MOBILE IP

Charles E. Perkins **Andrew F. Myles**

T.J. Watson Research Center, IBM **Macquarie University, Australia**

This paper specifies a new mobile host protocol, Mobile IP (MIP), to allow packets between mobile hosts or mobile networks and other hosts (including fixed hosts) to be delivered transparently along close to optimal routes.

MIP introduces a new IP option to encode MIP routing data into IP packets. Only a small amount of overhead is introduced to each packet and no changes are needed to layers above the IP layer.

1 Introduction

Computers are no longer large, expensive, and non-portable. Recent developments make it possible to buy inexpensive notebook and palm top computers that are both portable and extremely powerful. As a result, people are not constrained to using their computers in a single location.

At the same time, computer networking is becoming indispensable to the average user. The ability to access information and communicate globally has amplified the utility of computers many times over but users increasingly want to attach to the network, with the same level of service, wherever they happen to be working.

A typical user might usually connect to the network at their desk but occasionally will take their computer to a conference room in another part of the building, or even at another site. However, today's networking protocols, including the TCP/IP and the OSI protocol suites, have been designed under the tacit assumption that computers are always attached to the network at a single physical location. Today, host migration is assumed to occur so rarely that it can be handled manually.

For example, consider the process of host migration in the above scenario using the IP protocol. If the user's desk and the conference room have direct access to the same IP subnet then the migration process is trivial. Otherwise, the user must acquire a new IP address. After acquiring a new address from the appropriate local authority, which may not be an easy task, numerous configuration files on the migrating machine, on various name servers and on other machines (that use the original IP address to identify the migrating machine) need to be modified.

This migration process does achieve host migration but only after a slow, error prone configuration procedure that a typical user does not have the skills or desire to carry out. In addition, the computer then has a completely different identity and so all existing network applications must be restarted. Clearly, new protocols are needed to ensure that host migration can be achieved transparently to the user.

Ease of migration will become even more important as wireless network interfaces become widely available. Once the user is unconstrained by cable it is likely that frequent network migration will become common. Host migration may not even be under the user's control. For example, if a mobile host operates in overlapping wireless cells it may migrate from cell to cell based on dynamic factors such as load and noise. The new protocols will not only have to be transparent but also efficient so that they can handle rapid network migration.

Increasingly, computer networks are being interconnected into a single global network. The Internet is a good example of a growing global network. Users already have access to network facilities anywhere in the world, and not just in the local area. The new protocols should therefore be designed to operate both locally and in the wide area.

This paper specifies a new mobile host protocol, Mobile IP (MIP), that allows hosts to migrate transparently and rapidly around both the local and the wide area network in an IP environment. The protocol is based on the best aspects of various protocols proposed by IBM, Sony, and Columbia University. MIP modifies the existing IP (Internet Protocol) so that mobility can occur transparently to higher layers.

A comparison of all the proposals protocol may be found in [5]. It is easily concluded that none of the proposals are ideal; MIP suggests a new set of engineering compromises building on the previous proposals.

The Sony proposal [6] for example, is timer based and its success depends on the selection of its timer parameters, which requires difficult compromises. In addition, it requires complex packet processing in routers and for communications between stationary hosts and mobile hosts.

The Columbia University proposal [2] works very well in the local area but has poor wide area facilities (in particular, wide area routing).

The original IBM proposal [3] was designed assuming that loose source routing has been implemented correctly in existing networks. This assumption has since been shown to be incorrect. Even if loose source routing were implemented correctly on a wide scale it still has some problems related to UDP communications, security, an inability to handle base stations crashing, and inefficient handling by existing routers.

2 Infrastructure

A mobile host (MH) is the entity that is allowed to migrate around the IP network. It is assigned a constant IP address, known as its home address, on a home network. The home network may correspond to a physical subnet or a virtual subnet.

When a MH connects to the network at a new location it identifies, using a beaconing protocol, an Internet Access Point (IAP) connected to the local subnet that agrees to act as its agent. This access point also becomes the MH's default router. The address of the IAP is known as the MH's visiting address; the combination of a MH's home address and visiting address is known as a binding. The set of MHs using a particular IAP must be in the access zone of the IAP.

