Scalable and Robust Internetwork Routing for Mobile Hosts

David B. Johnson

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891

dbj@cs.cmu.edu

Abstract

This paper describes a new protocol for transparently routing packets to mobile hosts operating in a large internetwork. The protocol, called the Mobile Host Routing Protocol (MHRP), allows any host to become mobile at any time, yet there is no penalty for a host being "mobile capable," since the protocol automatically uses only the standard internetwork routing mechanisms and adds no overhead when a mobile host is currently connected to its home network. The paper concentrates on the design of MHRP as it applies to the Internet using IP. Mobile hosts use only their "home" IP addresses, regardless of their current location in the Internet. No changes are required in stationary hosts that communicate with mobile hosts, and no changes are required in mobile hosts above the IP level. MHRP introduces several new features to provide better robustness for routing to mobile hosts, and provides better scalability to very large numbers of mobile hosts than previous mobile host protocols.

1. Introduction

Mobile hosts such as notebook and palmtop computers and portable workstations are now widely available and affordable, and the distinction between desktop workstations and portable computers is beginning to disappear in terms of both features and computational power. A mobile host may be in use continuously, through a wireless network interface, as the host is carried from one location to another; or it may simply be disconnected from the network at its current location, temporarily moved to a new location, and reconnected to the network through either a wireless or conventional wired network interface.

However, current internetworking protocols, including IP [10], ISO CLNP [14], NetWare IPX [20], and AppleTalk [15], require mobile hosts to change their network addresses when moving to a new network, making host movement inconvenient and error prone. The new address must be edited into various configuration files, and currently running network applications must usually be restarted. The use of hierarchical addressing and routing schemes in internetworking protocols reduces the size of the routing tables that must be maintained at each router and exchanged between routers, and simplifies the routing decisions at each router, but prevents packets addressed to a mobile host from reaching that host when it is currently away from its "home" location.

For example, IP addresses are composed of a *network number*, identifying the network to which the host is attached, and a *host number*, identifying the particular host within that network. IP expects to be able to route a packet to a host based on the network number contained in the host's IP address. If a host changes its point of connection to the Internet and moves to a new network, IP packets destined for it will no longer reach it correctly.

This paper describes the *Mobile Host Routing Protocol (MHRP)* [5], a protocol for transparently supporting the routing of packets to mobile hosts operating in a large internetwork. The protocol requires no changes to non-mobile hosts or to backbone routers, and requires no changes to mobile hosts above the network level. MHRP is designed to provide efficient and robust operation and to scale well to very large numbers of mobile hosts. The protocol allows any host to become a "mobile host" simply by moving away from its home network. A mobile host is assigned a permanent network address in the same way as any other host, and always uses only its home address. There is no penalty for a host being "mobile capable," since the protocol automatically uses only the standard internetwork routing mechanisms and adds no overhead when a host is currently connected to its home network.

This paper concentrates on the design of MHRP as it applies to the Internet using IP [10]. Section 2 describes the MHRP infrastructure. Section 3 describes the protocol used when a mobile host moves to a new network, and Section 4 describes the mechanism used for IP packet routing and delivery to mobile hosts. The features introduced in MHRP for robustness are discussed in Section 5. Section 6 describes several examples of the use of MHRP. Section 7 compares MHRP with previous mobile host protocols, and Section 8 presents conclusions.

This research was supported in part by the Advanced Research Projects Agency under Contract DABT63-93-C-0054. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. government.

2. Infrastructure

Each mobile host owned by some organization is assigned a permanent IP address by that organization within one of the IP networks belonging to that organization. This network is called the "home" network for the mobile host, and the mobile host will use this IP address whether attached to this home network or currently attached to some "foreign" network.

In order for mobile hosts to be able to visit some foreign network, a *foreign agent* must be present on that network that is willing to allow mobile hosts to connect to it. The foreign agent maintains a list recording the IP address of each visiting mobile host currently connected to that network (for which it has agreed to serve as the foreign agent), and forwards arriving IP packets addressed to one of those mobile hosts directly to the mobile host. The "location" of a mobile host is represented by the IP address of its current foreign agent. The foreign agent is normally the router that connects this network to the rest of the Internet, but may also be a separate support host on that network. The connection of a mobile host to its current foreign agent may be either wired or wireless.

In order for hosts on some network to become mobile hosts (to leave their home network and connect to some foreign network), a *home agent* must be present on their home network. The home agent maintains a database recording the IP address of the foreign agent currently serving each mobile host for which this is the home network. When a mobile host moves to a new network and connects to a new foreign agent, it must notify its home agent. For mobile hosts not currently connected to their home network, the home agent must also arrange to intercept any packets arriving on the home network addressed to those hosts. Each packet intercepted by the home agent is forwarded to the foreign agent currently serving the destination mobile host, which then delivers the packet locally to that host.

For example, consider the sample internetwork illustrated in Figure 1. Host *M* is a mobile host, with an IP address within network *B*. Network *B* is thus called *M*'s "home" network. However, *M* is currently connected to network *D*, a wireless network connected to network *C* through router *R4*. Routers *R1*, *R2*, and *R3* connect networks *A*, *B*, and *C*, respectively, to a backbone network. Suppose *R2* is the home agent for mobile host *M*, and *R4* is the current foreign agent serving *M*. If host *S* sends an IP packet to *M* using *M*'s IP address, the standard IP routing algorithms will deliver the packet to *M*'s home network, where it will be intercepted by *R2*. Using the address of *M*'s current foreign agent recorded in its location database, *R2* forwards the packet to *R4*, which then delivers the packet locally to *M*.

