Horizontal” versus “Vertical” Handoffs

We define a *horizontal handoff* as a handoff between base stations that are using the same kind of wireless network interface. This is the most common definition of handoff. We also define a new type of handoff, a *vertical handoff*, as occurring between base stations that are using different wireless network interfaces. This naming convention follows from the overlay network structure, with networks with increasing cell sizes at higher levels in the hierarchy (Figure 2). We divide vertical handoffs into two categories: an *upward vertical handoff* is a handoff to a network with a larger cell size, and a *downward vertical handoff* is a handoff to a network with a smaller cell size. These two cases are not necessarily symmetric: the handoff from a lower overlay (one with higher bandwidth per unit volume) to a higher overlay (one with lower bandwidth per unit volume) usually does not appear the same to a Mobile Host (MH), in terms of connectivity, as a handoff from a higher to lower overlay.

There are some important differences between horizontal and vertical handoffs that affect our strategy for implementing vertical handoffs. These are:

- Many network interfaces have an inherent diversity that arises because they operate at different frequencies. For example, the room-size overlay may use infrared frequencies, the building-size overlay network may use radio frequencies, and the wide-area data system may use yet different radio frequencies. Another way in which diversity exists is in the spread spectrum techniques of different devices. Some devices may use Direct Sequence Spread Spectrum, (DSSS), while others may use Frequency Hopping Spread Spectrum (FHSS). Some of our the optimizations to reduce handoff latency will take advantage of this diversity.
- In a single-overlay network, a MH is ideally within range of a single base station at a time. The MH is usually within range of multiple base stations only during a handoff. In a multiple-overlay network, a mobile device can be within range of several base stations simultaneously for long periods of time.
- In a single-overlay network, the choice of “best” base station is usually obvious: the mobile chooses the base station with the largest signal strength, perhaps incorporating some amount of thresholding and hysteresis. In a multiple-overlay network, the choice of the “best” network cannot usually be determined by factors such as signal strength, because the networks have such




Figure 2. Horizontal vs. Vertical Handoffs

varying characteristics. For example, an in-building RF network with a low signal strength may still yield better performance than a wide-area data network with a high signal strength.

2.3 Primary Challenges

The primary technical challenges in the design of a seamless vertical handoff system are:

- *Low Latency Handoff: make the switch between networks as seamless as possible for disruption-intolerant applications and with as little data loss as possible.*

Our goal is to allow a typical user to use fully-interactive multimedia communication tools across all of these networks, even though the networks provide different levels of service. For example, the Infrared network may support full-motion video and high-quality audio, the RF network may support a lower frame-rate video and lower quality audio, and the wide-area network may support only audio.

- *Power Savings: minimize the power drain due to multiple simultaneously active network interfaces.*

The simplest approach to managing multiple wireless network interfaces is to keep all of them on all of the time. Measurements of commercially available wireless network interfaces [SGHK96] show that keeping an IBM Infrared and WaveLan RF interface on all of the time consumes approximately 1.5 watts. This power drain is approximately 20% of a total power drain of a typical laptop computer [FZ94]. At

these levels of power consumption, effective management of the network interfaces is crucial.

- *Bandwidth Overhead: minimize the amount of additional network traffic used to implement the overlay structure.*

Implementing vertical handoffs in wireless overlays requires bandwidth overheads in the form of beacon packets and handoff messages that are necessary to provide service to roaming users, and we want to minimize these costs while also providing minimal disruption for transitions between networks.

- *Discover the right time to perform handoffs in a wireless channel that is difficult to predict and characterize.*

Ideally, our system should keep a user connected to the lowest overlay network (where the bandwidth per unit area is largest) for as long as possible until it is absolutely necessary to move to a higher overlay. The challenge is that it is difficult to predict when the mobile will become disconnected from the current overlay based only on channel characteristics. In Appendix A, we show why it is difficult to predict higher-level network characteristics such as packet error rate from lower-level characteristics as signal strength and quality.

- *Work with commercially available devices over which we do not have direct control.*

We must depend on existing networking technologies to have a full range of wireless networks at our disposal. Although we assume that we can control certain room-size and building-size overlays, we also assume that the wide-area data overlay is owned and administered by a third party and that we cannot directly control the overlay's infrastructure. This is an important consideration because it limits modifications we can make to support vertical handoffs. For example, for those networks we cannot control, it is not possible to change the code running in the those network's base stations or gateways to support our specialized handoff schemes.

There are many inherent trade-offs in meeting these challenges: reducing power consumption by keeping network interfaces off when not in use increases handoff latency. Similarly, zero-latency handoff could be achieved by simply sending and receiving data across all Network Interfaces (NIs) at all times, but this an inordinate waste of bandwidth and power.

3. The Basic Handoff System

In this section we describe our wireless testbed and the basic system used to implement vertical handoffs.




Figure 3. Overview of the Handoff System

Handoffs are built on top of the mobile routing capabilities of Mobile IP. The infrastructure we use is similar to the one described in [Ses95] and the Mobile IP specification [Per95]. Mobile Hosts (MHs) connect to a wired infrastructure via Base Stations (BSs). A Home Agent (HA) performs the same functions as in Mobile IP. The difference is that the care-of address is not a unicast address but an encapsulating multicast address. The MH is responsible for initiating handoffs between base stations and between networks. A small group of Base Stations are selected by the MH to listen on this multicast address for packets encapsulated and sent by the HA. One of the BSs is selected by the MH to be a *forwarding* base station; it decapsulates and forwards the packets it receives on the multicast address to the mobile. The other BSs are *buffering* base stations; they hold a small number of packets from the HA in a circular buffer. When the MH initiates a handoff, it instructs the old base station to move from forwarding to buffering mode. The new base station then forwards the buffered packets that the mobile has not yet received to it. For networks where the base station infrastructure is not under our control, the Home Agent acts as the base station to the Mobile Host; the HA sends separate unicast packets to the care-of address of the MH's wide-area data interface.

