1995

An Architecture for A Camputs-Scale Wireless Mobile Internet

Douglas E. Comer
Purdue University, comer@cs.purdue.edu

John c. Lin

Vincent F. Russo

Report Number:

95-058

AN ARCHITECTURE FOR A CAMPUS-SCALE WIRELESS MOBILE INTERNET

Douglas E. Comer
John C. Lin
Vincent F. Russo

Department of Computer Science
Purdue University
West Lafayette, IN 47907

CSD-TR 95-058
September 1995

An Architecture For A Campus-Scale Wireless Mobile Internet

Purdue Technical Report CSD-TR 95-058
Douglas E. Comer, John C. Lin and Vincent F. Russo
Department of Computer Sciences
Purdue University
West Lafayette, Indiana 47907, U.S.A.
{comer, lin, russo}@cs.purdue.edu

Abstract

This paper proposes a network architecture to support mobile communication in a large university campus. The design allows each mobile host to maintain network connectivity while roaming freely in the campus. Called *Crosspoint*, the design combines wireless local-area network (LAN) technology with high-speed Asynchronous Transfer Mode (ATM) switching technology. The combination provides a wireless communication system with sufficient aggregate bandwidth to handle massive, synchronized movements of mobile hosts. Furthermore, the design requires no modification to conventional network software on mobile hosts, stationary hosts, or existing IP routers. We have implemented a prototype of the Crosspoint design. This paper also describes the basic protocols used in the implementation.

1 Introduction

Recent advances in personal computing and wireless LAN technologies have resulted in affordable laptop and palmtop computers with wireless networking capability. A computer with a wireless LAN adaptor can communicate directly with other wireless computers in the same wireless LAN while roaming. To communicate with computers outside the wireless LAN, a roaming wireless computer (or a mobile host) uses a nearby *base station*. A base station is normally a stationary computer with a wireless interface and a connection to conventional network facilities using terrestrial links. In particular, a base station that connects to the global TCP/IP Internet can provide a mobile host with access to other computers at sites around the world.

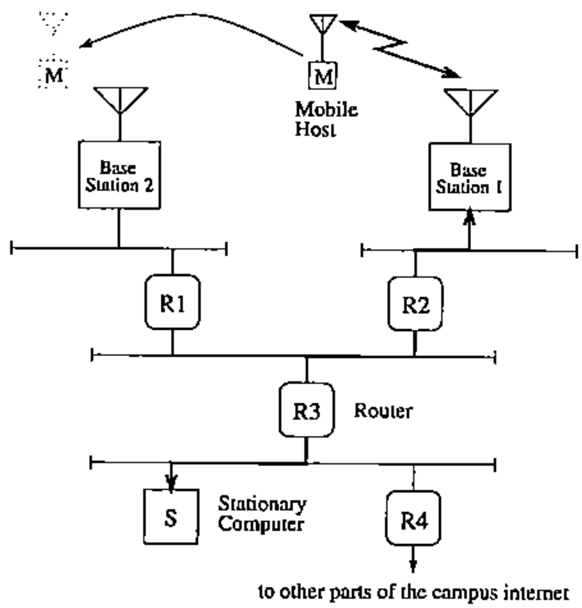

Figure 1: An example internet supporting mobile communication

A base station can provide network access for a group of mobiles that are within the *area* covered by its wireless interface. Because a base station can only cover a limited area, multiple base stations are needed to provide coverage for a large area. Attaching multiple base stations to an internet introduces routing problems that result when a mobile host migrates from base station to base station. As Figure 1 illustrates, a mobile host, M, uses base station 1 to communicate with a stationary computer, S. Base station 1 forwards packets from M to router $R2$; the packets then travels through $R3$ and S. Packets from S to M traverse the same path in the reverse direction. To maintain connectivity as M migrates to base station 2's area, $R3$ must change its next-hop route from $R2$ to $R1$, and $R1$ must change its next-hop route to base station 2. In fact, all the routers on the campus internet must change routing entries that correspond to M because M may communicate with any host reachable from the campus internet. Note that routers exchange routing information to update routing entries. Furthermore, packets that carry routing information compete with data packets for network bandwidth.

