To appear in *Proceedings of the Ninth Annual IEEE Workshop on Computer Communications*, October 1994.

Scalable Support for Transparent Mobile Host Internetworking

David B. Johnson

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213-3891

dbj@cs.cmu.edu

Abstract

This paper considers the problem of providing transparent support for very large numbers of mobile hosts within a large internetwork such as the Internet. The availability of powerful mobile computing devices and wireless networking products and services is increasing dramatically, but internetworking protocols such as IP used in the Internet cannot currently support host movement. To address this need, the Internet Engineering Task Force (IETF) is currently developing protocols for mobile hosts in the Internet. This paper reviews the current state of that effort and discusses its scalability to very large numbers of mobile hosts.

1. Introduction

The global Internet is growing at a tremendous rate. There are now about 3.5 million hosts connected to the Internet, and this number is doubling approximately every year. The average time between new networks connecting to the Internet is about 10 minutes. Initiatives such as the National Information Infrastructure and the increasing commercial uses of the Internet are likely to create even faster growth in the future.

At the same time, portable computing devices such as laptop and palmtop computers are becoming widely available at very affordable prices, and many new wireless networking products and services are becoming available based on technologies such as spread-spectrum radio, infrared, cellular, and satellite. Mobile computers today often are as capable as many home or office desktop computers and workstations, featuring powerful CPUs, large main memories, hundreds of megabytes of disk space, multimedia sound capabilities, and color displays. High-speed local area wireless networks are commonly available with speeds up to 2 megabits per second, and wide-area wireless networks are available that provide metropolitan or even nationwide service.

With these dramatic increases in portability and ease of network access, it becomes natural for users to expect to be able to access the Internet at any time and from anywhere, and to remain connected and continue to use the network as they move about. However, internetworking protocols such as IP [15] used in the Internet cannot currently support host mobility. A mobile user, today, must usually at least modify a number of configuration files and restart all network connections when connecting to the Internet at a different point or through a different network, making host movement difficult, time consuming, and error prone.

To address this need in the Internet, the Mobile IP Working Group of the Internet Engineering Task Force (IETF) has been working over the past few years to develop standard protocols to support mobile hosts operating in the Internet [4, 5, 8, 9, 12, 13, 14, 17, 21, 22, 24]. Section 2 of this paper describes the general problem of mobility management and packet routing to mobile hosts in a large internetwork. Section 3 gives a summary of the current state of the basic IETF Mobile IP protocol, and Section 4 describes extensions to this protocol also being developed within the IETF for optimizing packet routing to mobile hosts. Section 5 discusses the scalability of this work to very large numbers of mobile hosts, and Section 6 presents conclusions.

This research was supported in part by the Wireless Initiative of the Information Networking Institute at Carnegie Mellon University.

2. Problem Analysis

2.1. Internetwork Routing

In order to provide scalable routing support, internetworking protocols such as IP [15], ISO CLNP [19], NetWare IPX [23], and AppleTalk [20], use *hierarchical* addressing and routing schemes. For example, in IP, the network address of a host is divided into two levels of hierarchy: a *network number* identifying the network to which the host is connected, and a *host number* identifying the particular host within that network. Routers within the Internet know (and care) only how to route packets based on the network number of the destination address in each packet; once the packet reaches that network, it is then delivered to the correct individual host on that network.

Aggregating the routing decision at each level of the hierarchy in this way reduces the size of the routing tables that each router must maintain, reduces the size of the routing updates that routers must exchange, and simplifies the routing decisions at each router. Hierarchical addressing and routing has proven to be essential to keep up with the exponential growth of the Internet, in particular. The original two-level hierarchy of Internet addressing in IP has already been transparently extended at the bottom with *subnetting* [11] and at the top through use of *CIDR* [3]. In the IETF's "IPng" effort to develop the next generation of the IP protocol, support for many more levels of hierarchy than in the present version of IP is an explicit design goal [10].

