Abstract

In this article we explore various network layer concepts that play a crucial role in the design of mobile networking systems. We show that mobility is essentially an address translation problem and is best resolved at the network layer. We describe services that must be supported at the network layer to carry out the task of address translation. Using these service primitives as building blocks, we describe a network-layer architecture which enables smooth integration of mobile end systems within the existing Internet. A summary of some of the key Mobile IP proposals is presented, and it is shown that each proposal can be viewed as a special case of the architecture outlined in this article.

Network Layer Mobility: An Architecture and Survey

Pravin Bhagwat, Charles Perkins, and Satish Tripathi

Mobile end systems frequently change their point of attachment to the network. In such an environment, in order for mobile devices to run without disruption, an internetworking infrastructure is needed. In addition, a common networking protocol is required which can support network-wide mobility. Mobile devices also need to communicate with the existing pool of information servers and file servers, which means that internetworking solutions for connecting stationary and mobile systems are also required. Unfortunately, the Internet Protocol (IP), which forms the fabric of the current worldwide data communication network, falls short of meeting this demand. The current Internet suite of protocols — Transmission Control Protocol/Internet Protocol (TCP/IP) — were designed under the assumption that end systems are stationary. If, during an active network session, one end of the connection moves, the network session is broken. Naturally, all networking services layered on top of TCP/IP are also disrupted when end systems become mobile. There are two approaches for solving this problem. One is to completely redesign internetworking protocols with the specific goal of supporting mobile end systems. The other approach is to provide additional services, which make mobile internetworking possible, at the network layer in a backward-compatible manner. The first approach, though an interesting possibility from a research viewpoint, is infeasible because it would require radical changes to the currently deployed networking infrastructure. The latter approach is the focus of our investigation.

To ensure interoperability with the existing infrastructure, the handling of mobility should be completely transparent to the protocols and applications running on stationary hosts. In other words, from a stationary end system's perspective, a mobile host should appear like any other stationary host connected to the Internet. This means that the same naming and addressing conventions, those originally developed for stationary hosts, must apply to mobile hosts. In addition, any changes in a mobile's network attachment point should be completely hidden from the protocols and applications running on stationary hosts.

In this article we explore various network-layer concepts that play a crucial role in the design of mobile networking systems. We show that mobility is essentially an *address translation* problem and is best resolved at the network layer. We have identified the fundamental services that must be supported at the network layer to carry out the task of address translation. Using these service primitives as building blocks, we describe a network-layer architecture which enables smooth integration of mobile end systems within the existing Internet. In the second half of this article, we present a summary of some of the key mobile IP proposals and show that each proposal can be viewed as a special case of the proposed architecture. It is worth pointing out that our objective is not to propose a new protocol for supporting mobility, but to highlight various design choices and the trade-offs involved in the design of internetworked mobile systems.

Internet Naming and Addressing

The Internet is a large collection of networks which share the same address space and interoperate using a common set of protocols, such as TCP/IP [1, 2]. A fundamental concept of the Internet architecture is that each host[1] has a unique network address, by which it is reachable from other hosts in the network. Data are carried in the form of packets which contain source and destination addresses. To communicate with another host, a source need only know the address of the destination. Internet routers cooperate to carry packets from a source to a destination node.

Internet routers maintain a view of network topology in the form of routing tables. These tables are consulted when making packet-routing decisions. The process of routing involves inspecting the destination address contained in the packet and, based on the contents of the routing table, determining the next-hop router to which the packet should be relayed. Each router along the path from a source to a destination node repeats this process until the packet is finally delivered to the destination host.

If host addresses are treated as *flat identifiers*, routers will be required to maintain routing information on a per-host basis. Obviously, this is not feasible given the large number of hosts (over 10 million!) connected to the Internet. A natural solution is to impose a hierarchy on the address structure. Hierarchical addressing is essential if the routing architecture is to be scalable. The Internet, for example, deploys a multilevel hierarchical addressing scheme [3].

[1] *In Internet jargon, "host" means an end system connected to the Internet.*

1070-9916/96/$05.00 © 1996 IEEE

Internet Addressing

Each host in the Internet is assigned a unique 32-bit internet address (also known as an IP address) which consists of two parts: network-id and host-id. IP addresses are commonly represented using dotted notation, where each octet is represented as a decimal number and dots are used as octet separators.

Under the current Internet addressing scheme, routers only need to maintain network topology information at the granularity of individual networks. This means only the network part of the destination address is used in making the routing decision. Though hierarchical addressing makes routing simple and manageable, as a natural consequence it puts certain restrictions on address usage. A hierarchical address can only be used within the domain of its definition. For example, an Internet address is only meaningful as long as the host using it remains connected to the network denoted by the network-id part of the address. When the host moves to a new network, it must be allocated a new address which is derived from the address space of the new network. In order for the Internet routing to work, *a mobile host must be associated with a new address when it moves.*

Naming

Hosts are also identified in the network by their *host names*. Names are user-defined aliases (strings of characters) which are used to denote hosts. An important distinction between names and addresses is that addresses may be protocol-specific (e.g., an IP address, CLNP address, IPX address, XNS address), but names are not. Names provide a way for applications to make reference to network entities without having to know anything about the underlying network protocol in use. This is useful because users find names easier to use and remember than cumbersome network addresses.

Though applications refer to end systems by names, when packets are transported through the Internet each must contain an IP address of a destination node. This is because Internet routers do not understand names; they can only interpret addresses. Therefore, a translation mechanism is required for mapping host names to addresses. To accommodate a large, rapidly expanding set of names, a decentralized naming mechanism called the "domain name system" (DNS) was deployed in the Internet. DNS stores name-to-address mappings in a distributed data structure. Finding the address of the host is essentially a directory lookup operation (Fig. 1). When two hosts on the Internet need to communicate with each other, the source node performs a DNS lookup to obtain the destination node's address and then initiates a connection setup procedure. During connection setup, each end of the connection learns about the address of the other end. As long as the connection is active no additional DNS lookups are performed, because name-to-address binding is assumed to be static and is not expected to change during a connection lifetime.

The Mobility Problem

To illustrate why host mobility poses a problem at the network layer, it is important to emphasize the distinction between the concepts of *name* and *address*. A name is a location-independent identifier of a host. An address, on the other hand, reflects a host's point of attachment to the network. For hosts that remain static throughout their lifetime, both names and addresses can be used interchangeably; but for a mobile host, an address cannot be used as a unique identifier, because it must change with the location of the host. The name is the only location-independent identifier that can be used to refer to mobile hosts.

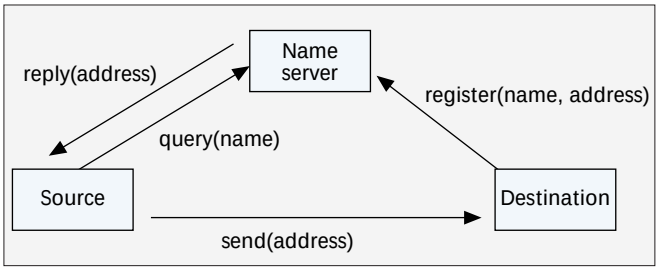

Figure 1. *DNS-based name-to-address resolution.*

Mobility Problem: Directory Service View

In networks where hosts are static, name-to-address bindings never change. Host mobility makes this binding a function of time. Therefore, network-layer mechanisms are required for resolving names into addresses and tracking the location of hosts as they move. The DNS, which provides name-to-address translation service in the Internet today, could be enhanced to meet the additional demands. However, this task is made difficult by many hurdles:

- Historically, the DNS had no provision to handle dynamic updates. This is because it was originally designed to provide name lookup service for stationary hosts only.
- The DNS design attempts to optimize the *access* cost, not the *update* cost. Server replication and client caching provide significant performance gains for access-only systems, but result in very poor performance when updates are performed. In a mobile environment, both updates and accesses are likely.
- DNS clients cache DNS records to reduce latency for future accesses and to reduce load on the name servers. There is no callback mechanism generally available from servers to clients in case cache entries become invalid.

