Reducing Router-Crossings in a Mobile Intranet *

Rohit Dube Ibrahim Korpeoglu Satish K. Tripathi
Mobile Computing and Multimedia Laboratory
Institute for Advanced Computing Studies and Department of Computer Science
University of Maryland
College Park, MD 20742
{*rohit,korpe,tripathi*}*@cs.umd.edu*

CS-TR-3735, UMIACS-TR-97-1

Abstract

Current general purpose mobility solutions like Mobile-IP involve multiple *router-crossings* even when the mobile host moves within an intranet from one subnet of a router to another. An environment consisting of a large number of mobile hosts would congest the router causing hosts to experience high latency and jitter. This paper presents a mechanism to eliminate multiple router-crossings in a mobile intranet, which reduces the load on the routers and the hand-off and data latency at the mobile hosts.

1 Introduction

With the increasing popularity of the web and web based applications, traffic from hosts to and from the Internet is going up steadily. This has increased the load on routers connecting campus and building intranets to the Internet Service Providers (ISPs) causing them to become the primary bottleneck in the Internet today.

Availability of Mobile-IP ([Per96a]) implementations is popularizing the use of laptops as *internet-enabled* mobile hosts. Because Mobile-IP relies on some static hosts acting as Home and Foreign Agents to tunnel traffic to and from a mobile host's current location, a data packet making its way to a mobile host crosses the router twice, exasperating the router and increasing the latency and jitter seen by the mobile host. For example in figure 1, a packet destined for the mobile host (MH) from a web server in the Internet, gets routed from the ISP router to the campus-router, which in turn routes the packet to the home subnet of MH. As MH has moved away to a foreign subnet, the packet is picked up by the Home Agent (HA), encapsulated and tunneled to the Foreign Agent (FA) on the current subnet of MH. FA decapsulates the packet and passes it onto MH. Thus the packet not only traverses the campus-router twice, it also traverses the protocol stack up and then down on both HA and FA. Since the mobile host moves between two different subnets on the same router, these traversals can be avoided by short-circuiting at the router.

In a campus or building environment it is very likely that movement of mobile hosts would be restricted to subnets on a single or a small group of routers under the control of one administrative authority. We use this observation to design a mechanism to eliminate the

*This work is supported in part by NSF (grant CCR 9318933), IBM, and Novell.

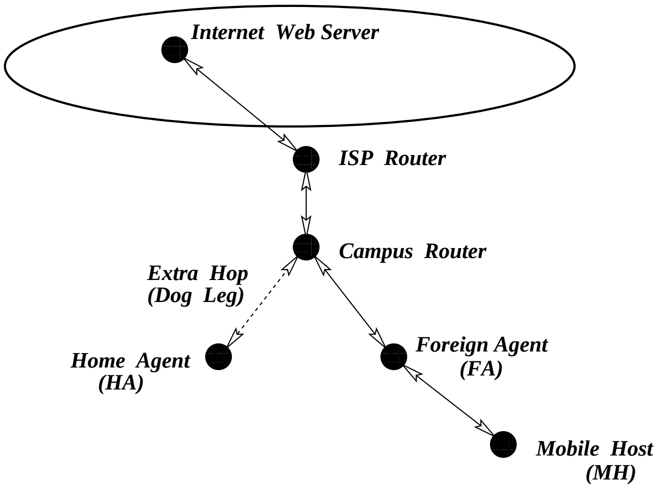

Figure 1: Multiple Router Crossings

stack traversals on HA and FA and the duplicate router-crossings on the campus-router, in the common case of a mobile host moving between subnets of the same router. This is done by co-locating the Home and Foreign Agents of all the subnets of a router, into a single entity on the router. We then extend this mechanism to multiple routers under one administrative domain. As we'll see later in the paper, this technique reduces hand-off and data latency seen by applications running on mobile hosts besides reducing the load on the router.

The rest of the paper is organized as follows: section 2 presents related work, section 3 describes our architecture and section 4 discusses our implementation experience. Section 5 analytically compares our approach with conventional solutions. Finally, section 6 presents our conclusion and discusses some ideas for future research.

