

- [42] F. A. Tobagi and L. Kleinrock. Packet Switching in Radio Channels: Part II - the Hidden Terminal Problem in Carrier Sense Multiple-Access Modes and the Busy-Tone Solution. *IEEE Transactions on Communications*, 23(12):1417–1433, 1975.
- [43] Bruce Tuch. Development of WaveLAN, and ISM Band Wireless LAN. *AT&T Technical Journal*, 72(4):27–37, July/August 1993.
- [44] 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.

- [19] 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.
- [20] J. Ioannidis and G. Q. Maguire Jr. The Design and Implementation of a Mobile Internetworking Architecture. In *Proceedings of USENIX Winter '93 Conference*, pages 491–502, Jan. 1993.
- [21] ITU-T. Draft Recommendation I.150: B-ISDN ATM Functional Characteristics. *ITU Study Group XVIII*, June 1992.
- [22] ITU-T. Draft Recommendation I.361: B-ISDN ATM Layer Specification. *ITU Study Group XVIII*, June 1992.
- [23] ITU-T. Draft Recommendation I.363: B-ISDN AAL Specification. *ITU Study Group XVIII*, Jan. 1993.
- [24] David B. Johnson and David A. Maltz. Protocols for Adaptive Wireless and Mobile Networking. *IEEE Personal Communications*, 3(1):34–42, Feb. 1996.
- [25] Phil Karn. MACA - A New Channel Access Method for Packet Radio. In *Proceedings of the ARRL 9th Computer Networking Conference*, Sept. 1990.
- [26] R. Kohno, R. Meidan, and B. Milstein. Spread Spectrum Access Methods for Wireless Communications. *IEEE Communications Magazine*, pages 58–67, Jan. 1995.
- [27] M. T. Le, F. Burghardt, S. Seshan, and J. Rabaey. InfoNet: the Networking Infrastructure of InfoPad. In *Proceedings of Compton 1995*, March 1995.
- [28] John C. Lin. *An Architecture For A Campus-Scale Wireless Mobile Network*. PhD thesis, Purdue University, Dec. 1996. <ftp://gwen.cs.purdue.edu/pub/lin/thesis.ps.Z>.
- [29] Shankar Narayanaswamy et al. Application and Network Support for InfoPad. *IEEE Personal Communications*, March 1996.
- [30] C. Perkins and P. Bhagwat. A Mobile Networking System Based on Internet Protocol. *IEEE Personal Communications*, First Quarter 1994.
- [31] Charles Perkins. RFC-2002: IP Mobility Support. *Request For Comments*, Oct. 1996.
- [32] 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.
- [33] D. Plummer. RFC-826: Ethernet Address Resolution Protocol. *Request For Comments*, Nov. 1982.
- [34] J. Postel. RFC-792: Internet Control Message Protocol. *Request For Comments*, Sept. 1981.
- [35] Jon Postel. RFC-925: Multi-LAN Address Resolution. *Request For Comments*, Oct. 1984.
- [36] Martin De Prycker. *Asynchronous Transfer Mode: Solution for Broadband ISDN*. Prentice-Hall, Englewood Cliffs, New Jersey, third edition, 1995.
- [37] Theodore S. Rappaport. *Wireless Communications: Principle and Practice*. Prentice-Hall, Englewood Cliffs, New Jersey, 1996.
- [38] C. Sunshine and J. Postel. IEN-135: Addressing Mobile Hosts in the ARPA Internet Environment. *Internet Engineering Notes*, March 1980.
- [39] Lucent Technologies. WaveLAN/ISA Network Adaptor Card. <ftp://ftp.wavelan.com/pub/pdf-file/isa/fs-isa.pdf>, 1997.
- [40] F. Teraoka, K. Uehara, H. Sunahara, and J. Murai. VIP: A Protocol Providing Host Mobility. *Communications of the ACM*, 37(8):67–75, Aug. 1994.
- [41] 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.

9 Acknowledgments

The authors would like to thank the people who contributed to the Crosspoint project, including Dr. John Steel, Harold Brown, Lebin Cheng, Karthik Prabhakar, David Starkey, and David Stevens. We are grateful for the support from the facilities staff of the department, especially Dan Trinkle and Tim Korb.

References

- [1] D. Allen. Hidden Terminal Problems in Wireless LAN's. *IEEE 802.11 Working Group Paper*, 1993.
- [2] M. G. Baker, X. Zhao, S. Cheshire, and J. Stone. Supporting Mobility in MosquitoNet. In *Proceedings of the 1996 USENIX Technical Conference*, Jan. 1996.
- [3] Mary Baker. Changing Communication Environments in MosquitoNet. In *Proceedings of the IEEE Workshop on Mobile and Computing Systems and Applications*, pages 64–68, Dec. 1994.
- [4] H. Balakrishnan, S. Seshan, and Randy H. Katz. Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks. In *ACM Wireless Networks*, Dec. 1995.
- [5] Pravin Bhagwat and Charles Perkins. A Mobile Networking System based on Internet Protocol (IP). In *Proceedings of Mobile and Location Independent Computing*, pages 69–82, Aug. 1993.
- [6] T. Blackwell et al. Secure Short-Cut Routing for Mobile IP. In *Proceedings of USENIX Summer 1994 Conference*, pages 305–316, June 1994.
- [7] R. Braden and J. Postel. RFC-1009: Requirements for Internet Gateways. *Request For Comments*, June 1987.
- [8] R. Caceres and L. Iftode. Improving the Performance of Reliable Transport Protocols in Mobile Computing Environments. *IEEE Journal on Selected Areas in Communications*, Oct. 1994.
- [9] Vint Cerf. IEN-110: Internet Addressing and Naming in a Tactical Environment. *Internet Engineering Notes*, August 1979.
- [10] Stuart Cheshire and Mary Baker. Experiences with a Wireless Network in MosquitoNet. In *Proceedings of the IEEE Hot Interconnects Symposium '95*, Aug. 1995.
- [11] Douglas E. Comer. *Operating System Design: The Xinu Approach*. Prentice-Hall, Englewood Cliffs, New Jersey, 1984.
- [12] Douglas E. Comer. *Internetworking with TCP/IP: Principles, Protocols, and Architecture, Vol. I*. Prentice-Hall, Englewood Cliffs, New Jersey, third edition, 1995.
- [13] AT&T Wireless Communications and Networking Division. Data Manual: WaveLAN Air Interface. *Document No.: 407-0024785 (Draft)*, June 1995.
- [14] CDPD Consortium. Cellular Digital Packet Data System Specification (Release 1.0). July 1993.
- [15] S. Deering. RFC-1112: Host Extension for IP Multicasting. *Request For Comments*, Aug. 1989.
- [16] The ATM Forum. *ATM User-Network Interface Specification Version 3.0*. Prentice-Hall, Englewood Cliffs, New Jersey, 1993.
- [17] K. S. Gilhousen et al. On the Capacity of a Cellular CDMA System. *IEEE Transactions on Vehicular Technology*, 40(2):303–312, May 1991.
- [18] Alex Hills and David B. Johnson. A Wireless Data Network Infrastructure at Carnegie Mellon University. *IEEE Personal Communications*, 3(1):56–63, Feb. 1996.

computing area.

