Abstract

Next-generation wireless network standards are currently being defined. The access network architectures have several specialized components tailored for their respective wireless link technologies, even though the services provided by these different wireless networks are fairly similar. In this article we propose a homogeneous IP-based network as a common access network for the different wireless technologies. The IP-based access network uses the Internet standard, Mobile IP, to support macro-mobility of mobile hosts, and HAWAII to support micro-mobility and paging functionality of current wireless networks. We also illustrate how the proposed IP-based solution can interwork with existing infrastructure so that deployment can be incremental.

IP-Based Access Network Infrastructure for Next-Generation Wireless Data Networks

Ramachandran Ramjee, Thomas F. La Porta, Luca Salgarelli, Sandra Thuel, and Kannan Varadhan, Bell Labs, Lucent Technologies
Li Li, Cornell University

Traffic on the Internet is growing exponentially due to an increased subscriber base and new applications. Wide area wireless networks are also experiencing rapid growth in terms of subscribers. Currently, there are many efforts underway to provide data services on wireless access networks.

The Internet Protocol (IP) is the dominant internetworking protocol in operation today. The logical choice for a networking protocol for wireless data networks is also IP for several reasons. First, by using an IP-based network, applications written for wired data networks can operate on wireless networks. Second, to defray cost, integrated wireless and wireline networks can be built and managed. Third, advances on IP technology, such as IP telephony and quality of service (QoS), may be directly applied to the wireless networks. This will enable wireless networks based on IP to provide voice service as well as data services, thus allowing them to tap into the vast subscriber base of cellular voice customers.

We believe all mobility-related functionality should be handled at the IP (network) layer. This enables the deployment of a homogeneous, IP-based wireless access network that is independent of the different wireless interfaces. Only wireless-link-specific processing is relegated to the base stations. We achieve this by extending the IP layer software running in routers and base stations in the access network.

We adopt a domain-based division of the IP mobility protocols. One of the motivations for our domain-based approach hinges on the assumption that most mobility is local to a domain. In particular, most user mobility is typically contained within an administrative domain of the network. Since an administrative domain is under the control of a single authority, it is possible to incorporate special support for mobility in the infrastructure. This domain-based management approach is similar to the division of existing routing protocols into intradomain routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) and interdomain routing protocols such as Border Gateway Protocol (BGP). Another recent solution for mobility management, Cellular IP [1], also adopts a domain-based approach. However, while network elements in the Cellular IP domain are specialized for mobility management, our solution augments regular IP routers with mobility support so that these routers can also be used to route other wired IP traffic as well.

In this article we present the design of an IP-based access network infrastructure for next-generation wireless networks. In the access network we use the Internet standard, Mobile IP, as the interdomain protocol to support macro-mobility; we use an extension of Mobile IP, called HAWAII, as the intradomain protocol for supporting micro-mobility and paging functionalities. The HAWAII protocol results in less disruption to user traffic during handoff and fewer updates to the home agent than Mobile IP. HAWAII's paging support allows for efficient battery consumption at the mobile host. Furthermore, since HAWAII allows mobile hosts to retain their network address while moving within a domain, QoS support is simplified. Finally, we also illustrate how the proposed IP-based solution can interwork with existing wireless infrastructure for incremental deployment.

The remainder of the article is organized as follows. We overview the architectures of emerging wide-area wireless data networks, illustrating architectures from both European and North American standards. We also present an IP-based wireless access network architecture using Mobile IP and HAWAII, an enhancement of Mobile IP. We compare the performance of HAWAII with Mobile IP. We illustrate how our IP-based architecture can be part of the next-generation wide-area wireless data network architecture. We then present our conclusions.

Emerging Wireless Data Network Architectures

Here we will overview two wireless data network architectures currently being defined.

GPRS

General Packet Radio Service (GPRS) [2] is being defined by the European Telecommunications Standards Institute (ETSI) to provide packet data service using Global System for Mobile Communications (GSM) cellular networks. A high-level dia-

1070-9916/00/$10.00 © 2000 IEEE

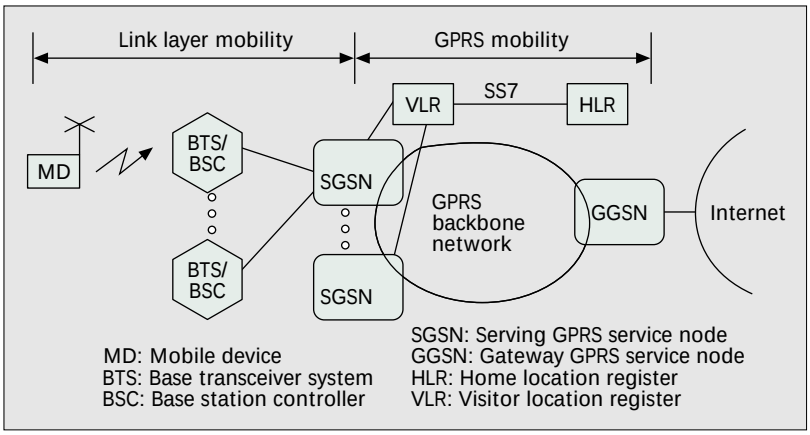

Figure 1. *GPRS network architecture.*

gram of a GPRS network is shown in Fig. 1. GPRS uses a combination of link-layer and newly defined higher-layer techniques for mobility management.

On the air interface, GPRS supports registration, authentication, paging, and handoff (called *cell reselection*), as well as procedures for channel access to transmit data packets. GPRS allows the mobile host to operate in two distinct states: an active state where the network knows the location of the mobile host's current base station, and a standby state where the network knows only the approximate location of the user, such as a set of base stations, called the *paging area*, in which the mobile host resides. One of the motivations for defining the standby state is to reduce the host's battery power consumption by allowing the mobile host to only notify the network when it moves out of the paging area. If data packets for a mobile host in standby state arrive at the wireless access network, the serving GPRS service node (SGSN) pages the mobile host in its paging area to determine the mobile host's current base station before delivering the data packets.

In the backbone network, GPRS defines a new tunneling protocol built on top of an IP network, called the Generic Tunneling Protocol (GTP), to handle device mobility, and support registration and authentication procedures. Data packets flowing through the tunnel are encapsulated with an outer GTP/UDP/IP header. This adds 48 bytes of header overhead to each data packet, which is substantial for voice-over-IP applications that transmit data packets with a small payload. GPRS also defines a QoS profile for each user with attributes for precedence, delay, reliability, and peak and mean throughput classes. However, the drawback of defining GPRS-specific QoS support mechanisms is that advances in IP QoS support such as integrated [3] and differentiated [4] services may not be directly applicable.