A design for a distributed location directory service for mobile hosts was proposed by Awerbuch and Peleg in [4]. They formally proved an important theoretical result which established that a system cannot optimize both access and update operations.[2] Using the concept of *regional directories* (a type of cache), they proposed a distributed directory layout which guarantees that the communication overhead of access and update operations is within a poly-logarithmic factor of the lower bound.

As far as the Internet is concerned, distributed-directory-service-based solutions do not appear very attractive because they cannot be deployed without changing existing host software. The current size of the Internet makes any such change to host software almost impossible to achieve. Hence, an alternative solution method is required.

Mobility Problem: Internet View

When host names were originally deployed, it was implicitly assumed that the *name*-to-*address* binding remained static. Instead of referring to hosts by name, protocols were developed that referred to hosts through their addresses. A standard example is a TCP connection which is identified by a 4-tuple:

< *source IP address, source TCP port,*
destination IP address, destination TCP port >

If neither host moves, all components of the connection identifier will remain fixed; thus, a continuous TCP session can be maintained between the two hosts. If either end of the connection moves, we run into the following problem:

2 In their paper they use the terms `Find` *and* `Move` *to denote these operations.*

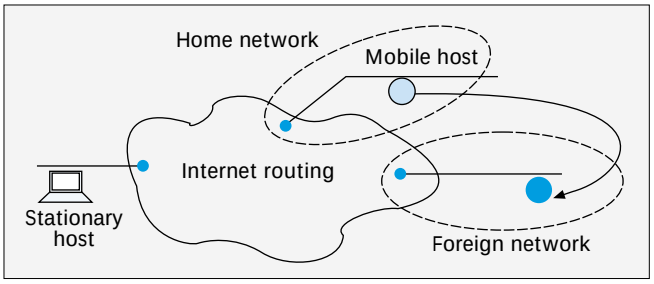

Figure 2. *Illustration of terms.*

- If the mobile host acquires a new IP address, then its associated TCP connection identifier also changes. This causes all TCP connections involving the mobile host to be broken.
- If the mobile host retains its address, then the routing system cannot forward packets to its new locations.

The fundamental problem is that in the Internet architecture, an IP address serves dual purposes. From the transport- and application-layer perspective, it serves as an *endpoint identifier*, and at the network layer the same IP address is used as a *routing directive*. This problem is not specific to the Internet architecture; in fact, all other contemporary connectionless network architectures, such as open systems interconnection (OSI), IPX, and XNS, suffer from this problem. Because our objective is to ensure that connections are not broken when hosts move, we can say that *in order to retain transport-layer sessions, a mobile host's address must be preserved regardless of its point of attachment to the network.*

An immediate consequence of this choice is that we cannot rely on the existing addressing paradigm for delivering packets to a mobile host's new location. A solution might be to keep per-mobile-host routing information at all routers, but this completely breaks the hierarchical model of routing, causing unbounded growth in the size of routing tables. Thus, the problem of supporting mobile hosts within the Internet is not just keeping track of hosts. A mechanism has to be designed for forwarding packets to mobile hosts without modifying and compromising the scalable nature of the Internet routing mechanism.

Network-Layer Solution Architecture

In this section we describe a network-layer architecture that allows smooth integration of mobile end systems within the Internet. Our objective is to highlight and analyze the essential aspects of providing mobility extensions in any connectionless network where routers depend on addresses stored in the packet. The specific details involved in designing a mobile networking system will be discussed later. For ease of exposition, we will first introduce a few definitions (Fig. 2).

Mobile Host — An internet host is called a *mobile host* if it frequently changes its point of attachment to the network. A change in the attachment point can happen while one or more transport-layer sessions involving the mobile host are in progress. It is assumed that the rate of change of location is slower than the time it takes for the mobile routing protocols to take into account the mobile host's new location.

Home Address — Like any other internet host, a mobile host is also assigned an IP address which is referred to as its *home address*. A standard 32-bit address is allocated using the same guidelines that apply to stationary hosts. When the DNS is queried with a mobile host's name, it returns the home address of the mobile host.

Home Network — Within each administrative domain, network administrators usually reserve one or more subnetwork(s) for mobile hosts. The home address of a mobile host is allocated from the address space of one of these subnetworks, referred to as the *home network* in the subsequent discussion. The terms *home address* and *home network* could also apply to stationary hosts. The only difference is that stationary hosts always remain connected to their home network, while mobile hosts sometimes may not be found at their home networks.

Foreign Network — Any connected segment of an Internet, other than the home network of a mobile host, to which the mobile host is allowed to attach is referred to as a *foreign network*.

Notice that the above definitions are relative to a mobile host. The same network could operate as both a home and foreign network, depending on which mobile host is connected to it. As long as a mobile host remains connected to its home network, existing internet routing mechanism are sufficient to route packets up to its current location. It is only when it moves to a different network that additional mechanisms are required. If a mobile host moves within its home network (e.g., detaches from one Ethernet point and attaches through another Ethernet point), it does not constitute a move from the network-layer point of view. A collection of link-layer networks, which are interconnected through bridges, is called a "layer 2 segment." Existing link-layer bridging mechanisms are capable of routing packets up to end systems as long as they remain connected to the same layer 2 segment. Within a layer 2 segment, a packet can be delivered solely on the basis of the destination node's link layer address; the network layer routing is not required.

In the previous section, we made two crucial observations:

- The home address of a mobile host cannot be used for routing packets to its current location (except when it is attached to its home network).
- A mobile host's address must be preserved in order to retain all active transport connections involving the mobile host.

These are conflicting requirements. From the first observation, when a host moves a new address, reflecting its new point of attachment to the network, must be used for the purpose of routing. The second observation says just the opposite: the original address must be preserved to retain all active network sessions.

Two-Tier Addressing

We introduce the concept of *two-tier addressing* to resolve the problem associated with the dual uses of an internet address. Our solution involves associating two internet addresses with each mobile host (Fig. 3). The first component of the address reflects the mobile's point of attachment to the network while the second component denotes its home address. The first address component serves as a *routing directive*. It changes whenever a mobile host moves to a new location. The second component of the address serves as an *end-point identifier*. It remains static throughout the lifetime of a mobile host. The purpose of two-tier addressing is to decouple the dual role of an internet address into two disjoint, well-defined functions.

The concept of two-tier addressing is illustrated in Fig. 3. Packets destined to mobile hosts contain the destination address in the two-tier format. The Internet routing system only looks at the first component of the address and routes those packets to the point where the mobile host is attached. At this point, the first address component is discarded. Only the second address component, the home address of the mobile host, is used in subsequent protocol processing. From an end host's perspective, this means it notices no difference when it is attached to its home from when it is located in a

foreign network. In other words, the mobile host remains *virtually* connected to its home. Packets that originate from the mobile host and are destined to the stationary host (S) do not require any special handling, since the Internet routing system can deliver those packets based on their destination addresses. If S is also mobile, then the same two-tier addressing mechanism can be used to route packets to its current location.

It is important to note that two-tier addressing is only a logical concept. Its realization does not necessarily require carrying two addresses in the destination address field of the network-layer packets. In fact, doing so would require changes in the existing packet formats, necessitating changes to host and router software. It is desirable to support the two-tier addressing method using the existing mechanisms available in the Internet. The following sections describe how this goal can be achieved.

