

Comparison of Mobile Host Protocols for IP

Andrew Myles & David Skellern,

School of Mathematics, Physics, Computing, Electronics,
Macquarie University 2109,
Sydney, Australia.

andrewm@mpce.mq.edu.au

14 April 1993

Abstract

Host mobility is becoming an increasingly important feature with the recent arrival of notebook and palm top computers, the development of wireless network interfaces and the implementation of the global network. This paper describes and compares three proposals from Sony, IBM and Columbia University for mobile host protocols (MHP) that are compatible with the TCP/IP protocol suite. A set of basic requirements for a MHP are also suggested and it is observed that none of the three proposals entirely satisfies these requirements. Each proposal has faults in their implementation of mobile network layer functionality. Moreover, it is noted they do not address problems that must be solved in both higher and lower layers.

Key Words: Mobile Protocols, Wireless Networks, Networks, IP

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 globally communicate and access information has amplified the utility of computers many times over. Users increasingly want to attach to the network, with the same level of service, wherever they happen to be working.

A typical scenario might be a user that is usually connected to the network at their desk but occasionally would like to use their computer in a conference room in another part of the building or even at another site. The problem is that 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. Any host migration that does occur is assumed to be rare enough that it can be handled manually.

For example, consider the process of host migration in the above scenario using the TCP/IP protocol suite. If the user's desk and the conference room have direct

access to the same IP subnet then the migration process is trivial. In situations where this is not the case the only solution is for the user to 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. Even worse, the computer now 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. Often host migration will not even be under the user's control. For example, if a mobile host is 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. This trend will continue with the proliferation of metropolitan and wide area network services including SMDS, Frame Relay and ultimately B-ISDN. 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 in both the local and wide area.

The first goal of this paper is to evaluate and compare three leading proposals for mobile host protocols from Columbia University [Ioannidis, 92a], [Ioannidis, 92d], Sony [Teraoka, 91], [Teraoka, 92a], [Teraoka, 92b] and IBM [Perkins, 92a] that are compatible with the TCP/IP protocol suite. After discussing the meaning of mobility in Section 2, each of the proposals is briefly described in Section 3. Section 4 evaluates and compares the proposals in the context of a purely network layer solution to mobility. Section 5 discusses the reasons a network layer solution is insufficient and why mobility must be implemented at all levels.

2 Mobility

The primary aim of a *mobile host protocol* (MHP) is to provide host mobility. A host is defined to be mobile if, as it migrates around the local or wide area network, users transparently experience the same standard of network service as when the host remains attached to the network at a fixed location.

An essential requirement of mobility is thus *operational transparency* (ie, users do not need to perform any special actions due to host migration). Operational transparency can only be achieved by providing mechanisms to detect migration and to perform the appropriate actions to ensure continuing network services from all hosts to the *mobile host's* (MH) new location. A protocol that requires a host to be restarted every time it migrates is not mobile but merely *portable*.

However, operational transparency by itself does not necessarily lead to a useful service since it makes no guarantees with respect to the quality of service. In particular, operationally transparent migration can be achieved at the cost of a much lower level of performance.

If a MH user is running network applications at a particular location, their performance should be similar to that of the same applications running on a fixed host (FH). If the MH migrates to a new location the applications should continue to operate in exactly the same way and with similar performance as the same applications running on a FH at the new location. The use of a MHP should not affect the performance experienced by users of FHs. An additional requirement for mobility is thus *performance transparency*.

Factors that ensure a MHP has performance transparency include optimum routing of packets to and from MHs, efficient and robust migration procedures and efficient use of network resources such as transmission and processing bandwidth.

A MHP must also be feasible to implement in the real world. Relevant factors include cost, host impact and infrastructure requirements. Practical and market considerations also demand that a MHP be backward compatible with existing network protocols.

3 Mobile IP

According to [Postel, 81] the Internet Protocol (IP), in addition to fragmentation and reassembly, is responsible for providing, "the functions necessary to deliver a package of bits (an Internet datagram) from a source to a destination over an interconnected system of networks". Although the existing IP, IPv4, assumes that hosts are fixed, the above definition implicitly designates IP responsibility for routing datagrams to and from MHs.

Three current proposals, from Columbia University, Sony and IBM add functionality to IPv4 for routing to and from MHs in an attempt to provide mobility. Compatibility with IPv4 is a major criterion in all three proposals.

IPv4 assigns each host an address that is used by the higher layers to uniquely identify source and destination hosts. However, an IP address has a dual purpose. By its division into network and host parts, it also contains location information. An IP address, in its role as an identifier, must be constant during migration to maintain the host's identity. Therefore, the location information for a MH must be contained elsewhere and propagated through the network as required.