This hierarchy, however, prevents packets from being routed correctly to a mobile host while away from its home network, since a host's address logically encodes its location. Without special handling for mobility, packets addressed to a mobile host will be routed by the Internet only to the mobile host's home network. This is true of any hierarchical routing and addressing scheme, whether provider-based or geographical.

2.2. Location Registry

It is important to be able to support packet routing to mobile hosts from existing correspondent hosts that have not been modified to support mobility. Given the very large number of hosts already deployed within the Internet, it seems quite likely that some will not be upgraded to support mobility for some time. Furthermore, some existing hosts may never be upgraded, for example because the organizations owning some hosts may lack the interest or resources to upgrade, or because the original vendor no longer offers support for particular products owned by some customers. The ability to support unmodified correspondent hosts also allows any correspondent host to communicate with any other host without being concerned whether or not it is currently mobile.

It therefore becomes logical to provide basic mobility support for a mobile host through a location registry recording the mobile host's current location, that can be accessed through the mobile host's home network. An unmodified correspondent host (or one that simply does not know that a particular mobile host is in fact mobile) will send IP packets for that mobile host in the same way as all IP packets are sent today. Such packets will thus reach the mobile host's home network, where they may be intercepted by some mobility support agent and forwarded to the mobile host's current location.

Requiring the sender to instead query the location registry before sending is incompatible with the goals of supporting existing unmodified correspondent hosts and of not requiring the sender to be aware of whether a particular destination host is currently mobile before sending. Accessing the location registry through the mobile host's home network also avoids any requirement for changes to the basic routing algorithms of the Internet, and allows each organization owning some network to manage this functionality for all of its own mobile hosts with this home network, improving scalability and easing manageability.

2.3. Packet Tunneling

Some mechanism is needed to cause a packet for a mobile host to be routed to that host's current location rather than (only) to its home network. In order to avoid distributing routing information for a mobile host throughout the Internet so that the new routing decision could be made at each hop, it must be possible to modify each packet for a mobile host in such a way that the routing infrastructure of the Internet will route the modified packet to a location identified in the packet. This type of packet forwarding is known as *tunneling*.

In tunneling a packet from one node to another, only these two nodes (the two endpoints of the tunnel) need know that tunneling is taking place. Routers between the node tunneling the packet and the new destination node to which the packet is tunneled simply route the packet at each hop in the same way as any ordinary IP packet. There is thus no need to modify existing routers, such as within the Internet backbone, nor to modify existing Internet routing algorithms.

2.4. Caching and Consistency

The mechanisms suggested above allow packets for a mobile host to be sent to it at its current location, but support forwarding only through an agent on the mobile host's home network. For example, if a mobile host, say MH1, is visiting some network, even packets from a correspondent host on this same network must be routed through the Internet to this agent on MH1's home network, only to then be tunneled back to the original network for delivery to MH1. If the correspondent host in this example is actually another mobile host, say MH2, then packets from MH1 to MH2 must likewise be routed through some agent on MH2's home network and back to the original network for delivery to MH2. This indirect routing places unnecessary overhead on the Internet, on each mobile host's home network, and on the agent providing forwarding service from each home network. Such indirect routing may also significantly increase the latency in packet delivery to a mobile host.

Correspondent hosts that have been modified to support mobility should be able to learn the current location of a mobile host with which they are communicating, and to then use this location to tunnel their own future packets directly to the mobile host. By caching this location, the expense of discovering this location can be avoided on each individual packet sent to the mobile host. However, this caching creates the problem of *cache consistency* when the mobile host then moves to a new location, since the correspondent host's cache will still point to the old location. In order to support smooth handoff from one location to another, the protocol must be able to update correspondent host's caches and should provide some support for packets that may be tunneled based on a temporarily out-of-date cache.

3. The Basic Mobile IP Protocol

This section provides an overview of the current state of the basic IETF Mobile IP protocol. The protocol provides transparent routing of packets to a mobile host and requires no modification to existing routers or correspondent hosts. No support is provided, however, for caching a mobile host's location at correspondent hosts or for allowing correspondent hosts to tunnel packets directly to a mobile host's current location. These extensions are being produced within the IETF as a separate set of modifications to this basic protocol, and are discussed in Section 4.

