

Towards Integrating Wireless LANs with Wireless WANs using Mobile-IP

Sumi Helal and Choonhwa Lee

Computer and Information Science and Engineering Department
University of Florida, Gainesville, FL 32611, {[helal,chl](mailto:helal,chl@cise.ufl.edu)}@cise.ufl.edu

Yongguang Zhang

HRL Laboratories, LLC
3011 Malibu Canyon Rd, Malibu, CA 90265
ygz@hrl.com

Abstract

To allow a seamless integration between wireless LANs and Wireless WANs, we developed a full stack adaptation model and a simple subnet architecture that superimposes Mobile-IP on cellular-type wireless LANs. The idea is to use Mobile IP as an integrative layer atop different LAN/WAN networks. While Mobile-IP is widely used in wireless WANs, it is not known how well it performs under a wireless LAN environment, against native MAC-level handoff. Through experimentation using 802.11 W-LAN, we found that under practical values of handoff frequencies, the performance of Mobile IP based W-LAN handoff is almost identical to the performance of W-LAN handoff. Further performance studies show the suitability of Mobile-IP as an integrative layer in this architecture.

I. INTRODUCTION

Recent advances in portable computers and wireless LAN/WAN technologies have engendered two new paradigms of computing known as *nomadic* and *mobile* computing. Untethered users with wireless-capable portable computers are either *nomadic* (e.g., within an office, a building, or a pedestrian outdoor area), or *mobile* (e.g., in a taxicab, a train, or even an airplane). Nomadic computing utilizes wireless LAN networking technology such as the IEEE 802.11, which is a low mobility (typically, up to 5 M/h), high bit rate (ranging from 2 to 25 Mbps) network. Mobile computing, on the other hand, utilizes wireless switched or packet data networks such as GSM, CDPD, and iDEN, which are high mobility (up to 60 M/h), low bit rate (ranging from 9.6Kbps to 40Kbps) networks.

Despite recent advances in the achieved bandwidth of all types of wireless networking technology, indoor networks continue to provide much higher bandwidth than outdoor networks. This bandwidth gap is expected to either persist or widen in the near future. What is needed for truly ubiquitous connectivity is an adaptability infrastructure that will allow users to roam freely while transparently switching between wireless LAN/WAN

networks without disrupting the applications. Bridging wireless WAN and wireless LAN will fuse in the economical advantages and the high bit rate of wireless LAN with the ubiquity and ad-hoc mobility allowed by wireless WAN.

Integrating wireless LAN/WAN is in the mainstream of a desired and expected evolution of future wireless networks. As the first tele-services evolved from voice services delivered via wire-line terminals (handsets), to cordless phones, to cellular phones, wireless data tele-services will similarly evolve with network deployment. We have already witnessed voice services evolving to build into basic data services such as CDPD (Cellular Digital Packet Data). When basic data services become ubiquitous, they will ultimately be required to move toward higher rate data services, and then evolve to bandwidth on demand. Another testimony on network evolution is the cordless phone, which evolved into the Japanese PHS system offering higher data bit rate and greater mobility. Guided by this history of network evolution, we highly anticipate the evolution and convergence of wireless LANs and wireless WANs into an integrated networking service.

To realize this vision, we have developed a Full Stack Adaptation (FSA) concept. It is based on an architecture that integrates *horizontal* (base stations utilize the same wireless networking technology) and *vertical* (base stations use different technologies) handoff, while allowing applications to participate fully in the handoff process. In this article, we focus on integrating Mobile IP into the FSA architecture. We first introduce the FSA architecture and then present ideas and experimental validation on how to superimpose Mobile IP on Wireless LANs, as a step towards a full integration of wireless LAN/WAN. We implemented and experimentally evaluated a Subnet Architecture (SA) on top of wireless LAN coverage cells, using the Mobile IP protocol. We describe the SA architecture and present experimental results demonstrating its feasibility in an actual implementation using IEEE 802.11 Wireless LANs and Mobile IP.

II. THE FULL STACK ADAPTATION PROJECT

Since no single wireless network technology meets the ideal of high bandwidth, low latency, universal availability, and low cost, it is inevitable that, for the foreseeable future, multiple network interfaces will be required for ubiquitously connected mobile hosts. The Full Stack Adaptation Project (FSA)¹ is based on an architecture that integrates *horizontal* (base stations utilize the same wireless networking technology) and *vertical* (base stations use different technologies) handoff, while allowing applications to participate fully in the handoff process.