The base stations send out periodic beacons similar to Mobile IP foreign agent advertisements. The MH

listens to these packets and decides which base station should be forwarding packets for the mobile, which base stations should be buffering packets in anticipation of a handoff, and which base stations should be a members of the multicast group assigned for a single mobile.

In the baseline horizontal handoff scheme [Ses95], the relative signal strength of these beacon packets was compared and the base station with the highest was chosen as the forwarding base station. For network interfaces that support channel quality measurements, we keep this feature. For devices that do not, we treat all base stations equally. This is not much of a detriment, however, due to the nature of overlay networks. From a room-sized network, the most logical transition is not to another room-sized network but to a building-size network. The wide-area data networks usually handle mobility transparently, so the MH is not required to make handoff decisions. In our system, the MH need only make horizontal handoff decisions for the building-size networks.

Upward vertical handoffs are initiated when several beacons on the currently connected network are not received. The MH then decides that the current network is not reachable and hands over to the next higher network, first turning on the upper network's device if necessary. *Downward* vertical handoffs are initiated when several beacons in a row are heard from a lower overlay's NI. The MH decides that the mobile is now within range of the lower overlay's NI and switches to the lower overlay.

3.1 Implementation Details

A more detailed breakdown of the base station and mobile handoff system components is shown in Figure 4. The handoff decisions are made by a user-level program running on the MH, the Handoff Controller (HC). A user-level program, `beacond`, sends out the beacon packets on each wireless subnet. The forwarding and buffering of multicast packets from the home agent is handled in the kernel networking code by a subroutine `BS_mip_ctrl()`. This subroutine consults a translation table for each packet and determines whether to forward or buffer the packet. This translation table is manipulated by a user-level program, `encapd`. `Encapd` manipulates the translation table in response to encapsulation requests from the Handoff Controller. On the MH, there is a similar kernel-level subroutine, `MH_mip_ctrl()`, that consults a structure that controls the filtering and counting of duplicate packets over multiple network interfaces. This filtering and counting process is used in the Packet and Header Doublecasting schemes of Section 6. The HC is responsible for manipulating this kernel level structure from user level. The function `MH_mip_ctrl()` can also notify the HC of important handoff-related events such as a disconnection from an overlay. These events will be described in more detail in Section 6. A detailed description of the




Figure 4. Detailed breakdown of mobile and base station components

interfaces can be found in Appendix B.

In our system, the MH has primary responsibility for initiating handoffs. However, the MH can take *advice* from an external source about the choice of network or base station to connect to. Possible external sources include:

- A *user-visible control panel* that allows the user to specify specific constraints about which networks to use.
- A *subnet manager* that uses judicious advice messages to avoid cell hotspots and increase the utilization of sparsely populated cells. For example, it may be advantageous to switch some users to a higher overlay network if the cell that they are currently using is congested or close to capacity.

The HC can also export state messages to the external sources to allow them to determine the state of the handoff controller.

The entire handoff system consists of approximately 4000 lines of user-level code and 1200 lines of kernel-level code. Of that, approximately 1200 lines of user-level code and 200 lines of kernel-level code are

devoted specifically to vertical handoffs.

4. Description of Metrics

In this section, we describe the variables and metrics that we use to quantify the performance and overhead of our handoff system.

4.1 Notation

We use the following symbols:

S_H = the size of a IP+Ethernet packet header

S_B = size of a beacon packet.

S_M = size of a mobile-initiated handoff message (in bytes).

S_D = size of a user's data packet (in bytes).

L_U = latency of the upper network interface (in ms)

L_L = latency of the lower network interface (in ms)

B_U = bandwidth of the upper network interface (in kilobits/sec).

B_L = bandwidth of the lower network interface (in kilobits/sec).

P_L = power consumption of the lower interface (in mW).

P_U = power consumption of the upper interface (in mW).

N_B = spacing between beacon packets (in ms).

N_D = spacing between user data packets (in ms).

T_B = threshold number of beacons packets (not) heard before initiating a handoff.

T_D = threshold number of data packets heard on the new interface before initiating a handoff to that new interface.

D = the length of power-saving duty cycle for NIs that are in sleep mode (in seconds).

4.2 Handoff Latency, Power, and Bandwidth Overhead

We define the handoff latency L as the amount of time from when the mobile is disconnected from the old base station to when the mobile receives the first packet from the new base station. We break down the

latency required to complete a vertical handoff into the following components:

$$L = L_D + L_P + L_N + L_F$$

- L_D is the component of latency where the mobile discovers that it must hand off to a new wireless overlay. This could be to an upper overlay as a result of moving out of range of the current overlay: for example, moving out of a room or moving out of a building to outside. This could also be to a lower overlay as a result of moving back into coverage of a lower overlay: for example, moving back into a room or building. In the basic system, this is largely a function of the beaconing frequency. If the beacon spacing is large, then L_D will be larger because the mobile will take longer to discover that it can no longer hear the current overlay. Also note that because of the overlapping nature of cells, in horizontal and downward vertical handoffs, this component of latency is not visible to the user as a disconnection. The mobile is still connected to the old base station during the time where it discovers that it can hear the new base station.
- L_P is the component of latency where the mobile must power on the upper or lower network interface. This includes any network registration time. This component of latency may or may not be visible to the user depending on whether or not the device was already on at the time the handoff occurred. Ideally, with the mechanisms described in Section 6.1 we can predict when the user is likely to hand off and can make this component of latency invisible to the user.
- L_N is the component of latency where the mobile must inform the new base station to start forwarding any data to the mobile. This is usually a function of the network latency.
- L_F is the component of latency where the base station sends the first data packet across the new network to the mobile. This component is a function of the network latency and bandwidth.

Some of these parts of the latency may be large and not be controlled. Many wide-area wireless networks such as Metricom [Met96] and CDPD [GW96] have a network registration process that must occur before a device can be connected. This can increase the value of L_P if the device must register with the network before it can be used. Wide-area wireless networks also have a much larger latency than local-area wireless networks, and this can affect the value of L_F .