The overheads of propagating routing updates are especially apparent in a large university campus where 50,000 mobile hosts occupy in a small geographic area. More important, movements of mobiles at a university are massive and synchronized — a large percentage of the population migrates to new locations during each change of class. Without a careful design, the campus internet may experience network congestion when most students attempt to communicate from new locations. The situation becomes worse because congestion can cause delay or loss of routing updates, making data packets to follow nonoptimum paths. Diverse capabilities in students' mobile computers also complicate the design. In particular, some students will choose mobile computers that do not have a preemptive, multitasking operating system. We seek a design that accepts such systems as first-class members.

Many researchers have considered the problems of supporting mobile communication in an internet environment [1, 2, 5, 6, 11, 16, 17, 18, 19]. However, none of the designs meets the requirements stated above. For example, to avoid changing routes in the internet, researchers [18, 19] propose assigning two IP addresses to a mobile host's wireless interface. One static address is used to form transport connections; the other address is dynamically assigned each time when a mobile host enters a new base station's area. However, doing so requires substantial modification to network software, which is designed to use a single, static address. Furthermore, encapsulation will be needed because transport protocols often fix endpoint addresses when connections are formed. Another solution uses IP-tunneling to forward datagrams among base stations [6]. Such solution introduces unnecessary overheads and waste network bandwidth. Furthermore, the solution requires each mobile to execute a backgroud process that determines with which base station to associate.

In this paper, we propose a new design. Called *Crosspoint*, the design can handle the problem of synchronized route changes for a large university campus. Furthermore, the design requires no modification or addition to the network software on mobile hosts, stationary hosts, or existing routers. We will describe our design in the following sections.

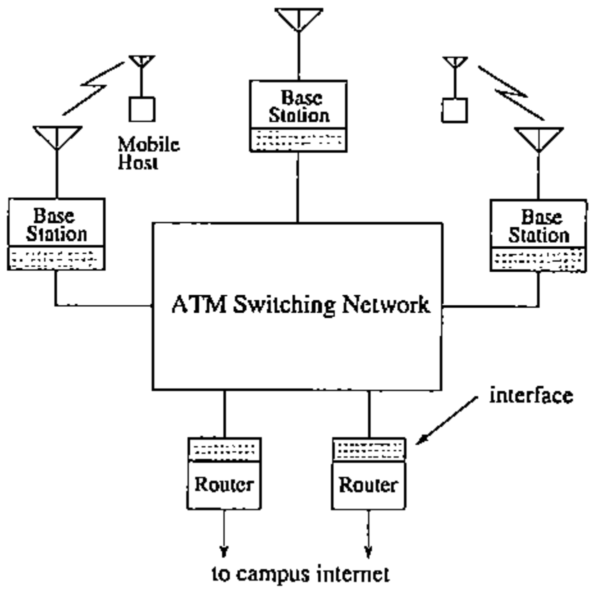

Figure 2: The proposed architecture

The remainder of this paper is organized as follows. Section 2 describes the Crosspoint architecture. Section 3 describes the basic protocols. Section 4 uses an example to illustrate how the architecture and protocols fit together to provide seamless mobile communication. Section 5 presents the current status of the Crosspoint project. Finally, section 6 summarizes the paper and discusses future work.

2 Proposed Architecture

Physical Interconnect. The proposed architecture uses a dedicated Asynchronous Transfer Mode (ATM) switching network [7, 8, 15] to interconnect base stations and special purpose routers. Figure 2 illustrates an example configuration. In the figure, base stations provide wireless access for the mobile hosts; routers with connections to the campus internet provide mobile hosts with access to the global Internet. Processors (i.e., base stations or routers) attached to the ATM network use ATM's high-speed switching fabric to transport mobile hosts' datagrams and to exchange control information. The switching fabric provides a high-bandwidth, low-delay interconnect among processors. Because an ATM fabric has sufficient capacity to handle the traffic from many processors, additional base stations and routers can be added as needed. It is feasible, for example, to scale the architecture to many base stations per building on a large campus.