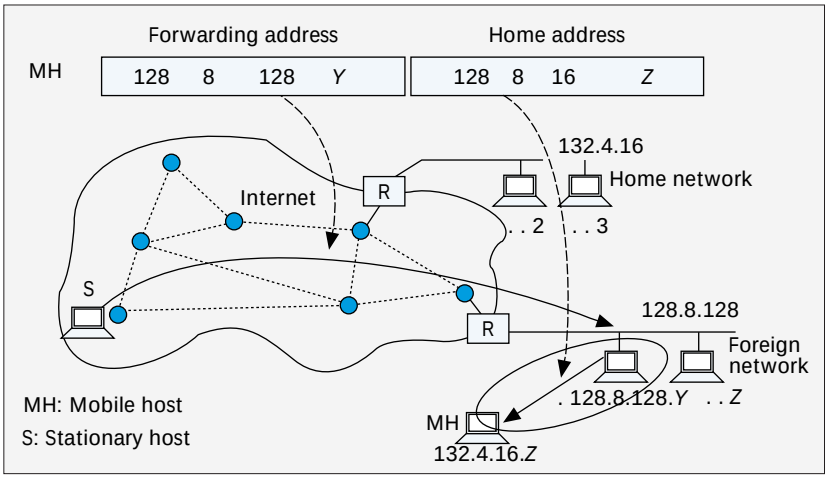

Figure 3. *Two-tier addressing for mobile hosts.*

Architecture Components

Forwarding Agent — When away from its home network, a mobile host can attach to the Internet through a foreign network. For the purpose of forwarding datagrams to its new location, an address derived from the address space of the foreign network must be used. Packets destined to the mobile host contain the address of a *forwarding agent* (FA) in the forwarding address subfield of the two-tier address. An FA provides an access point through which mobile hosts can attach to the network. It receives packets on behalf of mobile hosts, and forwards them to appropriate mobile hosts after necessary protocol processing.

Conceptually, the processing at the FA involves stripping the forwarding address part of the two-tier address and exposing the home address of the mobile host. Once the packet arrives at the FA, the forwarding address is no longer required in the subsequent protocol processing. When a packet arrives at the FA, it contains the address of the FA in its destination address field. The FA essentially maps the contents of the destination address field (the forwarding address) to the home address of the associated mobile host. We use the notation g to denote this mapping function:

g(forwarding address) $\rightarrow$ (home address)

An FA should be able to relay packets to the mobile host on the basis of its home address. This is easy if the FA and the mobile host are directly connected (normally over a wireless link). Otherwise, the routing protocol operating in the foreign network should advertise host specific routing information within the foreign network to facilitate routing of these packets to mobile hosts. Normally, we would expect a wireless base station to operate as an FA, in which case the mobile host and FA would be directly connected to each other over a wireless link.

A mechanism is required so that mobile hosts can discover the address of an FA when they connect to a foreign network. Similarly, a mechanism is required so that the FA can determine the identities of all mobile hosts that require its service. The simplest way to achieve this is through a route advertisement and a registration protocol. Forwarding agents periodically advertise their presence in the foreign network. Beaconing, the periodic broadcast of messages over the wireless medium, is the most commonly used method. Mobile hosts can listen to broadcasts, determine the identity (address) of the nearest FA, and initiate a registration sequence.

Location Directory (LD) — The component in the architecture that records the association between the home and forwarding addresses of a mobile host is called a *location directory* (LD). The LD contains the most up-to-date mapping between a mobile host and its associated FA. Mobile Hosts are required to send updates to the LD whenever they move to a new location.

Because the number of mobile hosts is expected to be very large, a centralized realization of the LD is deemed infeasible. A policy for distributing LD components should take many factors into consideration, such as the cost of access, ease of locating LD components, and security and ownership of location information. Because the LD will be accessed very frequently, a good distribution method should exploit the locality of access patterns and provide uniform load balancing among all LD components. Given a model for the LD access pattern, the LD distribution can be formulated as an optimization problem [5]. Unfortunately, these mathematical results [5–7] cannot be applied directly in the Internet. The primary reason is that in the Internet factors such as ease of location, security, and ownership take precedence over any cost optimization considerations.

A feasible distribution scheme in the Internet is the *owner-maintains* rule. According to this scheme, the LD entries for mobile hosts are maintained at their respective home networks. Advantages of this scheme are:

- Some agent on each home network is responsible for maintaining, securing, authenticating, and distributing LD information for its mobile hosts. This policy fits well within the Internet philosophy of autonomous operation.
- No special mechanisms are required to locate the LD components. It is important to point out that in a distributed scheme, in order for a source to send a query to the right LD component, the source is required to know the address of the LD component in advance. Under the *owner-maintains* rule, a source simply sends a query addressed to the mobile host. The packet is delivered to the home network by normal internet routing, where it is intercepted by the home router and subsequently relayed to the correct LD component.

This is certainly not the only possible distribution scheme. Later in this article we will discuss other options while reviewing various mobile IP proposals.

Address Translation Agent — Hosts that need to communicate with a mobile host insert the mobile's home address in the destination address field of all packets they issue. At some point during the routing process this address should be replaced by the address of the FA associated with the mobile host. The entity which performs this operation is called an *address translation agent* (ATA). The process of address translation involves querying the LD, obtaining the FA address,

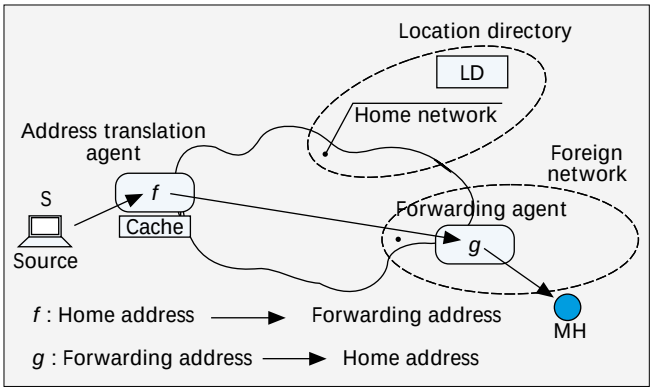

Figure 4. *Packet forwarding model.*

and subsequently making use of this address in forwarding packets to the correct location of the mobile host. The address translation function is:

f(home address) → (forwarding address)

From a two-tier addressing perspective, an ATA initializes the forwarding address part of the destination address. In an actual implementation this could be achieved by prefixing the original destination address of the packet with the FA's address. This operation can be performed at the source host; the only problem is that the function *f* cannot be computed without making changes to the existing host software of millions of hosts.

For performance reasons, an ATA may decide to cache frequently used LD entries in making forwarding decisions. Querying the LD before each address translation operation could be prohibitively expensive, particularly so when the ATA and LD are geographically separated. Caching, however, introduces a new requirement in the architecture: maintaining consistency between the LD and its cached entries throughout the Internet.

Location Update Protocol

Keeping the LD up to date in the face of frequently changing host locations is crucial. Keeping cached LD entries consistent with the master LD is an equally important consideration. Inconsistencies could make mobile hosts inaccessible and even cause the formation of routing loops in some cases. The purpose of the *location update protocol* (LUP) is to provide reliable mechanisms for keeping the LD and its cached copies consistent at all times.

To a large extent, the choice of LUP depends on the caching policy used. Together, they determine the scalability and routing characteristics of a mobility solution. In systems that do not permit LD caching, ATAs must be collocated with the LD, since issuing an LD query for each packet an ATA forwards is prohibitively expensive. In such systems, packets addressed to mobile hosts first travel all the way to the home network before any address translation (function *f*) is performed. Clearly, the paths that packets follow are nonoptimal in this case. Caching improves the routing efficiency of a mobile networking system because packets do not have to travel to home networks before being forwarded toward the FAs associated with the destinations. At the same time, caching makes the system more complex and vulnerable to security attacks. If cache updates are not properly authenticated, it is possible to redirect packets away from a mobile host and cause denial of service.