Any host or router may also function as a *cache agent* by caching the location of one or more mobile hosts. Caching the location of a mobile host enables a host or router to forward packets for that mobile host directly to its foreign agent, without going through the host's home network and home agent. For example, in Figure 1, if *S* implements the MHRP protocol and chooses to function as a cache agent, it could optimize its own communication to *M*.

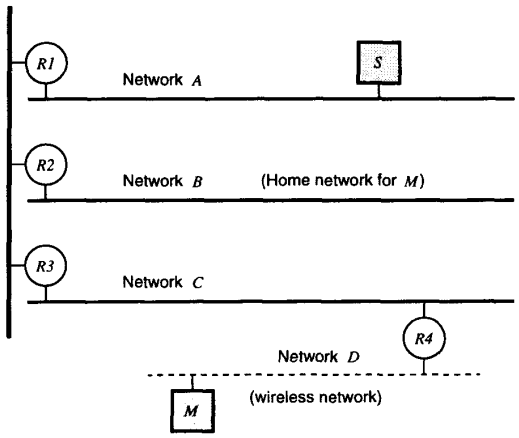

Figure 1 An example section of an internetwork

The cache maintained by a cache agent is only an optimization to improve routing of subsequent packets to a mobile host, and the contents of the (finite) cache space provided by any cache agent may be maintained by any local cache replacement policy. The consistency of each cache entry with respect to the true location of that mobile host as recorded by its home agent is maintained as needed by MHRP. In particular, when a mobile host moves to a new foreign agent, the location for that mobile host currently cached by any cache agents becomes out-of-date, since these cache entries still point to the old foreign agent. However, if out-of-date location information from some cache agent is used to forward a packet, MHRP will in turn forward the packet to the correct new foreign agent for that mobile host, and all out-of-date cache entries used in forwarding that packet will automatically be updated for use in forwarding subsequent packets to the mobile host.

The functionality of a foreign agent, home agent, and cache agent may be provided by separate hosts or routers on a network, or may be combined in different ways on one or more hosts or routers. For example, any node functioning as a home agent, foreign agent, or mobile host should generally also function as a cache agent. A single router on some network providing the functionality of both a home agent and a foreign agent would allow hosts with that network as their home network to become mobile and to connect to the Internet on other networks, as well as allowing hosts from other networks to connect to that network as visiting mobile hosts. Although not required, all other Internet hosts should also support functioning as a cache agent, in order to optimize their own communication with mobile hosts.

In case no foreign agent is available on some foreign network, a mobile host may also be able to serve as its own foreign agent, if it is able to obtain a temporary IP address within that foreign network. The temporary IP address would be used *only* as the address of that mobile host's foreign agent, and all packets intended for the mobile host would be tunneled to that address in the same way as for

other mobile hosts' foreign agents; the mobile host itself would continue to use only its home IP address. Providing this support for a mobile host serving as its own foreign agent is optional, and any methods for obtaining such a temporary IP address are beyond the scope of this paper.

Each organization manages its own home agent (or agents) to support the routing of IP packets to the mobile hosts owned by that organization. For example, if that organization requires increased reliability of service for its own mobile hosts, it can replicate the home agent function on several support hosts on its own network, although these hosts must cooperate to provide a consistent view of the database recording the current location of each of that home network's mobile hosts. Likewise, additional links connecting the home network to the Internet may be installed to provide continuous connectivity to the home agent in case one of the links to the Internet is temporarily down. Such decisions can be made independently by each organization, and directly benefit that organization's mobile hosts.

3. Moving a Mobile Host

A mobile host may move from one network to another at any time. Mobile hosts notice their own movement and identify a new foreign agent (or their own home agent) through an *agent discovery* protocol, similar to the Internet's ICMP router discovery protocol [2]. Foreign agents and home agents periodically multicast an agent advertisement message on their local networks; mobile hosts may wait to hear the next periodic advertisement message, or may optionally multicast an agent solicitation message when attempting to find a new agent. Mobile hosts realize that they have returned to their home network when they hear an advertisement from their own home agent.

A mobile host may explicitly disconnect from its current foreign agent (or from its home agent) before moving, in cases of planned disconnection. In many cases, though, such as for continuously moving hosts connected through a wireless interface, it may not be possible for a mobile host to explicitly disconnect before moving. For example, such a mobile host might be moved out of range of the transceiver at its old foreign agent at any time simply by being carried physically too far from it. Once it is within range of a new foreign agent, it may reconnect to the network through this new foreign agent and implicitly disconnect from its old foreign agent at the same time.

When a mobile host disconnects from its current network, it first notifies its home agent, and then notifies its old foreign agent from which it is disconnecting. As a special case, if the host is disconnecting from its home network, only its home agent is notified. When a mobile host reconnects to a new network, it must first notify its new foreign agent, and then notify it's home home agent (and its old foreign agent if the host did not explicitly disconnect from its old network earlier). As a special case, if the mobile host is reconnecting to its home network, only its home agent is notified. The mobile host registers a special foreign agent address of zero with its home agent when reconnecting to its home network.

At the home agent, the notifications of a mobile host disconnecting from or reconnecting to the network are used to maintain a record of the current location of the mobile host. The record is maintained in a database giving, for each mobile host for which this is the home network, the IP address of the current foreign agent for that mobile host. The database may be maintained in the memory of the home agent, but for reliability, should also be recorded on disk to survive any crashes and subsequent reboots of the home agent.

When the old foreign agent receives notification that a mobile host is disconnecting from it, it removes the mobile host from its list of locally visiting mobile hosts. The old foreign agent may optionally also cache the IP address of the new foreign agent for this mobile host, if it is capable of also functioning as a cache agent. This cache entry thus becomes a "forwarding pointer" to the new location for the mobile host, although this cache entry and is treated in the same way as any other cache entry maintained by a cache agent. Such a "forwarding pointer" may be useful in maintaining connectivity to a frequently moving mobile host during periods in which that host's home agent may be temporarily inaccessible.