Addressing and Routing. To accommodate 50,000 mobile hosts, a class-B IP address space is used for the wireless interfaces. Like a stationary host, once the wireless interface of a mobile host is configured with an IP address selected from the address space, the address remains fixed. As a mobile host migrates from base station to base station, base stations cooperate to track the mobile host. Collectively, base stations provide the appearance of a single, seamless LAN to which mobile hosts attach. Campus routers perceive the mobile network as a single subnetwork interconnected to the campus internet using Crosspoint routers; all datagrams destined for the mobile hosts are forwarded to the Crosspoint routers. The Crosspoint routers then use the ATM hardware to deliver the datagrams to base stations (see next subsection for details). Consequently, IP routing remains unchanged as a mobile host migrates from base station to base station. Therefore, propagating route changes throughout the campus internet is unnecessary.

Default Router Address. A single IP address selected from the class-B space is reserved for the wireless interface of every base station. A mobile host installs the reserved address as its default router address so that the mobile host can use a nearby base station to communicate with hosts that it cannot reach directly. Because all the base stations use the same IP address, a mobile host need not change its default router when migrating from one base station to another.

2.1 Virtual Circuit Management and Datagram Forwarding

ATM is a connection-oriented switching technology. A host attached to an ATM network uses *virtual circuits* to communicate with other hosts connected to the same network. Before communication takes place, a virtual circuit must be established between the communicating ATM hosts. Furthermore, ATM guarantees service quality for each established virtual circuit. For example, an ATM host can specify the needed bandwidth and the priority for each individual circuit.

In Crosspoint, a processor attached to the ATM network maintains a virtual circuit to each other processor over which it forwards datagrams. In addition, each pair of processors uses a second high-priority virtual circuit for sending control information. Thus, a processor will have two virtual circuits open to each other processor. To deliver a datagram to a mobile host, a processor chooses among the virtual circuits used for data, selecting the circuit that leads to the correct base station. Using two separate circuits ensures that data traffic does not compete with control traffic for network bandwidth. More important, packets containing routing updates have higher priority than data packets. In fact, each routing update requires only a single ATM cell and can be transported without adaptation. Thus, routing updates are efficient and do not propagate beyond machines that attach directly to the ATM network.

To process control and data packets at high speed, each processor connected to the ATM network includes a special interface, shown as a shaded box in Figure 2. The interfaces handle local routing and route changes that result when a mobile host migrates to a new base station. In particular, the interface implements an address-to-circuit binding for selecting an outgoing virtual circuit, given a destination address (i.e., a mobile's address). When a mobile host migrates to a new base station, the interfaces cooperate to handle all route changes. When routing information arrives at the interface, the interface automatically updates its address-to-circuit binding and begins using the new binding. In addition, the special interface executes the protocol modules that allow a mobile host to maintain transport connectivity while roaming. We describe the protocols next.

3 Basic Protocols

Because a base station's radio signal can only cover a limited area, a roaming host that is about to move out of the current base station's area must establish a radio link with a new base station, or a disruption of network connectivity will occur. By careful placement of each base station, the combined areas of all base stations can cover the entire campus, allowing a mobile host to maintain radio contact with a base station at all times. To provide total coverage, a base station's area may overlap with the areas of multiple base stations. Furthermore, a base station may receive radio signals from neighboring base stations.

3.1 Overlapping Areas

Overlapping areas present a challenge to protocol design. When situated in an overlapping area, a mobile host's radio signals may reach multiple base stations, and signals from multiple base stations may reach the mobile host. To avoid confusion, base stations can use a coding scheme at the radio level to distinguish each other's area [4, 10, 12]. However, doing so requires a mobile host to monitor various radio signals from multiple base stations and to reconfigure the interface when switching to a new base station. Our design does not require a mobile host to have such capabilities. In fact, all base stations and mobile hosts in Crosspoint use the same coded radio signal. Protocols at the base stations ensure the following invariants:

- At any time, only one base station handles a mobile host's communication requests.
- Base stations do not communicate with each other over the wireless interface.

The following subsections describe how the protocols ensure the invariants in the presence of overlapping areas.

3.2 Mobile Host Detection

Observe that computers emit packets when they try to communicate with other computers. Mobile hosts are no exception. Furthermore, mobile hosts tend to initiate communications with stationary server computers that are stable and contain resources. Thus, a base station can use packets (or frames) emitted from mobile hosts to detect their presence[1]. By listening in promiscuous mode, each base station monitors the activities of all the mobile hosts in its area.