We define the power overhead P as the amount of power that must be consumed to initiate a handoff. This is a function of the number and type of wireless interfaces that are on.

We define the bandwidth overhead B as the number of bytes sent per second by the base station that are

Type	User-visible Bandwidth	Cell Diameter	Latency	Registration Time (95% Conf Interval)	Power Consumption (mW)
Infrared (IBM Infrared)	850 kb/sec	7 meters	2-5ms	6.7 ms (5.7-7.8 ms)	349.6
In-Building RF (915 Mhz/2.4Ghz WaveLan)	1.6 Mb/sec	100 meters	2-5 ms	110.4 ms (93.8-127.0 ms)	1148.6 (915) 1318.8 (2.4)
Wide-Area Data (Metri-com)[Met96]	40 kb/sec	1 km	≈100 ms	7.6 sec (6.3-8.9 sec)	346.9

TABLE 1. Bandwidths, Latencies, and Registration Times for Our Networks

not actual data packets. These are usually beacon packets or other messaging that the mobile uses to initiate a handoff.

4.3 Power Usage

As previously mentioned in Section 2.3, power management of multiple wireless devices is important. Table 1 shows the steady state power consumption of the network interfaces when they are in an idle state. Our system handles power control in the following way: all network interfaces for overlays higher than the current network interface will be kept off by default. They will be turned on when geographic hints indicate that a handoff may be likely. This will allow use when possible to hide the value of L_P from the user. The NI immediately below the current NI will be put into a power saving low duty cycle sleep state where it will wake up every D seconds and listen for beacons on the lower interface for a short amount of time. This will increase the value of L_D for downward vertical handoffs, but the mobile will still be connected to the upper overlay for this period of time. and there will be no user-visible disruption.

5. Results for the Base System

5.1 Measurement Testbed

Our testbed consists of IBM ThinkPads and Intel-based PCs running a modified version of BSD/OS 2.0.1. Table 1 shows the specific wireless networks that we use along with typical bandwidths, latencies, and registration times. We use the IBM Infrared Wireless LAN network as our room-size network, the AT&T WaveLan as our building-size network, and the Metricom Ricochet System as our wide-area data network. The registration times were measured by sending a stream of UDP packets to a mobile host, turning on the network interface, and marking the time between the network interface was turned on and when the first data packet was received by the mobile.

Type	Ld (up/down)	Ln(up/down)	Lf(up/down)	P	B
Basic	$N_B(T_B-1)+N_B/2$	L_U+S_M/B_U	L_U+S_D/B_U	P_L	$N_B S_B$
	$D/2+ N_B(T_B)+N_B/2$	L_L+S_M/B_D	L_L+S_D/B_L		

TABLE 2. Predicted latency and cost for the basic system

Transition	Ld+Ln	95% Conf (Ld+Ln)	Lf	95% Conf (Lf)	Total	B
Infrared→ WaveLan	2.5 sec	1.85 sec-3.25 sec	8.35 ms	7.19ms-9.51ms	2.56 sec	64
WaveLan→ Infrared	3.5 sec	N/A	20.34 ms	4.634 ms-36.05 ms	3.520 sec	64
WaveLan→Metricom	2.79 sec	2.7 sec-2.99sec	125.96 ms	109.17ms-142.17ms	2.92 sec	64
Metricom→WaveLan	3.8 sec	N/A	134 ms	111.2ms-156.8ms	3.93 sec	64


TABLE 3. Breakdown of handoff latency for the basic system

5.2 Measurement Methodology

We measured the latency of handoffs by sending a continuous stream of UDP packets to the MH. For the Infrared to WaveLan transitions, this stream was limited to 500 kilobits/sec. For the Wavelan to Metricom transitions, the stream was limited to 50 kilobits/sec. The handoff was initiated by having the MH listen to a socket and turn the lower interface off and on in response to external messages. An observer machine was running a packet sniffer and the resulting packet trace was post-processed to determine when the external messages triggered the turn-on and turn-off of the interface, when the MH send the decapsulation requests to the BSs, and when the first packets arrived over the new interface to the MH.

5.3 Predicted Performance

Table 2 shows the analytical derivations for the values of L, P, and B as a function of the variables in Section 4. For upward handoffs in the basic system, the MH must wait for approximately T_B beacons to determine that the current overlay is no longer reachable. The $N_B/2$ term accounts for the fact that the overlay may be lost anywhere between two beacon times: on average this happens midway between two beacons. For downward handoffs, an additional $D/2$ seconds must be spent waiting for the lower interface to come out of its power saving state and hear the lower overlay's beacons. The mobile must then notify the upper base station to start forwarding new packets: this takes L_U+S_M/B_U seconds for the upward handoffs and L_D+S_M/B_D seconds for downward handoffs. Finally, the new base station must forward the first data packet to the mobile: this takes L_U+S_D/B_U seconds for upward handoffs and L_U+S_D/B_U seconds for downward handoffs. The steady state power consumption of this scheme is only P_L mW, because only a single interface needs to be on a time to trigger a handoff. The steady-state bandwidth overhead is from the beacon messages: this consumes approximately N_B*S_B bytes per second.