Unlike the IETF Mobile IP scheme, which is designed mainly to solve the *macro* mobility management problem [31], the Crosspoint approach focuses on solving the *micro* mobility management problem. That is, the IETF Mobile IP scheme focuses on providing mobility support for mobiles that move from one internet to another, whereas Crosspoint provides mobility support for wireless mobiles that roam in an area covered by multiple wireless interfaces. The two schemes are complementary to each other. For example, the Mobile IP scheme can be used to support mobile computing among Crosspoint networks. From the perspective of Mobile IP, a Crosspoint network is a subnet attached to the Internet using one or more Crosspoint routers. The Crosspoint routers can serve as the home agents for local mobiles as well as the foreign agents for visiting mobiles. Further research is needed to investigate the impact of supporting Mobile IP on the design and implementation of Crosspoint.

Another area for future research is to extend the current design to support IP multicasting. Initial investigation indicates that in addition to using the Internet Group Management Protocol (IGMP) [15] to track active groups, base stations need a protocol to determine the multicast groups of which a mobile is a member. For example, during the ownership transfer for a mobile, the previous owner can inform the new owner about the multicast groups in which the mobile is participating. Thus, the new owner can determine whether to join the groups on behalf of the mobile. Alternatively, the new owner can send an IGMP membership query message to the mobile immediately after capturing the mobile. Once the group membership information is available, Crosspoint processors must cooperate to route multicast traffic. Whether point-to-multipoint virtual circuits are suitable for transport multicast traffic warrants further study.

Finally, we are investigating schemes that can improve TCP performance. Experiments performed on the prototype shows that TCP does not perform well in a wireless environment, where packet loss rate is higher than a wired environment. Researchers have proposed various schemes to improve TCP performance in a wireless environment. For example, one scheme [8] exploits the fast-retransmit mechanism. However, the scheme requires a mobile host to use a modified version of TCP/IP. Another scheme installs a *snoop* module [4] at each base station. The snoop module caches unacknowledged data for mobiles and performs local retransmissions across the wireless link to mobiles. The snoop approach is more suitable for Crosspoint because it allows base stations to handle the details of improving TCP throughput, without requiring modifications to the TCP/IP software on mobiles. Further research is needed to determine whether the snoop approach works well in Crosspoint. In particular, the handoff protocol may require modification, in order to transfer the state information maintained by the snoop module to the new owner.

support mobile hosts [31]. The IETF approach can be further refined to support optimal routing [6, 24] by installing additional software on nonmobile hosts. A mobile can also become a forwarding agent for itself when visiting a foreign network.

A number of institutions are building wireless data network system for mobile computing research. Carnegie Mellon University is building a wireless networking infrastructure called *Wireless Andrew* [18]. Currently, Wireless Andrew consists of two types of wireless systems: a wide area system using 19.2 Kbits/second Cellular Digital Packet Data (CDPD) [14] service and a local area system using 2 Mbits/second Lucent WaveLAN technology [43]. Wireless Andrew provides roaming support between the CDPD network and the WaveLAN network.

At the University of California at Berkeley, researchers are building a wireless network system, *InfoNet* [27], for supporting the InfoPad project [29]. A mobile computer in InfoPad is a portable multimedia terminal called *Pad*. Each Pad uses two wireless links to communicate with a *Gateway*, which provides Pads with access to a backbone network. The link from a Pad to a Gateway is a Proxim radio link with a capacity of 244 Kbits/second; the link from a Gateway to a Pad is a Plessey radio link with a capacity of 700 Kbits/second [27]. Each Gateway supports a small geographical region, called *picocell*, with a typical radius of 10 meters. Server processes use protocols to support seamless migration of a Pad from one picocell to another. The servers make the Pad appear to be a stationary terminal attached to the backbone network.

The *MosquitoNet* project [3] at Stanford University is investigating operating system and application issues in mobile and wireless computing. Researchers have built a testbed that consists of a wireless network and a collection of wired networks. The wireless network uses the Ricochet micro-cellular data network service provided by Metricom. The service uses pole-top radio units to provide wireless access for mobiles. Each radio unit offers a raw data rate of 100 Kbits/second. Ricochet uses Metricom proprietary routing protocols to provide roaming service to mobiles [10]. MosquitoNet uses a scheme similar to the IETF Mobile IP scheme to support transparent host migration among the wireless network and the wired networks, with emphasis on not using foreign agents [2].

8 Conclusion and Future Work

We have described the design and implementation of a campus-sized wireless network called Crosspoint. The architectural design uses a scalable switching fabric that provides high-speed interconnections among base stations and Crosspoint routers. The protocol design focuses on supporting seamless wireless mobile communication without modifying the networking software and hardware of each mobile. From the perspective of a mobile, the Crosspoint network is a single wireless LAN that covers the entire campus. A prototype implementation of Crosspoint has been installed in the Computer Science building and used daily by students and staff. The prototype also serves as a testbed for future research in wireless mobile

configured with a Crosspoint IP address and the default router address, it can access the Internet via the base stations using vendor supplied networking software and driver code without any modifications. Reference [28] documents a set of measurements performed on the prototype, such as handoff latency, the impact of handoff on TCP throughput, and signal strength variation observed by base stations when a mobile travels through an overlapping area.

7 Related Work

In the early days of Internet, researchers [9] had discovered the weakness of IP's addressing and routing scheme in supporting mobile hosts². In Internet Engineering Notes (IEN) 135 [38], Sunshine and Postel proposed using loose source routing and *forwarders* to support mobile hosts. They also proposed reserving a single network ID for the mobiles. A mobile host uses the same address regardless of the physical network to which the host attaches. A global database maintains the bindings of mobile-to-forwarder. A new message format was designed to query and update the location of a mobile. When communicating with a mobile, a host is responsible for obtaining the address of the mobile's forwarder and creating the loose source routing option in the IP header.

Researchers [5, 30] at IBM corporation also proposed using loose source routing to support mobile computing. Unlike the IEN-135 approach, the IBM approach does not restrict all mobiles to use the same network ID. Each mobile has a *home network*. The network ID of a mobile's IP address identifies the home network of the mobile. Each home network has at least one *mobile router (MR)* that forwards datagrams for a mobile that is away from its home network.

Researchers [41] at SONY corporation proposed an approach that assigns two IP addresses to a mobile: one permanent address, called a *virtual IP (VIP)* address that is used to identify the mobile, and one temporary IP (TIP) address for forwarding purpose. Like the IBM approach, the SONY approach also uses routers, called *primary resolvers* [40], at home networks to intercept datagrams and forwards datagrams for mobiles.