2 Related Work

The Mobile-IP specification [Per96a] allows mobile hosts to move between subnets by maintaining a forwarding pointer at the mobile host's Home Agent. Ordinarily, every time the host changes its subnet (and hence its Foreign Agent) a registration request is sent back to the Home Agent. All data packets are then tunneled from the Home Agent to the new Foreign Agent. The base Mobile-IP protocol suffers from two performance problems: high *handoff latency* due to the registration messages exchanged between the Foreign and the Home Agents and high *data latency* due to the indirect path taken by data packets as described previously. The indirect path also increases congestion at the already overloaded routers.

Caceres and Padmanabhan [CP96] describe a method by which wireless machines moving between base-stations on the same subnet use proxy and gratuitous arps [Ste94] to quietly accomplish a hand-off without going through the Home Agent. For movement between subnets in the same administrative domain, a hierarchy of Foreign Agents similar to one described in [Per96b] is suggested, where the hand-off latency following a move is decreased by using hierarchical Foreign Agents which shield the remote Home Agent from the knowledge of a local move.

Blackwell et al [B+94], Myles [MJP95], Johnson and Perkins [JP96], [PJ96] cache addresses of Foreign Agents on correspondent hosts to tunnel packets directly to a mobile host's Foreign

Agent (i.e. without going through the Home Agent). Hosts which do not implement the *FA-cache* protocol have to take the longer route to reach the mobile hosts. In any case, data packets incur the extra hop through the Foreign Agent, even if the mobile host was directly visible to the router. Bhagwat and Perkins [BP93] use IP's loose source route option to achieve the same, but suffers from the disadvantage of slower and sometimes incorrect processing of the options on the intermediate routers. Perkins and Luo [PL95] use explicit assignment of new care-of IP addresses, local to the current point of attachment to effect mobility as well as as a direct data path. However, this requires the availability of DHCP servers and forces mobile hosts to implement a Foreign Agent.

In this paper, we describe a mechanism which reduces both handoff and data latency for the common case of movement restricted to a campus. The reduced latencies follow from the co-location of the Home Agent and the Foreign Agent and their placement on the router connecting the LANs to the Internet. Unlike the approaces discussed above, our implementation requires minimal support from mobile hosts and none from any static hosts. The burden of supporting mobility lies mostly on the routers which are the nodes worst effected by the sub-optimal routes.

3 Mobile Intranet Architecture

We observe that handoff latencies stem from the registration packet exchange between FA and HA (figure 1). This exchange can be removed if FA and HA were co-located. Since FA is necessarily on the foreign subnet and HA is necessarily on the home subnet, a co-located FA and HA can exist only on a node which is both on the home and the foreign subnets. In most LAN configurations, there is only one such entity: the router. Besides reducing the handoff latency, co-locating HA and FA and placing them on the router has the effect of reducing data latencies as any packets destined for MH can be routed directly onto the MHs current subnet. Taken together, this amounts to reducing multiple router-crossings and stack traversals mentioned earlier, for *all* packets.

In the following sections we describe the architecture in detail. The addressing scheme used by the architecture is discussed first, followed by the protocol operation.

3.1 The Addressing Scheme

In figure 2, correspondent hosts CH1, CH2, CH3 and CH4 are on subnets X.Y.A, X.Y.B, X.Y.C and X.Y.D respectively. All the mobile hosts (eg MH1) serviced by CR1 are on a virtual subnet X.Y.M Similarly, mobile hosts (eg MH2) serviced by CR2 are on virtual subnet X.Y.N. In the absence of host routes, R routes packets destined for subnet M to CR1 and those for subnet N to CR2.

3.2 Mobility within a Router

MH1 which is physically on subnet A, converses with other machines on virtual subnet M directly, using MAC addresses. If the target mobile host is on the same physical subnet as MH1, then packets can be exchanged over the wire without any additional support. However if the target is on a different physical subnet, then CR1 acts as a bridge and relays packets back and forth. This is achieved by making CR1 proxy arp for mobile hosts on subnet M.