3.1. Infrastructure

Each mobile host is assigned a unique *home address* in the same way as any other Internet host, within its *home network*. Hosts communicating with a mobile host are known as *correspondent hosts* and may, themselves, be either mobile or stationary. In sending an IP packet to a mobile host, a correspondent host always addresses the packet to the mobile host's home address, regardless of the mobile host's current location.

Each mobile host must have a *home agent* on its home network that maintains a registry of the mobile host's current location. This location is identified as a *care-of address*, and the association between a mobile host's home address and its current care-of address is called a *mobility binding*, or simply a *binding*. Each time the mobile host establishes a new care-of address, it must *register* the new binding with its home agent so that the home agent always knows the current binding of each mobile host that it serves. A home agent may handle any number of mobile hosts that share a common home network.

A mobile host, when connecting to a network away from its home network, may be assigned a care-of address in one of two ways. Normally, the mobile host will attempt to discover a *foreign agent* within the network being visited, using an *agent discovery* protocol. The mobile host then *registers* with the foreign agent, and the IP address of the foreign agent is used as the mobile host's care-of address. The foreign agent acts as a local forwarder for packets arriving for the mobile host and for all other locally visiting mobile host's registered with this foreign agent. Alternatively, if the mobile host can obtain a temporary

local address within the network being visited (such as through DHCP [2]), the mobile host may use this temporary address as its care-of address.

While a mobile host is away from its home network, a mobile host's home agent acts to forward all packets for the mobile host to its current location for delivery locally to the mobile host. Packets addressed to the mobile host that appear on the mobile host's home network must be intercepted by the mobile host's home agent, for example using "proxy" ARP [16] or through cooperation with the local routing protocol in use on the home network.

For each such packet intercepted, the home agent tunnels the packet to the mobile host's current care-of address. If the care-of address is provided by a foreign agent, the foreign agent removes any tunneling headers from the packet and delivers the packet locally to the mobile host by transmitting it over the local network on which the mobile host is registered. If the mobile host is using a locally obtained temporary address as a care-of address, the tunneled packet is delivered directly to the mobile host.

Home agents and foreign agents may be provided by separate nodes on any network, or a single node may implement the functionality of both a home agent (for its own mobile hosts) and a foreign agent (for other visiting mobile hosts). Similarly, either function or both may be provided by any of the IP routers on a network, or they may be provided by separate support hosts on that network.

3.2. Agent Discovery

The *agent discovery* protocol operates as a compatible extension of the existing *ICMP router discovery* protocol [1]. It provides a means for a mobile host to detect when it has moved from one network to another, and for it to detect when it has returned home. When moving into a new foreign network, the agent discovery protocol also provides a means for a mobile host to discover a suitable foreign agent in this new network with which to register.

On some networks, depending on the particular type of network, additional link-layer support may be available to assist in some or all of the purposes of the agent discovery protocol. A new protocol must be defined for agent discovery, however, at least for use on networks for which no link-layer support is available. By defining a new protocol, mobile hosts are also provided with a a common method for agent discovery that can operate in the same way over all types of networks. If additional link-layer support is available, it can optionally be used by mobile hosts that support it to assist in agent discovery.

Home agents and foreign agents periodically advertise their presence by multicasting an *agent advertisement* message on each network to which they are connected and for which they are configured to provide service. Mobile hosts listen for agent advertisement messages to determine which home agents or foreign agents are on the network to which they are currently connected. If a mobile host receives an advertisement from its own home agent, it deduces that it has returned home and registers directly with its home agent. Otherwise, the mobile host chooses whether to retain its current registration or to register with a new foreign agent from among those it knows of.

While at home or registered with a foreign agent, a mobile host expects to continue to receive periodic advertisements from its home agent or from its current foreign agent, respectively. If it fails to receive a number of consecutive expected advertisements, the mobile host may deduce either that it has moved or that its home agent or current foreign agent has failed. If the mobile host has recently received other advertisements, it may attempt registration with one of those foreign agents. Otherwise, the mobile host may multicast an *agent solicitation* message onto its current network, which should be answered by an agent advertisement message from each home agent or foreign agent receiving the solicitation on this network.