Packet Forwarding Operation

With the inclusion of address translation agents and forwarding agents, the operation of packet forwarding can be easily illustrated. Figure 4 illustrates how packets from a stationary host (S) are routed to a mobile host (MH). S sends out packets which are addressed to the home address of the MH. These are intercepted by an ATA which maps (using function *f*) the original destination of the packet to the address of the forwarding agent. Once these packets arrive at the forwarding agent, the FA remaps (using function *g*) the destination to the home address of the mobile host and delivers them to the mobile host. Along the path from the source to the destination, packets twice undergo an *address translation* operation. The end result of this translation process, the function *gof*, is an identity mapping, which means that the whole process of address translation is completely transparent to hosts located at both ends of the path, which communicate as if they were stationary. The transport-layer protocols and the applications running on stationary as well as mobile hosts operate without any modifications whatsoever. This property of the solution architecture is termed *transport-layer transparency*.

The proposed architecture preserves transport-layer transparency regardless of where and how in the network the LD, ATAs, and FAs are distributed. This flexibility enables us to capture the design choices made in other mobile IP proposals. In the next section, we'll show that each one of these proposals can be viewed as a special case of the proposed architecture.

Address Translation Mechanisms

So far we have described how various components of the architecture cooperate with each other to perform the necessary address translation operations, but the actual mechanisms were not mentioned. Within the Internet there are two possible ways of doing it: using either *encapsulation* or *loose source routing* (LSR). A brief description of both follows.

Encapsulation — In the encapsulation method a new header is appended at the beginning of the original datagram (Fig. 5). The outer header contains the address of the forwarding agent, while the inner header contains the home address of the mobile host. Since the Internet routing system only looks at the outer datagram header, it routes this packet to the forwarding agent. The forwarding agent strips the outer datagram header and delivers the original datagram locally to the mobile host.

Loose Source Routing — Loose source routing is an option supported in IP which can also be used to perform address translation operation.[3] Using IP's source-routing option, an ATA can cause datagrams addressed to a mobile host's home address to be routed via an FA. Figure 6 illustrates how this is done. An LSR option is used to specify a list of addresses. The Internet routing system routes the datagram containing the LSR option to each address, one by one, in the sequence in which they appear in the list. The current destination is kept in the destination address field of the datagram header, and a pointer points to the address to be visited next in the sequence. When the datagram arrives at the current destination, the contents of the destination address field are swapped with the address pointed to by the next-hop pointer, and the pointer is advanced to the next address in the list. This process is repeated until the datagram is delivered to the address that occurred last in the original list of addresses included in the LSR option. At this point the next-hop pointer in the LSR option points past the last address.

As a natural consequence of LSR option processing, the

[3] *Originally it was included in IP not for this purpose, but to help debug network problems.*

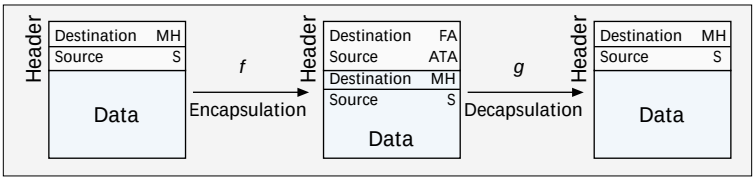

Figure 5. *Illustration of encapsulation and decapsulation.*

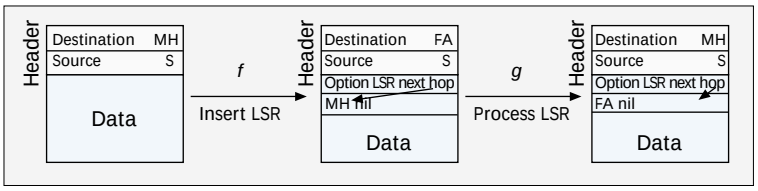

Figure 6. *Using loose source routing to perform address translation.*

path a packet follows (the list of addresses visited en route) is automatically recorded in the packet. The destination can reverse this list and send a reply back to the source along the reverse path. In [8, 9] authors show how this feature is used to design a mobile networking scheme that collocates the ATA with the source and the FA with the destination.

In this section we showed how components of the proposed architecture mutually cooperate to overlay a packet-forwarding service on top of an existing routing infrastructure. It is important to point out that the ATA and FA only represent functions that must be supported, not machines that must be deployed in the network. In fact, the proposed architecture allows a great deal of flexibility in placement of these functions in the network. This flexibility allows us to experiment with various design alternatives and devise a solution for a specific target environment.

Mapping to Candidate Mobile IP Proposals

Over the past several years, many proposals have been made for supporting host mobility on datagram-based internetworks. A vast majority of these proposals have been designed to be compatible with today's TCP/IP-based Internet. The candidate proposals differ widely in terms of the specific components they propose to add to the Internet, the mechanisms they use for address translation, and the policy they use for managing location updates. In this section, we will show that all mobile IP proposals can be viewed as a special case of our proposed network architecture.

In our model, the ATA and FA represent the two basic functions that must be supported by any proposal that supports mobility. We will demonstrate this fact by explaining the operation of each mobile IP proposal in terms of these two functional entities. Basically, all proposals attempt to provide an address translation service through deployment of some additional entities in the network. They differ only in terms of their choice of where they locate these functions, the specific LUP they use, and whether they use encapsulation or source routing to effect address translation. Below we present a short summary of related mobile IP proposals, with a short note following each proposal outlining how its operation can be captured by our proposed solution architecture.

Columbia Scheme

The scheme proposed by Ioannidis *et al.* [10, 11] is designed primarily to support mobility within a campus environment. Mobile hosts are allocated addresses from a subnetwork which is reserved for use by wireless hosts. A group of cooperating mobile support routers (MSRs) advertise reachability to the wireless subnet. MSRs provide an access point through which mobile hosts can connect to the campus backbone, and are also responsible for forwarding traffic to and from mobile hosts. Each mobile host, regardless of its location within a campus, is always reachable via one of the MSRs. When a host sends a packet to a mobile host, it first gets delivered to the MSR closest to the source host. This MSR either delivers the packet (if the destination mobile host lies in its wireless cell) or forwards it to the MSR responsible for the destination mobile host. If an MSR does not know which MSR is currently responsible for a destination, it sends a `WHO_HAS` query to all MSRs in the campus and awaits a reply message from the responsible MSR. When sending a packet to the destination, an MSR encapsulates the packet and delivers it to the target MSR. Upon receiving this packet, the target MSR strips the encapsulation header and relays the original packet to the mobile host.

In the Columbia proposal (Fig. 7), an MSR performs both *encapsulation* and *decapsulation* operations, meaning that both functions, *f* and *g*, are collocated at the MSR. For packets addressed to mobile hosts in its coverage area, an MSR acts like an FA. For packets addressed to other mobile hosts it acts like an ATA. Each MSR maintains a table of mobile hosts in its wireless cell. These tables together constitute the segment of the LD that is associated with mobile hosts on the campus network. This LD distribution scheme can also be thought of as a distributed realization of the owner-maintains rule. Recall that in the owner-maintains rule, the segment of the LD was collocated with the home router. An MSR in the Columbia scheme is a distributed realization of the home router. As a result, the table of mobile hosts maintained at an MSR constitutes a distributed segment of the LD that is required to be maintained at the home router.