Unlike other approaches [5, 6, 11] that require each base station to broadcast a *beacon* periodically and each mobile host to process the beacon and determine with which base station to associate, mobile hosts in Crosspoint do not participate in supporting seamless mobile communication. In fact, mobiles are completely unaware of the existence of multiple base stations around them. The network software transmits and receives packets as if a mobile host is associated with a single wireless LAN. Therefore, a mobile host can use conventional network software without modification or addition. Base stations cooperate to handle all the details required to support seamless mobile communication. Conceptually, the base stations provide an illusion of a single wireless LAN that covers the entire campus.

3.3 Initial Capture

When a mobile host powers on and initiates communication, it emits a frame (e.g., an Address Resolution Protocol (ARP) [13] frame to query its default router's hardware address) that allows nearby base stations to detect its presence. Because a mobile host may be situated in an overlapping area, multiple base stations may receive the frame. To ensure that only one base station captures a new mobile host, base stations uses an initial capture protocol. The protocol randomly assigns a new mobile to one of the base stations that simultaneously detect the mobile host's presence. The following example explains how the

[1] A base station can use an ICMP [14] echo request to elicit transmission from a mobile when necessary.

protocol works.

Assume that a mobile host, M, powers on and emits a frame that is received by a nearby base station, B. B immediately sends a CAPTURE message to the neighboring base stations over the control circuits[2]. In addition, B starts a *capture timer* that expires after $T_{capture}$ seconds. In the CAPTURE message, B includes a locally generated random number. The random number acts as a *bid* for M. If no other base station has detected M, B has no competitor and captures M after the capture timer expires. If there are other base stations that also detect M, they behave exactly the same as B: each computes a random number, sends a CAPTURE message to the neighboring base stations, and starts the capture timer. In this way, all the base stations that have detected M inform each other using the CAPTURE messages. A base station cancels the capture timer (loses its bid for the mobile) when an incoming CAPTURE message contains a random number that is greater than the number it had generated. Eventually, the base station that generates the largest random number captures M after the capture timer expires[3]. The winning base station immediately broadcasts a message to inform all the other processors that it *owns* M.

The situation in which a base station detects a mobile host that is already captured by another base station is handled by the handoff protocol, which we describe next.

3.4 Handoff

Handoff refers to a transfer of mobile ownership; a base station hands off a mobile host to a nearby base station. It is to use radio signal strength to determine when a handoff should occur: if a new base station can maintain a better radio contact with a mobile host, the base station that currently owns the mobile hands off ownership to the new base station. However, such scheme requires hardware support, and we do not assume such support is available. Instead, our design uses software. Each base station includes a module that uses the following two criteria to determine when to hand off a mobile host:

- A base station that receives a packet from a new mobile host consults the current owner before capturing the mobile.
- The base station that currently owns a mobile is given priority in maintaining the ownership.

Because our design assumes that a base station does not have hardware support to measure the radio link quality to a mobile host, the protocol software uses frame reception to approximate link quality. For example, if a mobile emits a frame that is received by both the owning base station and a nearby base station, the protocol deduces two equally good links, regardless of the actual quality of each individual link. Similarly, if the nearby base station receives the frame but the owning base station misses it, the protocol deduces that the nearby base station has a better link than the owner.

Like the initial capture protocol, the handoff protocol uses ATM control circuits to exchange control messages. When a base station receives a frame from a mobile that it does not own, the

[2] A base station can use an ATM point-to-multipoint circuit if the ATM hardware supports it.

[3] The base station with highest address captures the mobile host when there is a tie in the random numbers.

base station sends a message across a control circuit to the mobile's owner. The protocol requires each base station to maintain a timestamp for each mobile host. Whenever an owning base station receives a frame transmitted by a mobile, the base station updates its timestamp for the mobile. When a message arrives from another base station that has received a frame from the mobile, the owning base station uses the timestamp information to determine whether to hand off the mobile host to the other base station or retain ownership.