At Columbia University, researchers [19, 20] proposed an approach that uses packet forwarding and IP tunneling to support mobile hosts. The Columbia approach uses *Mobile Support Routers (MSRs)* to support a campus-sized wireless mobile internet. Each MSR supports wireless communication for a geographical area called a *cell*. MSRs use IP tunnels to turn the physically disjoint cells into a coherent *virtual* subnet of the campus internet.

The Internet Engineering Task Force (IETF) Mobile IP Working Group uses a forwarding scheme similar to the Packet Forwarding Method described in [44] and the forwarding scheme of the Columbia approach. The IETF approach uses *home agents*, *foreign agents*, and IP tunnels between the two types of agents to

²At that time, airborne packet radios were mobile hosts.

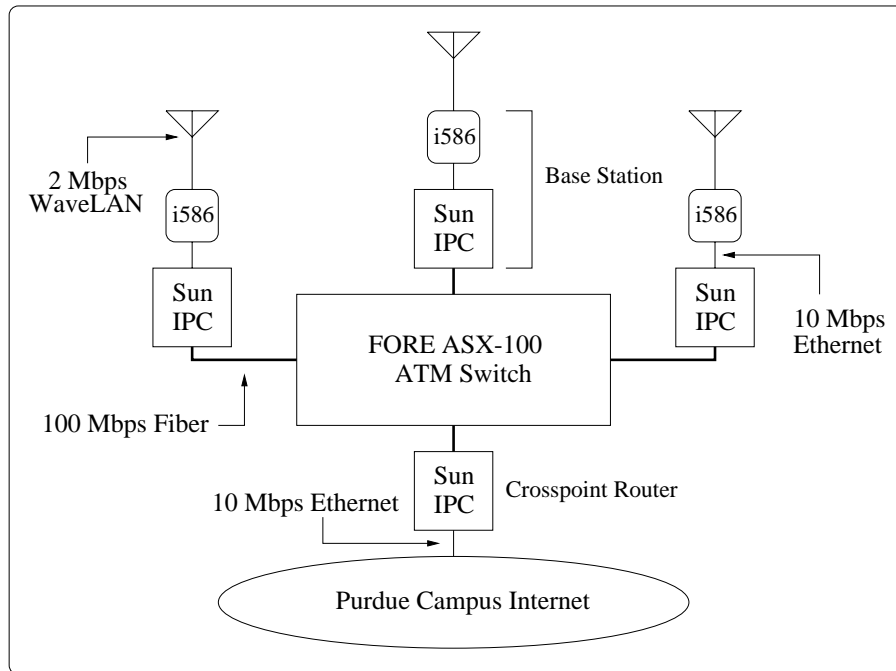


Figure 6: Illustration of the prototype Crosspoint network

inefficiency.

6 Prototype Implementation

A prototype implementation of Crosspoint has been working since 1995. As Figure 6 illustrates, the prototype consists of an ATM switch, a Crosspoint router, and three base stations.

The Crosspoint router is a SPARC IPC workstation. Each base station consists of two computers: a SPARC IPC workstation with an ATM interface and a 120 MHz Pentium PC with a 2 Mbps WaveLAN [39, 43] interface; the two computers communicate using a 10 Mbps Ethernet. The connection between each IPC workstation and the ATM switch is a 100 Mbps fiber link. Each IPC workstation uses two point-to-point permanent virtual circuits (PVCs) to communicate with each of the other IPC workstations. One PVC uses AAL5 to carry IP datagrams, and the other uses raw ATM cells to carry control messages. The prototype system attaches to the campus internet using a 10 Mbps Ethernet interface.

All the IPC workstations run version 4.1.3 of the SunOS operating system. All the Pentium PCs run the Xinu operating system [11]. The core of the Crosspoint software resides in the SunOS kernel. The driver for the wireless interface and the rest of the Crosspoint software reside in Xinu.

The three base stations were placed on each floor of the Computer Science building. Students use notebook PCs equipped with WaveLAN network adaptors to access the network. Once a mobile PC is

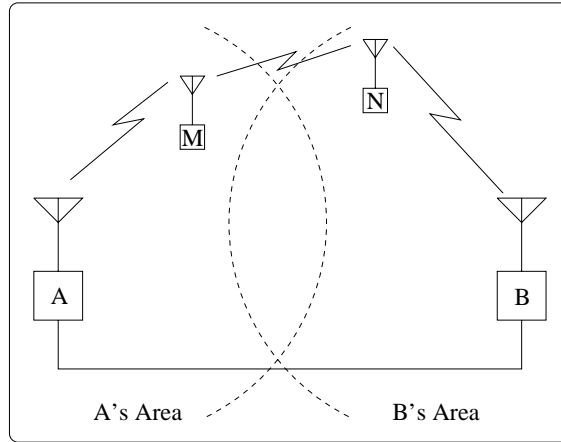


Figure 5: A network configuration that illustrates the hidden terminal situation. M is in-range of both A and N . N is out-of-range of A . A and N are said to be hidden terminals to each other.

the request. The binding of N 's IP to physical address on M depends on which ARP reply arrives *later*, because the binding carried in the later reply supersedes that carried in the former one. If the proxy-ARP from A arrives later, M will communicate with N via A , instead of directly with N . Fortunately, the hidden terminal situation does not impede the communication between the two mobiles; M can still communicate with N via base stations A and B .

5.5 ARP Cache and Direct Communication

Although direct communication between two mobiles is more efficient, cached ARP entry can prevent the mobiles from communicating with each other if they move out-of-range of each other. Consider the example in Figure 5 again. Suppose mobiles M and N communicate with each other directly. Thus, each maintains the other's IP-to-physical address binding in the ARP cache. When the mobiles move out-of-range of each other, the cached bindings make the communication impossible. Moreover, nearby base stations, A and B , cannot help because the destination physical addresses of the frames emitted from the mobiles are not the default router physical address.

Decreasing mobiles' ARP cache timeout value reduces the probability of stale ARP entries impeding communication between mobiles. However, it increases ARP broadcast frequency and requires additional software or configuration on mobiles. A scheme that requires no change on ARP cache timeout is to enforce the following policy: two mobiles always communicate via one or more base stations. Base stations can use address binding techniques such as proxy-ARP to implement the policy. However, the policy introduces inefficiency: two mobiles that are in-range of each other cannot communicate directly if there are base stations around them. Observations so far have shown that two mobiles that are in-range of each other rarely communicate. If the observation holds, the simplicity of this approach may outweigh the introduced

are intended for mobiles in the neighbor's area, thus creating a forwarding loop.

5.3 Indirect Communication via Base Stations

Two mobiles that are out-of-range of each other cannot communicate directly. However, if each is in-range of a base station, the two can communicate via the base stations. The base stations use a technique known as *proxy-ARP* [7, 12, 35] to make the communication possible. Consider the example illustrated in Figure 4.