The FSA architecture allows a two-way interplay between applications and network connectivity on vertical handoffs. On the one hand, an application may provide modes that allow a user to specify indirectly the relative importance of high bandwidth, low latency, or low cost. An example is a teleconferencing application that has “inattentive” and “attentive” modes. A user paying casual attention to a conference (“inattentive mode”) may not require the best possible transmission quality, or might have audio

¹ Joint project between University of Florida and University of New Orleans.

squelched. But the user may suddenly begin to pay strict attention to the broadcast (“attentive mode”), which might require more bandwidth. The FSA architecture allows applications to provide recommendations to the handoff mechanism that can result in substantial savings. On the other hand, some vertical handoffs are required—to maintain connectivity, an alternate interface must be chosen if a mobile host moves out of the range of service for a particular network interface, or if the signal strength of one network decreases. Allowing applications to register interest in imminent or completed vertical handoffs allows the development of novel reactions to changes in quality of service. For example, an application requiring high bandwidth might alert the mobile user that additional movement may result in reduced service (a handoff is imminent, and the only other available interfaces provide inferior service). Highly adaptive applications might choose to hoard data if high bandwidth becomes available. They might choose to take a checkpoint to ward against loss of data in the face of decreasing battery strength and/or the likelihood of disconnection. Or the application might choose to rely more heavily on cached data because of reduced bandwidth caused by a vertical handoff.

A detailed schematic of the FSA architecture is depicted in Figure 1. Three network interfaces are being supported including Fast Ethernet, 802.11 wireless LAN, and *iDEN* wireless Packet data over PPP/SLIP. In order to integrate vertical and horizontal handoff, and to provide a fixed IP solution under multiple networks, the Mobile-IP protocol is fitted into the FSA architecture. This allows handoff to occur across cells of the same network type, or across cells belonging to different networks, with the goal of disrupting applications to a minimum degree. The Mobile-IP layer is largely unaware of vertical handoffs, other than registration of a new “care of” address for the mobile host. This is important in order to maintain compatibility with a variety of network interfaces that support Mobile-IP.




Figure 1. Adaptation Stack Architecture

The FSA architecture propagates information in both directions in the stack. Applications make their quality of service requirements known, and these requirements are used to guide vertical handoff decisions. In turn, information about the effects of mandated vertical handoffs is trickled upward to the application adaptation layer. The vertical LAN/WAN handoff layer monitors network characteristics, available power, and application requirements in order to intelligently perform vertical handoffs between available interfaces. Our application adaptation focus is on the use of *thin clients* as an alternative to specialized proxies. The details of the vertical handoff layer and the application adaptation layers under the FSA architecture are outside the scope of this paper and are not discussed.

III. INTEGRATING MOBILE-IP WITH CELLULAR-TYPE WIRELESS LAN

Our goal is to integrate Mobile IP into the FSA architecture with support to three network types: fast Ethernet, iDEN, and 802.11 Wireless LANs. The iDEN network is based on Mobile IP and therefore imposes no problems to solve. Ethernet is straightforward as

well, leaving 802.11 Wireless LAN to be the network interface needy of integration with Mobile IP.

III-A. Background on Mobile IP

The IETF Mobile IP Working Group has been working to specify a mechanism in IPv4, called Mobile IP, which is to accommodate node mobility within the Internet [1]. Mobile IP enables a mobile computer to move to a new network without changing its IP address or disrupting communication connections. The IETF Mobile IP architecture introduces new functional entities called *home agent* and *foreign agent* which cooperate to allow a *mobile node* to move without changing its IP address. Each mobile node is associated with a home network on which a designated host acts as its home agent. When mobile node is away from its home network, home agent is responsible for intercepting and forwarding packets destined to it.