The handoff algorithm can be explained by an example. Assume that: base station B owns a mobile host, M; M emits a frame, and the frame is received by nearby base stations. If it receives the frame, B updates the timestamp that corresponds to M and forwards the frame. Other base stations that receive the frame buffer the frame, send a HANDOFF message to B, and start a handoff timer that expires after $T_{handoff}$ seconds. When it receives a HANDOFF message, B computes Δt, the difference between the time at which the HANDOFF arrives and B's timestamp for M. If Δt is greater than $T_{hthresh}$, B answers with a HANDOFF_ACK message, allowing the sending station to capture M; otherwise, B answers with a HANDOFF_NACK message, denying the sending station's handoff request. After it sends a HANDOFF_ACK, B denies subsequent handoff requests (i.e., requests from other base stations that have received a frame from M). Value $T_{hthresh}$ is normally a fraction of a second (e.g., 50 ms). Because the HANDOFF message traverses the ATM network with little delay, B deduces that M is still in its area when Δt is less than $T_{hthresh}$. If Δt is no less than $T_{hthresh}$, B deduces that it missed the packet that causes a base station to send the HANDOFF message and allows the station to capture M.

A base station that receives a HANDOFF_ACK cancels the handoff timer, forwards the buffered frames, and broadcasts a control message to declare the ownership of the captured mobile. A HANDOFF_NACK causes a base station to cancel the handoff timer and discard the buffered frames. When a mobile is situated in an overlapping area, the handoff timer allows a nonowner base station to capture the mobile in case the current owner station fails. Expiration of the handoff timer triggers a base station to use the initial capture protocol to capture a mobile.

When a mobile stays in an overlapping area, multiple base stations may receive frames from the mobile. If each frame from the mobile causes multiple base stations to send HANDOFF messages to the owner base station, the owner station may be overwhelmed with HANDOFF requests. To avoid excessive messages, base stations bound the rate at which they send HANDOFF messages. Specifically, a base station imposes a delay of at least T_{hdelay} seconds between successive HANDOFF messages.

3.5 Avoiding Wireless Communications Among Base Stations

A mobile host uses a nearby base station to communicate with hosts that it cannot reach directly. That is, the base station serves as the default router for the mobile host. To allow frame transmission across physical medium, IP uses ARP to bind the default router address to a hardware address. The binding is then stored in the ARP cache with a predefined lifetime. An ARP cache improves ef-

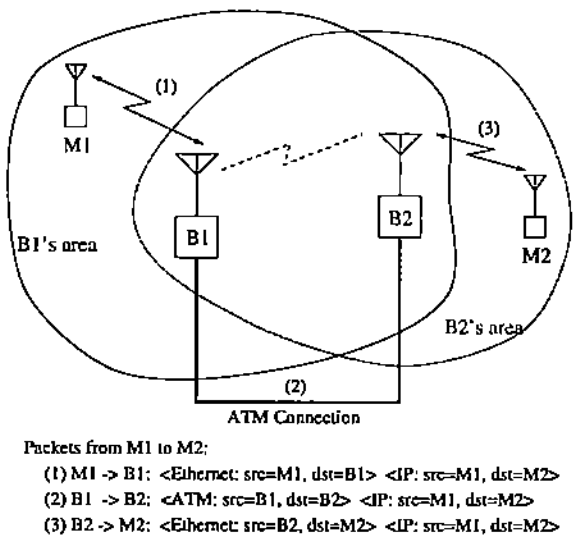

Figure 3: A network configuration that illustrates the need to prohibit communications among base stations over the wireless interface

ficiency by eliminating unnecessary ARP broadcasts. However, a cache can introduce binding errors when one IP address can map to various hardware addresses. The default router address is an example of such IP address. In particular, when a mobile host enters a new base station's area, the mobile's ARP cache will maintain the binding of the default router address to the previous base station's hardware address. Fortunately, an incorrect binding does not prevent the new base station from receiving frames from the mobile host, because base stations monitor network activities in promiscuous mode. The protocol software in a base station disregards a frame's destination address and uses IP addresses to identify the originating host and the intended recipient.