The major difference among the three MHP proposals is in how the location information is represented and propagated. In the following sub-sections the techniques used by each proposal are described and the differences highlighted. The reader may obtain further details by referring to the appropriate draft RFC ([Teraoka, 92a], [Ioannidis, 92a], [Perkins, 92a]) for each proposal.

3.1 Sony MHP

The Sony MHP identifies a MH using two 32 bit, IP style addresses; the *virtual IP* address (VIP) and the *temporary IP* address (TIP). The VIP is a permanent address that is used by the higher protocol layers. The network part of the VIP identifies the MH's home network.

When a MH is connected to the network it uses special local facilities (such as DHCP [Droms, 92]) to acquire a TIP on the local subnet. This address represents the current network location of the MH. If the MH is connected to its home network then the VIP and TIP are equal. A MH always informs a gateway attached to its home network of its current TIP.

The division of the location and identification information into two addresses requires extra space within packet headers. The Sony MHP defines a new IP option to carry the identification information (source and destination VIPs) while the address fields in the packet header carry the location information (source and destination TIPs). The VIP option, shown in Figure 1, also carries packet type, hold time and time stamp information.

Option Type	Option Length	VIP _{ver}	VIP _{type}	VIP _{hold}
source VIP address (VIP _{srcAddr})				
destination VIP address (VIP _{dstAddr})				
source address time stamp (VIP _{srcTS})				
destination address time stamp (VIP _{dstTS})				

Figure 1. VIP Option Format

The Sony MHP defines an efficient method of distributing the VIP to TIP mapping called the *Propagating Cache Method*. Routers and hosts that participate in the Sony MHP learn and cache a MH's mapping, in an *Address Mapping Table* (AMT), by snooping on the header and option fields of passing packets. In this way the information is distributed only to those parts of the network that might need to route packets to a particular MH.

An AMT entry is held until it is either invalidated, obsoleted, timed out or updated. It may be invalidated by a special management packet, *VipDelAmt*, that is transmitted by the home gateway when a mapping changes. The invalidation process is not guaranteed and when invalid entries are detected a *VipErrObs* packet is issued to obsolete the erroneous cache entries. Cache entries are timed out *VIP_{hold}* after they are either created or updated and are updated continuously by passing packets.

The operation of the Sony MHP can be best understood by considering an example where MH1 wants to send a packet to MH2. MH1 knows the VIP of MH2 but may not know its TIP (ie, its current location). In this case, it sets the IP address field of the packet equal to the VIP and transmits the packet. The packet will then be routed towards MH2's home gateway using existing IP routing facilities.

If an intermediate Sony router has a more recent mapping in its AMT cache (as measured by comparing the AMT and packet time stamps) it replaces the destination IP address field. The packet then propagates towards MH2's current location. The packet header may be modified a number of times in this way as the packet propagates through the network.

In the case where no intermediate Sony router knows the current mapping for MH2 or when there are no intermediate Sony routers the packet will eventually arrive at MH2's home gateway. If MH2 is attached to its home network, the packet will be delivered. If not, the home gateway will modify the destination IP address field to reflect MH2's current location and retransmit the packet.

Subsequently, when MH1 receives a packet back from MH2, it will store the mapping contained within the packet in its own AMT. All future packets from MH1 to MH2 will then be sent directly to MH2's current location through the most efficient route.

Intermediate Sony routers will also cache mappings for MH1 and MH2, which will assist in the efficient delivery of packets from other hosts to these MHs.

Non-Sony routers ignore the VIP option containing the VIP addresses and process the packet normally.

Now suppose MH2 migrates. On arrival at its new location it will acquire a new location address (TIP) and will inform its home gateway using the *VipConn* management packet. The home gateway tries to invalidate existing cache entries for MH2 throughout the network by transmitting a *VipDelAmt* packet. This packet should propagate to all Sony routers and hosts that contain a mapping for MH2. However, if obsolete cache entries still exist they will either be detected and invalidated during normal operation or they will timeout. After migration packets will be routed to MH2 using the process already described for initially locating MH2.

A further feature of the Sony MHP is its definition of a compatibility mode to ensure that MHs can communicate with FHs (ie hosts using existing protocols). The compatibility mode swaps the source IP and VIP in the packet so that a FH always receives the VIP as the source address.

3.2 Columbia MHP

The fundamental concept in the Columbia MHP is its definition of a virtual mobile subnet spread across a small number of real subnets. The mobile subnet exists wherever one of a small number of cooperating Mobile Support Routers (MSR) are located. The MSRs provide a gateway between the real subnets and the mobile subnet.