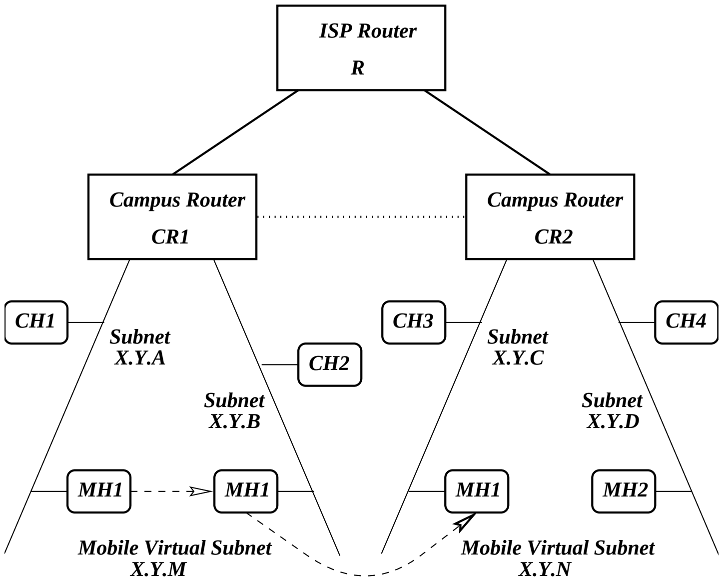

Figure 2: The Intranet Architecture

The default route on MH1 points to CR1. Since MH1 and CH1 are on different subnets, traffic between them has to go through CR1. This is wasteful but can only be optimized by modifying the stack on CH1 to send ethernet frames directly to MH1 (a concept similar to the FA-cache discussed in section 2). We do not implement such a cache and therefore do not discuss this further in the paper.

When MH1 moves from subnet A to subnet B, it broadcasts a *greet* message. CR1 picks up the message and updates the outgoing interface for the routing table entry of MH1 to point to the B subnet. It then sends back an acknowledgment message to MH1 indicating that the handoff has been completed. CR1 also sends out a gratuitous arp on subnet A annulling the arp entries on the hosts belonging to subnet M. Thereafter CR1 proxy arps for MH1 on the A subnet.

Now, consider data flowing from CH4 to MH1. A data packet first goes to R (ignore the dotted line between CR1 and CR2 for now). R looks up the route for MH1 and since it does not have a host route for MH1, it forwards the packet to CR1. CR1 looks up its route table and realizes that the next hop interface for MH1 is the interface connected to subnet B. It therefore sends the packet out on the wire from where MH1 receives it. No packets are sent to MH1s former location: subnet A. The return path to CH4 is straight-forward and follows the usual internet routing mechanism.

3.3 Mobility across Routers

If MH1 now moves to subnet C, it sends out a greet message as before. CR2 picks up this message and realizes that it is from a host on the M mobile subnet owned by CR1. CR2 sends a message to R, which creates a host route pointing to CR2 for MH1. R then sends a message to CR1, which results in CR1 updating the routing table entry corresponding to MH1 to point to R. CR1 sends out a gratuitous arp on the B net to anull any cached arp entries. R sends back an acknowledgment to CR2 which then creates a route table entry for MH1 pointing to the C subnet. Finally CR2 sends an acknowledgment to MH1.

If MH1 moves to subnet D, the protocol exchange is similar to the move from subnet A to subnet B. No node beyond CR2 need be involved, and the handoff finishes quickly.

Notice, that we have so far assumed a tree relationship amongst the routers. Routers are often put on a high-speed backbone within a campus for better performance. If such a backbone exists (the dotted line in figure 2 between CR1 and CR2), the extra hops from CR1 to R and then to CR2 can be avoided for traffic contained within the campus: traffic from CH1 to MH1 for example, when MH1 is on subnets C or D.

Mobility between routers across campuses can be handled by either extending the hierarchy beyond the ISP router or by using the regular version of Mobile-IP, with the campus routers acting as Home Agents. The former would suffer from administrative problems and the later from optimization problems.

3.4 Wireless Mobile Hosts

In the above sections, we do not explicitly discuss the physical medium and treat all machines as if they were on a regular ethernet LAN. This is not a problem for wireless mobile hosts as they act like ethernet connected machines with the base-station acting as a bridge. Mobility between connection points on the same subnet (which is not an issue for regular ethernet) is handled as in [CP96]: base-stations proxy for the mobile hosts and forward layer 2 frames to and from them appropriately. Since the nature of the physical link is invisible to the network layer, mobility between base-stations on different subnets can be handled just as with regular ethernet as described previously.