When the new foreign agent receives notification that a mobile host has connected to it, it creates an entry for that host in its list of locally visiting mobile hosts. When the foreign agent receives an IP packet addressed to one of the hosts in this list, the foreign agent transmits the packet over its local network to that mobile host. The method used by the foreign agent to learn the local physical network address corresponding to the visiting mobile host is specific to the particular type of local network involved. For example, the physical network address may be saved from the connection notification message when the mobile host connected to this foreign agent, or a dynamic address resolution protocol such as ARP [7] may be used to learn the physical network address when needed.

When the home agent receives notification that a mobile host is disconnecting from its home network, the home agent must arrange to intercept all subsequent packets transmitted on this network to the mobile host. For example, the home agent may broadcast an ARP "reply" message [7] on the local network (perhaps retransmitted a few times for reliability), in order to update the address resolution cache of any other hosts on that network so that they now believe that the physical network address corresponding to the disconnecting mobile host is the physical network address of the home agent itself. When the mobile host subsequently reconnects to its home network, the mobile host broadcasts a similar ARP "reply" message to the local network, in order to cause other hosts on the same network to update their ARP cache with the real physical network address for the mobile host, rather than the physical address of the home agent that they may still have in their caches. While a mobile host is disconnected from its home network, the home agent also answers ARP requests for the mobile host with "proxy" ARP [12].

As described above, a home agent must be present on each IP network (or subnet) having hosts that may become mobile, and a foreign agent must likewise be present on each

IP network (or subnet) to which mobile hosts may connect. The home agent can then intercept IP packets for a mobile host being transmitted over that network, and the foreign agent can transmit packets directly over that network to locally visiting mobile hosts.

It may also be possible to support an entire routing domain with one (or more) home agents or foreign agents by selectively using host-specific IP routes. When a mobile host disconnects from its home network, its home agent could begin advertising network reachability to that specific host. Such host-specific routes would be advertised only while the mobile host was disconnected from its home network, and would not be propagated outside that routing domain. Likewise, when a mobile host connects to some foreign network, the mobile host could begin advertising a host-specific route for itself, allowing a foreign agent (in that same routing domain) to deliver arriving packets to it. This routing would be advertised only while the mobile host was connected to this foreign network, and would not be propagated outside that routing domain.

4. Forwarding Packets to a Mobile Host

4.1. The MHRP Encapsulation Protocol

In order to forward IP packets destined for a mobile host to the foreign agent currently serving that host, MHRP introduces a new *encapsulation* protocol. The home agent or initial cache agent handling the packet transforms it into a new IP packet, addressed to the foreign agent, by adding a new header (the *MHRP header*) to the packet between the IP header and any existing transport-level header such as TCP [11] or UDP [8], as illustrated in Figure 2. Once the MHRP header is added, the packet uses only normal IP routing for delivery to the foreign agent. This use of encapsulation is known as *tunneling*.

Once received by the foreign agent, the MHRP header is removed from the packet, the original IP header is reconstructed, and the packet is transmitted by the foreign agent directly to the locally visiting mobile host. Unlike typical encapsulation protocols, MHRP does *not* add a complete new IP (or link-level) header to the packet, but rather only

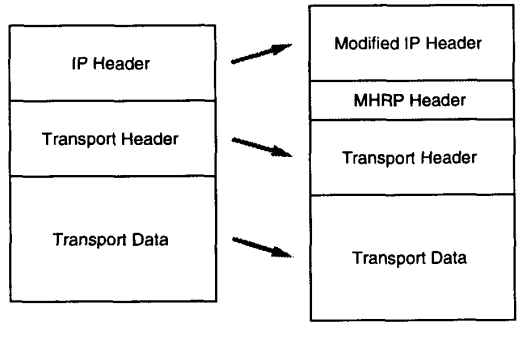

Figure 2 Building the MHRP header in a packet for tunneling to a mobile host

modifies the necessary fields in the existing IP header; this results in a significant savings in space overhead in the packet, and avoids complications in deciding whether or not to copy various parts of the existing IP header such as any IP options to the new header.

4.2. Adding the MHRP Header to a Packet

The MHRP header is illustrated in Figure 3. To add the MHRP header to a packet, the home agent or initial cache agent handling the packet performs the following steps:

- The original IP protocol number is copied from the IP header into the MHRP header, and is replaced in the IP header by the IP protocol number indicating MHRP.
- The original IP destination address (the mobile host) is copied from the IP header into the MHRP header, and is replaced in the IP header by the address of the foreign agent.
- Finally, unless the MHRP header is being built by the original sender of the packet, the original IP source address is copied from the IP header into the MHRP header, and is replaced in the IP header by the address of the cache agent or home agent building the MHRP header.

The shaded portions of the packet in Figure 3 are not modified in adding the MHPR header to the packet.

The MHRP header may be built by the original sending host itself, if it is currently also functioning as a cache agent and has a cache entry for the location of the mobile host. In this case, the list of previous IP source addresses in the MHRP header is empty, and the length of the constructed MHRP header is only 8 octets.

The MHRP header may also be built by the sender's first-hop router or by any other router that forwards the packet, if this router is currently functioning as a cache agent and has a cache entry for the destination IP address. If no cache agents are encountered, the packet will be routed at each hop according to the normal IP routing algorithms and will eventually reach the mobile host's home network, where the packet will be intercepted by the mobile host's home agent. The home agent will then build the MHRP header and tunnel the packet to the foreign agent. If the MHRP header is built by any host or router other than the original sender, the list of previous IP source addresses in the MHRP header contains a single entry (the address of the original sender), and the length of the constructed MHRP header is 12 octets.