3.3. Registration

Much of the basic IETF Mobile IP protocol deals with the issue of registration with a foreign agent and with a mobile host's home agent. When establishing service with a new foreign agent, a mobile host must register with that foreign agent, and must also register with its home agent to inform it of its new care-of address. When establishing a new temporarily assigned local IP address as a care-of address, a mobile host must likewise register with its home agent to inform it of this new address. Finally, when a mobile host returns to its home network it must register with its home

agent to inform it that it no longer is using a care-of address.

Each registration with a home agent or foreign agent has associated with it a *service lifetime* period, negotiated during the registration. After this lifetime period expires, the mobile host's registration is deleted. In order to maintain continued service from its home agent or foreign agent, the mobile host must re-register within this period. Either lifetime period may be set to infinity, in which case no re-registration is necessary.

All registrations with a mobile host's home agent must be authenticated. Although any authentication algorithm shared by both nodes may be used, the protocol defines a standard authentication algorithm using the MD5 message-digest algorithm [18], based on a secret shared between the mobile host and its home agent. The MD5 hash over the the shared secret and the important fields of the message is included in each registration message or reply, allowing the receiver to verify the source of the message and the fact that none of the important fields in the message have been changed since the message was sent. Since only the mobile host and its home agent know the shared secret, no other node can modify the message or can create a forged message. Administration of the shared secret is fairly simple, since both the mobile host and its home agent are owned by the same organization (both are assigned IP addresses in the home network owned by that organization). Manual configuration of the shared secret may be performed any time the mobile host is at home, while other administration of these nodes is being performed.

3.4. Tunneling

The protocol allows the use of any tunneling method shared between a mobile host's home agent and its current foreign agent (or the mobile host itself when a temporary local IP address is being used). During registration with its home agent, a list of supported tunneling methods is communicated to the home agent. For each packet later tunneled to the mobile host, the home agent may use any of these supported methods.

The protocol requires support for "IP in IP" encapsulation for tunneling. In this method, to tunnel an IP packet, a new IP header is wrapped around the existing packet; the source address in the new IP header is set to the address of the node tunneling the packet (the home agent), and the destination address is set to the mobile host's care-of address. This type of encapsulation may be used for tunneling any packet, but the overhead for this method is the addition of an entire new IP header (20 bytes) to the packet,

Support is also recommended for a more efficient "minimal" encapsulation protocol [5, 6, 7], which adds only 8 or 12 bytes to each packet. Here, only the modified fields of the original IP header are copied into a new forwarding header added to the packet between the original IP header and any transport-level header such as TCP or UDP. The fields in the original IP header are then replaced such that the source address is set to the address of the node tunneling the packet (the home agent), and the destination address is set to the mobile host's care-of address. This type of encapsulation adds less overhead to each packet, but it cannot be used with packets that have already been fragmented by IP, since the small forwarding header does not include the fields needed to represent that the original packet is a fragment rather than a whole IP packet.

4. Route Optimization

The basic IETF Mobile IP protocol fulfills its primary goal of providing transparent packet routing to mobile hosts operating in the Internet. However, all packets for a mobile host away from home must be routed through the mobile host's home network and home agent, severely limiting the performance transparency of the protocol and creating a significant bottleneck to potential scalability.

As suggested in Section 2, what is needed is the ability for correspondent hosts to be able to cache the location of a mobile host and to then send packets directly to the mobile host at its current location. We call this functionality *route optimization*, and together with Andrew Myles of Macquarie University and Charles Perkins of IBM, I have been working particularly to develop this functionality within the IETF protocol. This section provides an overview of the current state of the protocol extensions for route optimization.