The details of the beaconing protocol by which a MH may acquire and reconfirm an IAP as its agent are not defined in this paper; but a registration and re-registration process similar to that specified by Columbia is one reasonable choice.

Whenever a MH successfully negotiates with a new IAP to act as its agent, it increments a time stamp. The time stamp is passed to the IAP using the beaconing protocol, and uniquely identifies the binding between the MH and the IAP.

The binding of a MH, and its associated time stamp, may be included in any data packet that the MH transmits (using a MIP option) and in two special management packets (MIP Binding Reply and a MIP Binding Update).

Any MIP entity may request a MH's current binding and time stamp by sending the MH a MIP Binding Request. The target MH will reply with a MIP Binding Reply. A MH may notify other MIP entities of its current binding and time stamp by transmitting either a MIP Binding Reply or a MIP Binding Update. The only difference between these management packets is that the MIP entity receiving a MIP Binding Update must acknowledge it with a MIP Binding Reply containing the same binding and time stamp. Note that all source bindings have time stamps. In what follows, the presence of a time stamp in such packets must be understood even if it is not explicitly mentioned.

Each IAP maintains a list of MHs, for which it is currently acting as an agent, along with the time stamp for each binding. These bindings are knows as local bindings. An IAP may also cache non-local bindings and associated time stamps for other MHs. An IAP may learn about a binding and its associated time stamp from a MIP Binding Reply, a MIP Binding Update, the source binding in a packet containing MIP routing information or the beaconing protocol.

An IAP may require a new local binding to be authenticated before using it. If an IAP requires a new non-local binding to be authenticated it may send a MIP Binding Request to the target MH requesting authentication.

Bindings for a MH may be deleted if they are not reconfirmed within a certain period. Local bindings may be reconfirmed by the beaconing protocol. Non-local bindings may be reconfirmed if the same binding and time stamp for a

MH is received in a MIP Binding Reply, a MIP Binding Update or from the source binding in a packet containing MIP routing information.

If a MH migrates and the IAP is not notified that it is no longer acting as the MH's agent, the local binding time out mechanism ensures that an IAP does not misroute packets forever. The non-local binding time out mechanism ensures that a IAP does not misroute packets forever if some other IAP that is the subject of a binding crashes or becomes inaccessible.

If an IAP receives a new binding, and the binding's time stamp is greater than the previous time stamp the IAP may update its corresponding binding entry after obtaining any desired authentication. In all other cases, the new binding is ignored, although the IAP may use a MIP Binding Request or the beaconing protocol to resolve any inconsistency.

When a MH migrates and acquires a new agent, the MH ensures that its previous IAP agent discovers this fact by sending it a MIP Binding Update containing its new binding. The MIP Binding Update is retransmitted until it is acknowledged with a MIP Binding Reply or until the MH can be sure that its previous IAP has timed out its local binding for the MH.

A MH also maintains non-local bindings. It uses exactly the same mechanisms as an IAP to create, modify and time out non-local bindings.

When a MH acquires a new IAP agent it also notifies its location directory (LD) of its new binding by sending the LD a MIP Binding Update (with an incremented time stamp) until it is acknowledged with a MIP Binding Reply.

Even after the initial acknowledgment, the MH should continue to send MIP Binding Replys or Binding Updates to its LD. This process ensures that the LD always has a current binding (along with its associated time stamp) in its directory for all the MHs it serves, even after a LD crash. The frequency of retransmission is a configurable parameter that represents a trade off between reliability and bandwidth use.

A LD is responsible for making bindings for its MHs available to a home redirector (HR) that serves the MH's home networks. The HR can keep these bindings permanently. The protocol between the LD and HR is not specified in this paper.

An HR appears to the rest of the network to be a router. As a router it may use existing routing protocols (e.g., RIP) to advertise reachability to the home networks it serves. HRs can tunnel received packets to the current location of any of its active MHs.