4 Implementation

We implemented the first rung of the intranet hierarchy (single router case) using the testbed shown in figure 3. The router (graf) is an Intel pentium machine running the 4.4BSD *ip forwarding* code. MH is an Intel 486 (mobile host) with an IP address on the virtual mobile subnet. CH1 and CH2 are correspondent hosts on physical subnets 46 and 126 respectively. All hosts run BSD/OS 2.1 [MBKQ96] and are connected via 10Mbps ethernet. The following subsections give a brief sketch of the implementation issues and present experimental results which corroborate our approach.

4.1 Handoff

The mobile hosts connect to the network using PCMCIA ethernet cards. These hosts implement a *trigger* protocol which is activated whenever the ethernet card is re-inserted into the PCMCIA slot or the RJ45 jack (or T-connector) is re-connected to the ethernet segment. The trigger protocol sends a greet message with its IP address (on the virtual subnet), its MAC address and the subnet to which it was last connected. The router handling the subnet receives the message, does the route table modifications (described previously) and sends back an acknowledgment message to the mobile host indicating that the handoff has completed. It is possible to implement handoff using periodically broadcast *beacons* from the router, but this was left out to prevent the router from generating any more messages than it ordinarily does.

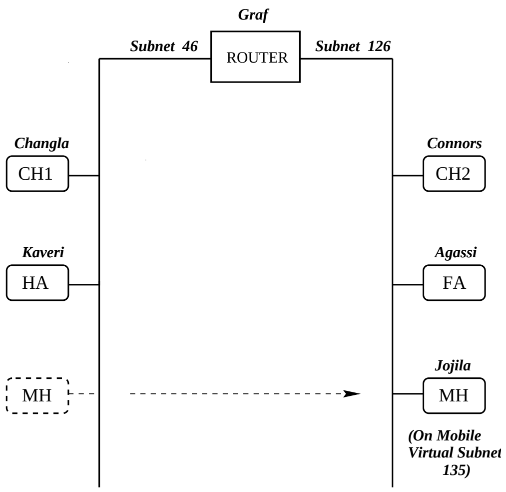

Figure 3: Experimental Setup

Case	Communication Direction	Conventional Routing Path	Reduced Routing Path
1	CH1 $\longrightarrow$ MH (at subnet 46)	CH1-MH	CH1-R-MH
2	CH2 $\longrightarrow$ MH (at subnet 46)	CH2-R-HA-MH	CH2-R-MH
3	CH1 $\longrightarrow$ MH (at subnet 126)	CH1-HA-R-FA-MH	CH1-R-MH
4	CH2 $\longrightarrow$ MH (at subnet 126)	CH2-R-HA-R-FA-MH	CH2-R-MH

Table 1: Data Path Combinations

4.2 Data Path and the Router

For our prototype, we implement the additional router functionality in user space using the Berkeley Packet Filter (bpf) [MJ93], [WS94]. A user level daemon on the router opens up a bpf device and uses it to sniff ethernet frames for and from hosts on the virtual subnet. If the daemon gets a frame containing a greet message from a mobile host, it does the appropriate processing and updates its routing table. Other packets from a mobile host are handled by the regular forwarding code on the router. Data packet for mobile hosts are processed as follows: the daemon consults its routing tables and a) drops the packet if it originated at a mobile host and is intended for a mobile host on the same physical subnet (in this case the mobile host would get the packet directly), and b) forwards the packet through the appropriate interface listed in the route table entry for all other cases.

4.3 Experimental Results

We measure the latencies observed at a mobile host using the testbed described above, and compare the results obtained using conventional routing and our *reduced* routing protocol. Since we implemented the reduced protocol over ethernet LANs, handoff latency (which would have been important for wireless hosts) is not relevant. Therefore, only data latency between correspondent hosts and mobile hosts is measured with microsecond accuracy.