Whenever the mobile node is away from its home, it registers its location (*care-of address*) with its home agent. Mobile IP can use two different types of care-of address: a *foreign agent care-of address* (an address of a foreign agent with which the mobile node is registered), and a *co-located care-of address* (an externally obtained local address which the mobile node has associated with one of its own network interfaces). The co-located care-of address can be dynamically acquired by the mobile node through Dynamic Host Configuration Protocol (DHCP) [8] on the local network. Depending on the type of its care-of address, the mobile node will register either directly with its home agent, or through a foreign agent, which forwards the registration to the home agent.

After a successful registration, the home agent intercepts packets arriving for the mobile node on its home network. Then, the home agent encapsulates the packet and sends it to the mobile node using the care-of address. In the case of foreign agent care-of address, when a foreign agent receives the encapsulated packet, it decapsulates and delivers the packet to the visiting mobile node. If a co-located care-of-address is used, the home agent sends the encapsulated packet to the mobile node directly, and the mobile node does the decapsulation itself.

I-A.1 Triangle Routing

The basic IETF Mobile IP specification has a routing anomaly, known as triangle routing, which is inherent due to its tunneling mechanism. In triangular routing, all packets sent to a mobile node must be sent the mobile node's home network and then forwarded to the node's current location. This two-step routing between correspondent host and mobile node causes increased network load and delay. As an effort to avoid this inefficient routing problem, there has been a research on route optimization [7] and extensions have been made to the basic IETF Mobile IP protocol to remedy the non-optimal routes.

I-A.2 DHCP

On wireless LANs, mobility support can be improved to a certain extent through the use of dynamic IP address allocation provided by DHCP. When a mobile node visits a

foreign network running DHCP server, it requests the use of an address for some period of time. The server returns an available IP address out of a pool of addresses, which is not already allocated on the local network. If the IP address is configured successfully, the mobile node is able to communicate by using network resources directly. The DHCP-based wireless LAN technology enables a new mobile node to be connected to some foreign network, while roaming around without any wire-line connection. DHCP-based wireless LANs intended for public access, commercial services already began deployment [12].

Real mobility support, however, can't be achieved through DHCP-based wireless LANs in that the network connection can be maintained only within wireless LAN boundary. A mobile node can't move to another LAN with its active network connection maintained, because the allocated IP address is valid only within the local network. To the contrary, in Mobile IP based wireless LAN (which is the approach taken in this project), a mobile node can move to other networks while maintaining network connections. It is Mobile IP that makes this mobility transparent to TCP and upper layers.

III-B. Mobile-IP based Wireless LAN

The simplest solution is to impose a one-to-one mapping between wireless cells and subnetworks. We refer to this solution as the Subnet Architecture or SA. In this architecture (depicted in Figure 2-a), each cell is assigned a unique subnet address. When a mobile node crosses from one cell to the next, it moves into a different subnet, and the Mobile IP network-level handoff is initiated, immediately following completion of the MAC-level handoff. The handoff results in a handshake between the new foreign agent and the home agent of the mobile node. On completion of the handoff, messages destined to the mobile node are forwarded by the home agent to the subnet where the mobile node is visiting, and, eventually, to the mobile node at the foreign network. The implementation of this architecture requires the detection of the MAC-level handoff occurrence on the mobile node, and subsequently initiating a Mobile IP handoff between the two networks.




Figure 2. Mobile IP based Wireless LAN Subnet Architectures

A generalization of the SA architecture is a many-to-one mapping of cells to subnets (Figure 2-b), in which only MAC-level handoff is initiated by the mobile node in response to any *intra-subnet* (yet inter-cell) mobility. Access Points (APs) belonging to different cells of the same subnet form a logical network with the same logical network name (e.g., Service Set Identity (SSID) in IEEE 802.11 LAN terms). When a mobile node moves across subnets, both MAC-level and Mobile IP level handoff protocols are initiated by the mobile node. This generalized architecture allows for higher flexibility in the design of the subnet based on administrative or work affinity instead of spatial affinity. For example, a team of researchers may occupy non-contiguous offices or may be separated by a concrete wall. In these circumstances, the many-to-one SA architecture allows one subnet to be assigned to the same team, whereas a one-to-one mapping architecture results in multiple subnets to be assigned to the research team.

One important advantage of the subnet architecture is increased security. Without Mobile IP, the entire roaming domain of the W-LAN is a broadcast domain where any mobile node can receive other nodes' packets (encrypted or not). In the subnet architecture, packets destined to mobile nodes of one subnet are not routed to any other mobile nodes outside the destination subnet. In a sense, this enables true multicast in the roaming domain, leading to a naturally more secure wireless internetworking.