4.3. The "Location Update" Message

MHRP sends a "location update" message to report the current location of a mobile host to specific hosts or routers that appear to need this information. A location update message is sent only when a specific need to update some host or router is identified. The recipient of a location update message, if it is capable of functioning as a cache agent, may cache the IP address of the foreign agent reported in the message.

Any intermediate router that forwards a location update message may also cache the address contained in the

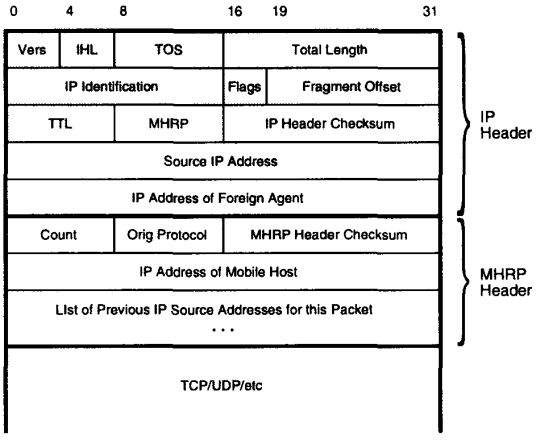

Figure 3 The MHRP header in an IP packet

message, if the router is capable of functioning as a cache agent. For example, a network of hosts that do not implement MHRP could be supported by a common first-hop router that is willing to function as a cache agent and caches the location of any mobile hosts with which these hosts correspond. However, in order to function as a cache agent, a router must examine each packet that it forwards, checking for location update messages and updating or adding to its cache. Routers should thus support a configuration option to enable or disable the capability to become a cache agent, avoiding the overhead of examining each packet forwarded except when needed. Routers that do not support MHRP, of course, also would experience no overhead from forwarding location update messages.

In the Internet, the location update message is defined as a new type of ICMP message [9]. The location update message includes the IP address of the mobile host and the IP address of the foreign agent currently serving the mobile host. The location update message is sent as an ICMP message due to its similarity with the existing ICMP redirect message type, and also to aid in backwards compatibility with hosts that do not implement MHRP. Any such host receiving a location update message will simply ignore the message, since any ICMP messages of unknown type must be silently discarded by all hosts [1].

Since not all hosts will support MHRP, any host or router that sends location update messages must provide some mechanism for limiting the rate at which it sends these messages to any single IP address. For example, a list could be maintained giving the IP addresses to which updates have been sent and the time at which an update was last sent to each address. This stored time on each list entry could also be used to implement LRU replacement of the entries within the list. This requirement to avoid location update message flooding is similar to the mechanism already required by hosts to limit the rate at which ARP requests are sent to any single IP address [1].

In an implementation, any host or router acting as a cache agent for a mobile host may record the IP address of that mobile host's foreign agent in the same table that it uses already to handle the existing host-specific ICMP redirect message type [9]. Before transmitting an IP packet, a host or router must currently search this table to find the correct first-hop router address to use for this destination IP address. By saving the IP address of a mobile host's foreign agent in this same table (with a different type field on the table entry), the correct foreign agent address for sending to a mobile host can be found in the cache with little or no additional cost. This table can also be used by a foreign agent to store a table entry for each locally visiting mobile host currently connected to that foreign agent, and to thus recognize that a packet that it is routing must be transmitted locally to a visiting mobile host.

4.4. Foreign Agent Processing of Tunneled Packets

Once received by the foreign agent, the packet is processed by the MHRP protocol module on the foreign agent. If the destination mobile host IP address is found in the foreign agent's list of locally visiting mobile hosts, the MHRP header is removed from the packet and the original IP header is reconstructed. The foreign agent then transmits the packet over the last hop to the directly connected mobile host.

If the mobile host has moved to a new foreign agent (or returned to its home network) since last connected to this foreign agent, the mobile host will not be found in the foreign agent's list of locally visiting mobile hosts. If this foreign agent is also functioning as a cache agent and has cached the location for this mobile host, the packet is tunneled by this (old) foreign agent to the new foreign agent by modifying the IP header and MHRP header to transform the packet into a new IP packet, addressed from the old foreign agent (the cache agent) to the new foreign agent. If, instead, the old foreign agent has no cached location information for this mobile host (either because the new location was not cached when the mobile host moved, or because that cache entry has subsequently been reused for some other mobile host) the packet is tunneled instead by the old foreign agent to the mobile host's *home* IP address, where it will be intercepted by the mobile host's home agent.

In re-tunneling the packet to the new foreign agent or to the mobile host's home agent, the old foreign agent modifies the IP header and MHRP header by the following steps:

- The current IP source address from the IP header is appended to a list of previous IP source addresses for this packet, maintained in the MHRP header. The size of the MHRP header in the packet thus is increased by 4 bytes.
- The IP source address in the IP header is replaced by the current IP destination address from the IP header (the foreign agent's own address).
- If tunneling to the new foreign agent, the IP destination address in the IP header is set to the address of the new foreign agent; otherwise, the packet is tunneled to the mobile host's home agent by setting the IP destination address in the IP header to the IP address of the mobile host, obtained from the MHRP header.

The packet is then forwarded to its new IP destination address using only normal IP routing. For tunneling to the mobile host's home address, the packet is intercepted by the mobile host's home agent in the same way as other packets transmitted on the home network addressed to the mobile host, as described in Section 3.

Any finite maximum length of the list of previous IP source addresses in the MHRP header may be imposed. When a host or router that is tunneling a packet to a mobile host attempts to add a new address to the list, if the list is already at the maximum length allowed by this implementation and there is thus no room in the list to add this new address, the following steps are performed:

- A location update message is sent to each address currently in the list. The current foreign agent address cached by this node (to which it will tunnel the packet itself) is reported in the update message as the address to which future packets for this mobile host are to be tunneled.
- The list is of previous IP source addresses in the MHRP header is truncated and reset to empty.
- The new address is added to the list as the single entry in the list.