If the source MH, or any IAP that receives a packet, cannot route it to the destination MH's current location (e.g., because it does not have a binding for the MH), the packet destination IP address is set to the destination MH address. Existing routing techniques then forward the packet to a HR that serves the destination MH's home network. In the absence of errors, a HR always has a binding for all the active MHs in the home networks it serves and thus can always tunnel the packet to the destination MH's known location.

A mobile support router (MSR) is similar to an IAP except that a MSR does not keep local bindings. A MSR is expected to be an existing router with additional MIP capabilities. In particular, a MSR keeps a cache of non-local bindings which it can use to tunnel packets to a destination MH's current known location. In some ways, a MSR is similar to a Sony router[6].

MSR's establish and maintain their cache of non-local binding entries in the same way as an IAP or a MH. However, a MSR may also snoop on MIP Binding Updates, MIP Binding Replys and the source bindings of passing data packets to create and maintain its non-local bindings. A MSR may also use MIP Binding Requests to actively discover a MH's current binding.

The functions of a MSR are to optimize communications between MHs, especially when the network is segmented and a MH's LD and/or HR's are not accessible, and to reduce the effect of triangle routing when a SH sends packet to a MH. Whereas a Sony router is an integral part of the MH protocol defined by Sony, a MSR in MIP is only an optimization.

A SH never includes a source binding in its outgoing packets, and a MSR or HR never inserts one when tunnelling a packet from a SH to a destination MH. This ensures that MIP entities receiving the packet do not learn that the SH has a binding to the MSR. The use of a binding from the SH to the MSR might cause non-optimal routes to be used from the MH to the SH.

Although MHs, IAPs, LDs and HRs are specified separately in this paper, there is no reason that various combinations cannot be the implemented within the same machine. Figure 1 summarizes the relationships of the various MIP entities.

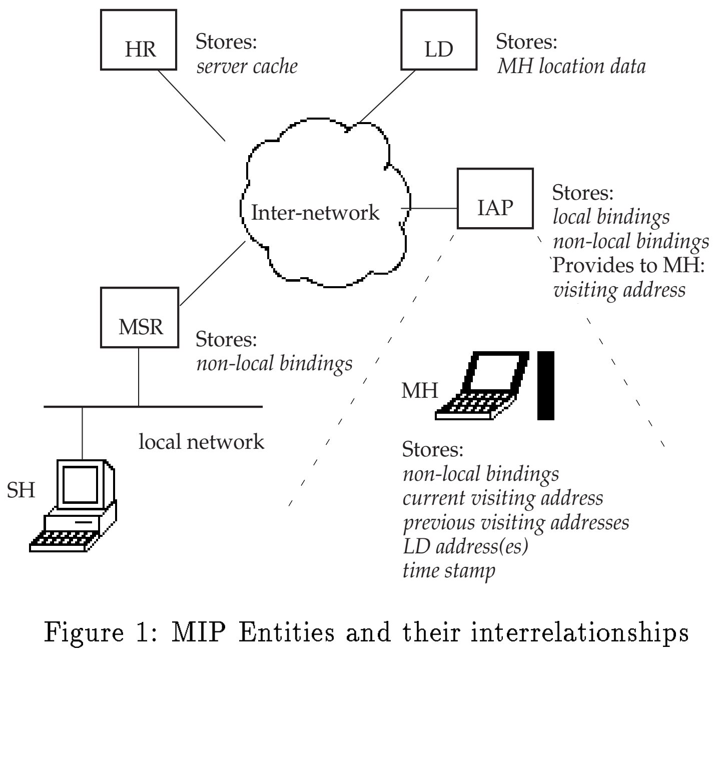

Figure 1: MIP Entities and their interrelationships

3 Option Syntax and Semantics

This section describes the syntax and semantics of the new IP option used by MIP. The MIP option (illustrated in Figure 2), consists of the following fields:

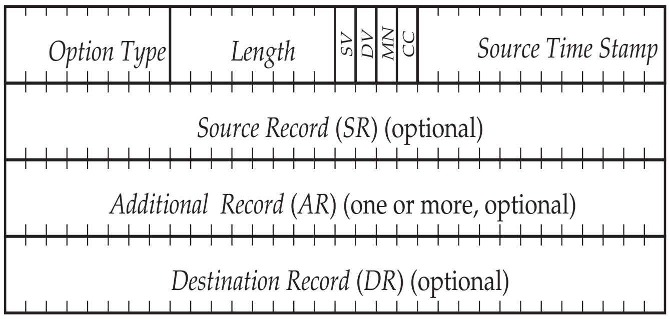

Figure 2: MIP Option

Option Type The option type field is one byte and has semantics as specified in the Internet Protocol [4].