4.1. Location Caching

Any node may optimize its own communication with mobile hosts by maintaining a *location cache* in which

it caches the binding of one or more mobile hosts. When sending a packet to a mobile host, if the sender has a location cache entry for this mobile host, it may tunnel the packet directly to the care-of address indicated in the cached binding. Likewise, a router when forwarding a packet may tunnel the packet directly to the destination mobile host's care-of address if the router has an entry in its location cache for the destination IP address of the packet.

In the absence of any location cache entry, packets destined for a mobile host will be routed to the mobile host's home network in the same way as any other IP packet, and are then tunneled to the mobile host's current care-of address by the mobile host's home agent. This is the only routing mechanism supported by the basic Mobile IP protocol. With route optimization, though, as a side effect of this indirect routing of a packet to a mobile host, the original sender of the packet is informed of the mobile host's current binding, giving the sender an opportunity to cache the binding.

A node may create a location cache entry for a mobile host only when it has received and authenticated the mobile host's binding. Likewise, a node may update an existing location cache entry for a mobile host, such as after the mobile host has moved to a new foreign agent, only when it has received and authenticated the mobile host's new binding.

Optimal routing of packets from a correspondent host can be achieved if the correspondent host implements a location cache. A router implementing a location cache can also provide routing assistance for packets that it forwards from correspondent hosts that do not implement the Mobile IP route optimization protocol. For example, a network of nodes that do not implement route optimization could be supported by a common first-hop router that maintains a location cache.

A location cache will, by necessity, have a finite size. Any node implementing a location cache may manage the space in its cache using any local cache replacement policy. If a packet is sent to a destination for which the cache entry has been dropped from the cache, the packet will be routed normally through the mobile host's home network and will be tunneled to the mobile host's care-of address by its home agent. As when a location cache entry is initially created, this indirect routing to the mobile host will result in the original sender of the packet being informed of the mobile host's current binding, allowing it to add this entry again to its location cache.

4.2. Foreign Agent Handoff

When a mobile host moves and registers with a new foreign agent, the basic Mobile IP protocol notifies the mobile host's previous foreign agent that it should no longer serve the mobile host. Route optimization extends this notification to allow a mobile host's previous foreign agent to be reliably notified of the mobile host's new binding. When registering its new care-of address with its home agent, the home agent may be requested to reliably send a *binding update* message (for which an acknowledgement is expected) to the mobile host's previous foreign agent. The previous foreign agent may then create a location cache entry for the mobile host, serving as a "forwarding pointer" to the new care-of address. The previous foreign agent may then use this cached binding to tunnel any subsequently arriving packets for the mobile host to its new location.

Such a location cache entry at a mobile host's previous foreign agent is treated in the same way as any other location cache entry. In particular, this location cache entry may be deleted from the cache at any time. Suppose a node (such as this previous foreign agent) receives some packet that has been tunneled to this node, but this node is unable to deliver the packet locally to the destination mobile host (it is not the mobile host itself, and it does not believe that it is currently serving as a foreign agent for this mobile host). In this case, the node tunnels the packet to the mobile host's home agent, from which the packet will be re-tunneled to the mobile host's current location.

4.3. Location Cache Updates

When a mobile host's home agent intercepts a packet from the home network and tunnels it to the mobile host, the home agent may deduce that the original sender of the packet has no binding in its location cache for the destination mobile host. In this case, the home agent sends a *binding update* packet to the sender, informing it of the mobile host's current binding. No acknowledgement for this binding update is needed,

since any future packets intercepted by the home agent from this sender for the mobile host will serve to cause a retransmission of the update.

Similarly, when the foreign agent serving a mobile host receives a packet for the mobile host, if it appears to the foreign agent that the original sender of the packet either has no location cache entry or has an out-of-date location cache entry for this mobile host, the foreign agent notifies the mobile host's home agent to send a binding update packet to the sender. When a mobile host is using a temporary local IP address as its own care-of address, if it receives a packet for which it appears that the original sender of the packet either has no location cache entry or has an out-of-date location cache entry for this mobile host, the mobile host notifies its home agent to send a binding update to the sender. As in the case above, no acknowledgement is needed for this binding update packet.