MSRs acquire LD cache entries on a need-to-know basis by sending a multicast `WHO_HAS` query to all MSRs in the campus. The response to this query is generated by the MSR that possesses the primary copy (in other words, the MSR responsible for the destination mobile host). The LUP uses a *lazy-update* approach. When a mobile host moves, only the primary copy and the previous copy of the LD entry is updated. Cached entries are assumed to be correct by default. When cached entries turn stale, the first packet that is forwarded using the stale entry generates an error message from the old MSR, causing the source MSR to flush its cache and then multicast a `WHO_HAS` message.

Because functions *f* and *g* are required to be supported only in new entities (MSRs) added to the system, the Columbia proposal can operate without requiring any modifi-

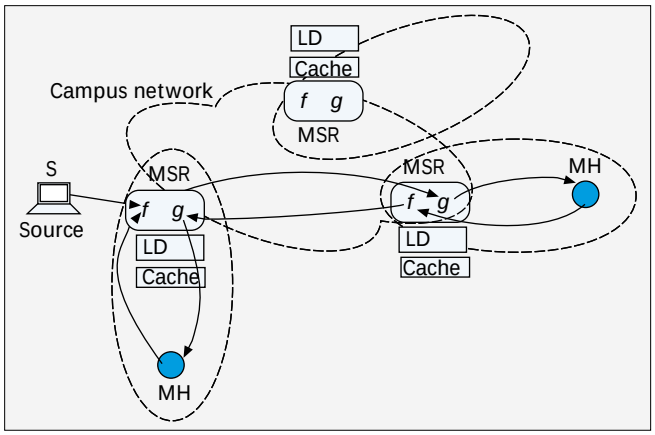

Figure 7. *Mapping to the Columbia proposal.*

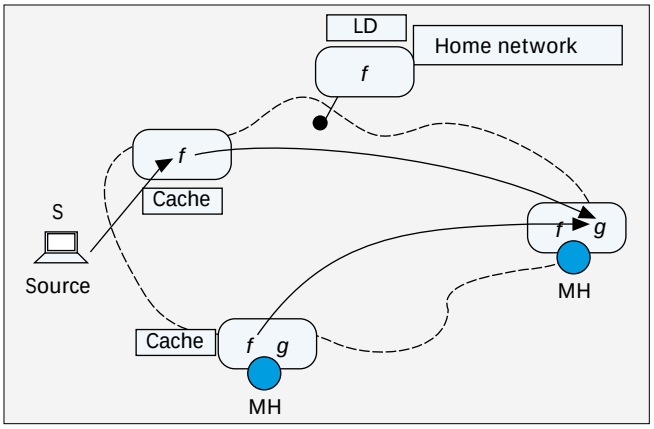

Figure 8. *Mapping to the Sony proposal.*

cations to the existing host and router software. This proposal presents a good combination of design choices for handling mobility within a campus environment; however, it does not scale well for use with the global Internet.

Sony Scheme

In Sony's proposal [12–14] (Fig. 8), a mobile host is assigned a new temporary address when it is attached to a new network. The router of the home network is notified of this new address through a special control message. Packets addressed to the mobile host, in addition to carrying its home address, can also carry its temporary address. Packets originating from a mobile host that is away from its home network always carry both home and temporary addresses in the source address field. Routers that forward these packets can examine the source addresses and cache the mapping (home to temporary) in their address mapping tables (AMTs). A source includes both addresses in all outgoing packets if it already has an AMT entry for the target host; otherwise, packets are forwarded to the home address. If a transit router has an AMT cache entry for the destination, it can intercept the packet and forward it to its correct location. If none of the transit routers have a cache entry, the home router is eventually responsible for forwarding the datagram.

When a host moves to a new location, all AMT cache entries are invalidated through a special `disconnect` control message which is broadcast in the network. Since this message of invalidation is not reliable, there is also a timeout associated with all AMT cache entries, which, on expiration, causes AMT entries to be purged.

This method requires modifications to routers and host software and has problems interoperating with the existing hosts because it also requires modifications to IP packet formats.

The Sony proposal collocates the forwarding agent function, *g*, with mobile hosts. In other words, it requires each mobile host to act as its own forwarding agent. The advantage is that packets can be directly tunneled to the mobile host, without intervention from a forwarding agent. This is useful, particularly for wired mobile hosts, which may at times connect to foreign networks that have no forwarding agents attached. The approach of collocating *g* with the mobile host has a disadvantage. It doubles the address space requirement for mobile hosts because, in addition to a home address, a temporary address is also required for operation. In some environments, this may be a serious problem.

In Sony's proposal, the home router acts as an ATA, and it also maintains the LD for mobile hosts that have been assigned addresses on the home network. To avoid routing each packet via the home router, the Sony proposal allows flexibility to collocate *f* with internet routers. Since LD cache entries are carried in the source address field of the Virtual IP protocol,[4] routers can acquire them just by inspecting the source address of packets they relay. Distributing LD caches across the Internet improves routing performance; however, it makes updates very costly. Sony's proposal, therefore, has a scalability problem. When a host moves to a new location, it is required to send a broadcast in the network to purge all cached LD entries.

LSR Scheme

In contrast with other proposals, which are encapsulation-based, the LSR proposal [8, 9, 15, 16] (Fig. 9) is based on the use of an existing IP option, LSR. The LSR scheme also allows each mobile host to retain its home address regardless of its current location. Associated with each home network is a mobile router (MR), which is responsible for advertising reachability to the home network, and for keeping track of the current location of each mobile host that has been assigned an address on that network. In a foreign network, mobile hosts attach to the Internet via wireless base stations known as mobile access stations (MASs). When a mobile host walks into the wireless cell of an MAS, it informs its MR of the internet address of the current MAS. The MR records this information in its routing table, and also informs the previously recorded MAS that the mobile host has moved out of its wireless cell. The packets sent to the mobile host first arrive at the MR by the normal routing process. To forward a packet to the mobile host's current location, the MR inserts an LSR option in the packet, specifying the current MAS as a transit router. The inserted LSR option causes this packet to be routed to the mobile host via the MAS. When the mobile host sends a reply to the source, it also inserts the LSR option in all outgoing packets, again specifying the current MAS as a transit router. When the stationary host receives this packet, it will reverse the recorded route, and insert it in all outgoing packets that are sent to the mobile host. Thus, subsequent packets originating from the stationary host will be automatically routed along an optimal path.

Notice that this proposal relies on the end host's ability to perform route reversal. Unfortunately, the vast majority of hosts in the Internet either do not perform correct route reversal or, in some cases, even drop LSR packets due to the security risk involved. Another problem is that packets carrying the LSR option receive poor service from IP routers. Most router vendors optimize their forwarding loop for the common case of a simple IP header. When a packet with options

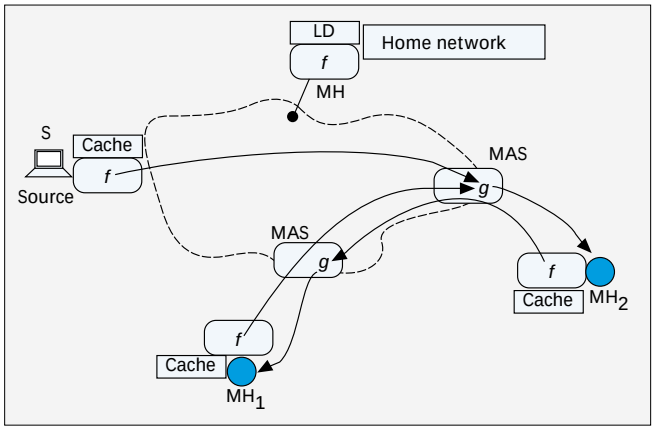

Figure 9. *Mapping to the LSR scheme.*

[4] *The modified IP protocol*

is received, it is bumped into a low-priority queue. Due to these limitations, the LSR proposal was not accepted as a candidate for further consideration within the Internet Engineering Task Force (IETF).