A MH is allocated a constant address on the mobile subnet. This means higher layer protocols have an unchanging view of a MH's identity. The fact that the mobile subnet is spread across a number of real subnets gives the MH its mobility.

Consider the example network shown in Figure 2, which illustrates a typical office network with a number of Ethernet segments (IP subnets 2.x, 3.x, 4.x) connected to a backbone (1.x) via routers. Cooperating MSRs are attached to each segment defining a mobile subnet (5.x) that is effectively spread across the three Ethernet segments.

Also shown is MH1 (5.1) and FH1 (2.3). MH1 is physically connected to segment 4 and has already registered with MSR 4 having identified it as its local MSR by listening for and responding to a *Beacon* packet that each MSR transmits at regular intervals.

Now suppose FH1 wants to transmit a packet to MH1. As the destination is on a different subnet the packet will initially be routed using existing routing protocols to the nearest router that advertises connectivity to the mobile subnet (probably MSR2 in this case). If MSR2 does not know the current location of MH1, it sends a *WhoHas* packet to each of the other MSRs¹ asking who has MH1 registered with it.

MSR4 will reply with an *IHave* packet which MSR2 caches as the current location of MH1. MSR2 then tunnels any packets addressed to MH1, using an encapsulation protocol, to MSR4, which de-

¹ Either using multicasting or by sending a packet to each of the other cooperating MSR's.

encapsulates the packet and delivers it to MH1. If MSR2 already had MH1's current location cached from previous activity, it does not need to perform the discovery process.

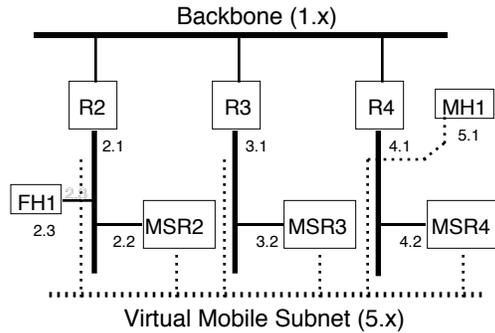


Figure 2. Example Columbia network.

Now suppose MH1 migrates to segment 3 and registers with MSR3. MSR3 will send an acknowledged *FwdPtr* packet to MSR4 informing it of MH1's new location. If MSR2 subsequently tunnels a packet destined for MH1 to MSR4, MSR4 will tunnel it to MSR3 and send a *Redirect* packet to MSR2 informing it of MH1's new location. Future packets from FH1 will then be tunneled by MSR2 directly to MSR3. Even if MH1 forgets where it was previously registered, cache time-outs ensure that invalid cache entries do not cause packets to be misrouted indefinitely.

The MHP, as described above, is not appropriate for wide area mobility. The process of identifying a MH's current location does not scale well since too many *WhoHas* packets are required to locate a single MH. It also requires a very large number of cooperating MSR's to be spread around the world. Recognising this, the Columbia MHP includes the following feature to provide wide area mobility.

When a MH migrates to a subnet where no MSR that handles its mobile subnet is attached, the Columbia MHP allows the MH to effectively become its own MSR. A MH in this mode of operation is known as a *popup*.

When a *popup* MH attaches to a foreign subnet it acquires a foreign address using local facilities (similar to a Sony MH). The MH then registers with a home MSR informing it of its foreign address. The MSR subsequently replies to any received *WhoHas* requests for the MH and tunnels any packets addressed to the MH to the MH's foreign address, where the MH performs its own de-encapsulation. Any packets that the MH transmits are tunneled back to the home MSR for forwarding.

3.3 IBM MHP

The IBM MHP is unique among the three MHP's discussed in that it relies on TCP and UDP layer facilities to assist in the provision of mobility. When TCP receives a packet with a loose source routing, (LSSR²) option in the IP header, [Braden, 89] specifies

² It is recommended that the reader refer to [jp] for a proper understanding of the operation of LSSR's.

that any packet sent in reply should employ the same LSSR in reverse. Similarly, when UDP receives a packet with a LSSR option, [Braden, 89] specifies that UDP should pass the LSSR to the application which should use the reverse LSSR for any reply. These facilities allow existing FHs to participate in routing to MHs.

Whenever a MH connects to the network it must register with a Base Station (BAS) that is also attached to the local subnet. A MH must also notify its Mobile Router (MR) of the address of the BAS (ie the MH's current location). A MR advertises connectivity to the group of MHs that it is configured to handle using existing routing protocols (eg RIP).

Now suppose a FH wants to transmit a packet to a MH. The FH transmits the packet normally and the packet is routed using existing routing protocols to the MH's MR. On receiving the packet the MR looks up the current location of the MH and inserts a LSSR in the packet with the MH's current BAS as the first hop.