In Fig. 1, the air interface protocols from the mobile device are terminated at the base terminal station and base station controllers (BTS/BSCs) (shown as a single box for simplicity). The GTP tunnels extend between the two GPRS gateway routers: the SGSN terminates one end of the tunnel and directs packets to the proper BTS/BSC using link layer protocols; the gateway GSN (GGSN) terminates the other end of the GPRS tunnel and is a gateway to the Internet. As a device moves between SGSNs, new GTP tunnels are established to manage mobility. As a device moves between BTSs/BSCs on a single SGSN, handoffs are handled at the link layer.[1]

GPRS reuses the same infrastructure deployed for GSM in order to support authentication, registration, and roaming. In particular, each SGSN is connected to a visitor location register (VLR), which holds a temporary database of the users currently attached to it. A permanent database of registered users is kept in the home location register (HLR), together with a pointer to their current VLRs. Whenever a new user has to be authenticated, the VLR contacts the user's HLR. The HLR replies to the VLR with authentication information which is composed of a set of random challenges and their corresponding responses, obtained with the use of a secret key that the HLR shares with the user. By sending the challenges to the user and comparing its responses with those obtained from the HLR, the VLR performs user authentication. Similarly, for cyphering between the SGSN and the user, the HLR can send to the VLR encryption keys, obtained from the same secret key known only to the user and HLR.

CDMA

Code-division multiple access (CDMA) networks use a combination of link-layer and IP-layer techniques to manage mobility. A simplified view of a CDMA network is shown in Fig. 2. CDMA networks define an air interface that performs similar functions to GPRS networks: registration, authentication, paging, handoff, and channel access. The network protocols defined for CDMA data networks are based on IP.

One important difference between GPRS and CDMA networks is that in CDMA networks a mobile device may communicate with more than one base station during a soft handoff, thereby transmitting duplicate data frames and increasing the probability of the correct reception of user data. As shown in Fig. 2 the duplicate data frames are received by a special network element, called the *frame selector*, which forwards the data frame with the highest probability of being uncorrupted and discards the rest. These frames are received by an interworking function (IWF) that reassembles the frames into IP packets that are sent to the Internet. For roaming, the IWF may act as a Mobile IP foreign agent.

For movement between base stations attached to the same frame selector, mobility is managed by link-layer techniques. For mobility between frame selectors, mobility could be managed using Mobile IP, which is discussed later. Note that in

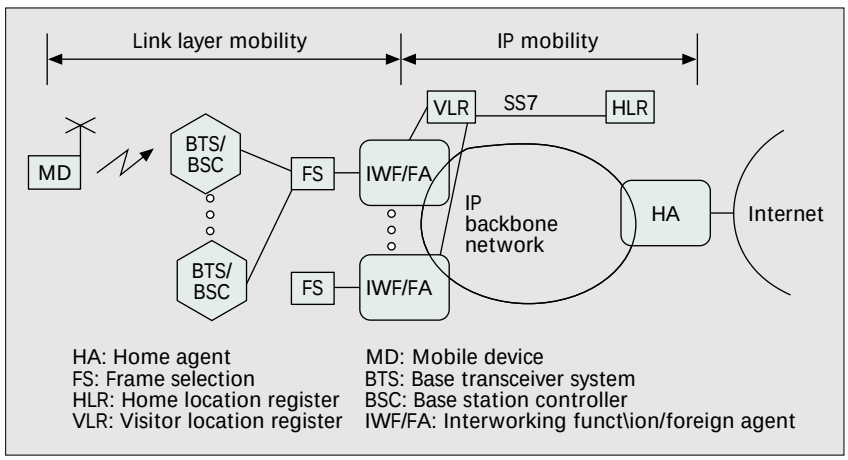

Figure 2. *CDMA network architecture.*

[1] *A recent proposal advocates the establishment of another GTP tunnel between the SGSN and BTS/BSC so that inter-BSC handoffs are handled similarly to inter-SGSN handoffs.*

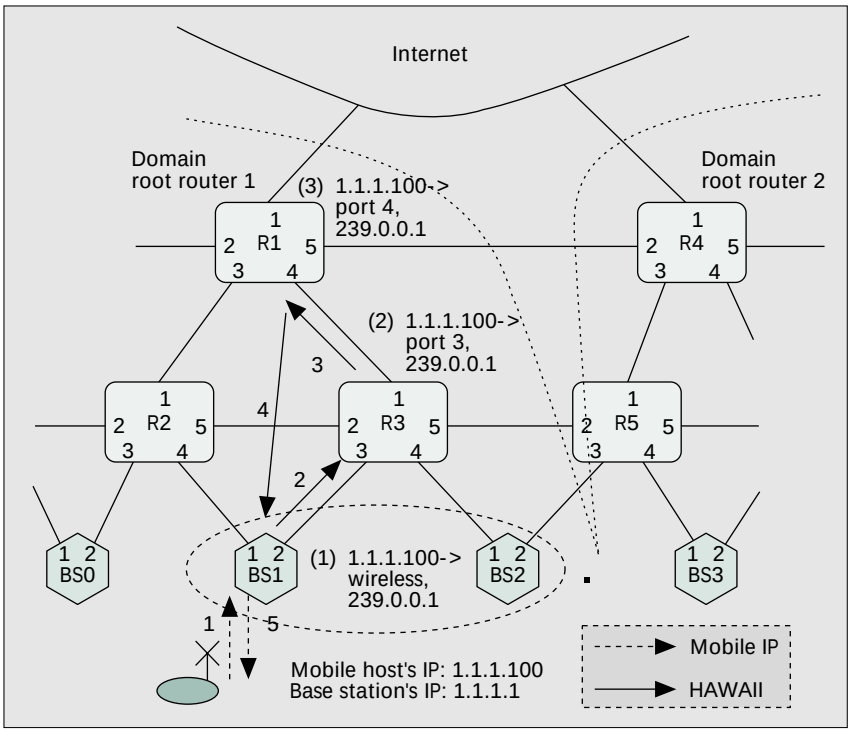

Figure 3. *Power-up.*

this case data packets between the home agent and the foreign agent are encapsulated using an IP-in- IP tunnel. This adds 20 bytes header overhead to each data packet.