In this proposal, the MR acts as an ATA and is also responsible for maintaining the LD. The MAS acts as an FA for mobile hosts that lie in its wireless cell. The key feature of this proposal is that it enables function *f* to be collocated with all internet hosts without requiring changes to host software. All internet hosts, when generating replies to packets that are received with the LSR option, are required to do the route reversal [17]. For TCP connections, the route reversal is performed by the protocol processing module and, in User Datagram Protocol (UDP) connections, this responsibility lies with the applications. From our reference architecture viewpoint, the process of route reversal corresponds to the task which an ATA is required to carry out. Thus, this scheme attempts to exploit mechanisms already available within IP to achieve collocation of the ATA with end hosts.

The LSR scheme maintains a distributed version of the LD, but no special protocol is required for distributing and managing LD cache entries. LD entries are automatically acquired through the incoming LSR option. Recall that packets which arrive at a stationary host already contain the address of the MAS. This, together with the source address of the packet, constitutes an LD cache entry. When a host starts a new session with a mobile host, it has no LD cache entry for the destination. Naturally, the first packet is routed to the destination via the MR. When the ACK for this packet arrives, it contains the LD cache entry in the incoming LSR option. This LD entry is maintained, on a per-session basis, only as long as the corresponding TCP session is alive. When the session terminates, the corresponding LD entry is purged. If the destination moves during an active session, the LD cache entry becomes inconsistent; however, it gets updated as soon as the next packet from the destination arrives at the source. This constitutes a pure on-demand cache update policy which has a good scaling property. Following a host's movement, only those LD cache entries which are in use are updated. Compared to Sony's proposal, which requires a message to be broadcast to the network, significantly fewer messages are exchanged. Naturally, an on-demand cache update policy lends itself to a scalable design, with respect to both the size of the network and the rate of host mobility.

Mobile IP Working Group Proposal

IETF has created a Mobile IP working group to come up with a proposal for near-term deployment within the Internet (Fig. 10). In this design [18], each mobile host retains its home address regardless of the mobile host's location. When the mobile host visits a foreign network, it is associated with a *care-of address*, which is an Internet address associated with the mobile host's current point of attachment. The care-of address identifies either the mobile host directly, if the address is acquired through Dynamic Host Configuration Protocol (DHCP) [19], or a foreign agent that is responsible for providing access to visiting mobile hosts. When away from home, the mobile host registers its care-of address with a home agent; the home agent is responsible for intercepting datagrams addressed to the mobile host's home address and tunneling (encapsulating) them to the associated care-of address.

In this scheme all datagrams addressed to a mobile host are always routed via the home agent. However, the packets in the reverse direction (i.e., those originating from the mobile host and addressed to a stationary host) are relayed along the shortest path by the Internet routing system. This gives rise to what is known as the triangle routing problem.

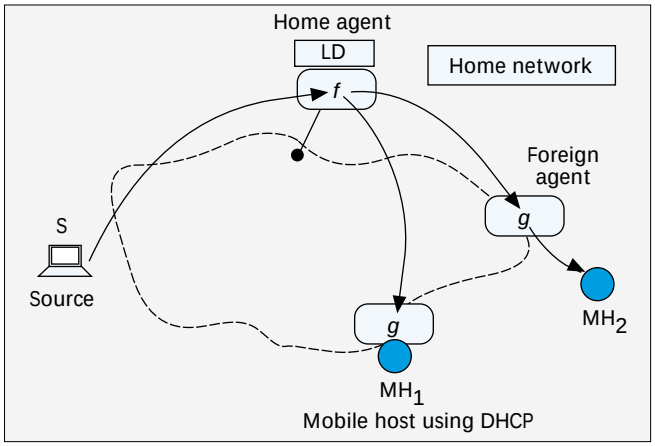

Figure 10. *Triangle routing: mobile IP proposal.*

Route optimization is possible if the location information is allowed to be cached; however, this proposal does not permit caching of LD entries because of security concerns. Currently, the Internet does not provide any secure mechanism for distributing cache entries. Any entity in the Internet can masquerade as a home agent and reroute traffic away from a mobile host just by redistributing fake cache entries. This proposal, therefore, takes the stand that routing based on cached location information is insecure, and the best possible defense against security attacks is to not use it at all. The cost of this choice is that routing is always nonoptimal.

When the mobile host arrives at a foreign network, it can listen for (or solicit) agent advertisements to determine whether a foreign agent is available. If so, the registration request to the Home Agent is sent via the foreign agent; otherwise, the mobile host must acquire a care-of address (through DHCP) and then register with the home agent.

The IETF Mobile IP proposal reflects a design choice that collocates *f* with the home agent and *g* with the foreign agent. This proposal also allows *g* to be collocated with the mobile host. This happens when the mobile host acquires a temporary address via DHCP or Point-to-Point Protocol (PPP). The LUP is very simple; the mobile host notifies the home agent whenever it moves to a new location. Since the LD entries are never cached, the question of maintaining consistency does not even arise.

Mobile IP with Route Optimization

Route optimization [20] is basically a protocol by which Internet hosts can learn the current care-of address for a mobile node — that is, they can create a valid binding (an LD cache entry) for the mobile node, and become ATAs (Fig. 11). Once an Internet host has a valid binding, the host can encapsulate packets and send them directly to the care-of address for the mobile node, just as the mobile node's home agent does in the basic Mobile IP specification. The correspondent host can also optionally use an abbreviated style of encapsulation called *minimal encapsulation* [21], which typically in this case adds 8 bytes to the original IP datagram.

Aside from the difficulty of changing existing Internet hosts to use new techniques to deal with mobility, route optimization faces the additional technical difficulty (and requirement) of enabling the recipient hosts to be sure that the location update information is authentic. The absence of authentication techniques would leave a home agent vulnerable to cooperating with a malicious host that wanted to hijack traffic destined for mobile hosts. Similarly, any other host accepting cache updates on the mobile node's location needs to be able to ensure that the updates are authentic. Providing a high enough degree of confidence in the authenticity of the

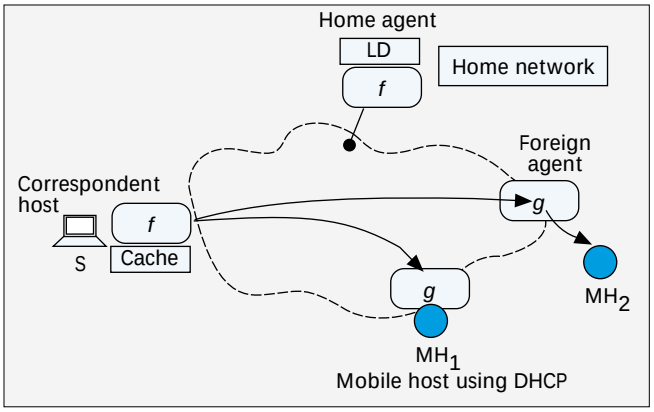

Figure 11. *Mobile IP with route optimization.*

location updates has been a driving factor in the design of the route optimization protocol. The goal, then, of route optimization is to enable the delivery of authentic binding updates (as needed) to arbitrary Internet hosts.

If a correspondent host has no binding for a mobile node, the home agent will receive packets from the correspondent host destined for the mobile node. In this case, the home agent is well placed to send a valid binding to the correspondent host. If the correspondent host has a stale (incorrect) binding for a mobile node, the situation is more difficult. Usually, the binding associates the mobile node to a care-of address offered by a foreign agent that no longer serves the mobile node. In this case, the foreign agent notifies the correspondent host (via a *binding warning* message) to request a new binding update from the home agent. Route optimization assumes that foreign agents typically have no security association with correspondent hosts, and thus cannot send authenticated binding updates directly to them.