Page 1

Figure 5. Breakdown of Handoff Latency

5.4 Measured Performance

Table 3 shows the measured results for the basic system. L_p is not included; we assume that the interface is already turned on. We use a an N_B of 1 second, and a T_B of 3; when more than three beacon times pass without hearing any beacons, the MH considers the current overlay unreachable and switches to the next higher overlay. Similarly, when the MH hears three beacons in a row from a previously disconnected overlay, the MH considers the overlay reachable again and switches back to the old overlay. The choice of three is a heuristic; for heavily loaded networks, beacons may be delayed or lost. Three beacons was the lowest choice of a beacon threshold that still accounts for this effect in our system. In general, the predicted values of L_D , L_N , and L_F agree well with the measured values. From Figure 5, we see that the handoff latency is dominated by L_D , from 2.5 to 3.8 seconds. Even for the wide-area data network (Metricom), which has radically different network characteristics, the handoff time is dominated by L_D . The other parts of the handoff latency take approximately as long of that described in [Ses95]. The difference between this system and the system described there is that in their system, the mobile has perfect knowledge about the relative quality of the connection to the base stations: it is easy for the mobile to determine which of two homogeneous base stations is better. In our system, we assume no knowledge about the rel-

ative quality of base stations: determining which of two heterogeneous wireless networks is better is impossible or difficult.

6. Optimizations

For some applications, such as non-interactive file transfers and WWW access, the total handoff latency from the basic system is acceptable. One of our goals, however, is to support interactive multimedia communication across these network interfaces, and for these applications, a latency of approximately three seconds is unacceptable. With this in mind, we examined several optimizations to the base strategy that allow us to reduce the value of L_D during handoff.

6.1 When to use them

The schemes described in this section are used in situations where the application indicates that a low handoff latency is important, such as real time interactive voice or video. In these cases, the total handoff latency should be less than 300ms if possible. These optimizations are not be used continuously, however. They are used only when the mobile is in a situation where it may hand off soon. *Note that this is not the same as determining that a mobile must hand off immediately.* Although we show in Appendix A that handoff is difficult to predict given only channel characteristics, alternative hints can be used to predict that a handoff is likely. These include

- *User input:* the user can also instruct the MH to be more aggressive about handoff, at the cost of more power consumption, by using these optimizations. When the user is likely to leave the building, she can put the MH in a more aggressive mode. The user can take the MH out of this state when not moving.
- *Received signal strength:* although signal strength does not map directly to packet loss, it can indicate the distance from the base station. When the MH notices that the signal strength is gradually decreasing, it can assume that the user is moving away from the base station, and when the signal strength is increasing, the MH can assume that the user is moving toward the base station. When the strongest base station that the MH can hear has a low signal strength that has been decreasing, the MH can assume that a vertical handoff may be likely and use some of these optimizations.
- *Geographic hints:* we can use traces to predict which cells are the gateways to a new overlay networks. For example, although the overlapping nature of wireless overlays means that a user can be potentially connected to multiple networks at once, the transitions between networks are a func-

tion of the building geography. For example, a vertical handoff is only possible from certain places in the building, and only certain cells cover these locations. For example, only 1 in-building RF cell is likely to cover the exit of an office building. The base stations covering these cells could add information in their beacon packets indicating that this cell is near the exit to a building, and that a vertical handoff to a wide-area network may be likely.

- *Handoff Frequency*: The MH can also track the frequency of handoffs and use these optimizations when more handoffs are occurring, indicating that movement out of this overlay's coverage is more likely. This approach has been suggested for switching high-tier and low-tier PCS systems [SJ91].

All of these optimizations have some additional cost in terms of power or overhead bandwidth, and lead to some reduction of latency.

6.2 Optimizations

The following are optimizations that we can make to reduce the handoff latency:

- **Fast Beaconing**: the MH can selectively instruct a subset of the base stations that are listening to the multicast group for the MH to transmit beacon packets at a higher frequency than once per second. This allows the mobile to determine more quickly that it has not heard several beacons from the base station and should hand off.
- **Packet Doublecasting**: the MH can place a into forwarding mode a subset of the base stations that are listening to the multicast group for the MH. This means that a copy of the packet will be transmitted over each network for each base station. In our scheme, only two base stations are placed in forwarding mode simultaneously; the current base station and the strongest base station of the next higher overlay. Duplicate packets are filtered out at the network layer by `MH_mip_ctrl()` at the MH, which also keeps track of which packets have been received by the mobile over which interface. When more than N_D consecutive packets are received on the new interface with none on the old interface, `MH_mip_ctrl()` notifies the HC that the old overlay is unreachable and the HC initiates a handoff to the new interface. This approach simply extends the multicast tree of packets destined for the mobile to the network layer of the mobile itself. Note that this approach does at the network layer what the IS-95 CDMA Cellular phone standard and the ARDIS wide-area data system do at the physical layer. In IS-95, multiple base stations send duplicate copies of the same data using the same CDMA codes. The MH's receiver is already

equipped to handle multiple time-shifted copies of the same waveform, and a MH simply moves into the cell of the new base station seamlessly. In ARDIS, multiple base stations transmit the same data at the physical layer to achieve better in-building penetration.

- **Header Doublecasting:** This approach takes advantage of the fact that in the Packet Doublecasting approach, the duplicate packets on the upper interface are used as the indicator of handoff. If the forwarding time from the upper overlay's base station is low, then the full packets do not have to be sent until the actual handoff occurs. In this approach, the MH can place a likely next base station into a "buffer but forward packet headers" mode. `MH_mip_ctrl()` keeps track of which packets or packet headers has been received by the mobile, and notifies the HC when more than N_D headers have been received on the new base station with no packets on the old base station. This approach has an advantage over the Packet Doublecasting approach in that less data is sent on the upper overlay.

Both doublecasting approaches have an advantage over the beaconing systems in that they use extra resources only when the MH is actively sending data. When the user is not sending data, no extra bandwidth is used. Another advantage of the doublecasting approaches is that the data is the only overhead for handoff. If fast beaconing were used, then the beacons would be competing with the user data for network resources. The advantage of the beaconing systems is that if multiple users in a cell want low-latency handoff, and broadcast is supported, they can benefit from the aggregation of beacon packets across all users.