In order to determine if a notification should be sent, the foreign agent serving the mobile host (or the mobile host itself if using its own care-of address) compares the original source address of the packet (inside the encapsulation used for tunneling the packet) to the current source address of the packet (the address of the node that tunneled the packet to this node). If the packet was tunneled to this node by the original sender, then the sender must have an up-to-date location cache entry pointing to this care-of address. Otherwise, a binding update is needed, unless the packet was tunneled to this node by the mobile host's home agent.

All binding update packets are sent by a mobile host's home agent, which is in complete control of which other nodes it allows to learn the mobile host's binding. If, for any local administrative reasons, the home agent wants to keep a particular mobile host's current binding private (from all or only some other nodes), it is not required to send a binding update that would otherwise be sent by the protocol.

Included in each binding update message sent by the home agent is an indication of the time remaining in the service lifetime associated with the mobile host's current registration. Any location cache entry established or updated in response to this binding update must be marked to be deleted after the expiration of this period. A node wanting to provide continued service with a particular location cache entry may attempt to reconfirm that binding before the expiration of this lifetime period. Location cache entry reconfirmation may be appropriate when the node has indications (such as an open transport-level connection to the mobile host) that the location cache entry is still needed. This reconfirmation is performed by the node actively requesting the mobile host's home agent to send a new binding update message to the node.

Each node must provide some mechanism to limit the rate at which it sends binding updates to to the same node about any given binding, and to limit the rate at which it sends notifications about the need to send binding updates for any given mobile host to the same home agent. Some nodes will not implement the route optimization extensions of the Mobile IP protocol, and those that do may be limited in the number of bindings they can cache or the speed with which they can process binding updates. After a small number of binding updates or notifications sent to the same node about some binding, the sending node must quickly decrease the maximum rate at which new updates or notifications for this binding are sent.

4.4. Authentication

All messages that add or change an entry in a location cache must be authenticated using the same type of authentication algorithm as is used in the basic Mobile IP protocol for registration with a mobile host's home agent (Section 3.3). This authentication verifies the source of the message and ensures that none of the important fields of the message have been changed since the message was sent.

In particular, a node receiving a binding update packet must verify the message's authentication before altering the the contents of its location cache in response to the message. This requirement for authentication covers all binding update packets: those sent to build or update a location cache entry in response to a packet routed indirectly to a mobile host, as well as those sent to notify a mobile host's previous foreign agent of its new binding. Without such authentication, a malicious node anywhere in the Internet could forge a binding update message, allowing it to arbitrarily intercept or redirect packets destined for any other node in the Internet.

In the basic Mobile IP protocol, only a mobile host's registration with its home agent must be authenticated, allowing the simple solution of a manually configured

secret shared between a mobile host and its home agent. For route optimization, a mobile host's home agent must be able to authenticate a binding update sent to any other node in the Internet. This form of general authentication is currently complicated by the lack of a standard key management or authentication protocol in the Internet, and by the lack of any generally available key distribution infrastructure. A number of restricted authentication schemes are possible in the short term, before the necessary protocols and infrastructure are available, but the details of such schemes are beyond the scope of this paper. The procedures described in Sections 4.2 and 4.3, however, have been designed to minimize the impact of these current restrictions, and may be modified when more general authentication mechanisms become available.

5. Protocol Scalability

The combination of the basic IETF Mobile IP protocol described in Section 3 and the extensions for route optimization described in Section 4 can provide highly scalable support for packet routing to large numbers of mobile hosts in the Internet. This section considers the different factors affecting the scalability of the protocol.

5.1. The Home Network

Each organization owning an IP network supports all mobile hosts for which this is the home network. This arrangement allows mobility support to scale as new organizations and new networks connect to the Internet, avoiding any centralized support bottleneck. Each organization thus also may control the level of expense or effort which they expend to support their own mobile hosts, and their own mobile hosts directly benefit from these expenditures.