CDMA networks use HLR/VLR mechanisms similar to GPRS[2] for supporting user authentication, registration, and roaming. In addition, CDMA networks use authentication procedures defined for Mobile IP.

An IP-Based Wireless Data Network

As illustrated earlier, while IP forms the basis of the backbone network in each architecture, there are still several specialized components and procedures for mobility support in the access network. We envision the next-generation wireless access network as a pure IP-based network, where base stations will be IP addressable entities. In the remainder of this section we describe our IP-based architecture assuming that the base station is IP addressable. We will see later how one may substitute the base station by the appropriate mobile host's next-hop IP entity when mapping the architecture to current wireless access networks.

Mobile IP

Mobile IP is the current standard for supporting macro-mobility in IP networks [5]. Mobile IP defines two entities to provide mobility support: a *home agent* and a *foreign agent*. The home agent is statically assigned to the mobile host based on the permanent home IP address of the mobile host. The foreign agent is assigned to the mobile host based on its current location. The foreign agent has associated with it an IP address called the *care-of* address. Packets destined for a mobile host are intercepted by the home agent, and tunneled, perhaps using IP inside IP, to the foreign agent using the care-of address. The foreign agent decapsulates the packets and forwards them directly to the mobile host. Therefore, the foreign agent is the IP entity closest to the mobile host. In wireless networks this will be a base station, or a router attached directly to a base station like the IWF of a CDMA network.

Mobile IP mandates the authentication of each signaling message, to prevent malicious users from setting up unauthorized tunnels. For authentication purposes, security associations exist between users and their home agents, home agents and foreign agents, and users and foreign agents. Such security associations can be statically configured, by distributing permanent keys to the interested parties. However, the static configuration of each security association can lead to severe scalability problems, particularly in wide-area networks. To avoid this, the Internet Engineering Task Force (IETF) is standardizing protocols and architectures to permit authentication, authorization, and accounting (AAA) [6] servers to distribute short-lived authentication keys to Mobile IP nodes. In this case security associations are configured dynamically by AAA servers, and last only for the duration of a single session.

Mobile IP provides a good framework for allowing users to roam outside their home networks without disruption to their applications. However, it was not designed specifically for wide-area wireless networks or to manage micro-mobility. As such, it has several limitations when applied to next-generation wireless network architectures where handoffs across base stations will be handled at the IP layer.[3]

First, Mobile IP treats all forms of mobility uniformly; therefore, a user moving a short distance, perhaps between two base stations, uses the same mechanisms as another user registering from a remote domain. This entails changing the IP care-of address of the mobile host and notifying the home agent of the movement. Because these movements may be frequent, the overhead of these notifications is a concern. Also, this may cause significant disruption (loss and delay) to user traffic as a handoff occurs. Furthermore, the tunneling of data packets results in nonoptimal routing and header overheads.

Second, Mobile IP does not support paging. Paging facilitates efficient power management at the mobile host by allowing the host to update the network less frequently at the cost of providing the network with only approximate location information. In Mobile IP the mobile host is expected to update the network on every move. This results in excessive battery power consumption, which is unacceptable for wide-area wireless devices.

Finally, there has recently been tremendous interest in supporting QoS in the Internet through the use of differentiated [4] and integrated services [3]. A mobile host using Mobile IP acquires a new care-of address on every handoff from one base station to another. This would trigger the establishment of new QoS reservations from the home agent to the mobile host even though most of the path between the home agent and the mobile host is unchanged, as is likely to be the case for local mobility within a domain.

In summary, while Mobile IP should be the basis for the mobility protocol in wide-area wireless data networks, it has

[2] *The actual protocol used is IS41 in CDMA vs. MAP for GSM.*

[3] *Note that the way Mobile IP is used in the current CDMA networks, handoffs across IWFs will not be as frequent and some of the limitations described may not apply.*

several limitations when applied to wide area wireless networks with high-mobility users that may require QoS. We therefore extend Mobile IP to address these limitations. The architecture of our IP-based wireless access network is described later.

HAWAII

We now illustrate the operation of our proposed IP-based mobility solution with an elaborate example. The details of the protocol can be found in [7–9]. The example is divided into the following four subsections, illustrating power-up, micro-mobility, paging, and macro-mobility functionalities. We then conclude with a discussion of security issues in HAWAII.

Power-Up — Our IP-based access network is segregated into a hierarchy of domains, loosely modeled on the autonomous system hierarchy used in the Internet. The gateway into each domain is called the *domain root router*. Each host is assumed to have an IP address and a home domain. While this address assignment can be static, we prefer that the mobile host be assigned a dynamic address through DHCP during power up. This results in better IP address utilization efficiency for the wireless access network. In this case, assuming that the domain in which the mobile host is powered up belongs to the mobile host's service provider, the domain becomes the host's *home domain*. Because mobile hosts typically act as clients, as they activate applications their servers will learn their IP addresses. Also, directory servers can be used to learn the dynamically assigned address of the host.

The use of a dynamic address for mobile hosts is similar to the dialup model of service provided by Internet service providers to fixed hosts. The difference is that the users in wireless networks are mobile, and the home domain is determined by where the host is powered up rather than which modem access number is dialed. Apart from requiring fewer IP addresses than static allocation of IP addresses, we will see below that this also results in optimal routing with no tunneling as long as the user does not move out of a domain while powered up.

In our architecture, when operating in a domain, the mobile host maintains the assigned IP address regardless of its location. In order to maintain IP routing between the domain root router and the mobile host, HAWAII establishes special paths as the mobile host moves. The algorithm used to update selected routers to establish or maintain connectivity with a mobile host is termed a *path setup scheme*. The path setup scheme for power-up is illustrated in Fig. 3.

The figure shows two HAWAII domains with border routers, domain root router 1 and 2. Two base stations, BS1 and BS2, are assumed to be part of a multicast group, 239.0.0.1; this is relevant to paging, which will be discussed later. The HAWAII forwarding entries are shown adjacent to the routers. These entries are prepended with a message number indicating which message was responsible for establishing the entry (a message number of zero indicates a pre-existing entry). The entry consists of an IP address with an outgoing interface number for forwarding packets destined to that IP address and a multicast address corresponding to the group of base stations to which the user is currently attached.