Disregarding a frame's destination address can create a forwarding problem, as Figure 3 illustrates. In the figure, two base stations, $B1$ and $B2$, are in range with each other. Mobile hosts, $M1$ and $M2$, communicate with each other via $B1$, $B2$, and the ATM connection between the two base stations. An IP packet from $M1$ to $M2$ requires three stages of transport through the path. First, an Ethernet frame carries the IP packet from $M1$ to $B1$. Second, the ATM hardware transports the packet from $B1$ to $B2$. Third, another Ethernet frame carries the packet from $B2$ to $M2$. Throughout the journey, the source and destination IP addresses of the datagram remain the same, while the source and destination hardware addresses change from stage to stage. Note that $B1$ can receive the frame transmitted by $B2$ because $B1$ and $B2$ are in range with each other. When $B1$ receives the frame intended for $M2$, it ignores the destination hardware address and processes the enclosed IP datagram. Because the datagram indicates the source is a local mobile host (i.e., $M1$) and the recipient is $M2$, $B1$ forwards the datagram to $B2$, thus creating a forwarding loop.

A solution to the problem is for a base station to check the source hardware address of each incoming frame. A base station discards any frame with source hardware address matches a neighboring base station's hardware address. For the solution to work, each base station must maintain a list of hardware addresses; each entry in the list corresponds to the wireless interface's hardware address of a neighboring base station that may cause the forwarding problem. The overhead of maintaining the list and checking each incoming frame against the list are major drawbacks of the solution. We describe a novel solution that avoids the drawbacks below.

The solution uses a single hardware address for all base stations (e.g., a multicast address), and then uses the hardware address to determine whether the frame is destined for a base station

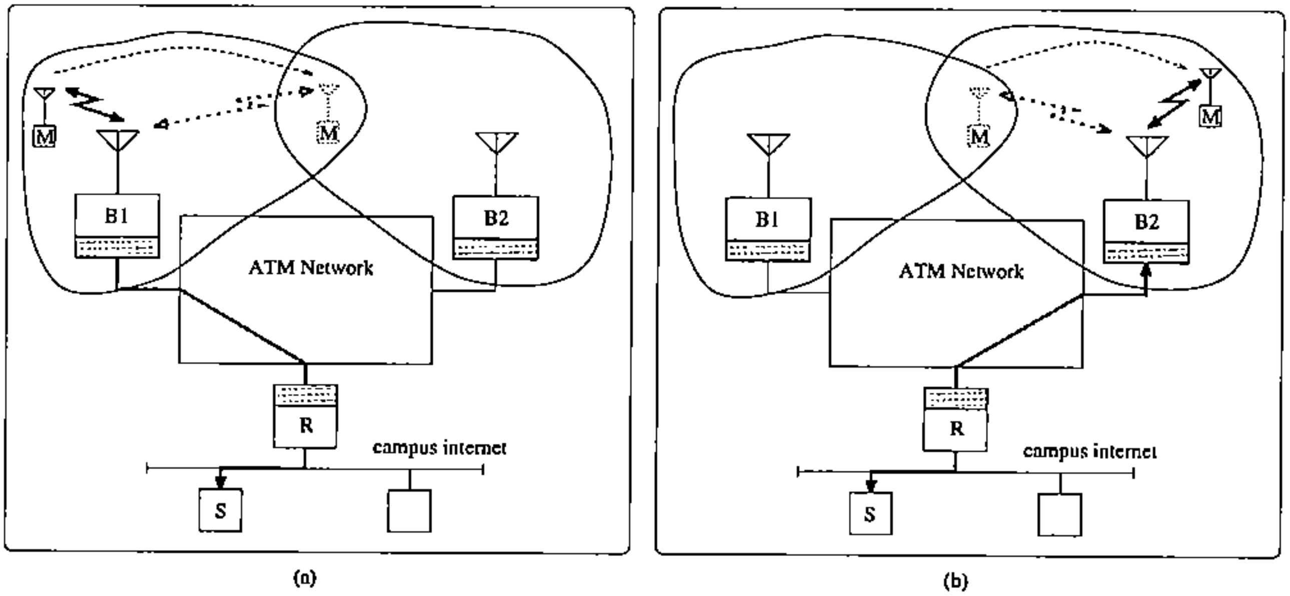

Figure 4: A roaming mobile host communicating with a stationary computer using Crosspoint

or for a mobile host. Because all base stations use the same hardware address, a mobile host always binds the default router address to the same value. Thus, all frames emitted from mobile hosts that are destined for any base station contain the same destination hardware address. By checking the destination address of each incoming frame, a base station can easily distinguish frames that are destined for base stations from those destined for mobile hosts. In the above example, $B1$ can discard frames from $B2$ that are destined for $M2$ because each of the frames contains $M2$'s hardware address as destination.