For example, an organization wanting to provide higher performance or more reliable access to the home agent for any of its mobile hosts may install higher bandwidth or additional links connecting their own home network to the Internet. The functionality of the home agent may also be replicated or distributed on multiple nodes on the home network; as long as a consistent view of the mobile hosts bindings is maintained, such arrangements are entirely at the option of the organization owning the network and need not affect other nodes within the Internet. The home agent functionality and the home network may be scaled to support any number of mobile hosts owned by this organization.

While a mobile host is at home, it is treated in the same way as any ordinary IP host, and no overhead is added to packets sent to it while at home. When the mobile host leaves home and registers a care-of address, its home agent begins tunneling packets for it, location cache entries are gradually created at different correspondent hosts or routers, and they then begin tunneling packets for the mobile host directly to its current location. As the mobile host moves from one care-of address to another, the location caches are updated as needed. When the mobile host later returns home, this same mechanism causes these location cache entries to be deleted; packets destined to this mobile host are then sent in the same way as any IP packets sent to an ordinary stationary host that has never moved.

It thus becomes feasible to upgrade all hosts, at any convenient time, to be "mobile capable," with no performance penalty [6]. Any mobile capable host could then become mobile at any future time as needed simply by leaving its home network and registering elsewhere. This property simplifies the installation of new hosts, as no decision need be made as to whether each host will need to be mobile at any future time.

5.2. The Foreign Network

Each network that allows mobile hosts to visit can control their own resource allocation using any local policies determined by the organization owning that particular network. For example, a foreign agent may limit the number of simultaneous visitors that it allows to register; if additional mobile hosts request registration, the foreign agent may return an error to each indicating that registration has been denied due to local resource allocation limits. Any organization may install additional or more powerful foreign agents or higher bandwidth local networks in order to provide any desired level of support for visiting users. Similarly, each organization may impose any administrative policies on the provision of service to visiting mobile hosts. For example, they may only allow mobile hosts for which prior billing arrangements have been established to register.

By deploying one or more foreign agents, the protocol places no new demands on IP address space allocation. Any organization wanting to provide service for visiting mobile hosts but not willing to deploy a foreign agent may support any number of visitors by reserving a portion of their local IP address space for dynamic allocation as care-of addresses for visiting mobile hosts.

5.3. Location Caches

The deployment and operation of a location cache in any node is only an optimization to the protocol, and no location caches are required. Each location cache may be limited to any size as needed or desired by any local administrative policies. Similarly, any local cache replacement policy may be used to manage the space within the cache.

If the location cache at some node is too small to be able to store a cached binding for each mobile host with which this node is actively communicating, the local cache replacement policy determines which entries are retained in the cache. For example, the use of LRU replacement will keep the most recently used entries in the cache. Other possible cache replacement policies might weight each entry by the number of times it had been recently accessed, or by some administratively assigned priority based on the a list of preferred hosts for which bindings should be cached. Such decisions are entirely local to the node (and organization) implementing the location cache.

5.4. Impact on the Network

No changes to the routing infrastructure of the Internet are required to support Mobile IP. By tunneling packets to a mobile host, all routers through which the tunneled packet must pass treat the packet exactly as any ordinary IP packet, using existing Internet routing algorithms. The Mobile IP protocol can thus be deployed incrementally, with each organization adding home agents or foreign agents as the need arises. All hosts and routers can be upgraded at any time, if desired, to support location caches.

By using route optimization, the overall overhead on the Internet can be minimized. Routing packets indirectly to a mobile host through the mobile host's home network and home agent places overhead on all links and nodes along this path, but route optimization allows this longer, indirect path to be avoided. Route optimization also reduces the resource demands on each home network, and avoids any possible performance bottleneck at the home network or at the home agent.

6. Conclusion

Recent increases in the availability of mobile computers and wireless networks provides the opportunity to integrate these technologies seamlessly into the Internet. Mobile users should be able to move about, remaining connecting to the Internet, utilizing the best available network connection at any time, whether wired or wireless. For example, a mobile host in its owner's office may be connected to an Ethernet, but when disconnected and carried away, it could transparently switch to a connection through a high-speed local area wireless network. While moving around within the building, the host could switch transparently from one wireless subnet to another, and when leaving the building, could again switch transparently to a wide-area wireless data service.