Since administrative constraints prevent us from hooking up our router to the rest of the campus network, we restrict our experiments to subnets directly connected to graf. Based on the relative placement of the correspondent and mobile hosts, four interesting combinations are identified and are listed in table 1. For each combination we measure the latency of data packets of size 256, 512, 768 and 1024 bytes, between the correspondent host and the mobile host. These latencies along with 95% confidence intervals for 60 measurements are shown in the graphs below. Note that only the route between the correspondent host and mobile host is of interest as the reverse path is optimal for both Mobile-IP and our reduced protocol.

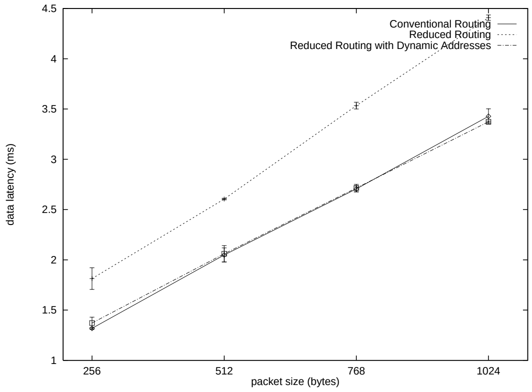

Figure 4: [Case 1] CH1 – MH (at subnet 46)

For conventional Mobile-IP, case 1 corresponds to MH at its home subnet communicating with a correspondent host belonging to the same physical subnet. There is no truly equivalent case for our architecture as there is no concept of a home subnet within a campus. Since MH and CH1 have different IP addresses, packets between them are forced to go through the router. The latencies are shown in figure 4 and are lower for Mobile-IP as expected. The longer route is purely an artifact of the mobile virtual subnet and can be eliminated by assigning an alias IP address on subnet 46 to MH. As shown in figure 4, the latencies observed with such a dynamic address assignment scheme are nearly the same as Mobile-IP.

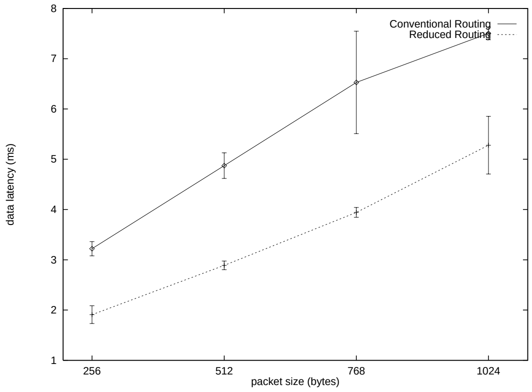

Figure 5: [Case 2] CH2 – MH (at subnet 46)

Case 2 does not arise in base Mobile-IP but may for the modified scheme [Mon96] where all traffic between MH and CH2 is tunneled though the Home Agent (HA). In the reduced mobility scheme, CH2 and MH1 communicate through the router without involving HA, saving the extra hop. The latencies are compared in figure 5 and are lower for the reduced protocol.

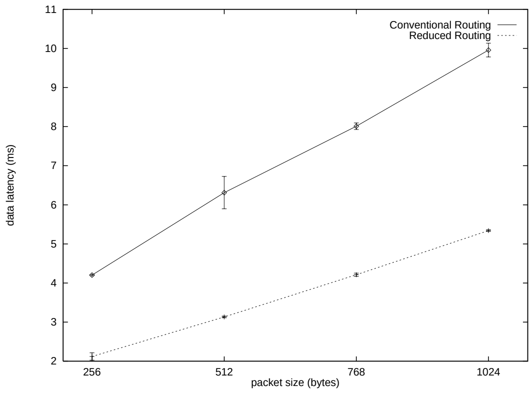

Figure 6: [Case 3] CH1 – MH (at subnet 126)

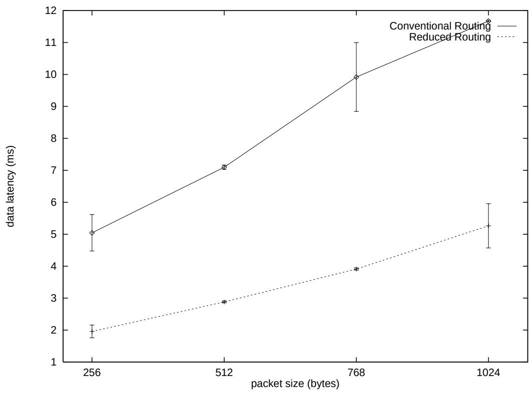

Figure 7: [Case 4] CH2 – MH (at subnet 126)

For conventional Mobile-IP, cases 3 and 4 correspond to a mobile host in a foreign subnet communicating with a correspondent host on its home and foreign subnets respectively. The former involves four hops from CH1 to MH and the later five from CH2 to MH, unlike our scheme which requires one hop each way. Figures 6 and 7 show the difference in the latencies. As expected, the reduced protocol exhibits the maximum improvement over Mobile-IP for these two cases due to the large difference in the number of hops.