4 Combining Architecture and Protocols

With the understanding of the proposed architecture and the basic protocols, we use an example to illustrate how they fit together to provide seamless mobile communication. Figure 4 illustrates a roaming mobile, M, communicating with a stationary server computer, S. Router R serves as the default router to the campus internet for base stations $B1$ and $B2$. That is, the base stations always forward datagrams destined for stationary hosts to R. When M initiates communication with S, $B1$ captures M using the initial capture protocol. After capturing M, $B1$ immediately sends a route update to R and $B2$ over the ATM control circuits. Using the route update, R modifies the address-to-circuit binding that corresponds to M, allowing datagrams from S to reach M (Figure 4 (a)). When M migrates to the area covered by both $B1$ and $B2$, the handoff modules on both stations use the control circuits to exchange protocol messages. After $B1$ allows $B2$ to capture M, $B2$ immediately sends a route update to R and $B1$. R updates M's address-to-circuit binding using the received information and begins forwarding datagrams destined for M to $B2$ (Figure 4 (b)). Thus, M maintains network connectivity with S during the migration. Furthermore, no routing update has been propagated to the campus internet when M changes base station. Consequently, routers and hosts on the campus internet are not affected by M's mobility.

5 Current Status

A prototype implementation of the Crosspoint design has been working since May of 1995. The hardware configuration consists of three base stations, one router, and a FORE Systems ATM switch. A base station consists of a PC with a wireless interface and a SPARC station with an ATM interface; the two processors communicate via an Ethernet. The ATM interface driver and Crosspoint protocol software are implemented on the SPARC processors running SunOS 4.1.3; the driver software for the wireless interface and part of the protocol modules are implemented on the PCs running Xinu operating system [3]. All the wireless interfaces are AIRLAN adaptors made by Solectek Corporation. A mobile PC runs Microsoft Windows 3.1 with TCP/IP support. A mobile PC only requires to configure its IP address and the default router address to access the wireless facility. Once configured, a mobile PC can communicate with hosts on the Internet and maintain network connectivity while roaming within the Computer Science building. Communications between two mobile PCs, situated in the same area or in different areas, are also supported.

6 Summary and Future Work

This paper proposes an architecture for a campus-scale wireless mobile internet. Central to the architectural design is an ATM switching network to which base stations and special purpose routers attach. A processor (e.g., a base station or a router) attached the ATM network uses two virtual circuits to communicate with each of the other processors attached to the same network: one circuit uses ATM Adaptation Layer 5 (AAL5) [9] to transport IP datagrams; the other high priority circuit uses raw ATM cells (without adaptation) to carry control information. The virtual circuits provide high-speed, low latency interconnections among processors to support seamless mobile communication.

The protocol design focuses on providing a mobile host with seamless mobile communication, without modifying the mobile's network software and hardware. A mobile host transmits and receives packets as if it is associated with a single wireless LAN, regardless of its location in the campus. Protocol software on a base station handles all the details required to support seamless mobile communication for a roaming mobile host. When a mobile host is situated in an area covered by multiple base stations, the protocol software ensures that only one base station handles a mobile's communication requests. In essence, the protocol software uses the high-speed ATM connections among processors to provide an illusion of a single wireless LAN that covers the entire campus. Routers and hosts on the campus internet perceive mobile hosts as stationary computers attached to a single subnet.