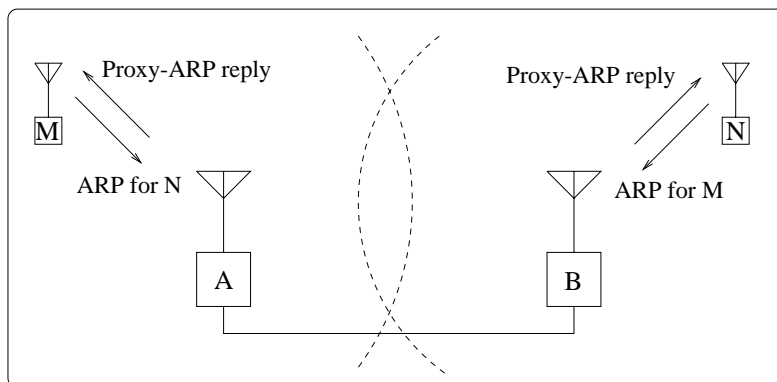


Figure 4: Illustration of how two base stations use proxy-ARP to allow two mobiles that are out-of-range of each other to communicate

In the figure, base station A is the owner of mobile M , and base station B is the owner of mobile N . When M initiates communication with N the first time, it broadcasts an ARP request to determine N 's physical address. Because M and N are out-of-range of each other, N cannot receive the ARP request. However, base station A can receive the ARP request. A knows N is reachable remotely, so it answers the ARP request on behalf of N . The proxy-ARP reply carries the binding of N 's IP address to A 's physical address. After processing the proxy-ARP reply, M sends datagrams destined for N to A , which in turn forwards them to B , and finally to N .

Likewise, when N sends the first datagram to M , it broadcasts an ARP request to determine M 's physical address. B answers the ARP request for M . Thus, N sends datagrams destined for M to B . When B receives the datagrams, it forwards the datagrams to A , which then forwards them to M . The communication between M and N is thus established.

5.4 The Hidden Terminal Situation

The proxy-ARP scheme described above can prevent two mobiles from communicating directly with each other if a situation known as *hidden terminal* [1, 25, 42] occurs. Consider the example in Figure 5.

Unlike the previous example, mobiles M and N are in-range of each other. When M broadcasts an ARP request to determine N 's physical address, both A and N can receive the request and will answer

R change the state for the mobile to reachable.

5 Address Binding and Communication Between Mobiles

The wireless interface used in Crosspoint supports direct communication between two mobiles [13]. Because mobiles in Crosspoint use unmodified network software and reside in a single IP network, a mobile always assumes direct communication with another mobile is possible. Communication difficulties arise when two mobiles that are out-of-range try to communicate directly. This section describes how base stations use IP-to-physical address binding techniques to allow two mobiles that are out-of-range of each other to communicate. It also discusses how address binding can facilitate or impede mobile communications.

5.1 The Default Router Physical Address

As explained in subsection 2.4, the wireless interface of every base station uses the same IP address, and the address is the default router address of every mobile. A mobile uses the default router address to communicate with its current owning base station. To send frames to its owner, a mobile uses ARP (Address Resolution Protocol) [33] to determine the owner's physical address. Once the owner's IP-to-physical address binding is available, it stores the binding in a cache (called *ARP cache*) to reduce the frequency of using ARP.

Unlike conventional network adaptors that filter incoming frames based on the destination physical address, the wireless interface of every Crosspoint base station accepts all incoming frames (i.e., it operates in promiscuous mode). Base stations use the *source* IP address carried in each incoming frame to determine whether it should forward or drop the frame. This approach allows a mobile to switch base stations without changing the default router's IP-to-physical address binding.

In fact, the wireless interface of every base station uses the *same* physical address. As a result, the IP-to-physical address binding of a mobile's default router remains the same regardless of which base station is the mobile's owner. Furthermore, datagrams sent from mobiles that are destined for base stations all have identical destination physical addresses. Thus, base stations can use the destination physical address to determine whether a mobile is sending frames to another mobile directly or to a base station.

5.2 Avoiding Wireless Communication Among Base Stations

Using the same physical address for the wireless interface of every base station has another nice property that frames emitted from base stations all have identical *source* physical addresses. A base station can use this property to ignore frames coming from neighboring base stations. Without ignoring frames from neighboring base stations, a base station will forward frames emitted from a neighboring base station that

processor resolves the situation by sending a control message to select one of the base stations as the owner.

4.4 Handling Outdated Routing Entries

When a base station captures a mobile, it broadcasts a route update message. Only the previous owner base station uses timeouts and retransmissions to ensure receiving the route update. For the other Crosspoint processors, the new owner does not check whether they have received the route update. If a processor misses the route update, its routing entry for the mobile becomes outdated.

To rectify outdated routing entries, Crosspoint processors use two schemes. First, base stations use broadcast to propagate a route update. Although broadcasting a route update does not guarantee that the route update will arrive at every receiver, each broadcast provides a new opportunity for a Crosspoint processor to update the routing entry for a mobile. Second, in response to a misrouted datagram, a processor uses a redirect control message to correct the sending processor's routing entry.

It is important that a redirect message for a mobile is originated from a processor that has reliable ownership information for the mobile. Observe that the set of processors that are likely to receive a misrouted datagram destined for a mobile is the set of the previous owners of the mobile. As mentioned earlier, the handoff protocol ensures that each of the previous owners maintains reliable ownership information for the mobile. Thus, a redirect message in response to a misrouted datagram will contain reliable information that the receiver can use to reach the correct owner.

Another form of a redirect message is a handoff NAK message. If a processor receives a handoff request for a mobile that the processor does not own or has handed off to another processor, the processor denies the request using a handoff NAK message. The NAK message carries the owner of the mobile back to the sender.

4.5 Handling Denial of Service

A lost route update can cause a more serious consequence: denial of service. To understand the cause, consider the following scenario. Suppose a previously unreachable mobile, M , becomes active and emits a frame. A base station captures M and broadcasts a route update for M . If a Crosspoint router, R , misses the route update, M 's state on R will remain unreachable. R will discard incoming datagrams destined for M . Consequently, M cannot communicate with hosts outside Crosspoint via router R .

Crosspoint processors use the following scheme to solve the service denial problem. Suppose a mobile emits a datagram, and the mobile's owning base station forwards the datagram to a Crosspoint processor. When the datagram arrives, the processor can infer that the mobile is reachable.

Applying the scheme to the example given earlier, router R checks each datagram received from the Crosspoint interconnect. If a datagram is originated from a mobile, and the mobile has a unreachable state,

to capture a mobile, the owner ensures that the new owner captures the mobile. The ownership transfer completes after the owner receives a route update for the mobile from the new owner. Thus, the previous owner always maintains a correct routing entry for the mobile. Maintaining correct routing entry allows the previous owner to redirect misrouted datagrams to the new owner and correct the routing entry of the processor that forwards the datagrams.