6.3 Performance

Table 4 shows the analytical definitions of L_D , L_N , L_F , P and B as a function of the variables described in Section 4. The formulas are very similar to those in Table 2. The fast beaconing system is identical to the basic system. In the header and packet doublecast systems, the MH must wait for T_D packets to arrive over the upper interface before the handoff is triggered. In the packet doublecast system, the notification and forwarding latencies are effectively zero: the mobile only has to change the packet filtering in the forwarding layer. The header doublecast scheme has the same notification and forwarding latencies as the beaconing system. The power consumption P of the doublecast schemes is more than the beaconing schemes, as both wireless interfaces must be on for the MH to make the handoff decision. The bandwidth overhead of the packet doublecast scheme is equal to the data rate at which the MH is receiving data ($N_D * S_D$). In the header doublecast scheme, the bandwidth proportional to the data rate at which the MH is

Type	Ld (up/down)	Ln(up/down)	Lf(up/down)	P	B
Fast Beacons	$N_B(T_B-1)+N_B/2$	L_U+S_M/B_U	L_U+S_D/B_U	P_L	$N_B S_B$
	$D/2+ N_B(T_B)+N_B/2$	L_L+S_M/B_L	L_L+S_D/B_L		
Packet Doublecasting	$N_D(T_D-1)+N_D/2$	0 0	0 0	P_L+P_U	$N_D S_D$
Header Doublecasting	$N_D(T_D-1)+N_D/2$	L_U+S_M/B_U L_L+S_M/B_L	L_U+S_D/B_U L_L+S_D/B_L	P_L+P_U	$N_D S_H$

TABLE 4. Analytical expressions for L and B for the various schemes

Transition	Ld+Ln	95% Conf Ld+La	Lf	95% conf Lf	Total L	B
Infrared→WaveLan	490 ms	256ms-723ms	7.02ms	3.6ms-10.4ms	501.1ms	310
WaveLan→Infrared	700ms	N/A	11.1ms	5.87ms-16.3ms	711.1ms	310
WaveLan→Metricom	511ms	457ms-607ms	Same as Basic	Same as Basic	636.96ms	310
Metricom→WaveLan	723 ms	N/A	Same as Basic	Same as Basic	857 ms	310

TABLE 5. Actual values of L and B for the Fast Beaconing System

Transition	Ld+Ln	95% Conf Ld+Ln	Lf	95% Conf Lf	Total L	B
Infrared→ WaveLan	202.4ms	131.3ms-243.46ms	0	0	202.4ms	65000
WaveLan→ Infrared	200 ms	N/S	0	0	200ms	65000
WaveLan→ Metricom	1599.7 ms	1470.4-1729.09ms	0	0	1599ms	6250
Metricom → WaveLan	2396.47ms	2186.13ms-2606.81ms	0	0	2396.47ms	6250

TABLE 6. Actual values of L, B, and P for the Packet Doublecasting System

receiving data, but only a small header of size S_H is sent on the upper overlay for every data packet of size S_D sent on the lower overlay.

6.3.1 Fast Beaconing

Table 5 shows the measured breakdown of handoff latency for the fast beaconing system. Beacons were sent out every 200 ms instead of every second in the original case. As in the basic system, the beacon threshold was set to 3 beacons. The measured values of L_D , L_F and B agree well with the analytical results. The latency has dropped by a factor of approximately 5 when compared to the basic system with a factor of 5 increase in bandwidth overhead.

6.3.2 Packet Doublecasting

Table 6 shows the measured breakdown of handoff latency for the packet doublecasting system. We used a packet threshold of 10 packets: 10 packets must be received by the mobile over 1 interface before the MH decides to switch to a new overlay. The choice of 10 is a heuristic: when packets are being sent over

Transition	Ld+Ln	95% Conf Ld+Ln	Lf	95% Conf Lf	Total	B
Infrared→ WaveLan	170.8ms	133.75ms-208.01ms	10.9ms	10.2ms-11.7ms	181ms	2075
WaveLan→ Infrared	170ms	N/A	11.7ms	9.1ms-24.3ms	181.7ms	2075
WaveLan → Metricom	Same as Packet	Same as Packet	Same as Basic	Same as Basic	1725.69ms	207.5
Metricom → WaveLan	Same as Packet	Same as Packet	Same as Basic	Same as Basic	2530.47ms	207.5

TABLE 7. Actual values of L and B for the Header Doublecast Scheme

multiple interfaces, ideally there would be a perfect interleaving of packets from the lower and upper interfaces. In practice, the interleaving is coarser grained: a burst of packets arrives over 1 interface, and then a burst over the other interface. The value of 10 was chosen to be larger than the largest burst of packets that we observed on heavily loaded networks. For the Infrared to WaveLan handoffs, this approach achieves a lower handoff latency than the basic system (approximately 200ms), but at a considerable cost. For the WaveLan network, this overhead of 65000 bytes/sec is approximately 32.5% of the network's maximum user-visible bandwidth of 1.6 Mb. For the WaveLan to Metricom handoffs, the latency is much larger than the approach that uses fast beaconing. The reason for this comes from the way in which vertical handoffs are initiated. In the fast beaconing system, the networks are considered independently: the presence or absence of beacons indicates whether or not to hand off. In the multicast approaches, this independence is not present because the networks are being compared relative to each other. Packets arrive over multiple network interfaces and must be considered together before a handoff decision can be made. For the Metricom network, the available bandwidth is sufficiently low that the amount of time it takes for the threshold number of packets to arrive is greater than the time it takes to independently consider the WaveLan.

6.3.3 Header Doublecasting

Table 7 shows the measured breakdown of handoff latency for the header doublecast scheme. As in the packet doublecasting scheme, we used a header threshold of 10 packet headers.

The table shows that for the WaveLan to Infrared handoffs, the header doublecasting scheme achieves a slightly lower latency as the full doublecasting scheme (171ms versus 200ms) with a dramatic decrease in bandwidth resources on the upper network. For the WaveLan network, this overhead is approximately 1% of the user-visible bandwidth of 1.6 Mb. For the WaveLan to Metricom handoff, the bandwidth overhead is dramatically decreased, but the value of L_D has not dropped equally. The reason for this is that the Metricom network is mainly latency bound: it can transmit approximately the same number of packets per

second regardless of their size. In addition, the value of L_F has now increased, as the packets must be forwarded from the Home Agent. This implies that the header and packet doublecasting approaches hold little advantage over the Beaconing approaches when used on low-bandwidth, high latency networks such as wide-area data networks.