The new address being added to the list in the MHRP header (copied from the IP source address in the IP header) identifies the node that tunneled this packet to the current host or router. Each address from the list of previous IP source addresses in the MHRP header identifies the source of a previous tunnel used in forwarding this packet, and thus identifies an out-of-date location cache. By sending a location update message to each of these nodes, each of these cache agents will then point more directly to the current foreign agent for this mobile host.

4.5. Handling of Returned ICMP Error Messages

In the Internet, an ICMP error message is returned to the sender of a packet to indicate any error (such as "destination unreachable") encountered in forwarding the packet [9]. Correctly handling returned ICMP error messages is difficult with any IP encapsulation protocol. With MHRP, any ICMP error messages returned in response to an IP packet sent to a mobile host must be reported back to the original sender of the packet causing the error. Furthermore, ICMP error messages include a portion of the original IP packet in error, and this returned packet must also be returned to the original sender in a form that makes sense to this sender.

When a packet is tunneled using MHRP, the addresses in the IP header are rewritten so that the IP destination address of the packet is the endpoint of the tunnel, and the IP source address of packet is the head of the tunnel (the home agent or cache agent tunneling the packet). If the mobile host has moved since being connected to the foreign agent at the endpoint of this tunnel, the packet will be tunneled again, either to the next foreign agent or to the mobile host's home agent. MHRP achieves the correct handling of returned ICMP error messages by causing the ICMP error message to travel back to the sender along the same set of tunnels that the original packet followed, and by reversing any changes made by MHRP to the original packet in the returned packet carried by the ICMP error message.

When initially sent by some host or router, an ICMP error message will be transmitted to the cache agent or home agent at the head of the most recent tunnel through which the packet was traveling when it encountered the error. This agent then reverses the changes it made to the original packet in the returned copy of the packet included in the ICMP message, and resends the modified ICMP message to the home agent or cache agent (or the original sending host) that sent the original message to this agent. The list of previous IP source addresses for a packet maintained in the packet's MHRP header is used to determine the address to which to resend the modified ICMP message.

This procedure also allows each cache agent or home agent along the path to process the error message locally. For example, if a "destination unreachable" message is returned, it might indicate that some router along this path to the the cached location for the destination mobile host is unreachable, not that the mobile host itself is unreachable; in this case, the cache agent may also delete its cache entry for this mobile host before resending the modified ICMP error message.

A returned ICMP error message may contain a copy of the entire original message [1], but might instead contain only the IP header and first 8 bytes of the the transport-level header and data of the original message [9, 1]. If the returned message contains at least the entire MHRP header and 8 bytes beyond the MHRP header (the first 8 bytes of the original transport-level header and data), the ICMP error message can be correctly forwarded back to the original sender. If, however, less of the original packet is returned in the ICMP message, little can be done by a cache agent beyond deleting its cache entry for this mobile host. The next packet from this sender to the same mobile host will then not be tunneled by this cache agent, and may thus follow a different path to its destination.

On the other hand, returned ICMP *reply* messages [9] such as "echo reply" cause no similar problems. Since the original IP packet containing the ICMP request message is reconstructed by the foreign agent before the packet is transmitted to the visiting mobile host, any ICMP reply message generated by the mobile host will be returned directly to the original sender.

5. Protocol Robustness

5.1. Cache Agent Consistency Maintenance

The list of previous IP source addresses for a packet maintained in the packet's MHRP header is used primarily for updating any out-of-date cache agents that were used in routing the packet.

When a packet reaches the correct foreign agent (a foreign agent for which the mobile host's IP address is found in the foreign agent's list of locally visiting mobile hosts), the IP header and MHRP header contain the following addresses:

- The IP destination address in the IP header is the address of this foreign agent.

- The IP source address in the IP header is the address of the last cache agent (or the home agent) used in routing the packet to that foreign agent. This cache agent points directly to the correct foreign agent and is not out-of-date.
- If the list of previous IP source addresses for this packet in the packet's MHRP header is not empty, the first address in this list is the address of the original sender of the packet. This sender either has no cache entry for this mobile host, or has a cache entry for this mobile host that is out-of-date.
- Each other address in the list of previous IP source addresses for this packet in the packet's MHRP header is the address of an out-of-date cache agent for this mobile host.

This foreign agent sends a location update message to each address in the list of previous IP source addresses in the MHRP header, causing each of them to update their cache entry to point to this foreign agent.

Likewise, if the packet is tunneled to the home agent, because the old foreign agent did not have a "forwarding pointer" cached for the new foreign agent, the IP header and MHRP header will contain the following addresses:

- The IP destination address in the IP header is the mobile host's IP address.
- The IP source address in the IP header is the address of the old foreign agent that tunneled this packet to the home agent.
- The first address in the list of previous IP source addresses for this packet in the packet's MHRP header is the address of the original sender of the packet. (The list cannot be empty in this case.) This sender either has no cache entry for this mobile host, or has a cache entry for this mobile host that is out-of-date.
- Each other address in the list of previous IP source addresses for this packet in the packet's MHRP header is the address of an out-of-date cache agent for this mobile host.

Once the packet is intercepted by the home agent on the mobile host's home network, the home agent sends a location update message to each address in the list of previous IP source addresses in the MHRP header, causing each to update its cache entry to point to the correct (new) foreign agent, as recorded in the home agent's database. The home agent also sends a location update message to the old foreign agent that tunneled the packet to the home agent (identified in the IP source address field in the IP header), allowing it to become a cache agent for this mobile host so that any subsequent packets arriving for this mobile host can be tunneled directly to its correct new foreign agent.