The mobile host first sends a Mobile IP registration message (1) to its current base station. The base station, BS1, identifies that the mobile host is powering up in its home domain based on parameters in the registration message. It then adds a forwarding entry for the mobile host and initiates a HAWAII power-up message (2) to router R3. Router R3 similarly adds a forwarding entry to forward packets for the mobile host toward base station BS1 and sends the message (3) to the domain root router, R1. Router R1 adds a forwarding entry toward Router R3 and sends an acknowledgment (4) back to the base station, which then sends a Mobile IP registration reply (5) to the mobile host.

At this time, packets destined for the mobile host's IP address, 1.1.1.100, arrive from the Internet to domain root router R1 based on the subnet portion of the IP address (1.1.1.0), and then get delivered to the mobile host through routers R1 and R3, and base station BS1 based on the host-based forwarding entries established by HAWAII. *Note that there is no tunneling involved in this case, and packets traverse the optimal route to the mobile host.*

Micro-Mobility — Now let us consider what happens when the mobile host is handed off from BS1 to BS2. The sequence of messages exchanged is illustrated in Fig. 4. Note that the mobile host maintains its IP address (1.1.1.100) since this movement is within a HAWAII domain.

As before, the mobile host sends a Mobile IP registration message (1) to its new base station, BS2, informing it that base station BS1 was its previous base station. BS2 initiates a HAWAII handoff message (2) toward BS1. BS1 adds an entry for the mobile host so that future packets are forwarded toward BS2. It then sends the HAWAII message (3) to R3. R3 changes its forwarding entry from port 3 to port 4 so that packets destined for the mobile host now travel to base station BS2. It then forwards the HAWAII message (4) to BS2, which updates its forwarding table and sends a Mobile IP registration reply to the mobile host.

Note that this particular way of updating routers and base stations is called the *forwarding path setup scheme* in HAWAII.

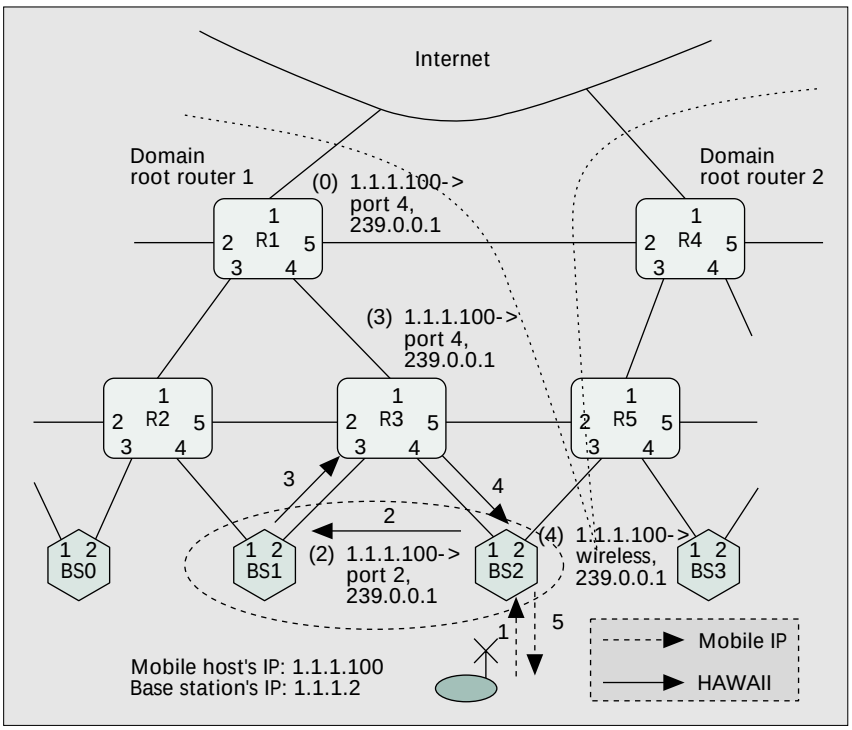

Figure 4. *Intradomain handoff.*

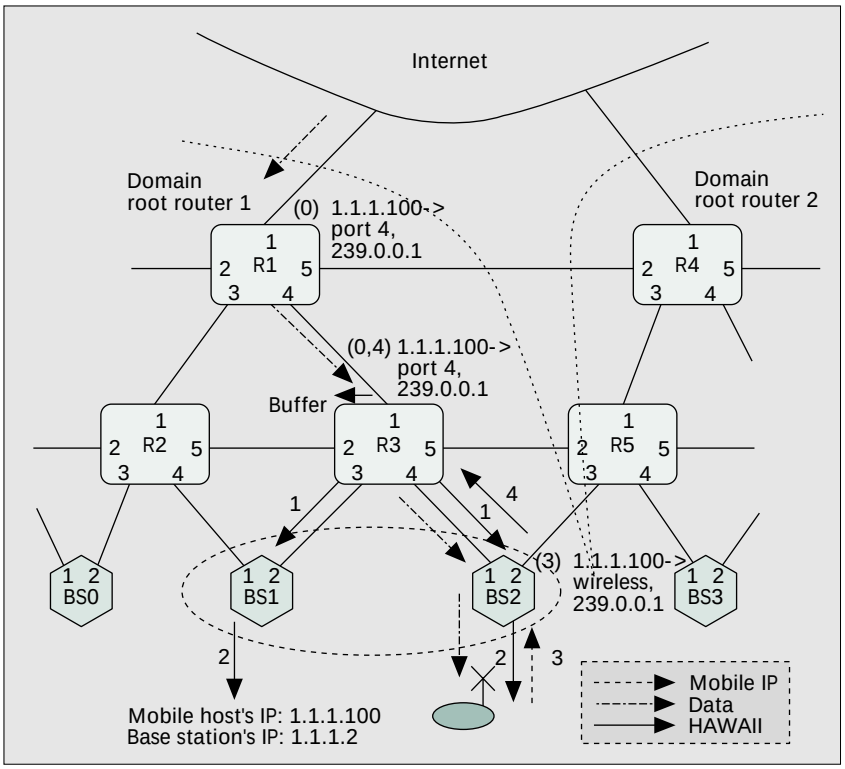

Figure 5. *Paging.*

This is useful in networks such as time-division multiple access (TDMA), where the mobile host cannot listen to two base stations simultaneously. In the case of networks such as CDMA where the mobile host could listen to multiple base stations simultaneously, it is possible to directly divert traffic from R3 instead of forwarding; this leads to a different algorithm for updating the routers and is called a *nonforwarding path setup scheme* in HAWAII. These schemes and other variations, including multicasting from R3 during handoff, are discussed in detail in [7]. *The advantage of custom tailoring these path setup schemes for different wireless networks is that disruption to user traffic can be minimized.* This is especially critical in next-generation wireless data networks where voice-over-IP and other multimedia traffic will likely be carried. Note that the path setup schemes can be analogous to the soft andoff functionality of current CDMA networks, albeit performed at the IP layer.