If the binding associates the mobile node to a care-of address that is stale or does not exist, then the correspondent host will have to purge its binding in response to an Internet Control Message Protocol (ICMP) message indicating that the care-of address is unreachable. The worst case occurs when the care-of address is reachable, but consumes the datagram with no ICMP error and still does not deliver it to the mobile node. In this case, the correspondent node will have to rely on timing out the binding for the mobile node according to its lifetime or, in more fortunate circumstances, purging the binding because of error conditions generated by higher-level protocols (such as TCP or the application protocol generating the datagrams). Fortunately, the latter case is likely to be rare, only occurring when a correctly functioning foreign agent hangs.

The route optimization proposal improves the basic Mobile IP design by collocating *f* with correspondent hosts. The location update protocol operates to maintain valid LD cache entries at those hosts which are likely to send packets to the mobile node. This constitutes an on-demand location update protocol. Binding update messages are authenticated in order to ensure that malicious hosts cannot disrupt traffic between correspondent hosts and mobile nodes using the route optimization protocol. Smooth handoffs between foreign agents are modeled as a special case of the route optimization techniques, and session keys are obtained to secure the necessary binding update messages in this common case.

IPv6 Mobility Proposal

A new version of IP has been designed, and is conventionally known as IPv6 (see [22] for protocol details). IPv6 has a 128-bit address space and a greatly improved scheme for handling options (Fig. 12). Part of the basic set of requirements for IPv6 is that it be designed to handle mobility well. In this section we will briefly describe a proposal which achieves this objective for IPv6. The existing version of IP will sometimes be referred to as IPv4 (IP version 4) to distinguish it from IPv6.

Because there is no substantial base of installed IPv6 systems yet, the IPv6 proposal is not constrained by compatibility requirements with existing systems. This is a huge advantage; for example, the main obstacle to designing an efficient protocol for route optimization with the base (IPv4) mobile IP specification is that most existing Internet hosts cannot be expected to successfully interpret the route optimization messages. If all IPv6 hosts support the mobile-IPv6 protocol described in this section, almost all traffic for mobile nodes will follow optimal routes. Also, since every IPv6 router can be assumed to support mobile IPv6, every network can be a home network, and any node on that network can roam the Internet.

IPv6 mobility operates by carefully sending binding updates (i.e., LD cache entries) to any host likely to need them. The binding update associates a care-of address (just as with IPv4) with the mobile node's home address — but with IPv6 no foreign agents are needed. This is because g can always be collocated with IPv6 hosts. Each mobile node receives packets at the care-of address it obtained via neighbor discovery at its current point of attachment; nevertheless, its correspondent hosts send datagrams addressed to the mobile node at its home address. However, before the correspondent host actually transmits the datagram, it places the care-of address in a *routing header*, which is analogous to a *loose source route* in IPv4 (see the section on the LSR scheme). In this way, a mobile node can move from one care-of address to another and still maintain all of its existing connections, which are associated with its home address, not some fleeting care-of address. However, as the mobile node moves from place to place, the correspondent host inserts different routing headers in the datagrams before transmitting them.

There is still a home agent (ATA) needed for IPv6, in case a correspondent IPv6 host does not have a binding for a mobile node; however, the home agent should not see many packets. Because the home agent tunnels packets to a mobile node by using encapsulation instead of a routing header, the mobile node can easily detect whenever its correspondent hosts do not have a binding for it.

In IPv6, the mobile node is always responsible for delivering binding updates to its correspondent hosts. If a mobile node moves to a new point of attachment, then any correspondent hosts that have recently sent packets to the mobile node should get binding updates. Most correspondent hosts with open TCP connections with the mobile node should get binding updates. Also, as just mentioned, any time a mobile node gets a datagram encapsulated by the home agent, the

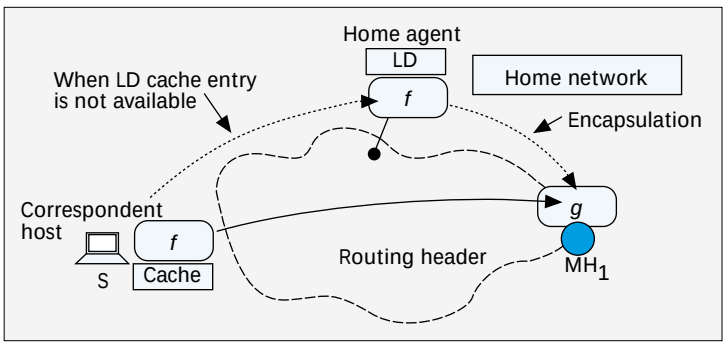

Figure 12. *IPv6 mobility proposal.*

Scheme	Address translation	Forwarding agent (*g*)	Location directory	Location update protocol
Columbia	Collocated with MSR	Collocated with MSR	Distributed among MSRs	Only primary copy modified; lazy-update policy used for updating cache entries
Sony	Collocated with all hosts and routers	Collocated with mobile hosts	LD maintained at home router; cache entries acquired by snooping a packet header	Only primary copy modified by the explicit connect message; cache entrees modified by broadcasting a `disconnect` message, or auto-flushed by a timeout message
LSR scheme	Collocated with all hosts and home routers	Collocated with mobile hosts	LD maintained at home router; cache entries acquired through incoming LSR option	Only primary copy modified; cache entries automatically updated when packets with new LSR option arrive; on-demand update policy, no broadcasts
IETF Mobile IP	Collocated with home routers	Collocated with foreign agent, or with mobile host if DHCP is used	LD maintained at home router only	Caching of LD entries not allowed; when a host moves, only primary copy modified
With route optimization	Collocated with correspondent hosts		Correspondent hosts able to cache LD entries	Home agent responsible for sending binding update message.
IPv6 proposal	Collocated with all hosts and home routers	Collocated with mobile hosts	LD maintained at home router; cache entries acquired through binding update messages issued by mobile hosts.	Mobile host responsible for updating the primary copy and all LD cache entries; on-demand update policy, no broadcasts

Table 1. *Functional comparison of mobile IP schemes.*

mobile node should certainly send a binding update to the source of that datagram. In IPv6, the mobile node that does the best job of sending out binding updates appropriately will receive the best performance from the Internet, and will place the least load on the Internet. Of course, any such binding update sent to a correspondent host should be authenticated to allow the correspondent host to trust the veracity of the update.

Because there are no foreign agents, we should consider the effects of moving from one point of attachment to another. With each such point of attachment, the mobile node will get a new care-of address; unfortunately, there are no foreign agents to help effect smooth handoffs from one point of attachment to the next. However, and especially in the case of wireless communications, there is no reason for the mobile node to halt operation at its previous care-of address. If the mobile node is still within range of the previous point of attachment, it can still receive packets at its previous care-of address. Because each IPv6 node is required to be able to handle multiple IPv6 addresses at each of its network interfaces (i.e, each IPv6 must have multihoming capability), each IPv6 mobile node must be prepared to handle multiple care-of addresses as needed. With this in mind, we expect that smooth handoffs will easily be processed by the mobile nodes themselves without any help from foreign agents. In the case of wired attachments, of course, smooth handoffs are simultaneously harder to visualize and more difficult to provide for by such multihoming techniques.