The testbed implementation and the experiments reported in this paper are based on the one-to-one mapping SA architecture. Our work can be easily extended to the many-to-one mapping subnet architecture.

IV. EXPERIMENTAL ANALYSIS

We have developed a testbed consisting of a home network and two foreign networks. To experiment with the one-to-one subnet configuration, we used two wireless cells, each mapped to a different IP subnet. While the mobile node continuously moves from one cell to the other, we monitored TCP/IP performance in order to quantify the effects and overheads caused by the composite MAC/Mobile IP handoff and to compare it with MAC-only handoff performance. Furthermore, we analyzed the extent of disruption the composite handoff can cause to active connections.

IV-A. The Testbed

As shown at Figure 3, our testbed consists of a home network-segment and two foreign network-segments that are plugged into a PC router ports. While the home network is a wired LAN, the foreign networks are IEEE 802.11 wireless LANs. As the HA, a SparcStation 10 workstation is connected to the home network via wired Ethernet card. The mobile node is an IBM ThinkPad 390 laptop equipped with BayStack 660 Wireless LAN PC card. Both the HA and the mobile node run Linux RedHat 5.2. We used SUNY's Mobile-IP implementation [3], which is one of a few Mobile-IP implementations based on Linux. At each foreign network, an Access Point (Nortel Networks BayStack 660) covers a cell and communicates with the mobile node in its range. The BayStack W-LAN is a 2 Mbps Direct Sequence Spread Spectrum (DSSS) radio technology that is IEEE 802.11 compliant. It provides a wireless coverage range of up to 300 feet in a standard office environment and up to 2,000 feet in an open environment.

Since two cells are overlapped, the mobile node can hear from both of APs. Also, notice that a subnet range covered by Mobile IP coincide with a range by the corresponding wireless LAN cell, since our testbed provides a one-to-one mapping of wireless LAN cells to Mobile IP cells. As shown in Figure 3, foreign agents are not supported in our testbed, which means that the mobile node decapsulates the tunneled packets by itself.




Figure 3. One-to-one mapping subnet architecture testbed

IV-B. Handoff Coercion and Propagation

In the one-to-one mapping Subnet Architecture, the handoff at the MAC layer means that the Mobile IP handoff should be performed immediately. For our experiment to measure the impact by both handoffs, we implement a user-level process, called *handoff controller*, in the mobile node. The handoff controller performs two major functions: *handoff coercion* by which we can trigger handoff each time handoff needs to be initiated at both layers during our experiment, and *handoff propagation* by which the handoff at MAC layer is propagated to Mobile IP layer.

We transferred a large file from the home agent, which is also acting as a correspondent node, to the mobile node, and vice versa. While transferring, the mobile node is switched between two APs, causing wireless LAN handoffs followed by Mobile IP handoffs. The handoff controller performs this coercion periodically, and at a given frequency. The handoff coercion at MAC layer is implemented via `ioctl(WLAN_BSSJOIN, ...)` system call supported by the Linux wireless LAN driver [11], whereas the Mobile IP handoff coercion is implemented by modifying Mobile IP implementation code [3]. This S/W controlled handoff coercion allows for precise control over the instant when handoffs start. In addition, it enables us to conduct our experiments without mobile node's physical movement.

As an 802.11 MAC handoff occurs, the mobile node, i.e., 802.11 LAN station, associates itself with a desired AP by exchanging Authentication request/response frame and Association request/response frame. In the course of the experiments, we observed

that under heavy traffic load, these handoff frames could get discarded when the transmission buffer of the wireless PC card becomes full. We remedied this by retrying the `ioctl(WLAN_BSSJOIN)` call.