Another important aspect of HAWAII and its path setup schemes is that they operate locally. In Fig. 4, note that only R3 and the two base stations are involved in processing the updates. R1 is unaffected by this movement of the mobile host since its forwarding entry pointing to R3 is unchanged. *This leads to much better scalability than an approach based on Mobile IP*. Performance results in [7] for typical network configuration show that HAWAII results in almost one-tenth the processing requirements of using a centralized approach based on Mobile IP.

Finally, *maintaining the IP address of the mobile host unchanged across movements within the same domain results in straightforward support for QoS*. In the case of using a reservation protocol such as RSVP, reservations need only be restored locally during handoff (at R3 and BS2); prior reservations at other routers such as R1 and the backbone routers can be maintained unchanged since the mobile host's IP address, used to identify flows, remains the same.

Paging — Recall that in GPRS, the paging functionality is performed in a centralized fashion by an SGSN and can be considered a link-layer function. In our architecture we use HAWAII to support IP-level paging in a distributed, scalable, and flexible manner.

Assume that the mobile host illustrated in Fig. 4 is idle and goes into standby state. Subsequently, the network only knows that the mobile host is present in one of the base stations in its paging area, denoted in this example by BS1 and BS2. In our architecture we use an administratively scoped IP multicast group address (239.0.0.1) to identify the set of base stations belonging to the same paging area. Assume that at this time IP packets destined for the mobile host arrive at the domain root router. The network then needs to page BS1 and BS2 to determine the exact location of the mobile host. The procedures involved in delivering the packet to the mobile host in this case is illustrated in Fig. 5.

The data packets first arrive at R1. Based on its forwarding entry, R1 notes that the mobile host is in standby state. However, R1 determines that it does not belong to the multicast tree of the paging area for the mobile host, 239.0.0.1 (we assume that the multicast tree connecting BS1 and BS2 is rooted at R3). Therefore, R1 forwards the data packets downstream on port 4 toward R3. R3 performs similar processing and identifies that it is part of the multicast tree for the paging area of the mobile host. It then buffers the data packets and initiates a HAWAII page request (1) to the multicast group address. BS1 and BS2, which belong to the multicast group, receive the page message and broadcast a page message (2) on their respective wireless interfaces utilizing the underlying link-layer technology. The mobile host, which happens to remain under BS2 in this example, receives the link-layer page message (e.g., by periodically scanning the broadcast paging channel) and sends a Mobile IP registration message (3) to BS2. This triggers a HAWAII path setup message from BS2 to the paging initiator (4), R3. Updated forwarding entries are also established at BS2 and R3. The buffered data packets (as well as any arriving packets) are then forwarded to the mobile host through BS2.

A complete description of the paging procedure can be found in [9]. The motivation behind the algorithm is to push the burden of paging to the base station and low-level routers, and away from the domain root router so that scalability is enhanced.

One of the benefits of performing paging at the IP layer is flexibility. For example, the fixed paging approach used in current cellular networks and presented in the example above allow only for a fixed set of base stations to belong to a paging area. In the IP-based approach, a paging area is determined by the composition of a multicast group. This enables other approaches such as hierarchical and per-user paging where different sets of base stations are paged for each user.

Macro-Mobility — Finally, we illustrate the situation where the mobile host moves between base stations connected to different HAWAII domains. In this case, the mobile host acquires a second address, the care-of address, in the new domain. We assume the collocated care-of address (CCOA) model of Mobile IP since the CCOA uniquely identifies the mobile host for QoS support. However, if necessary, a network-based care-of address model can also be incorporated. In this example, the mobile host acquires a CCOA of 2.2.2.200 from the new domain. The sequence of messages exchanged

for interdomain handoff is illustrated in Fig. 6.

The mobile host sends a Mobile IP registration message (1) to BS3. Based on the parameters in the message, the base station detects that this is an interdomain handoff. BS3 first initiates the HAWAII power-up procedure in this domain (denoted by messages 2–4) for establishing host-based entries for the 2.2.2.200 collocated address. It then sends a Mobile IP registration message (5) to the home agent of the mobile host, R1. In our architecture, the home agent is collocated with the domain root router. Thus, the home agent at R1 establishes a tunnel entry, which will tunnel packets destined to the mobile host IP address of 1.1.1.100 to its new collocated address. Upon receiving a reply to the Mobile IP registration message (6) from the home agent, BS3 sends an acknowledgment to the mobile host (7).

At this time, packets destined to the 1.1.1.100 address of the mobile host reach R1 and then get tunneled to the 2.2.2.200 address. The tunneled packets with the 2.2.2.200 address reach R4 based on the subnet portion of the address and then get forwarded through R5 and BS3 to the mobile host based on the HAWAII forwarding entries established. Subsequent handoffs by the mobile host in this new domain will be handled locally by HAWAII as described previously in the micromobility section. Thus, the home agent is updated only when the mobile host crosses a domain boundary, a much rarer occurrence, resulting in reduced handoff latencies and improved scalability.