7. Related Work

The idea of wireless data overlays has previously appeared in many places. “Wireless Overlay Networks” were first introduced in [KB96]. The concept of overlay networks was also introduced in the realm of high-tier and low-tier PCS systems [Cox95]. CDPD [GW96] is described as a data overlay on top of the cellular phone system.

There have been numerous papers dealing with handoff across homogeneous Cellular [SJ91], ATM [AN94], and Picocellular [GS94] networks. Seshan et al. [SBK96] [Ses95] [BSK95] implemented a system for low-latency horizontal handoffs. Our work expands upon theirs in that it handles multiple wireless networks and cases where the mobile device cannot use channel characteristics to trigger handoffs.

Recent work has addressed the problem of integration of multiple network interfaces in a single mobile. The MosquitoNet project at Stanford [BZCS96] has mobile devices equipped with Ethernet PCMCIA cards and Metricom modems. They trigger handoffs from one network to another based on the insertion and removal of Ethernet PCMCIA cards. Bhagwat [Bha95] also deals with the problem of multiple network interfaces, handling the routing aspects of multiple network interfaces as a special case of Mobile IP. Our work differs from theirs in that it focuses, quantitatively, on how to switch from one network interface to another in a manner that is completely transparent to the user.

8. Conclusions and Future Work

We have described additions to a horizontal handoff system to support the simultaneous operation of multiple wireless network interfaces. This vertical handoff system gives mobile devices the ability to roam freely in wireless overlay networks with seamless transitions between networks and with negligible interruption to applications. We argue that the addition of multiple network interfaces introduces inherent trade-offs between handoff latency, power management, and bandwidth overhead. We also argue that channel characteristics alone cannot be used to trigger vertical handoffs and any successful system must react to network conditions rather than predict them. Our schemes require no knowledge about specific channel characteristics and depend only on higher-order characteristics such as the presence or absence

of beacon and data packets. We present detailed measurements of handoff latencies and their costs in terms of network resources for a variety of different schemes. Results show that a simple scheme leads to a handoff latency that is seconds long and is dominated by the time it takes the mobile to discover that the current overlay is no longer reachable. Optimizations to this basic scheme can reduce this penalty to as low as 170ms, with only a 1.5% overhead on network resources. For transitions from room-size to wide-area data networks, the handoff latency from the basic system can be reduced to approximately 600ms with a small bandwidth overhead as a result of fast beaconing. Other optimizations either have a high cost in terms of bandwidth overhead or do not decrease handoff latency due largely to the latency-bound nature of the wide-area network being used.

Future directions for research are the following:

- Our working system does not yet use geographic hints to limit the use of the optimizations described in Section 6 to when a handoff is likely. We plan to add and analyze the effectiveness that simple hints such as cell connectivity have in predicting the likelihood of imminent handoffs, and add more sophisticated schemes if necessary.
- Although the MH can take advice about which network or base station to use, we have not examined the policies used to perform load balancing of mobile hosts across networks.
- The header and packet doublecasting optimizations we use depend on the fact that packets are being sent to base stations of different networks. Currently, these data flows are identical. For networks that have vastly different characteristics, this is not an ideal situation for a user who is receiving 500Kb of full-motion audio and video over an in-building RF network and is about to hand off to a wide-area data network. Similar to the approach of layered video dissemination in [MJV96], we are experimenting with the idea of *delivery classes* of traffic specified at the source and routing different subsets of delivery classes to different networks as a function of the network's characteristics.

9. Acknowledgments

Many thanks go to Randy Katz, my research advisor, who was a constant source of ideas and provided important direction for this work. Randy also provided many comments on initial versions of this master's report and greatly improved its presentation. Thanks go to Giao Nguyen, who developed the trace collection tools and collected some of the traces for the prediction results. His tools were invaluable in

allowing us to show that prediction is difficult. Thanks also go to Hari Balakrishnan, who helped debug some of the kernel optimizations made for faster handoffs. Srinivasan Seshan developed the horizontal handoff system and provided a stable and readable code base upon which to make the vertical handoff scheme work efficiently. Hari Balakrishnan and Elan Amir provided many helpful comments on early drafts of this paper that greatly increased the presentation of the material. This work is supported by DARPA contract DAAB07-95-C-D154 and grants from the California MICRO Program, Hughes Aircraft Corporation, Metricom, and AT&T.

References

- [All90] A. O. Allen. *Probability, Statistics and Queueing Theory with Computer Science Applications*. Academic Press, Inc., Boston, MA., 1990.
- [AN94] A. S. Acampora and M. Naghshineh. An Architecture and Methodology for Mobile-Executed Handoff in Cellular ATM. *IEEE Journal on Selected Areas in Communications*, 12(8):1365–1375, October 1994.
- [Bha95] P. Bhagwat. *A framework for integrating Mobile Hosts within the Internet*. PhD thesis, University of Maryland, December 1995.
- [BSK95] H. Balakrishnan, S. Seshan, and R.H. Katz. Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks. *ACM Wireless Networks*, 1(4), December 1995.
- [BZCS96] M. Baker, X. Zhao, S. Cheshire, and J. Stone. Supporting Mobility in MosquitoNet. In *Proc. of the 1996 USENIX Conference*, January 1996.
- [Cox95] D. C. Cox. Wireless personal communications: What is it? *IEEE Personal Communications*, pages 20–34, 1995.
- [FZ94] George H. Foreman and John Zahojaran. The challenges of mobile computing. Technical Report UW-CSE-93-11-03, University of Washington, March 1994.
- [GS94] R. Ghai and S. Singh. An Architecture and Communications Protocol for Picocellular Networks. *IEEE Personal Communications Magazine*, 1(3):36–46, 1994.
- [GW96] V. Garg and J. Wilkes. *Wireless and Personal Communications Systems*. Prentice-Hall, 1996. Chapter 8.
- [KB96] R. H. Katz and E. A. Brewer. The Case for Wireless Overlay Networks. In *Proc SPIE Multimedia and Networking Conference (MMNC'96)*, San Jose, CA, Jan 1996.
- [Met96] Metricom web page. <http://www.metricom.com>, 1996.
- [MJV96] Steven McCanne, Van Jacobson, and Martin Vetterli. Receiver-driven layered multicast. *Submitted to SIGCOMM '96*, 1996.
- [Per95] C. Perkins. IP Mobility Support Draft 12. IETF Mobile-IP Draft, 1995.