4.3 Resolving Simultaneous Capture

A *simultaneous capture* situation occurs when the owner of a mobile fails while more than one base station waiting for a handoff reply detects the failure and broadcasts a route update at about the same time. As a result, the routing entry for the mobile on all processors may become inconsistent.

To prevent the situation, each route update message contains an additional field used to carry a *bid*. Before broadcasting a route update for a mobile, a base station forms a bid by concatenating the signal strength to the mobile with a random number, and includes the bid in the route update. Thus, a base station that has a better signal to the mobile has a higher bid, and base stations that have the same signal strength to the mobile have equal chance to capture the mobile. Finally, base stations compare host IDs to break a tie.

After broadcasting the route update, the base station monitors incoming route update messages. If another route update for the same mobile arrives shortly (e.g., within 50 ms) after the broadcast, the base station infers that a simultaneous capture has occurred and compares the bid carried in the incoming route update to the local bid. If the local bid is larger, the base station starts a timer for recapturing the mobile. If another route update that carries a larger bid arrives, the base station cancels the timer, terminating the attempt to recapture the mobile. If all the incoming bids are less than the local bid, when the timer expires, the base station broadcasts the route update again to ensure that the routing entry for the mobile is consistent across all processors.

Because there is no guarantee that all the base stations that take part in the simultaneous capture will receive the route update from each other, another round of bidding may occur. In the worst case, two or more base stations end up becoming the owner of the same mobile. When that occurs, a Crosspoint processor can help rectify the inconsistency.

Consider the following scenario. Suppose two adjacent base stations, *A* and *B*, simultaneously capture mobile *M*, and each misses the route update from the other. Thus, each of the base stations regards itself as the owner of *M* and forwards datagrams for *M*. Because either base station can reach *M* directly, mobile *M* can still receive incoming datagrams. However, both base stations will forward a datagram emitted from *M*, thus generating duplicates. If both base stations forward the datagram to a common Crosspoint processor (e.g., a Crosspoint router), the processor will receive two datagrams originated from the same mobile one after another from two *different* base stations, an indication of a simultaneous capture. The

requests [34] to elicit a response from the mobile (*local search*). A response from the mobile indicates that the mobile is still reachable. If it still fails to receive frames from the mobile after sending a fixed number of ICMP echo messages, the owner requests neighboring base stations for assistance to locate the mobile (*neighborhood search*). The neighboring base stations also use ICMP echo requests to elicit frames from the mobile. The mobile is deemed unreachable after the local search and neighborhood search fail.

4 Reliability Considerations

The Crosspoint design assumes that the Crosspoint interconnect only provides *best-effort* packet delivery. That is, the network hardware does not guarantee that a packet sent across the interconnect will arrive at the destination processor intact. Besides possible loss of packets, Crosspoint processors may fail. Both packet loss and processor failure can cause inconsistent routing states on Crosspoint processors. This section discusses schemes and protocol extensions that minimize or rectify routing inconsistency.

4.1 Handling Failure During Handoff

Recall that the handoff protocol uses a request-reply type of interaction between two base stations. A base station that tries to capture a mobile sends a handoff request to the mobile's owning base station and awaits a reply. Because communication failures can occur, the non-owner base station starts a timer before sending the request to avoid waiting forever. If the reply arrives, it cancels the timer and proceeds. If the timer expires, it assumes that the request was lost, retransmits the request, and starts another timer. If the reply does not arrive after a preset number of retransmissions, the non-owner base station assumes that the owner base station has failed and captures the mobile.

On the owner side, if it permits the non-owner base station to capture the mobile, it replies with a handoff ACK, starts a timer, and awaits a route update message from the non-owner. While waiting, the owner base station retransmits the handoff ACK message in response to a handoff request from the non-owner, and rejects handoff requests from the other base stations. If a route update message does not arrive before the timer expired, the owner base station assumes that the non-owner has failed and broadcasts a route update to recapture the mobile. In addition, the owner base station starts the revalidation protocol to ensure that the mobile will emit a frame that allows other nearby base stations (if any) to detect it. Broadcasting the route update is necessary because the non-owner base station may have failed during propagation of route update.

4.2 Reliable Ownership Transfer

The protocol interaction described above illustrates an important protocol design decision: ownership transfer from one base station to another is reliable. Once an owner base station allows another base station

protocol to ensure that only one base station captures the mobile. Interestingly, the handoff protocol can also be used for the initial capture.

The idea is simple. Each Crosspoint processor uses a predefined algorithm to initialize its routing table. The algorithm ensures that the routing entry of each mobile is consistent across all Crosspoint processors. In particular, the owner of a given mobile is set to the same *Crosspoint router*. Thus, when a mobile initiates communication, each of the base stations that detect the mobile will send a handoff request to the same Crosspoint router. Because a Crosspoint router never owns a mobile and does not know how many handoff requests it will receive, it allows the sender of the first handoff request to capture the mobile. Once a base station captures the mobile and propagates the routing update, it will process future handoff requests based on signal strength comparison.

3.5 The Recapture Protocol

Recall that a routing entry contains both the ownership and reachability information about a mobile. If a mobile becomes unreachable, Crosspoint processors change the reachability information in the routing entry for the mobile, but maintain the ownership information. When a previously unreachable mobile initiates communication, each of the base stations that detect the mobile looks up the routing entry that corresponds to the mobile and sends a handoff request to the owner of the mobile. If it also detects the mobile, the owner recaptures the mobile and then processes each incoming handoff request. If it does not detect the mobile, the owner permits the sender of the first arriving handoff request to capture the mobile.

3.6 The Revalidation Protocol

Once a base station has captured a mobile, the base station needs to revalidate the mobile periodically for two reasons. First, the mobile may be turned off by its user. Without a mechanism for determining the mobile's status, subsequent datagrams to the mobile may end up wasting resources. Once it has determined that the mobile is no longer reachable, the owner can send an ICMP host unreachable message back to the host that tries to communicate with the mobile.

Second, the mobile may not emit a frame when it roams into the area of a new base station. Consequently, the new base station cannot detect the mobile. Meanwhile, datagrams are being forwarded to the mobile's owner, which no longer has a wireless link to the mobile. Note that the mobile may emit frames if it were able to receive the datagrams forwarded by the owner. And, the emitted frames are exactly what the new base station needed to detect the mobile. When such a situation occurs, the owner needs a way to detect the mobile's absence and request the assistance of the other base stations to locate the mobile.

The owner base station of a mobile initiates the protocol whenever it receives a datagram destined for the mobile, and the mobile has not emitted a frame for a predefined period. The owner uses ICMP echo

a reply. The request contains a weighted running average of the measured signal strength to the mobile¹. When the request arrives, the owner compares its signal strength measurement to that contained in the request. If the incoming measurement is greater than the owner's by a threshold, the owner permits the non-owner to capture the mobile by responding with a positive acknowledgment (or an ACK); otherwise, it denies the request using a negative acknowledgment (or a NAK). If a handoff ACK arrives, the non-owner base station captures the mobile and broadcasts a route update to inform the other Crosspoint processors about the ownership change.