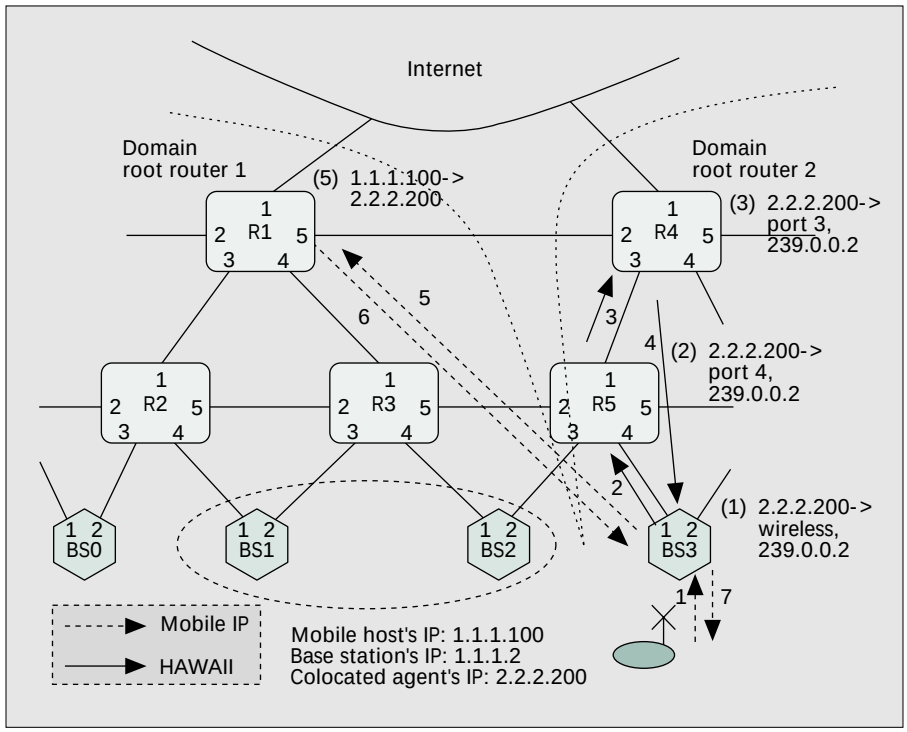

Figure 6. *Interdomain handoff.*

Security — HAWAII faces the same security concerns as any regionalized approach to IP mobility (e.g., [10]). Recall that HAWAII networks are organized hierarchically, and their performance advantage over basic Mobile IP is achieved through handoff procedures that involve only local nodes. Therefore, without involving the home agent, base stations must be able to:

- Verify the authenticity of the Mobile IP message coming from the mobile host
- Generate a reply that the mobile host will be able to verify as authentic

To achieve this goal, three separate security associations must be in place: the first between the base stations and the mobile host, the second between the base stations and the home agent, and the third between the home agent and the mobile host. For this purpose, an AAA [6] infrastructure can be used to distribute three sets of authentication keys, to the mobile host, to the base stations, and to the home agent. The basic Mobile IP authentication scheme must also be modified. In particular, in case of handoffs that do not involve the home agent, mobile hosts must be prepared to receive registration replies that contain only authentication information generated by base stations. At the same time, home agents must be prepared to receive surrogate registration requests, generated by base stations on behalf of mobile nodes. Such requests will not contain any authentication information generated by the hosts, only authentication data provided by the base stations.

Performance

We have implemented the HAWAII protocol on a testbed of PCs running the FreeBSD operating system. The processing of a HAWAII handoff message at a given node takes only 0.156 ms. The handoff latency for a mobile host connected to a network through Lucent's 2 Mb/s WaveLAN is approximately 5 ms (including 4 ms of latency on the wireless segment).

We next compare the performance of Mobile IP and HAWAII using simulation on a larger network with cross-traffic. We compare the disruption of HAWAII forwarding and nonforwarding schemes with basic Mobile IP as well as the Mobile IP route optimization (RO) [11] proposal. In the simulation the topology of the wireless access domain is a binary tree with three levels; at the lowest level, there are four base stations. The HA and the correspondent host are outside the domain, while the mobile user is handed off between the four BSs in the experiment. We simulate audio traffic to the mobile host from a correspondent host (and through the HA in case of Mobile IP) in the form of 160-byte UDP packets transmitted every 20 ms (64 kb/s).

In the case of an interactive audio application, a playout delay is typically used to overcome network jitter; if the packet arrives after its playout time, the packet is dropped. We are thus interested in *total packet loss* which includes both packets dropped due to late arrival as well as packets lost in the network due to handoff. In Fig. 7a we plot the total of dropped and lost packets per handoff (averaged over 100 or more handoffs) vs. playout delay for all four schemes. In this simulation topology, the propagation delay from correspondent host to mobile host is 25 ms for all the schemes except for basic Mobile IP which incurs an additional 100 ms delay due to routing through the HA.

In the case of basic Mobile IP, about 5 packets/handoff are lost in the network. This is because in our configuration the registration update from the mobile host takes about 100 ms (link delay of 50 ms and queuing delay of approximately 50 ms) to reach the HA. In this interval, about five packets are sent to the old BS and lost. Let us now compare the remaining three schemes. Consider a playout delay value of 100 ms in Fig. 7a. In this case, the Mobile IP RO scheme results in a