In a normal setting, handoffs are performed at the MAC layer as well as at Mobile IP layer, based on the received beacon signals and the mobility behavior of the mobile node. When handoff is determined to be needed at the MAC layer, the latter attempts to handoff by itself (without intervention from our handoff controller). If the handoff succeeds, our handoff controller gets notified, which in turn instructs the Mobile IP layer to perform its handoff. In this way, we propagate MAC layer handoff to Mobile IP layer, which is always the case with one-to-one mapping. Our handoff propagation mechanism is particularly suitable to the foreign network without foreign agent support in that it enables Mobile IP to start the handoff process immediately if needed, without any polling or additional delay at the IP layer. In our experiments –a laboratory setting, the MAC layer handoff is initiated by an explicit request from the handoff controller, because the two wireless LAN cells overlap significantly (both are located in the same laboratory room). Then, the resultant handoff notification and, in turn, Mobile IP handoff initiation are done the same way as in a normal setting, as described above.

V. EXPERIMENTAL RESULTS

For all experiments, we used *ftp* to transfer a large file from the correspondent node to the mobile node, while handoff is being coerced at a certain frequency. During the transfer, we tracked down TCP sequence number by using the “*tcpdump*” program. We measured *ftp* throughput and repeated our measurements for increased reliability of the results. The main two objectives of our experiments are:

- to quantify the overhead of the composite MAC/Mobile-IP handoffs and to compare it with MAC-only handoff, and
- to determine the range of acceptable handoff frequencies, within which degradation of TCP performance is acceptable.

V-A. Handoff Overhead

- IEEE 802.11 Handoff only Configuration.

In this experiment, two wireless cells form one IP network together and, therefore, two APs are assigned IP addresses belonging to the same network. This means only IEEE 802.11 handoff is needed. Mobile node movement is emulated by the handoff coercion mechanism described above. IEEE 802.11 Wireless LAN handoff is initiated by the handoff controller, whenever mobile node switches into the other AP. A large file is transferred via *ftp* between the correspondent node and mobile node to monitor the transfer performance at TCP layer. The purpose of this experiment is to measure the basic overhead associated with wireless LAN handoff.

We would like to mention that the effect of mobility on transport protocol behavior in wireless networks was first examined in [4], where TCP performance was

shown to degrade significantly due to the inappropriate reaction to the delays and packet loss incurred by host movement.




Figure 4. The Impact on TCP Performance due to Wireless LAN Handoff

Figure 4 shows the growth of TCP sequence number during a certain span of an ftp transfer, in the case of IEEE 802.11 handoff only. The sequence number ceases to increase for about .5 second near 7 second, 12 seconds, 17 seconds, and so on. These pauses are caused by the handoff every 5 seconds and amount to about 10% of total time. During these pauses, TCP transmits no new data and retransmits the unacknowledged packets.

- Composite IEEE 802.11 and Mobile IP Handoff

In this experiment, each of the two cells has its own IP network, as shown in Figure 3. First, wireless LAN handoff is coerced and on its completion, Mobile IP handoff is initiated. Notice that these periodical handoff coercion and propagation are handled by the handoff controller. In this setup, we ftp a large file from the correspondent node to the mobile node.




Figure 5. The Impact on TCP Performance of Mobile IP Handoff

Figure 5 shows how TCP reacts to the composite handoffs. There are long pauses in the growth of TCP sequence number around 387.5 second, 392.5 seconds, 397.5 seconds, and so on. These pauses are repeated every 5 seconds, which reflects exactly the fact that our handoff interval is 5 second for this experiment. The pause period, much longer than in the case of IEEE 802.11 handoff only, lasts up to about 40 % of total time. This is because both handoffs take much longer time to go through the following steps:

- Firstly, wireless LAN handoff is performed.
- Then, the mobile node gets a new IP address and reconfigures its network interface with this new address.
- Then, Mobile IP handoff is performed.

The longer the time spent in composite handoff, the more packets get lost and, in turn, more retransmissions are attempted at the TCP layer. In addition, by re-configuring itself, the network interface exacerbates IP packet loss and delays. It should be noted that these 3 steps cost more than in the case of wireless LAN handoff only, although the registration process overhead of Mobile IP is relatively small: *Register Request* and *Register Reply* messages are small UDP packets and the round trip time is hundreds of milliseconds in our testbed.

V-B. TCP Performance Evaluation

To experimentally evaluate the effect of composite handoff on TCP performance, we measured the time taken to ftp a large file (12,498,737 bytes, to be exact) from the correspondent node to the mobile node, while handoff is coerced at a certain frequency. This experiment is repeated in a series of different handoff frequencies and ftp throughput is measured for each experiment. We repeat these experiments in both cases (IEEE 802.11 handoff only and composite IEEE 802.11 handoff and Mobile IP handoff). We evaluate the impact of handoff by comparing the measured ftp throughput of the two cases with the baseline throughput involving no handoffs.