The current work in the IETF Mobile IP Working Group provides a good approach to reaching this vision of seamless transparent mobility. These protocols can efficiently scale to very large numbers of mobile hosts operating in a large internetwork. Such scalability will become crucial as the Internet continues its exponential growth, and as mobile users begin to account for a growing fraction of this population.

Acknowledgements

This paper has benefited greatly from discussions with many other participants in the Mobile IP Working Group of the Internet Engineering Task Force (IETF). I would particularly like to thank Andrew Myles and Charlie Perkins for their collaboration in our work within the IETF.

References

[1] Stephen E. Deering. ICMP router discovery messages. Internet Request For Comments RFC 1256, September 1991.

[2] Ralph Droms. Dynamic Host Configuration Protocol. Internet Request For Comments RFC 1541, October 1993.

[3] Vince Fuller, Tony Li, Jessica Yu, and Kannan Varadhan. Classless Inter-Domain Routing (CIDR): An address assignment and aggregation strategy. Internet Request For Comments RFC 1519, September 1993.

[4] John Ioannidis, Gerald Q. Maquire Jr, and Steve Deering. Protocols for supporting mobile IP hosts. Internet Draft, June 1992.

[5] David B. Johnson. Transparent Internet routing for IP mobile hosts. Internet Draft, July 1993.

[6] David B. Johnson. Ubiquitous mobile host internetworking. In *Proceedings of the Fourth Workshop on Workstation Operating Systems*, pages 85–90, October 1993.

[7] David B. Johnson. Scalable and robust internetwork routing for mobile hosts. In *Proceedings of the 14th International Conference on Distributed Computing Systems*, pages 2–11, June 1994.

[8] David B. Johnson, Andrew Myles, and Charles Perkins. Route optimization in Mobile IP. Internet Draft, July 1994.

[9] David B. Johnson, Andrew Myles, and Charles Perkins. The Internet Mobile Host Protocol (IMHP). Internet Draft, February 1994.

[10] Frank Kastenholz and Craig Partridge. Technical criteria for choosing IP:The next generation (IPng). Internet Draft, May 1994.

[11] J. Mogul and J. Postel. Internet standard subnetting procedure. Internet Request For Comments RFC 950, August 1985.

[12] Andrew Myles and Charles Perkins. Mobile IP (MIP). Internet Draft, September 1993.

[13] John Penners and Yakov Rekhter. Simple Mobile IP (SMIP). Internet Draft, September 1993.

[14] Charles Perkins and Yakov Rekhter. Support for mobility with connectionless network layer protocols (transport layer transparency). Internet Draft, January 1993.

[15] J. B. Postel, editor. Internet Protocol. Internet Request For Comments RFC 791, September 1981.

[16] J. B. Postel. Multi-LAN address resolution. Internet Request For Comments RFC 925, October 1984.

[17] Yakov Rekhter and Charles Perkins. Short-cut routing for mobile hosts. Internet Draft, July 1992.

[18] Ronald L. Rivest. The MD5 message-digest algorithm. Internet Request For Comments RFC 1321, April 1992.

[19] Marshall T. Rose. *The Open Book: A Practical Perspective on OSI*. Prentice Hall, Englewood Cliffs, NJ, 1990.

[20] Gursharan S. Sidhu, Richard F. Andrews, and Alan B. Oppenheimer. *Inside AppleTalk*. Addison Wesley, Reading, Massachusetts, 1990.

[21] W. A. Simpson, editor. IP mobility support. Internet Draft, September 1994.

[22] Fumio Teraoka and Keisuke Uehara. The virtual network protocol for host mobility. Internet Draft, April 1993.

[23] Paul Turner. NetWare communications processes. *NetWare Application Notes*, Novell Research, pages 25–81, September 1990.

[24] Hiromi Wada, Tatsuya Ohnishi, and Brian Marsh. Packet forwarding for mobile hosts. Internet Draft, November 1992.