5.2. Foreign Agent State Recovery

If a foreign agent "forgets" about a locally visiting mobile host, such as may happen when the foreign agent reboots, the mobile host will not be found in its list of locally visiting mobile hosts when a packet arrives for the host. In this case, the foreign agent will tunnel the packet to the home agent, as described in Section 5.1. The home agent then sends a location update message to each address in the list of previous IP source addresses present in the packet's MHRP header, as well as to the current IP source address in the packet's IP header. Each of these addresses identifies a host or router that has already handled this packet in routing it to the mobile host. The home agent compares these addresses to the current foreign agent address for the mobile host, as recorded in the home agent's database.

If a match is found in this comparison, the home agent discards the original packet, and the foreign agent will receive a location update message for this mobile host, identifying *itself* as the current foreign agent for the mobile host. The foreign agent could then simply add the mobile host back to its list of locally visiting mobile hosts, believing the home agent. Alternatively, the foreign agent could send a "query" message onto its local network to verify that the mobile host is actually connected to its network. For example, an ARP query message [7] could be used to elicit a reply from the mobile host indicating its presence.

To speed the state recovery of a foreign agent after it reboots, the foreign agent could also broadcast over its local network a query for all mobile hosts to initiate reconnection to it. Since such a broadcast cannot in general be guaranteed to reliably reach all locally visiting mobile hosts, the procedure described above based on the location update message and the list of previous IP source addresses in a packet's MHRP header is also necessary.

5.3. Robustness Against Routing Loops

No routing loops can be created by a correct implementation of this protocol. However, in a large internetwork such as the Internet, with many independent interoperating implementations of each protocol, some incorrect implementation could accidentally create a loop of cache agents. Such a loop would cause an arriving packet to be forwarded continuously around this loop until its IP "time-to-live" expired, potentially creating considerable congestion in the portion of the network involved in the loop.

The list of previous IP source addresses for a packet maintained in the packet's MHRP header may be used to easily detect any such loop that may be formed. When initiating the tunneling of a packet, the previous address in the IP source address field of the IP header is copied into the list in the packet's MHRP header, before being replaced in the IP header with the IP address of the host or router forwarding the packet. If the IP address of this node is already present in the list in the packet's MHRP header, then a forwarding loop exists involving the nodes identified in the list in the MHRP header; one pass around the loop has just been completed with the return of the packet to this node.

Any such loop detected can also easily be corrected using the list in the MHRP header. When a loop is detected by some node, that node sends a location update message to each address in the list, causing each of these cache agents to delete its cache entry for this mobile host. This, in effect, dissolves the loop. The original packet may then be discarded, or may be tunneled to the mobile host's IP address,

where it will be intercepted by the mobile host's home agent, as described in Section 4.

If the list in the MHRP header has been truncated because it reached its maximum allowable length (as described in Section 4.4), a loop may not be detected within a single cycle if the size of the loop is larger than the number of addresses allowed in the list. However, the size of the loop will contract during each cycle by a factor of the maximum list size. If the packet is still looping by the time the loop size contracts to be small enough to be recorded in the list in the MHRP header, the loop will then be detected and corrected. If the packet has been discarded by this time because its "time-to-live" field in its IP header has expired, the next packet will continue the loop contraction and detection procedure.

6. Examples

6.1. The Initial Packet to a Mobile Host

Suppose some host *S* is sending an IP packet to a mobile host *M*, as shown in Figure 1. *S* may not know (or need not know) that *M* is mobile. The packet is sent and routed in exactly the same way as any other IP packet, and thus reaches *M*'s home network. If *M* were currently connected to its home network, the packet would be delivered there directly to *M* with no extra overhead.

Figure 1 instead shows *M* connected to foreign network *D* through foreign agent *R4*. In this case, the packet is intercepted by *M*'s home agent, *R2*, which then tunnels the packet to *R4*. *R2* also returns a location update message to *S*. *S* or any router such as *R1* that sees this location update message may cache *M*'s location as a cache agent.

6.2. Subsequent Packets to a Mobile Host

Once a sending host such as *S* has cached the location of a mobile host to which it is sending packets, it may tunnel its own packets directly to that mobile host's foreign agent. This is expected to be the common case once MHPR becomes widely implemented in host software. A local network of hosts that do not yet support MHRP may also be supported by a single cache agent functioning in the IP router that connects that local network to the rest of the Internet. In this case, this router, such as *R1*, simply examines the packets that it forwards and tunnels any packets that are destined to addresses for which it has cached location information. This type of caching in intermediate routers may also be useful in supporting any hosts on a local network for which, for any reason, the protocol software running on those hosts cannot be modified to support MHRP.

6.3. When a Mobile Host Has Moved

When a mobile host moves to a new network, packets sent to that host must still reach it, and any cache agents that become out-of-date due to the host's movement must eventually be updated. Suppose mobile host *M* moves from *R4* to some new foreign agent, say *R5*. *M* notifies its home agent, *R2*, and its previous foreign agent, *R4*, of the move. The next packet that *S* sends to *M* will be tunneled first to *R4*, since this is the location for *M* that *S* has cached. If *R4* still has cached the new location of *M*, it will tunnel the packet to *R5*, where it will be delivered locally to *M*. *R5* will also send a location update message to *S*, reporting the new location of *M* at *R5*. On the other hand, if *R4* no longer has the cache entry giving the new location of *M*, it instead tunnels the packet to *M*'s home agent, which tunnels the packet on to the correct new foreign agent, *R5*, and returns a location update message to both *S* and *R4*.