- [SBK96] S. Seshan, H. Balakrishnan, and R. H. Katz. Handoffs in Cellular Wireless Networks: The Daedalus Implementation and Experience. *Kluwer Journal on Wireless Personal Communications*, 1996. To appear.
- [Ses95] S. Seshan. *Low Latency Handoffs in Cellular Data Networks*. PhD thesis, University of California at Berkeley, December 1995.
- [SGHK96] M. Stemm, P. Gauthier, D. Harada, and R. H Katz. Reducing Power Consumption of Network Interfaces in Hand-Held Devices, May 1996. Submitted for publication.
- [SJ91] Tekinay S. and B. Jabbari. Handover and Channel Assignment in Mobile Cellular Networks. *IEEE Communications Magazine*, 29(11):42–46, 1991.

Appendix A: Limitations of Signal Measurements

In this appendix, we show why the mechanisms to trigger horizontal handoffs in existing systems cannot be used to trigger vertical handoffs.

To measure the possible effectiveness of signal measurements, we collected packet-level traces of transfers to the mobile host while in its environment and communicating via the WaveLan interface. These packet-level traces include device-level measurements such as signal strength, quality, and silence level for each received packet as well as packet sequence numbers that we used to calculate the instantaneous packet loss rate. To show more formally that signal measurements are a poor indication of handoff, we calculate the correlation between each of the signal measurements from the WaveLan Device (signal strength, quality and silence level) and the instantaneous packet error rate. If the correlation between one of the signal measurements and the packet error is high, then the signal measurement can be used to predict when the mobile will be disconnected from the current wireless network. If the correlation is low, then device-level measurements are of little use and the MH must depend on higher-level measurements.

We calculated the sample correlation coefficient r of data sets $X=(x_1, x_2, \dots, x_n)$ and $Y=(y_1, y_2, \dots, y_n)$, defined by

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 \sum_{i=1}^n (y_i - \bar{y})^2}}$$

Metric	Min r	Max r	# times $\rho=0$?	Average r	Average $\rho=0$?
Strength	.01355	.38532	21	-.000687	YES
Quality	.01085	-.38464	19	.003266	YES
Silence Level	-.004168	-.61889	17	.007246	YES

TABLE 8. Summary of Statistical Analysis of Signal Measurements

and used the following algorithm (from [All90]) to test at the α level of significance the null hypothesis $H_0: r=0$, against the alternative hypothesis, $H_1: \rho \neq 0$:

1. Calculate the test statistic:

$$t_0 = \frac{r\sqrt{n-2}}{\sqrt{1-r^2}}$$

2. Accept or reject H_0 : if

$$|t_0| > t_{(n-2, \alpha/2)}$$

(where t is the student's t distribution), reject the null hypothesis and accept H_1 .

Note that this algorithm shows a stronger result than we are actually interested in. We would like to know if ρ is close to 1 rather than close to 0. If the hypothesis is not true, it does not indicate that prediction will not work; however, if the hypothesis H_0 is true, we can be sure that prediction based on signal measurements will not work.

A total of 22 traces were analyzed. All 22 traces had some periods where the mobile was disconnected and experienced a high error rate. We used an α of .05. The results are summarized in Table 8.

The first three columns are for each of the 22 trials considered separately. **Min |r|** shows the value of r where the absolute value of r was the closest to 0, and **Max |r|** shows the value of r where the absolute value of r was closest to 1. The third column shows the number of times where the hypothesis H_0 was satisfied. The last columns show the results when all 22 traces were combined and considered as a single trial. **Average r** shows the value of the sample correlation coefficient over all data points in all samples, and the final column shows whether or not the hypothesis H_0 was satisfied.

These results show that for the network interface we measured, signal level measurements are not a good indicator of impending disconnection. This means that for WaveLan interface, which is the only network

we use which supports channel measurements, our vertical handoff system must react to a disconnection by higher level means, such as a lack of packet reception, and react to these events rather than attempt to predict when a disconnection is likely.

Appendix B: Detailed Description of Component Interfaces

In this section, we describe in detail the interfaces between the components of the handoff system described in Section 3.1.

Beacon Packet Format:

This describes the format of the periodic beacon packets sent out by each base station. The actual messages are in the same order as the C data structure.

```
typedef enum {BEACON, BEACON_REQ, BEACON_REPLY} BMSGTypes;
/* BEACON -- a beacon from a base station */
/* BEACON_REQ -- please send me a beacon */
/* BEACON_REPLY -- a reply to a BEACON_REQ message */
typedef struct BSBeacon {
    struct in_addr wlAddr;    /* Wireless address of this base station */
    struct in_addr wiAddr;    /* Wired address of this base station */
} BSBeacon;

typedef struct BeaconReq {
    struct in_addr addr;      /* Please send a beacon to this address */
} BeaconReq;

typedef struct BeaconMsg {
    BMSGTypes type;          /* Is this a beacon, beacon request, or beacon reply */
    struct timeval curTime;   /* Time the beacon was sent */
    union {
        BSBeacon beacon;
        BeaconReq req;
    } data;
} BeaconMsg;
```

Encapsulation Request Format:

This is the format of the messages sent between the handoff controller and encapsd. The actual messages are in the same order as the C data structure.

```

typedef enum {UNENCAP_ADD, UNENCAP_ENA, UNENCAP_DIS, UNENCAP_DEL} BSRTypes;
/* These are the possible messages that the handoff controller can send */
/* UNENCAP_ADD -- Add an entry for this mobile */
/* UNENCAP_ENA -- Start forwarding for this mobile */
/* UNENCAP_DIS -- Stop forwarding for this mobile */
/* UNENCAP_DEL -- Delete an entry for this mobile */
typedef struct BSReq {
    BSRTypes type;
    struct in_addr bsAddr;      /* The base station that this message should corre-
respond to. (Allows bs->bs forwarding) */
    struct in_addr locAddr;    /* The local address of the mobile */
    struct in_addr realAddr;   /* The home address of the mobile */
    struct in_addr multiAddr; /* The multicast address used to send packets from
the home agent */
    struct mobileip_ids ids;   /* The last few packets seen by the mobile */
    int hdrOnly;              /* Whether to forward headers or full packets */
} BSReq;

```

Advice Message Format:

This is the format of the advice messages sent between external sources and the handoff controller.

```

/* These indicate what the message is: */
#define USE_THIS_NETWORK (0)      /* Use this network unless it is not avail-
able */
#define USE_ANY_NETWORK (1)      /* Cancels a USE_THIS_NETWORK or
DONT_USE_THIS_NETWORK message */
#define DONT_USE_THIS_NETWORK (2) /* Don't use this network unless it is the
only one available */
#define USE_THIS_BASESTATION (3) /* Use this base station unless it is not
available */
#define USE_ANY_BASESTATION (4)  /* Cancels a USE_THIS_BASESTATION or
DONT_USE_THIS_BASESTATION message */
#define DONT_USE_THIS_BASESTATION (5) /* Don't use this base station unless it is
the only one available */

/* If the top bit is set, it's a reply */
#define REPLY_MASK (0x80)

/* The message format is as follows: */
/* ----- */

```

```

/* | opcode | body |                               */
/* -----                               */
/*  1 byte  10 bytes(max)                       */

/* If a body field is an interface, it is a null-terminated string */
/* that is the interface name */
/* If a body field is an in_addr, it follows the same byte ordering */
/* conventions as in IP */

#define MAX_IF_SIZE (10)

typedef struct _AdviceBodyIf {
    char interface[MAX_IF_SIZE];
} AdviceBodyIf;

typedef struct _AdviceBodyBS {
    struct in_addr addr;
} AdviceBodyBS;

/* Here is a structure that holds an advice message */
typedef struct AdviceMsg {
    struct in_addr client;           /* Who this message is from */
    u_int opcode;                   /* What type of advice request */
    union {
        AdviceBodyIf interface;
        AdviceBodyBS bs;
    } body;
} AdviceMsg;

```

Handoff Controller State Format:

This is the format of the state messages sent between the Handoff Controller and External Sources.

```

/* These indicate what the message is: */

#define LIST_ALL_NETWORKS 0           /* Returns a list of all networks */
#define LIST_ACTIVE_NETWORK 1        /* Returns the currently connected network */
#define LIST_ALL_BASESTATIONS 2      /* Returns a list of all base stations */
#define LIST_ACTIVE_BASESTATION 3    /* Returns the currently connected base sta-

```

```

tion */
#define LIST_USED_NETWORKS 4          /* If a USE_THIS_NETWORK message has been
received for a network, return that network */
#define LIST_USED_BASESTATIONS 5      /* If a USE_THIS_BASESTATION message has been
received for a base station, return that base station */
#define LIST_DONTUSED_NETWORKS 6      /* If a DONT_USE_THIS_NETWORK message has
been received for a network, return that network */
#define LIST_DONTUSED_BASESTATIONS 7  /* If a DONT_USE_THIS_BASESTATION message
has been received for a base station, return that base station */

/* If the top bit is set, it's a reply */
#define REPLY_MASK (0x80)

/* So the message format is as follows: */
/* ----- */
/* | client | opcode | num_body | body1 | body2 | ... | bodyN | */
/* ----- */
/* 4 bytes  1 byte  1 byte  ? bytes ? bytes      ? bytes */

/* For "networks..." messages, a network is the Unix device name(efp, */
/*  irp, wlp, ...) as a null-terminated string. */
/* For "basestations..."messages, a body is an in_addr with bytes in network order */

typedef char String[100];

/* Here is a structure that can hold a message */
/* Only one of (body_ifs, body_addrs) is active at one time */
typedef struct _HandoffctlrMsg {
    struct in_addr client;
    u_char opcode;
    u_char numBody;
    String bodyIfs[100];
    struct in_addr bodyAddrs[100];
} HandoffctlrMsg;

```

Base Station Translation Table Modification Format:

Mobile Host Packet Filtering Request Format:

This is the format of the messages sent between the Handoff Controller and the kernel networking code on the MH side, and

encapd and the kernel networking code on the BS side. The kernel-level translation table is modified from user level via socket options. The same request structure is used for base station and mobile requests:

```
struct ip_mobreq {
    int type; /* Type of request */
    struct in_addr imr_ifaddr; /* Address of outgoing interface */
    struct in_multi imr_multi_addr; /* Multicast address of HA encaped packets */
    u_char imr_multicast_ttl; /* TTL for outgoing multicasts */
    struct in_addr imr_ins_addr; /* Address to insert for local packets */
    struct in_addr imr_ins_addr_new; /* New address to insert for local packets */
    struct mobileip_ids ids; /* IP ids of recently received packets */
    int handoffCtrlrPid; /* Proc.id of HC (used for notifications) */
    char blockedIf[10]; /* Interface for which to block packets */
    int burstThreshold; /* Number of packets over 1 interface before
notifying HC */
    int hdrOnly; /* Whether to forward full packets or only
headers */
};
```