The IPv6 mobility proposal reflects an ideal design choice which collocates *f* and *g* with all IPv6 hosts. The proposal employs a new address translation mechanism, the routing

Property	Columbia		Sony	LSR	IETF Mobile IP		IPv6
	Local	Popup				+ route-opt	
Optimal routing	Always	Never	Only if all routers are modified	Provided hosts support correct route reversal	Never	Always	Always
Address translation mechanism	Encapsulation		Encapsulation	LSR	Encapsulation		Routing header
Additional address space required	None	Double	Double	None	None, but required when using DHCP		Double
Scalability	Good		Poor	Good	Good		Excellent
Compatibility with IP	Total		Requires changes	Only if hosts and routers conform to standards	Total	Requires changes to hosts	Compatible with IPv6
Security	No authentication		Insecure	Insecure	Fully secure		Security via authenti cation header

Table 2. *Comparison of mobile IP schemes.*

header, which is functionally similar to the LSR option of IPv4. The location update protocol is based on an on-demand update policy, and mobile hosts are responsible for issuing all location updates. Overall, IPv6 looks very promising as an efficient and natural protocol for supporting mobility (see [23] for details about the protocol).

Summary

In this article, we first identified network-layer concepts that play a crucial role in the design of mobile networking systems. We showed that the process of *address translation* is fundamental to any mobility solution at the network layer. Our proposed network architecture employs three basic entities: the *address translation agent*, *forwarding agent*, and *location directory*, which cooperate with each other to carry out the operation of address translation. We showed that all candidate proposals for mobile IP can be visualized as a specific instance of our general architecture. We demonstrated this by showing a one-to-one mapping between the entities in our architecture and those required by the candidate proposals. The mappings represent a set of design choices (i.e., where in the network these entities are located) made in the candidate proposals. Tables 1 and 2 present a summary of our observations.

In addition to these design choices there are several other considerations, such as interoperability, backward compatibility, security, and authentication, which also play a crucial role in the design of a mobile networking system. Interested readers can refer to articles [11, 14, 24, 25] for an in-depth description of design and implementation issues.

References

[1] J. Postel, "Internet Protocol," RFC 791, Sept. 1981 (ftp://ds.internet.nic/rfc).
[2] J. Postel, "Transmission Control Protocol," RFC 793, Sept. 1981 (ftp://ds.internet.nic/rfc).
[3] P. S. Ford, Y. Rekhter, and H.-W. Braun, "Improving the Routing and Addressing of IP," *IEEE Network*, vol. 7, no. 3, May 1993, pp. 10–15.
[4] B. Awerbuch and D. Peled, "Concurrent On-Line Tracking on Mobile Users," *Proc. ACM SIGCOMM*, 1991, pp. 221–33.
[5] V. Anantharam *et al.*, "Optimization of a Database Hierarchy for Mobility Tracking in a Personal Communications Network," *Proc. Performance '93*, Sept. 1993.
[6] A. Bar-Noy, I. Kessler, and M. Sidi, "Mobile Users: To Update or Not to Update?" *Proc. IEEE INFOCOM*, Toronto, Canada, June 1994, pp. 570–76.
[7] A. Bar-Noy and I. Kessler, "Tracking Mobile Users in Wireless Communication Networks," *IEEE Trans. on Info. Theory*, Jan. 1994, pp. 45–65.
[8] C. Perkins and P. Bhagwat, "A Mobile Networking System Based on Internet Protocol," *IEEE Pers. Commun. Mag.*, vol. 1, no. 1, Feb. 1994, pp. 32–41.
[9] C. Perkins and Y. Rekhter, "Short-cut Routing for Mobile Hosts," expired Internet draft, July 1992.
[10] J. Ioannidis, D. Duchamp, and G. Q. Maguire Jr., "IP-Based Protocols for Mobile Internetworking," *Proc. ACM SIGCOMM,* 1991, pp. 235–45.
[11] J. Ioannidis and G. Q. Maguire Jr., "The Design and Implementation of a Mobile Internetworking Architecture," *Proc. Winter USENIX*, San Diego, CA, Jan. 1993, pp. 491–502.
[12] F. Teraoka, Y. Yokote, and M. Tokoro, "A Network Architecture Providing Host Migration Transparency," *Proc. ACM SIGCOMM*, Sept. 1991, pp. 209–20.
[13] F. Teraoka and M. Tokoro, "Host Migration Transparency in IP Networks," *Comp. Commun. Review*, Jan. 1993, pp. 45–65.
[14] F. Teraoka *et al.*, "VIP: A Protocol Providing Host Mobility," *Commun. ACM*, vol. 37, no. 8, Aug. 1994.
[15] P. Bhagwat and C. Perkins, "A Mobile Networking System Based on Internet Protocol (IP)," P*roc. USENIX Symp. on Mobile and Location Independent Comp.*, Cambridge, MA, Aug. 1993, pp. 69–82.
[16] D. B. Johnson, "Mobile Host Internetworking Using IP Loose Source Routing," Tech. Rep. CMU-CS-93-128, Carnegie Mellon Univ., Pittsburgh, PA, Feb. 1993.
[17] R. Braden, "Requirements for Internet Hosts — Communication Layers," RFC 1122, Oct. 1989 (ftp://ds.internet.nic/rfc).
[18] C. Perkins, draft-ietf-mobileip-protocol-15.txt;. draft RFC, work in progress, Feb. 1996 (ftp://ds.internet.nic/internet-drafts).
[19] R. Droms, "Dynamic Host Configuration Protocol," RFC1541, Oct. 1993. (ftp://ds.internet.nic/rfc)
[20] D. B. Johnson and C. Perkins, draft-ietf-mobileip-optim-03.txt; draft RFC, work in progress, Nov. 1995 (ftp://ds.internet.nic/internet-drafts).
[21] C. Perkins, draft-ietf-mobileip-minenc.txt; draft RFC, work in progress, Oct. 1995 (ftp://ds.internet.nic/internet-drafts).
[22] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," RFC 1883, Dec. 1995 (ftp://ds.internet.nic/rfc).
[23] C. Perkins and D. Johnson, draft-ietf-mobileip-ipv6-00.txt; draft RFC, work in progress, Jan. 1996 (ftp://ds.internet.nic/internet-drafts).
[24] H. Wada *et al.*, "Mobile Computing Environment Based on Internet Packet Forwarding," *Proc. Winter USENIX*, San Diego, CA, Jan. 1993, pp. 503–17.
[25] A. Myles and D. Skellern, "Comparing Four IP Based Mobile Host Protocols," *Proc. Joint Euro. Networking Conf.*, May 1993.

Biographies

PRAVIN BHAGWAT received the B. Tech. degree from the Indian Institute of Technology, Kanpur, India in 1990, and the M.S. and Ph.D. degrees in computer science from the University of Maryland, College Park, in 1992 and 1995, respectively, all in computer science. He was an IBM Graduate Fellow during his graduate studies. He joined IBM's Thomas J. Watson Research Center in 1995. Currently, he is a research staff member in the Wireless and Mobile Network Architecture group. His research interests include routing for mobile hosts, wireless networks, transport protocols, and distributed applications.

CHARLES PERKINS is a research staff member at IBM's T. J. Watson Research Center, investigating mobile and ad hoc networking, resource discovery, and automatic configuration for mobile computers. He is serving as the document editor for the Mobile IP working group of the Internet Engineering Task Force (IETF), and serves on the editorial boards for *ACM/IEEE Transactions on Networking*, *ACM Wireless Networks*, and *IEEE Personal Communications*. Charles holds a B.A. in mathematics and an M.E.E. degree from Rice University, and an M.A. in mathematics from Columbia University.

SATISH K. TRIPATHI is a professor in the Department of Computer Science at the University of Maryland, College Park. He served as the department chair from 1988 to 1995. Dr. Tripathi's research interests are in the areas of computer networks, mobile computing, ATM networks, and operating systems support for multimedia information. He has guest edited special issues of many journals and serves on the editorial boards of *Theoretical Computer Science*, ACM/Springer *Multimedia Systems*, and *IEEE Transactions on Computers*.