Figure 6. Relationship between Transfer Times and Handoff Frequencies

Figure 6 shows the measured ftp throughputs at several handoff frequencies. As expected, transfer time in the case involving both handoffs is generally longer than in the case of wireless LAN handoff only. In particular, the gap between the two cases is magnified with frequent handoff, as is at the 5 seconds of handoff frequency. But the transfer time dramatically decreases, as the handoff frequency increase. Figure 6 shows, however, that there is no significant difference at handoff frequencies over 30 seconds. This means that at such handoff frequencies, the composite handoff performance (in Mobile IP based W-LAN) is almost equivalent in performance to handoff in wireless LAN environments.

This result is of practical importance. To demonstrate, consider a typical Wireless LAN with an outdoor transmission range of up to 2000 feet, and a maximum vehicular speed of 5 miles/hour. If we assume that the mobile node moves straight across cells, the

needed minimal handoff frequency is 272 seconds (\gg 30 seconds, resulting in a TCP performance that is almost identical to W-LAN only handoff performance).

VI. CONCLUSION

We presented an architecture that aims at integrating wireless LANs and Wireless WANs, and showed the role of the Mobile-IP protocol as an integrative layer in that architecture. We also presented a simple subnet architecture that allows us to superimpose Mobile-IP on Wireless LANs. We carried out experiments with the objective of measuring and comparing 802.11 W-LAN handoff with the composite 802.11/Mobile-IP handoff in a real setting. The effect of the handoff on TCP performance was measured. We found that under practical values of handoff frequencies, the performance of Mobile IP based W-LAN handoff is almost identical to the performance of W-LAN handoff. The research ideas and measurements presented in this paper are only one step towards a full integration of nomadic and mobile computing.

Several enhancements to the basic Mobile IP protocol have recently been proposed and developed. One such enhancement is to reduce the handoff latency in order to minimize disruption to application due to handoff. In this paper, we focused our experiment only on the handoff of the basic Mobile IP. Further experiments using the enhanced Mobile IP are under investigation.

VII. REFERENCES

- [1] Charles E. Perkins, editor. IP mobility support, IETF RFC 2002, Oct. 1996.
- [2] IEEE Wireless Medium Access Control (MAC) and Physical layer (PHY) specifications. Draft Standard 802.11 D2.0, July 1995.
- [3] Abhijit Dixit and Vipul Gupta, Mobile-IP for Linux (ver 1.00), May 1996, Available at <http://anchor.cs.binghamton.edu/~mobileip/MIPv100/Doc/mip-doc.v100.ps>.
- [4] Ramon Caceres and Liviu Iftode, Improving the Performance of Reliable Transport Protocols in Mobile Computing Environments, IEEE Journal on Selected Areas in Communications, 13(5):850-857, June 1995.
- [5] Mary G. Baker, Xinhua Zhao, Stuart Cheshire and Jonathan Stone, Supporting Mobility in MosquitoNet, USENIX Winter 1996.
- [6] D. Johnson and D. Maltz, Protocols for Adaptive Wireless and Mobile Networking, IEEE Personal Communication, 3(1), Feb. 1996.
- [7] C. Perkins and D. Johnson, Mobility Support in IPv6, Proceedings of the Second Annual International Conference on Mobile Computing and Networking (MobiCom '96), Nov. 1996.
- [8] R. Droms, Dynamic Host Configuration Protocol, IETF RFC 2131, Mar. 1997.
- [9] H. Balakrishnan, S. Seshan and R. H. Katz, Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks, ACM Wireless Networks, 1(4): 469-482, Dec. 1995.
- [10] W. Richard Stevens, TCP/IP Illustrated, Volume 1: The Protocols. Addison-Wesley, Reading, Massachusetts, 1994.

- [11] AbsoluteValue Software, Inc., Linux Wireless LAN driver, <http://ftp.absoval.com/pub/linux-wlan>, 1999.
- [12] The MobileStar Network Home Page. <http://www.mobilestar.com>, 1999.