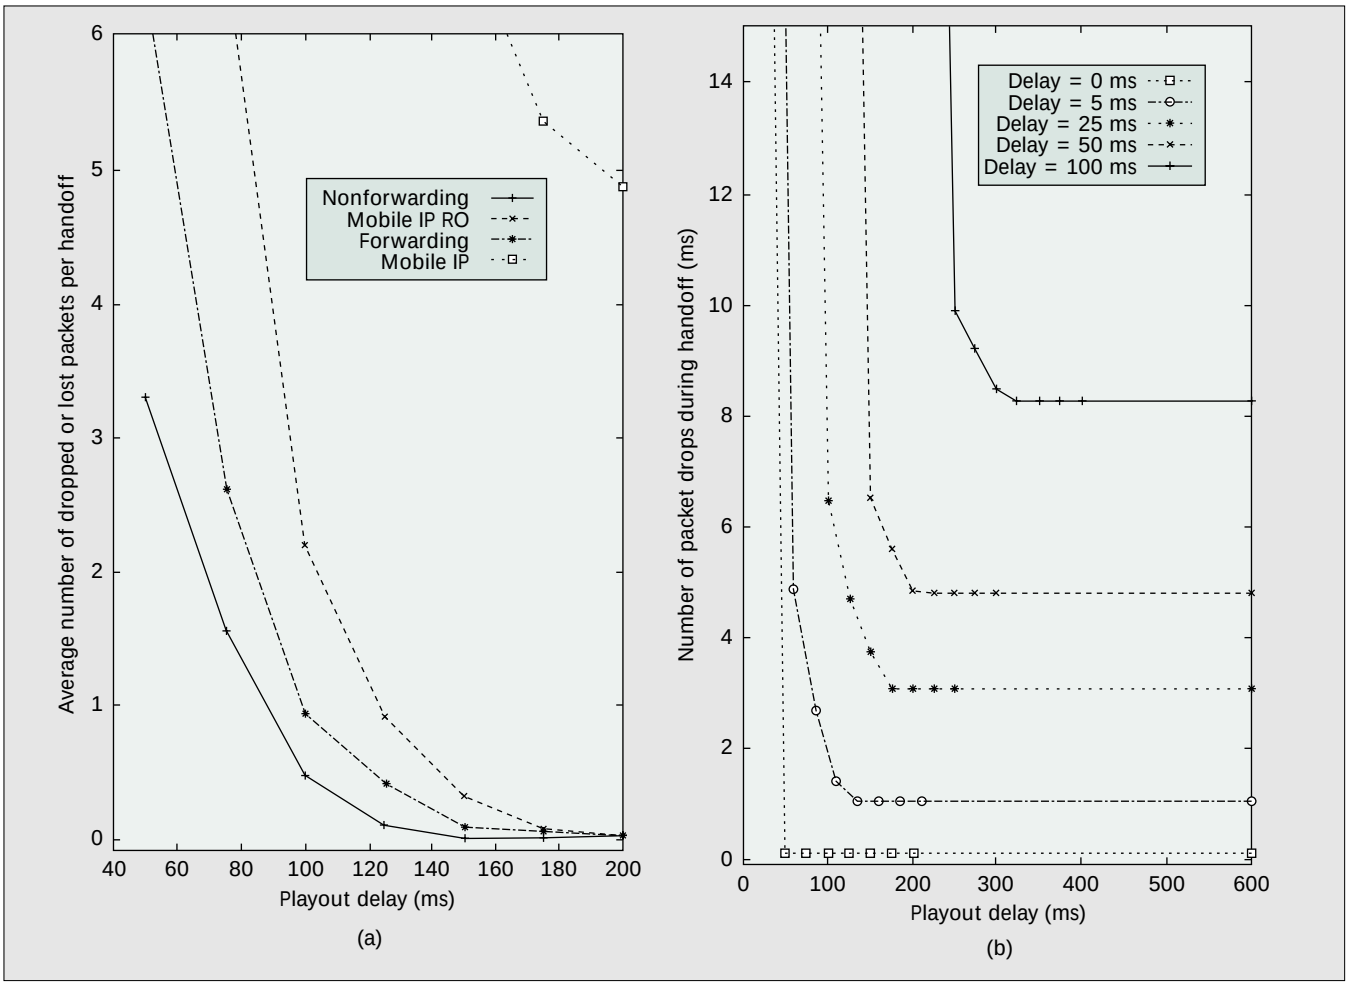

Figure 7. *Packet loss during two-hop handoff: a) all schemes; b) impact of delay to HA.*

total loss of about 3 packets/handoff, while the HAWAII schemes result in a total loss of less than 1 packet/handoff. This is because the HAWAII schemes switch over very quickly to the new route, while in Mobile IP RO the HA and then the correspondent host must be notified before packets use the new route. Of the HAWAII schemes, the nonforwarding scheme performs better than the forwarding one since the nonforwarding scheme is able to utilize the mobile host's ability to receive from multiple BSs.

We also examine the effect of the propagation delay to the HA on performance. The HAWAII schemes are unaffected since they operate locally. In the case of Mobile IP, as shown in Fig. 7b, when the delay to HA decreases, the performance approaches that of HAWAII. The same behavior is true for Mobile IP RO as well (not shown). Thus, by operating locally, the HAWAII schemes result in smaller disruption to interactive audio traffic than do the Mobile IP schemes.

A HAWAII-Based Next-Generation Wireless Data Network

Given the large installed base of wireless access networks, initial deployment of our architecture will have to interwork with current access networks. In this scenario, the BSC of a GPRS-based access network or the frame selection entity of a CDMA-based access network will likely be the mobile host's next-hop IP node. Here we describe two ways in which HAWAII can be used in wide area wireless networks. In both cases, HAWAII is used to enhance micro-mobility with an IP-based network.

A Pure HAWAII-Based Network

Figure 8 shows a pure HAWAII network at a high level. The HAWAII protocol runs in the network connecting the mobile device and domain router. In a GPRS network, the BSC/BTS

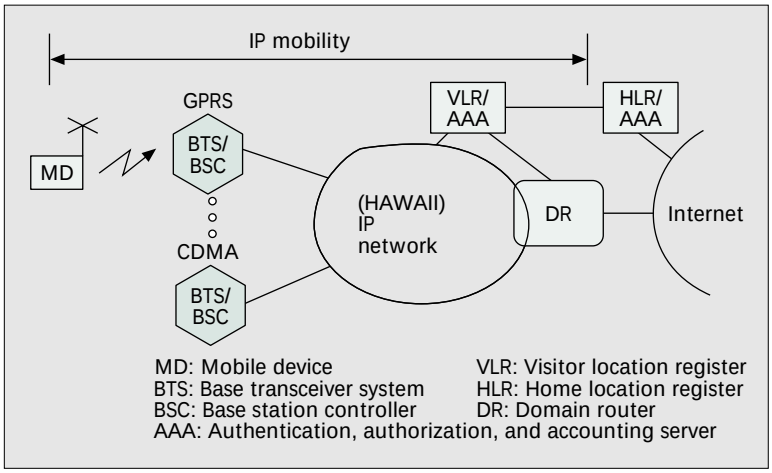

Figure 8. *The HAWAII-based architecture.*

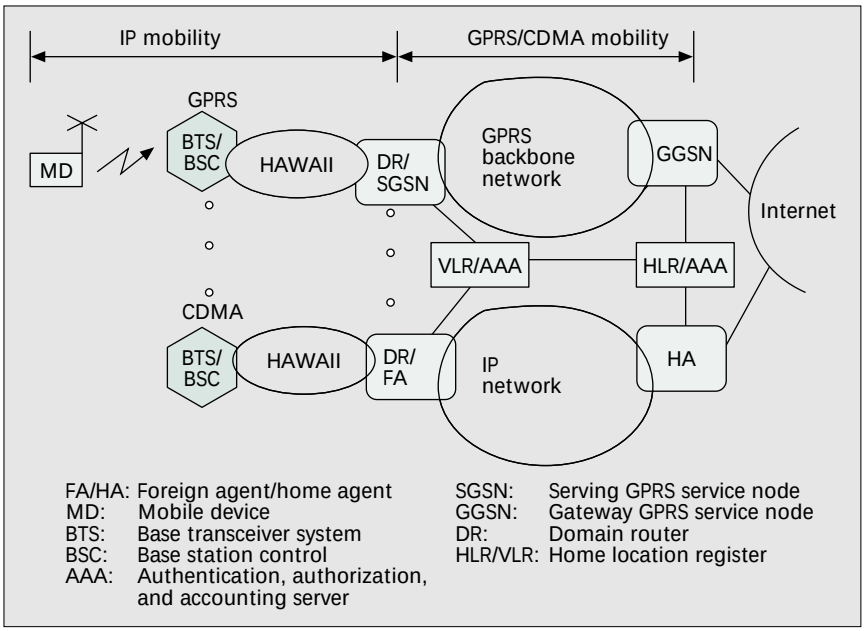

Figure 9. *HAWAII in the access network.*

act as a level 2 bridge to the air interface. In this case, mobility between BTSs attached to the same BSC is handled by link-layer techniques, and inter-BSC movement is handled by HAWAII. In a CDMA network, the frame selector would act as a bridge to the air interface in the same manner as the BSC.