We have implemented a prototype using the architecture and protocols described in this paper. The prototype allows us to experiment with ideas, to refine the design, and to investigate communication issues that are unique in a wireless environment. For example, we have observed that two mobiles can be in range with the same base station but out of range with each other. To make the communication between the two mobiles possible, the base station relays packets for the two mobiles. As the two mobiles move in range with each other, the base station should stop relaying and allow the two to communicate directly. Similarly, two mobiles in a base station's area can communicate directly then move out of range with each other (but still in range with the base station). Maintaining communication between the two mobiles requires the base station to relay packets for the two mobiles when they are out of range with each other. In both cases, the cached ARP bindings on each mobile host impede seamless communication between the two mobiles. We are experimenting with solutions for both cases.

7 Trademarks

Fore Systems is a trademark of Fore Systems, Incorporated. Microsoft is a registered trademark and Windows is a trademark of Microsoft Corporation. Sun and SunOS are trademarks of Sun Microsystems, Incorporated. SPARC is a registered trademark of SPARC International, Incorporated. AIRLAN is a trademark of Solectek Corporation.

References

[1] Kenneth G. Carlberg. A Routing Architecture That Supports Mobile End System. In *Proceedings of IEEE MILCOM 1992*, pages 6.3.1–6.3.6, Oct. 1992.

[2] D. Cohen, J. Postel, and R. Rom. IP Addressing And Routing In A Local Wireless Network. *IEEE INFOCOM*, pages 5A.3.1–5A.3.7, 1992.

[3] Douglas E. Comer. *Operating System Design: The Xinu Approach.* Prentice-Hall, Englewood Cliffs, New Jersey, 1984.

[4] K. S. Gilhousen et al. On the Capacity of a Cellular CDMA System. *IEEE Transactions on Vehicular Technology*, 40(2):303–312, May 1991.

[5] T. Blackwell et. al. Secure Short-Cut Routing for Mobile IP. In *Proceedings of USENIX Summer 1994 Conference*, pages 305–316, June 1994.

[6] J. Ioannidis, D. Duchamp, and G. Q. Maguire Jr. IP-based Protocols for Mobile Internetworking. In *Proceedings of ACM SIGCOMM '91*, pages 235–245, Sept. 1991.

[7] ITU-T. Draft Recommendation I.150: B-ISDN ATM Functional Characteristics. *ITU Study Group XVIII*, June 1992.

[8] ITU-T. Draft Recommendation I.361: B-ISDN ATM Layer Specification. *ITU Study Group XVIII*, June 1992.

[9] ITU-T. Draft Recommendation I.363: B-ISDN AAL Specification. *ITU Study Group XVIII*, Jan. 1993.

[10] R. Kohno, R. Meidan, and B. Milstein. Spread Spectrum Access Methods for Wireless Communications. *IEEE Communications Magazine*, pages 58–67, Jan. 1995.

[11] C. Perkins and P. Bhagwat. A Mobile Networking System Based on Internet Protocol. *IEEE Personnal Communications*, First Quarter 1994.

[12] R. L. Pickholtz, L. B. Milstein, and D. L. Schilling. Spread Spectrum for Mobile Communications. *IEEE Transactions on Vehicular Technology*, 40(2):313–321, May 1991.

[13] D. Plummer. RFC-826: Ethernet Address Resolution Protocol. *Request For Comments*, Nov. 1982. Internet Network Information Center.

[14] J. Postel. RFC-792: Internet Control Message Protocol. *Request For Comments*, Sept. 1981. Internet Network Information Center.

[15] Martin De Prycker. *Asynchronous Transfer Mode: Solution for Broadband ISDN*. Ellis Horwood, Chichester, England, second edition, 1993.

[16] Yakov Rekhter. An Architechture for Transport Layer Transparent Support for Mobility. *Journal of Internetworking: Research and Experience*, 4(4):223–246, Dec. 1993.

[17] C. Sunshine and J. Postel. IEN-135: Addressing Mobile Hosts in the ARPA Internet Environment. *Internet Engineering Notes*, March 1980.

[18] F. Teraoka, Y. Yokote, and M. Tokoro. A Network Architecture Providing Host Migration Transparency. In *Proceedings of ACM SIGCOMM '91*, pages 209–220, Sept. 1991.

[19] H. Wada, T. Yozawa, T. Ohnishi, and Y. Tanaka. Mobile Computing Environment Based on Internet Packet Forwarding. In *Proceedings of USENIX Winter '93 Conference*, pages 503–517, Jan. 1993.