Even though we implemented our protocol in user space, the improvements achieved are significant: latencies drop from 12 milliseconds to 5 milliseconds for packets of size 1024 bytes. This reduced latency coupled with reduced router load puts forth a strong case for in-kernel support of the mobility scheme.

5 Theoretical Analysis

Since financial constraints prevent us from doing an implementation oriented scalability study, we analytically compare the performance of the reduced protocol with conventional routing in the presence of multiple mobile hosts. The question we seek to answer is how does the data latency and the router load depend on the number of mobile hosts. Note that the goal is not to build a precise analytical model of a router, but to use a simple model to analyze how mobile hosts affect a router.

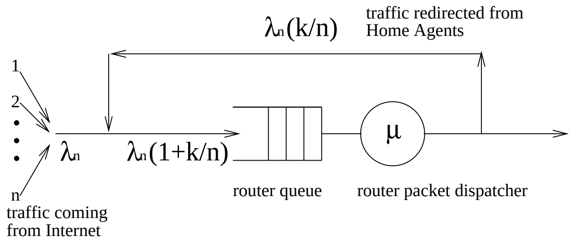

Figure 8: Router Model

We assume there are n hosts that are connected to the rest of the Internet via the router of interest. Let m of these n hosts be mobile hosts, and let k of these m mobile hosts be in foreign subnets ($m-k$ mobile hosts are in their home subnets). We make the simplifying assumption that each host sees the same traffic stream coming in through the router. Additionally, we assume that the packet interarrival times and packet lengths are exponentially distributed, so is the service time for each packet through the router. λ_n denotes the total average arrival rate of these n streams (packets/second), and μ denotes the average service rate (packets/second) of the router assuming a shared input buffer and FCFS service discipline.

If conventional Mobile-IP is used, the packets destined for mobile hosts away from their home subnets, first go to their Home Agents and then to their current subnet through the router. Therefore, we have feedback traffic as shown in figure 8, which is proportional to the number of mobile hosts away from their home subnets. For the reduced routing protocol there is no such feedback traffic, since the router is always aware of the current location of each mobile host (and is crossed only once for each packet).

By applying Little's formula [Tri82] to the model in figure 8, the average delay (D_{CR_s}) at the router for a packet routed using conventional routing, and destined for a static host or a mobile host in its home subnet turns out to be:

$$D_{CR_s} = \frac{1}{\mu - (1+\frac{k}{n})\lambda_n}, \quad 0 \le k \le n \tag{1}$$

The average delay (D_{CR_k}) for a packet destined for a mobile host in a foreign subnet is $2 \times D_{CR_s}$, since the router is crossed twice. (We are considering delay at the router only. The stack traversals at the Home and Foreign agents and the wire delay times make the total delay much worse.)

$$D_{CR_k} = \frac{2}{\mu - (1+\frac{k}{n})\lambda_n}, \quad 0 \le k \le n \tag{2}$$

For reduced routing there is no feedback traffic. Hence the average delay (D_{RR}) is the same for all hosts and is given by:

$$D_{RR} = \frac{1}{\mu - \lambda_n} \tag{3}$$

As an example, consider the testbed described in section 4. Since $k = 1$ and n is comparatively large, D_{CR_k} is twice D_{RR}. This is substantial improvement at the router which increases as k becomes more significant as compared to n.

Equations 1 and 2 show that the data latency in conventional routing increases rapidly with the ratio of mobile hosts away from their home subnets to the total number of hosts serviced by the router ($\frac{k}{n}$). When the total load λ_n increases, the delay becomes more sensitive to the ratio. However, for reduced routing, the data latency is insensitive to the presence of mobile hosts, as can be seen from equation 3, and is the same for all hosts.