Length The length field is one byte and has semantics as specified in Internet Protocol [4].

Source Valid (SV) The source valid bit indicates that the source record contains a valid address, the source time stamp field contains a valid time stamp and that the source MH binding in the MIP option may be used by other MIP entities to create or update location and agent cache entries.

Destination Valid (DV) The destination valid bit indicates that the destination record contains the address of the ultimate MH destination of the packet.

Mobile Network (MN) The mobile network bit is used to specify which MIP option is active when multiple MIP options are present in a packet. Multiple IP options are used to implement mobile networks in contrast to a single mobile hosts.

Source Record (SR) If the source valid bit is set then the source record exists and it contains the address of the IAP (or MSR) acting as an agent for the MH in the source IP address field. The source record, if it exists, always comes before the address record fields and the destination record.

Additional Record (AR) One or more additional records are used to hold the addresses of intermediate IAPs and MSRs that a packet passed through while being routed to its final MH destination. The number of additional records can be determined using the length of the MIP option and the existence of the source record and destination records. Additional record fields are used to aid in loop detection when inconsistent bindings are held by different entities..

Destination Record (DR) If the destination valid bit is set then the destination record exists and contains the address of the ultimate MH destination of the packet. The destination record, if it exists, always comes after the additional records and the source record.

4 Other Issues

4.1 Direct Communications

Many people will use their computers most often on their home networks. When doing so, they will expect the same level of service from their infrastructure as they receive today with, say, desktop workstations. In order to allow MHs to reside comfortably on their home networks, special care has to be taken with handling ARP requests from corresponding hosts.

A problem can arise if a MH, which has previously answered an ARP request, then moves away from its home network and leaves behind a stale entry in a SH's ARP cache. Since the way SHs handle ARP packets cannot be considered to be completely standardized, we suggest two methods for handling ARP requests for MH addresses, depending upon whether or not SHs can be relied upon to fulfil certain requirements in the way they handle their local ARP caches.

Define a gratuitous ARP as an ARP reply which is broadcast to all hosts in an Internet access zone, but not in response to any ARP request. Suppose all SHs listen to gratuitous ARPs, and use them to update any relevant entries in their local ARP caches whether or not the ARP reply was issued in response to an ARP request they had issued. Then a reasonably good solution is to make sure that a gratuitous ARP is issued on behalf of a MH whenever it migrates from one network to another. The gratuitous ARP will indicate that SHs should associate the IP address of the MH with the MAC address of the HR. The gratuitous ARP can be repeated as many times as needed to ensure that all the hosts on the home network do reliably receive it.

After the gratuitous ARP has been correctly processed by all hosts on the home network, packets sent to the MH's IP address (at the HR's MAC address) will subsequently be handled correctly as already specified by the MIP protocol. If the MH has moved away from its home network, then the gratuitous ARP will be issued by the HR, and subsequent packets destined for the MH will be sent through the HR just as desired.

More generally, in any Internet access zone, whenever a MH answers ARP requests and is being served by an IAP, that IAP must issue a gratuitous ARP on behalf of the MH after it moves. This gratuitous ARP will, again, associate the IP address of the MH with the MAC address of the IAP, and will be repeated by the IAP enough times to make sure that all the MHs within its access zone get the information. Subsequent packets destined for the MH will be directed to the IAP, which will then forward them using MIP.

If, on the other hand, SHs do not correctly process gratuitous ARPs, then the MHs must be configured to ignore all ARP requests. In this way, no SHs will ever store the MAC address of a MH within their ARP cache. On the home network at least, then the HR must perform proxy ARP duties on behalf of the MHs it serves. That way, all packets destined for MHs will end up at the HR, even it they could otherwise have been transmitted directly to the MHs on the same network. The HR must then forward the packets to the MH on behalf of the corresponding host. The HR is specified to perform the duties of an IAP for the MHs when they are connected to their home network, so the HR will have obtained the MAC address of the MHs using the beaconing protocol.