Suppose instead that mobile host *M* moves from *R4* to return to its home network. *M* informs its home agent, *R2*, of this move and registers with it a special foreign agent address of zero. *M* also tells its old foreign agent, *R4*, that it has returned to its home network, causing *R4* to delete *M* from its list of locally visiting mobile hosts; *R4* does not create a "forwarding pointer" cache entry for *M* in this case. The next packet that *S* sends to *M* will be tunneled to *R4* (from *S*'s cache entry), and then by *R4* to *M*'s home agent, as described above. Once received there by *M*, *M* will return a location update message to *S*, indicating that it is currently connected to its home network and that *S*'s cache entry for *M* should be deleted. Subsequent packets from *S* to *M* will thus be sent directly to *M* on its home network without involving MHRP.

7. Comparison to Previous Protocols

The problems of addressing and routing packets to mobile hosts on the Internet were first described by Sunshine and Postel [16]. They introduced the notion of a "forwarder" to which other hosts could send IP packets to be forwarded locally to a visiting mobile host, which is similar to MHRP's foreign agent. However, they described only a simple protocol requiring a global database, in which all mobile hosts would register the address of their current forwarder. Senders would query the global database for the correct forwarder host and use source routing to deliver the packet to the forwarder. After a mobile host has moved to a new location, the old forwarder will return a "host unreachable" message to the sender in response to any new packet arriving for the mobile host, and the sender must then consult the global database again to learn the new location of the mobile host and retransmit the packet.

An IP protocol for mobile hosts has been implemented by Ioannidis et al at Columbia University [3, 4] using an "IP-within-IP" (or IPIP) protocol to tunnel IP packets to the network to which a mobile host is currently connected. Their protocol is optimized for mobile host movement within its home "campus" and makes no provision for optimizing routing of packets to a mobile host when it moves outside its home campus. A set of Mobile Support Routers (or MSRs) on the home network advertise network reachability to all hosts on the home network, whether or not currently connected instead to some foreign network. If the mobile host is still within its home campus, the packet is delivered to one of the home MSRs, which then tunnels the packet to the correct MSR using IPIP. However, MSRs are required to cache the correct MSRs for other mobile hosts, and if not present in its cache, the original MSR must use a broadcast or multicast protocol among the other MSRs to find the one serving the destination mobile host. When a

mobile host connects to some foreign network, it must first obtain a temporary IP address within that network, and its MSR in its home network then tunnels its packets to this temporary address. All packets to a mobile host outside its home network must be routed first to its home MSR. Their protocol adds 24 bytes of overhaed to each packet sent to a mobile host, whereas MHRP normally adds only 8 bytes (or 12 bytes) to each packet.

Teraoka et al at Sony [19, 17, 18] have implemented a mobile host IP protocol in which all hosts have two addresses: a "Virtual IP" (or VIP) address that never changes, and a normal IP address that specifies the the host's current physical location and must be obtained as a temporary address when connecting to a new foreign network. All IP packets are modified to contain a VIP header, containing a host's VIP address, as well as a normal IP header, containing its temporary IP address, which is used for routing. The sender uses a cache to translate the destination VIP address to its current physical IP address. If it has no entry in its cache for this host, the packet is sent with the IP address initially set the same as the VIP address, which may cause the packet to travel as far as the mobile host's home network router, where the correct IP address is filled in and the packet is resent to the correct destination. Other routers in the Internet also cache the location of mobile hosts by remembering the source IP and VIP addresses of packets that they forward. When a mobile host moves to a new network, a flooding protocol is used to remove most of these cache entries for the host in other routers, but some may remain due to the way in which the flooding is propagated. Such an obsolete cache entry might cause a packet to be delivered to an incorrect host. An incorrect receiver discards the packet and returns an error message to the sender, which will then retransmit the original packet. The error message will also cause the cache entries at the routers through which it passes to be removed. The overhead added to each packet for the VIP header is 28 bytes.

Another mobile IP proposal has been made by Wada et al at Matsushita [21]. Each mobile host operates in one of two modes. As with the Columbia and Sony proposals, each mobile host must obtain a temporary IP address in the foreign network that it is visiting. In "forwarding mode," all packets to that mobile host are routed through a Packet Forwarding Server (or PFS) on that mobile host's home network, and are then tunneled to the mobile host's temporary IP address, using a protocol that they call the Internet Packet Transmission Protocol (or IPTP). Optimization of the routing to avoid going through the home network is not possible in forwarding mode. In "autonomous mode," though, senders can cache the temporary IP address of the mobile host and tunnel their own packets directly there. The overhead added to each packet with their protocol is 40 bytes, since a new IP header must be added, as well as a separate IPTP header.

Perkins and Rekhter at IBM [13, 6] have produced several proposals for mobile IP using the IP "loose source route and record" (LSRR) option [10]. Their proposals depend on the specific defined semantics of the LSRR option. Each mobile host registers with a "base station" in the foreign network being visited, similar to Sunshine and Postel's forwarders [16] and MHRP's foreign agents. All packets sent by a mobile host are sent through the mobile host's base station and include an LSRR option, such that when received at the packet's destination, the recorded route will indicate the correct path back to the mobile host through its base station. Hosts receiving a packet containing an LSRR option are supposed to save and reverse the recorded route for use in sending return packets. However, many existing implementations of the LSRR option either do not record the route correctly in the packet, or do not correctly reverse or save the recorded route when receiving a packet. Also, after moving, packets for a mobile host continue to go to the host's old location until some application on that host needs to send a normal IP packet to that destination. Their protocol normally adds only 8 bytes to each packet sent to a mobile host, although 8 bytes must also be added to each packet sent *from* a mobile host.