To support authentication and roaming, an HLR/VLR infrastructure could be used for authenticating access to the air interface, while an AAA-based infrastructure should be used to support authentication of Mobile IP and HAWAII transactions. Alternatively, the HLR could be augmented with an AAA interface, offering integrated support for GPRS/CDMA and Mobile IP/HAWAII authentication.

A Partial HAWAII-Based Network

Figure 9 shows a scenario in which HAWAII is used for mobility management in the access network while remaining connected to a backbone network using different networking technologies. The authentication infrastructure remains the same as in the pure HAWAII network example earlier.

For GSM, the backbone network would be a GPRS network. The HAWAII domain router is connected to the SGSN. GPRS protocols, such as GTP, would be used for registration, authentication, and high-level mobility. HAWAII would be used for micro-mobility and paging support. In this scenario, the benefits of HAWAII's micro-mobility management are realized for transporting user data, and the GPRS protocols are used to access databases, allowing GPRS infrastructure and management to be reused.

In a CDMA network, the backbone network would be basic Mobile IP. The foreign agent would be located at the domain root router. In this way, again micro-mobility would be handled by HAWAII and roaming by Mobile IP.

Conclusion

In this article we presented a homogeneous IP-based wireless access network architecture that supports different wide-area wireless technologies. This IP-based network uses the Internet standard, Mobile IP, to support macro-mobility of mobile hosts, and HAWAII to support the micro-mobility and paging functionality of current wireless networks. We also illustrated how the proposed IP-based solution can interwork with existing infrastructure so that incremental deployment can be achieved.

References

[1] A. Valko, A. Campbell, and J. Gomez, "Cellular IP," Internet draft, Nov. 1998.
[2] "Digital Cellular Telecommunication System, General Packet Radio Service, Service Description — Stage 2, GSM 03.60 v. 6.0," ETSI, 1998.
[3] S. Shenker, C. Partridge, and R. Guerin, "Specification of Guaranteed Quality of Service," RFC 2212, Sept. 1997.
[4] K. Nichols and S. Blake, "Differentiated Services Operational Model and Definitions," Internet draft, Feb 1998.
[5] C. E. Perkins, "IP Mobility Support," RFC 2002, Oct. 1996.
[6] P. Calhoun and C. E. Perkins, "DIAMETER Mobile IP Extensions," Internet draft, Nov. 1998; work in progress.
[7] R. Ramjee *et al.*, "HAWAII: A Domain-based Approach for Supporing Mobility in Wide-area Wireless Networks," *Proc. Int'l. Conf. Network Protocols*, Nov. 1999.
[8] R. Ramjee *et al.*, "IP Micro-Mobility Support using HAWAII," Internet draft, June 1999; work in progress.
[9] R. Ramjee, T. La Porta, and L. Li, "Paging Support for IP Mobility Using HAWAII," Internet draft, June 1999; work in progress.
[10] E. Gustafsson, A. Jonsson, and C. Perkins, "Mobile IP Regional Tunnel Management," Internet draft, Aug. 1999; work in progress.
[11] C. E. Perkins, and D.B. Johnson, "Route Optimization in Mobile IP," Internet draft, Nov. 1997.

Biographies

RAMACHANDRAN RAMJEE (ramjee@bell-labs.com) received his B.Tech. in computer science and engineering from the Indian Institute of Technology, Madras, in 1992, and his M.S. and Ph.D. in computer science from the University of Massachusetts, Amherst, in 1994 and 1997, respectively. He has been a member of technical staff at Bell Labs, Lucent Technologies since 1996. His research interests are signaling, mobility management, and QoS issues in wireless and high-speed networks.

THOMAS F. LA PORTA received his B.S.E.E. and M.S.E.E. degrees from The Cooper Union, New York, New York, in 1986 and 1987, respectively, and his Ph.D. degree in electrical engineering from Columbia University, New York, in 1992. He joined AT&T Bell Laboratories in 1986. He is currently head of the Networking Techniques Research Department in Bell Laboratories, Lucent Technologies where has has worked on various projects in wireless and mobile networking for the past several years. He received the Bell Labs Distinguished Technical Staff Award in 1996, and an Eta Kappa Nu Outstanding Young Electrical Engineer Award in 1996. His research interests include mobility management algorithms, signaling protocols and architectures for wireless and broadband networks, and protocol design. He was Editor-in-Chief of *IEEE Personal Communications* and is a technical editor on *ACM/Baltzer Journal of Mobile Networking and Applications*.

LUCA SALGARELLI received his Laurea (Dr.Eng. degree) in electrical engineering at Polytechnic of Milan, and his Master's degree (M.Phil.) in computer science at CEFRIEL, Milan, Italy. He is currently a member of technical staff at Bell Laboratories, Lucent Technologies. His interests are in designing, developing, and evaluating protocols and software architectures for IP-based networks.

SANDRA THUEL received her B.S. in computer engineering from the University of Puerto Rico, Mayaguez, in 1986, and her M.S. and Ph.D. in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, Pennsylvania, in 1988 and 1993, respectively. She has been a member of technical staff at Bell Laboratories, Lucent Technologies since 1993. Her research interests are QoS, scheduling and resource management, and issues regarding the convergence of voice and data networks.

KANNAN VARADHAN obtained his Ph.D. from the University of Southern California in 1998. He is interested in robustness issues in protocol design; more recently, he has become fascinated by studying routing protocol behavior in the Internet, and its observable behavior at the edges of the network and consequent effects on end-to-end protocols.

LI LI received his B.E. degree in automation from Beijing Polytechnic University in 1993 and his M.E in electrical engineering from the Institute of Automation, Chinese Academy of Science in 1996. He is currently a doctoral candidate in the Computer Science Department at Cornell University. He is interested in mobile computing, computer networking, and distributed systems. He is currently working in the area of mobile computing.