Sometimes SHs do not correctly process gratuitous ARPs, but nevertheless have a sufficiently small timeout value for entries in their ARP caches. Then it may be preferable to allow the MHs to answer ARP requests even though there will be a break in communications whenever the MH moves to a new access zone.

Direct communications between MHs and SHs connected to the same physical home network or the same physical network is an area that needs further study, to increase the flexibility of configurations on the Home Network while still allowing the mobile hosts to operate at maximum performance.

4.2 Cache Time Outs

Bindings have time outs on their entries. An entry is deleted unless it is reconfirmed during the time out period. When the entry is reconfirmed the time out period is reset.

Without a timeout an IAP might route a packet incorrectly so that it is lost for ever or until a higher layer timeout occurs. It is desirable that when a higher layer timeout occurs it informs MIP so that it can delete the corresponding cache entry.

The cache time out periods used determine the maximum disruption time if either an IAP crashes or a MH forgets to notify an old IAP. As both these events are expected to be rare, the time out may easily be long. A short time out period may cause non optimum routing to occur more often, because packets might need to be routed via the HR when IAPs, MHs and MSRs delete cache entries.

MIP avoids problems related to frequent cache time outs to certain extent by allowing cache entries to be reconfirmed (and thus the time out reset) by the source MH binding in incoming MIP packets.

4.3 Use of IP Options

MIP has a number of problems related to its use of an IP option to carry MIP routing information.

Experimentation has shown that not all hosts can accept IP options. Although this is contrary to standards as defined by the relevant RFCs, this means that these hosts would be unable to communicate with MIP mobile hosts as any packet the MIP mobile host transmitted containing a MIP option would be discarded.

It is also known that IP options are typically handled by routers on the Internet using a slow code path. When experiments were run sending large numbers packets containing IP options, many routers quickly became very congested. The only way to fix this problem is for routers to be re-implemented to handle IP options better.

4.4 Security

MIP also has a more serious problem related to security. As described, it assumes that an entity may immediately believe any binding it receives either in a management packet or in a source binding. This allows malicious hosts to send entities fake bindings and thus intercept packets or deny service to a mobile host.

The only way around this problem is for an entity receiving a binding to only treat it as being advisory and to authenticate the binding before use. One mechanism, called simple authentication, by which this can be achieved is described in [1] in relation to another mobile host protocol called IMHP.

5 Conclusion

MIP was designed at a time when security and authentication issues were not well understood and their importance was not recognised in a mobile environment. MIP also relied on the ability of the existing infrastructure to fully support IP options to transport mobility information around the network. Given these basic assumptions, MIP provides a very good basis for a mobile host protocol that satisfies most of the requirements proposed by [5].

However, given the realities of the Internet and the increasing importance of security and authentication, MIP was only a first step in the development of a mobile host protocol. Indeed, many of its concepts and its basic architecture have already been included in a new mobile host protocol, IMHP [1], that specifically provides authentication and security facilities.

References

[1] A. Myles C. Perkins D. Johnson. The internet mobile host protocol. In *Proceedings of the 2nd IEEE Int'l Conference on Universal Wireless Access*, pages 197–202, April 1994.

[2] J. Ioannidis G. Maguire Jr. The Design and Implementation of a Mobile Internetworking Architecture. In *Proceedings of Winter USENIX*, pages 491–502, San Diego, CA, Jan 1993.

[3] Y. Rekhter C. Perkins. Loose Source Routing for Mobile Hosts. Internet Draft, July 1992.

[4] J. Postel. Internet Protocol. RFC 791, Sep 1981.

[5] A. Myles D. Skellern. Comparing Four IP Based Mobile Host Protocols. *Computer Networks and ISDN Systems*, pages 349–356, Nov 1993.

[6] F. Teraoka M. Tokoro. Host Migration Transparency in IP Networks. *Computer Communication Review*, pages 45–65, Jan 1993.