3.3.1 Handling Multiple Incoming Handoff Requests

Because a mobile may stay in an overlapping area, the mobile's owner may receive handoff requests from multiple base stations. Because it does not know in advance how many handoff requests will arrive, the owner does not wait for all requests to arrive then process them. Instead, the owner processes each request as the request arrives. Once it has permitted a base station to capture the mobile, the owner denies subsequent handoff requests. The owner includes the new owner's ID in each handoff NAK message to inform the requesting base station that subsequent handoff requests should direct to the new owner.

3.3.2 Reducing the Frequency of Sending Handoff Requests

If every frame emitted from a mobile causes a handoff request sent to the mobile's owner, the owner may be overwhelmed when the mobile is situated in an overlapping area and emits many frames in a short interval. A base station uses two schemes to reduce the frequency of sending a handoff request.

First, if the signal strength to a mobile is weak (e.g., less than a threshold), a base station does not send a handoff request for the mobile, because the probability of receiving an ACK is low, and even if the base station receives an ACK, communication with the mobile may be impossible. We have observed that a base station's antenna can be more sensitive than a mobile's. As a result, the situation where a base station can receive a frame from a mobile, but the mobile cannot receive a frame from the base station can occur. In such circumstances, a base station can use the signal threshold to minimize the effect of asymmetric reception.

Second, a base station imposes a fixed time interval between two successive handoff requests. Thus, the rate at which a base station sends handoff requests is bounded.

3.4 The Initial Capture Protocol

When all the Crosspoint processors initialize the first time, no processor has a valid routing table. When a mobile initiates communication, all the base stations that detect the mobile must use the initial capture

¹Base stations use a received frame to obtain a sample of signal strength. Recent samples weigh more than old samples. The running average becomes invalid if no new frame arrives within a predefined time interval.

3.1 Routing

Routing within Crosspoint involves delivering an IP datagram destined for an active mobile to the base station that currently owns the mobile. Each Crosspoint processor maintains a routing table, which contains one entry for each mobile. A routing entry for a mobile contains both the ownership and reachability information of the mobile. When a datagram destined for a mobile arrives at a Crosspoint processor, the processor retrieves the routing entry that corresponds to the mobile and uses the entry to determine whether the mobile is reachable. If the mobile is unreachable, the processor drops the datagram and sends an ICMP host unreachable message [34] back to the sender. If the mobile is reachable, the processor uses the routing entry to determine the data circuit that leads to the owner base station, and sends the datagram over the circuit to the owner. When the datagram arrives, the owner base station performs the same routing table lookup, learns that the mobile is reachable locally, and delivers the datagram over the wireless interface to the mobile.

Clearly, maintaining correct routing entry for each mobile is essential for achieving optimal routing. The remainder of this section describes how Crosspoint uses protocol messages to establish and maintain routing entry for each mobile, without considering reliability issues such as processor failure or message loss. The reliability issues are the topic of the next section.

3.2 Protocol Overview

The goal of the protocol design is to ensure that at any time, exactly one base station handles a mobile's communication requests. Overlapping areas present a challenge to achieving the goal. When a mobile emits a frame, one or more base stations may receive the frame. Furthermore, a base station that receives the frame has no knowledge of which other base stations have received the same frame. The *initial capture protocol* ensures that only one base captures the mobile when the mobile initiates communication the first time. As the mobile roams, the *handoff protocol* ensures that the ownership of the mobile is passed from one base station to the next. The owner of the mobile uses the *revalidation protocol* to determine whether the mobile is still reachable. Finally, the *recapture protocol* ensures that only one base station captures the mobile when it was previously determined unreachable and initiates communication.

3.3 The Handoff Protocol

A base station uses the handoff protocol to negotiate and transfer the ownership of a mobile. The protocol follows a request-reply interaction between a base station that tries to capture a mobile and the mobile's owner. To capture the mobile, the non-owner base station sends a handoff request to the owner and awaits

However, 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.

Overlapping areas allow a mobile to change from one base station's area to another without losing radio contact. However, overlapping areas increase system design complexity. When a mobile emits a frame in an overlapping area, multiple base stations can receive the frame. To avoid confusion, base stations can use a coding scheme at the radio level [17, 26, 32] or at the frame level [13] to distinguish each other's area. However, the scheme requires a mobile to reconfigure the wireless interface when switching to a new base station. Our approach does not require a mobile to have such capability. Instead, all base stations and mobiles in Crosspoint use the same coded radio signal and frame. Base stations exchange control messages over the Crosspoint interconnect to ensure that only one station handles a mobile's communication requests at any time. A base station discards frames emitted from mobiles that it does not handle.

2.8 Handoff Overview

As a mobile migrates from the area of one base station to the area of another, the base station that currently handles the mobile (refer to as the *owner* of the mobile) must transfer the ownership to the next base station, so that the new base station can start handling the mobile. The process of transferring the ownership of a mobile from one base station to another is known as *handoff*. In a wireless environment, radio signal strength is commonly used as an indicator to determine when to hand off a mobile [37]. The owning base station hands off a mobile to a nearby station that can maintain better radio contact with the mobile.

In Crosspoint, because mobiles do not participate in handoff, base stations use frames emitted from mobiles as a hint for handoff. Each base station monitors the signal strength to a mobile using frames emitted from the mobile. When a base station receives a frame from a mobile that it does not own, the base station forms a handoff request that includes the measurement of the signal strength to the mobile and delivers the request to the mobile's owner. The owner uses a handoff algorithm to process the handoff request. The output of the handoff algorithm is a boolean value that instructs the owner to deny or accept the request. The next section will describe the handoff protocol in detail.

3 Crosspoint Routing and Protocols

Multiple base stations are needed to provide wireless coverage for the entire campus. Therefore, as a mobile roams the campus, base stations must cooperate to transfer the ownership of the mobile from one base station to another. This section describes how Crosspoint processors use protocols to determine the ownership of a mobile and to transfer the ownership of a mobile from one base station to the next. The ownership information is the routing information that Crosspoint processors use to forward datagrams for mobiles.

2.4 The Default Router Address

A single IP address selected from the class-B space is reserved for the wireless interface of every base station. A base station assigns the IP address to its wireless interface during initialization. Each mobile host installs the reserved address as its default router address. Thus, a mobile can use a nearby base station to communicate with nonmobile hosts. Because all base stations use the same IP address, a mobile host need not change its default router when switching from one base station to another.