6 Conclusions and Future Work

Ingress and egress routers are the ideal locations for placing mobility functionality as every packet to and from an end host is seen at these routers. In this paper, we described a mechanism which places such functionality at the ingress router to a campus supporting large scale mobility within its subnets. Our mobile intranet architecture improves both handoff and data latency observed at a host when it moves from one subnet to another within the campus/intranet. The routes obtained are optimal and the same as those achieved by the FA-cache approaches, but at a greatly reduced administrative cost as now a campus no longer has to rely on hosts implementing the FA-cache.

Cheshire and Baker [CB96] survey several approaches towards providing mobility, the fundamental assumption being that mobility should be transparent to routers. Based on our implementation experience and experimental results, we argue that making routers explicitly aware of mobility is a more scalable and manageable approach. The elimination of multiple router-crossings reduces the traffic flowing through the router and justifies the placement of additional functionality on them.

Having verified the feasibility of our approach with an implementation, we would like to carry out a simulation based scalability analysis for multiple campuses. Optimizing for mobility at a foreign campus, and security which has not been discussed in this paper would then become an issue and would have to be dealt with in a scalable way. Appropriate fixes to the protocols and the architecture would have to be made at that point. We would also like to extend our implementation to multiple routers and to wireless hosts, so that a complete and functional mobile infrastructure can be installed in the campus.

Acknowledgments

We thank NSF (grant CCR 9318933), Novell, and IBM for supporting this work. We also thank Sambit Sahu and Cynthia Rais for their comments on earlier drafts of this paper.

References

[B+94] T. Blackwell et al. Secure Short-Cut Routing for Mobile IP. In *USENIX Summer Technical Conference Proceedings*, June 1994.

[BP93] P. Bhagwat and C.E. Perkins. A Mobile Networking System based on Internet Protocol (IP). In *USENIX Symposium on Mobile and Location Independent Computing Proceedings*, August 1993.

[CB96] S. Cheshire and M. Baker. Internet Mobility 4x4. In *SIGCOMM Conference Proceedings*, August 1996.

[CP96] R. Caceres and V.N. Padmanabhan. Fast and Scalable Handoffs for Wireless Internetworks. In *MOBICOM Conference Proceedings*, November 1996.

[JP96] D.B. Johnson and C.E. Perkins. Route Optimization in Mobile-IP, November 1996. Internet Draft, Work in Progress.

[MBKQ96] M.K. McKusick, K. Bostic, M.J. Karels, and J.R. Quarterman. *The Design and Implementation of the 4.4BSD Operating System.* Addison-Wesley, 1996.

[MJ93] S. McCanne and V. Jacobson. The BSD Packet Filter: A New Architecture for User-level Packet Capture. In *USENIX Winter Technical Conference Proceedings*, January 1993.

[MJP95] A. Myles, D.B. Johnson, and C.E. Perkins. A Mobile Host Protocol Supporting route Optimization and Authentication. *IEEE JSAC*, 13(5), June 1995.

[Mon96] G. Montenegro. Bi-directional Tunneling for Mobile-IP, September 1996. Internet Draft, Work in Progress.

[Per96a] C.E. Perkins, editor. *IP Mobility Support.* Internet Engineering Task Force, October 1996. RFC 2002.

[Per96b] C.E. Perkins. Mobile-IP Local Registration with Hierarchical Foreign Agents, February 1996. Internet Draft, Work in Progress.

[PJ96] C.E. Perkins and D.B. Johnson. Mobility Support in IPv6. In *MOBICOM Conference Proceedings*, November 1996.

[PL95] C.E. Perkins and K. Luo. Using DHCP with computers that move. *Wireless Networks*, 1(3), 1995.

[Ste94] W.R. Stevens. *TCP/IP Illustrated, Volume 1.* Addison-Wesley, 1994.

[Tri82] K.S. Trivedi. *Probability & Statistics with Reliability, Queuing, and Computer Science Applications.* Prentice-Hall, 1982.

[WS94] G.R. Wright and W.R. Stevens. *TCP/IP Illustrated, Volume 2.* Addison-Wesley, 1994.