MHRP introduced a number of new features to provide better robustness for routing packets to mobile hosts [5], as described in Section 5. Previous protocols either did not address these issues, or had much less effective methods of handling them. For example, relying on the IP "time-to-live" to break any routing loops created can still cause considerable congestion in the network when a loop does occur. Existing routing protocols (for non-mobile hosts) attempt to ensure that no routing loops can occur, but when tunneling is used to reroute packets, many new possibilities for routing loops are created. If the rate at which new packets are being generated and sent into the loop exceeds the rate at which packets expire within the loop as the time-to-live of each packet counts down to zero, the number of packets in the loop will continue to rise. If the loop persists for any length of time, this congestion could cause routers within the loop to crash and could lead to the collapse of a portion of the Internet. Similarly, previous protocols had only limited provision for cache consistency maintenance, and did not support the recovery of a failed foreign agent (or base station or forwarder).

Another important factor to consider in the design of an internetworking protocol for mobile hosts is the ability of the protocol to efficiently support very large numbers of mobile hosts. As the popularity of portable computers continues to increase, the number of mobile hosts that must be handled will grow rapidly. MHRP can support this rapid growth in the number of mobile hosts, for example, because no global database or global communication is required, each home agent only manages the location of its own mobile hosts, and the amount of state information that must be saved or cached by other hosts or routers is small. The ability of Sunshine and Postel's protocol [16] to scale to large numbers of mobile hosts is limited, since it relies on a global database; and the Sony [19, 17] and Columbia [3] protocols require forms of broadcasting or multicasting. In addition, the requirement of their protocols and Matsushita's protocol [21] for mobile hosts to obtain a new temporary IP address when visiting a foreign network places a limit on their scalability, since the available IP address space within any foreign network number is lim-

ited. The IBM proposals using the IP LSRR option [13, 6], appear to have the best scalability of these previous protocols, but are limited in scalability due to problems with the use of IP options in general: any IP packet containing an IP option requires extra processing at each router that forwards the packet and cannot use the "fast path" optimized code in the router, since each option must be examined by the router to determine if it affects its handling of the packet. If large numbers of mobile hosts begin using the IP LSRR option, many existing routers may not be able to handle the significantly increased load. The use of the LSRR option also limits the scalability of these protocols, due to the large number of incorrect implementations of this option already deployed within the Internet.

8. Conclusion

This paper has described the design of the Mobile Host Routing Protocol (MHRP) as it applies to the Internet using IP. A mobile host may move from one network to another at any time, while always using only its "home" IP address. By always using the home IP address for a mobile host, the current location of a mobile host — and even the fact that the host is mobile — remains transparent above the IP level. Any host may be configured to be a "mobile host" by simply running the appropriate software on it. There is no penalty for this configuration, since the protocol automatically uses only the standard IP routing mechanisms, adding no overhead to IP, when a mobile host is currently connected to its home network. MHRP introduces several new features to provide better robustness for routing to mobile hosts, and provides better scalability to very large numbers of mobile hosts than previous mobile host protocols. An implementation of MHRP within the Berkeley networking code is currently in progress.

Acknowledgements

The protocol presented in this paper has benefited from the helpful comments and criticisms of the author's fellow members of the Internet Engineering Task Force "Mobile IP" working group. The Working Group is currently developing a complete protocol specification for IP routing to mobile hosts.

References

[1] Robert T. Braden, editor. Requirements for Internet hosts — communication layers. Internet Request For Comments RFC 1122, October 1989.

[2] S. Deering. ICMP router discovery messages. Internet Request For Comments RFC 1256, September 1991.

[3] John Ioannidis, Dan Duchamp, and Gerald Q. Maguire Jr. IP-based protocols for mobile internetworking. In *Proceedings of the SIGCOMM '91 Conference: Communications Architectures & Protocols*, pages 235–245, September 1991.

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

[5] David B. Johnson. Transparent Internet routing for IP mobile hosts. Internet Draft, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, July 1993.

[6] Charles Perkins and Yakov Rekhter. Support for mobility with connectionless network layer protocols (transport layer transparency). Internet Draft, IBM T.J. Watson Research Center, Yorktown Heights, New York, January 1993.

[7] David C. Plummer. An Ethernet address resolution protocol: Or converting network protocol addresses to 48.bit Ethernet addresses for transmission on Ethernet hardware. Internet Request For Comments RFC 826, November 1982.

[8] J. B. Postel. User Datagram Protocol. Internet Request For Comments RFC 768, August 1980.

[9] J. B. Postel, editor. Internet Control Message Protocol. Internet Request For Comments RFC 792, September 1981.

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

[11] J. B. Postel, editor. Transmission Control Protocol. Internet Request For Comments RFC 793, September 1981.

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

[13] Yakov Rekhter and Charles Perkins. Short-cut routing for mobile hosts. Internet Draft, IBM T.J. Watson Research Center, Yorktown Heights, New York, July 1992.

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

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

[16] C. Sunshine and J. Postel. Addressing mobile hosts in the ARPA Internet environment. Internet Engineering Note IEN 135, March 1980.

[17] Fumio Teraoka, Kim Claffy, and Mario Tokoro. Design, implementation, and evaluation of Virtual Internet Protocol. In *Proceedings of the 12th International Conference on Distributed Computing Systems*, pages 170–177, June 1992.

[18] Fumio Teraoka and Keisuke Uehara. The virtual network protocol for host mobility. Internet Draft, Sony Computer Science Laboratory and University of Electro-Communications, April 1993.

[19] Fumio Teraoka, Yasuhiko Yokote, and Mario Tokoro. A network architecture providing host migration transparency. In *Proceedings of the SIGCOMM '91 Conference: Communications Architectures & Protocols*, pages 209–220, September 1991.

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

[21] Hiromi Wada, Tatsuya Ohnishi, and Brian Marsh. Packet forwarding for mobile hosts. Internet Draft, Matsushita Electric Industrial Co., Ltd., Osaka, Japan, November 1992.