2.5 Passive Mobiles

Once the networking software of a mobile is configured with its IP address and the default router address, the mobile can access the wireless infrastructure. Unlike other approaches [6, 19, 30] that require each base station to broadcast a *beacon* periodically, and each mobile to install an additional software module that processes the beacon and determines 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 of a mobile transmits and receives packets as if the mobile is associated with a single wireless LAN that covers the entire campus. Thus, a mobile can use conventional network software without modification or addition.

2.6 Mobile Host Detection

Because mobiles are passive, base stations cannot rely on mobiles to send special packets to announce their presence. Thus, a base station must devise a way to detect the presence of mobiles in its area. Observe that computers emit packets when they try to communicate with other computers. Mobile hosts are no exception. Furthermore, mobile hosts tend to initiate communication with nonmobile server computers that are stable and contain resources. By listening in *promiscuous mode*, a base station can monitor frames emitted from mobiles to detect their presence. Moreover, a base station can use an ICMP echo request [34] to elicit transmission from a given mobile when necessary.

2.7 Overlapping Areas

The Crosspoint design assumes that a base station's wireless interface can provide wireless coverage of a geographic region much smaller than a campus. Thus, multiple base stations are needed to provide the wireless coverage for the entire campus. When a roaming mobile is about to move out of the current base station's area, the mobile 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 to maintain radio contact with a base station at all times.

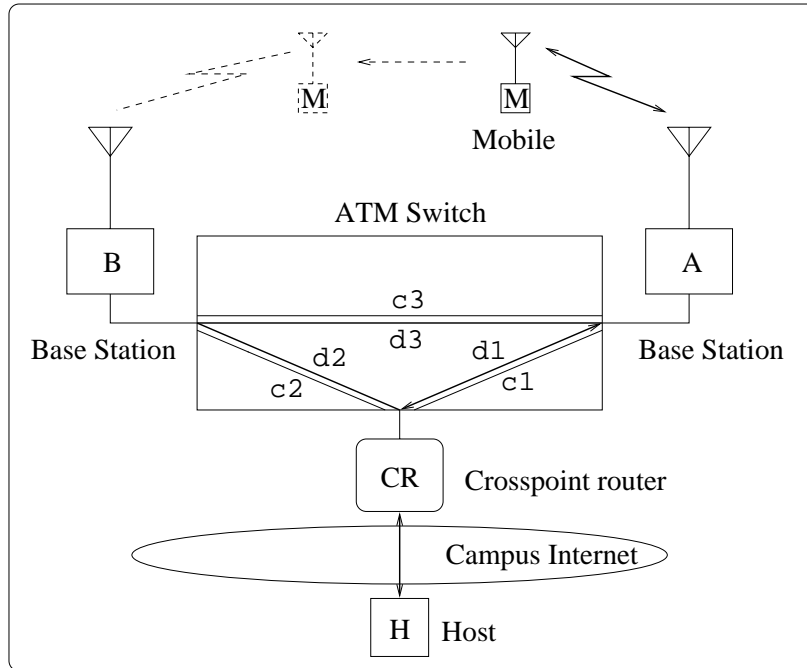


Figure 3: Mobile computing using a Crosspoint wireless network

datagrams pass through data circuit $d1$. When M migrates to the area of base station B , B detects the arrival of M and uses control circuit $c3$ to exchange control messages with A . Once A allows B to capture M , base station B will inform router CR , using control circuit $c2$, to forward the datagrams destined for M over data circuit $d2$, allowing seamless communication between mobile M and host H . Both M and H are unaware of the routing change. The Crosspoint protocol software handles all the details to support seamless mobile communication.

2.3 Addressing

To accommodate all the mobile hosts in a large campus, a class-B IP address space is reserved for the wireless interfaces. Each mobile has an IP address selected from the reserved address space. A mobile does not change address as it switches from one base station to another. As a mobile roams, base stations cooperate to track the mobile. Campus routers perceive the Crosspoint network as a single IP network interconnected to the campus internet using Crosspoint routers. Campus routers forward datagrams destined for mobiles to Crosspoint routers as if the datagrams are destined for nonmobile hosts. Consequently, as mobiles switch from base station to base station, it is unnecessary to propagate IP route changes throughout the campus internet.

2.1 Using ATM as the Crosspoint Interconnect

The current implementation uses a dedicated, high-speed ATM switching network as the Crosspoint interconnect. Because an ATM switch provides each attached processor with dedicated bandwidth, adding a new processor does not decrease the link capacity of others. Thus, additional base stations can be attached to provide greater wireless coverage. If the number of base stations grows beyond the size a switch can accommodate, additional switches can be added as needed. It is feasible, for example, to scale the architecture to many base stations per building on a large campus.

Each Crosspoint processor attached to the ATM network maintains two separate virtual circuits to each of the other processors: one data circuit and one control circuit. The data circuit uses AAL5 (ATM Adaptation Layer 5) [23] and is used for transporting IP datagrams. The control circuit uses raw ATM cells to carry control information. To ensure that control traffic has higher probability of being successfully delivered during congestion, the control circuit is assigned a higher priority than the data circuit. In addition, each control circuit is given a reserved bandwidth to guarantee that a large volume of data traffic does not impede control traffic.

Using two separate circuits ensures that IP datagrams will not be confused with control information because datagrams never travel on circuits used for control, and control packets never travel on circuits used for data. As a result, routing within Crosspoint is more efficient: a Crosspoint processor can transmit an IP datagram over a data circuit without prepending an extra header to mark the type of the transmitted packet. Furthermore, because data circuits use AAL5, which can carry data up to 65535 octets, datagrams from the wireless LANs or from the campus internet will travel through data circuits without being fragmented to smaller IP datagrams.

The disadvantage of maintaining two separate circuits from one Crosspoint processor to each of the other Crosspoint processors is that the number of virtual circuits increases quadratically proportional to the number of Crosspoint processors attached to the ATM network. In theory, ATM supports a maximum number of 16,777,216 (i.e., 2^{24}) virtual circuits at an end-user interface [16]. In practice, the maximum number of virtual circuits supported by an ATM switch depends on the capacity of the switch. For scaling beyond a campus-sized network, a hierarchical architectural design is needed. Alternatively, the IETF Mobile IP design [31] can be used to provide inter-campus mobility support.

2.2 Overview

We use the example network illustrated in Figure 3 to provide an overview on how the design can be used to support wireless mobile communication. In the figure, a Crosspoint router and two base stations are attached to an ATM switch. Mobile host M is communicating with host H , which resides in the campus internet. Base station A and Crosspoint router CR are forwarding IP datagrams between H and M . The

Section 6 describes a prototype implementation of Crosspoint. Section 7 describes related work. Finally, section 8 concludes the paper and discusses future work.

2 Crosspoint Wireless Mobile Internet

Crosspoint combines wireless LAN technology with high-speed switching technologies, such as Asynchronous Transfer Mode (ATM) [21, 22, 36]. The combination provides a wireless communication system with sufficient aggregate bandwidth to handle both data transfer and routing updates. Figure 2 illustrates the architectural design of Crosspoint.

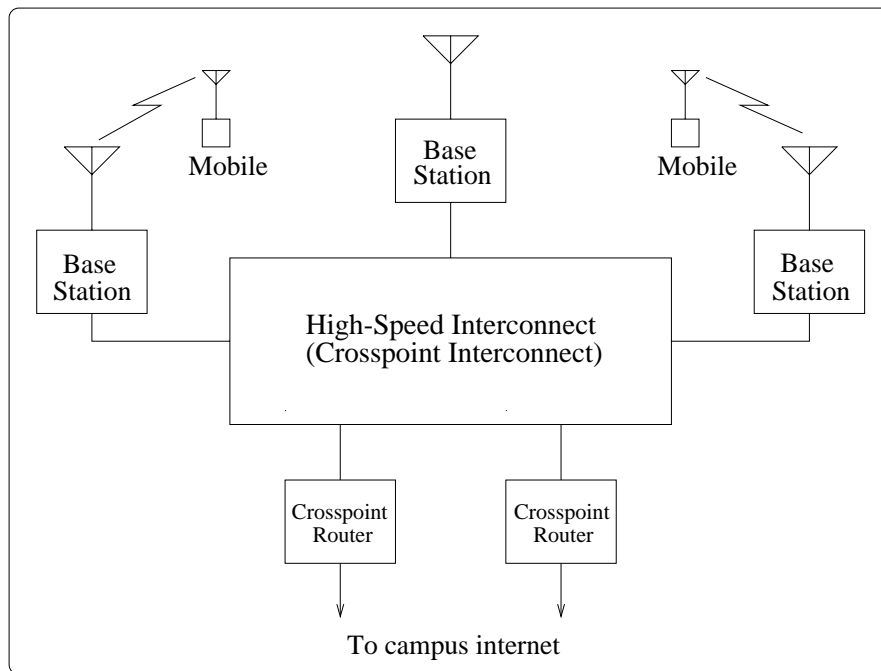


Figure 2: The architectural design of Crosspoint. A high-speed communication fabric interconnects all base stations and Crosspoint routers.

The design uses a scalable, high-speed communication fabric to interconnect all base stations and special purpose routers called *Crosspoint routers*. Base stations provide wireless access for mobile hosts (or mobiles). Crosspoint routers interconnect the Crosspoint network to the campus internet, allowing mobiles to communicate with hosts outside Crosspoint. The high-speed communication fabric provides high-bandwidth, low-latency communication channels among the attached *Crosspoint processors* (i.e., base stations and Crosspoint routers).

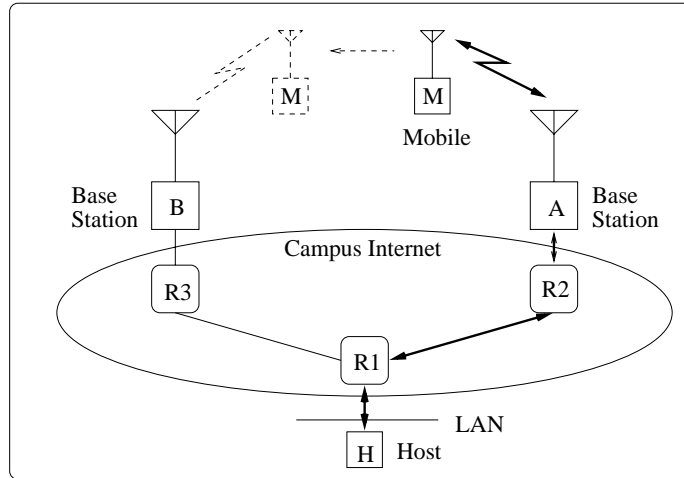


Figure 1: Illustration of a mobile computer communicating with a nonmobile computer in a campus internet that supports wireless mobile communication

M migrates to the coverage area of base station B , B must detect that M has arrived and then propagate a routing update message to allow packets destined for M to be forwarded to itself. To achieve optimal routing, B must propagate the routing update message to all the routers in the campus internet because M could be communicating with an arbitrary set of computers. Note that packets that carry the routing update message compete with data packets for network bandwidth.

The overhead of propagating routing updates is especially apparent in a large university campus where many (e.g., 50,000) mobile computers occupy 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 use their mobile computers to communicate from new locations, affecting not only the mobile computers, but also the nonmobile computers in the campus internet. The situation becomes worse when congestion causes delay or loss of routing updates, forcing data packets to follow nonoptimum paths.

This paper reports our experience in building a campus-sized wireless data communication system called *Crosspoint*. The system is capable of handling a large volume of routing update traffic and provides seamless wireless mobile internetworking without requiring modifications to the networking software on mobile computers, nonmobile computers, or routers in the existing Internet.

The remainder of this paper is organized as follows. Section 2 describes the architectural design of *Crosspoint*. Section 3 explains routing within *Crosspoint* and the protocols that maintain the routing table. Section 4 discusses schemes and protocol extensions that minimize or rectify routing inconsistency. Section 5 considers address binding issues and supporting communication between two mobile hosts.

Crosspoint: A Campus-Sized Wireless Mobile Network

Purdue Technical Report CSD-TR 95-058 (Revision, May 1999)
Douglas E. Comer, John C. Lin and Vincent F. Russo
Department of Computer Sciences
Purdue University
West Lafayette, Indiana 47907, U.S.A.

Abstract

This paper discusses a new approach to support wireless mobile internetworking on a university campus or similar environment. Called *Crosspoint*, the approach combines wireless local-area network technology with high-speed switching technology. The combination provides a wireless communication system with sufficient aggregate bandwidth to support a large volume of control and data traffic. Furthermore, the approach supports optimal routing to each mobile computer without requiring modification of the networking software on mobile computers, nonmobile computers, or routers in the existing Internet. This paper describes the design and implementation of Crosspoint. Through a prototype implementation, we have shown that the approach is feasible.

1 Introduction

Recent advances in personal computing and wireless local-area network (LAN) technologies have resulted in affordable laptop and palmtop computers with wireless networking capability. A portable computer with a wireless LAN adaptor can communicate directly with other wireless computers in the same wireless LAN. To communicate with computers that are far away, a wireless mobile computer uses a nearby *base station*. Normally, a base station is 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 wireless mobile computer with access to other computers at sites around the world.

The wireless interface of a base station can provide wireless coverage for a small geographical area (e.g., approximately 30 to 250 meters in diameter in an office environment). Mobile computers that reside within the area can use radio signals to communicate with the base station. Because a base station can provide wireless coverage for only 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 computer migrates from the area of one base station to the area of another. Consider the example internet illustrated in Figure 1.

In the figure, two base stations, A and B , attach to a campus internet. Mobile computer M is communicating with host H via base station A and two routers, R_1 and R_2 . To maintain network connectivity when