The packet is then forwarded to the MH via the BAS using existing routing facilities. Any return traffic from the MH to the FH specifies a LSSR with the BAS as the first hop. According to [Braden, 89], the returned LSSR will force subsequent packets from the FH to the MH to be routed directly from the FH to the MH via the BAS, bypassing the MR. This process is illustrated in Figure 3.

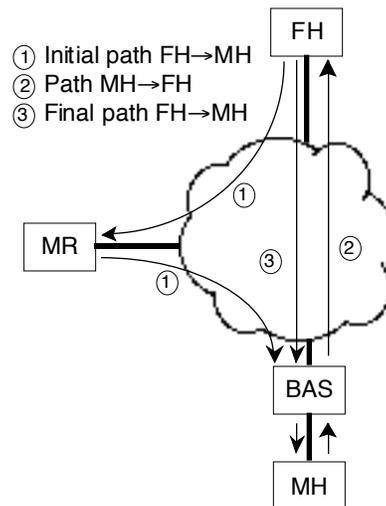


Figure 3. IBM MHP example.

When the MH migrates to another subnet it first registers with the new BAS, then notifies the MR of its new location and sends a message to the previous BAS telling it to delete its entry. Subsequent packets from the FH to the MH will be redirected by the original BAS back to the MR for correct routing until a return packet from the MH to the FH forces the FH to route the packets correctly via the new BAS.

Communication between two MHs is achieved using a similar process. The only difference is that the LSSRs will specify two hops, namely the BASs with which each MH is registered.

4 Comparison & Evaluation

This section evaluates and compares the three proposed MHPs on the basis of operational transparency, performance transparency and practicality in the context of a solution compatible with existing network protocols, as discussed in Section 2.

It should be noted that the three proposals discussed here are not the only proposals. At least two other serious TCP/IP based proposals have been made [Perkins, 93], [Wada, 92]. At the time of writing neither proposal was developed enough to fully include in this comparison. Work is also progressing to integrate mobility into the OSI architecture [Carlberg, 92]. Although the details of these proposals differ from those evaluated in this document, they use very similar techniques and have many features and problems in common with the three MHPs discussed.

4.1 Backward Compatibility

One of the main aims of the three MHPs is compatibility with existing network protocols. The investment in the existing network infrastructure is too large to discard simply for the sake of mobility.

The IBM MHP is based on the use of LSSRs, a facility not often used in the past. Although most routers process LSSRs correctly, it has been demonstrated [Ioannidis, 92b] that very few existing hosts implement LSSRs as specified by [Braden, 89].

In particular, most existing TCP/IP implementations upon receiving a packet with a LSSR option simply return it as a recorded route. Those few implementations that correctly reverse the LSSR, as required by [Braden, 89], do not track changes in the LSSR during a connection. Instead they always use the LSSR specified at connection set up.

An even greater problem occurs with UDP. [Braden, 89] requires that UDP pass received LSSRs to the application which is then supposed to reverse the LSSR for the reply. However, Unix sockets, for example, do not even have the ability to pass the appropriate information across the Application Programming Interface (API) and no known UDP application performs the correct LSSR reversal. Correct operation of the IBM proposal would therefore require changes to applications which is unacceptable in terms of backward compatibility.

It has been suggested [Deering, 92] that the IBM MHP interpretation of [Braden, 89] is incorrect as [Braden, 89] refers only to LSSR reversal on the reply to a request and not for future packets. A saving grace, however, for the IBM MHP is that communications can usually proceed, even if one end of a connection does not process LSSRs correctly, as packets to the MH will still be routed sub-optimally via the MR.

The ultimate success of the IBM MHP depends on future host software releases correctly implementing LSSRs. In the case of UDP applications this will require applications to be modified and new APIs to be specified. An alternative is to modify the UDP specification so that it handles LSSRs internally. This change has the undesirable side effect of making UDP a

state-based protocol. The usefulness of the IBM MHP must be judged on the likelihood and desirability of the appropriate modifications being implemented Internet-wide.

A lesser practical problem of backward compatibility effects both the IBM MHP and the Sony MHP. Some existing implementations, especially on PCs, drop any packet with any options [Myles, 92]. Until such implementations are fixed, communication between Sony MHs or IBM MHs and these implementations will be impossible. In contrast, the Columbia MHP has no specific backward compatibility problems.

A general problem that has been raised in relation to all three MHPs is an incompatibility with the developing multi-casting protocols. More study is needed to explore this issue.

4.2 Optimum Routing

A major factor in assessing the success of a MHP is its ability to transport packets through the network fabric using optimum or close-to-optimum routes. An optimum route can be defined as the route that existing routing protocols would use between two FHs. Sub-optimal routing results in wasted network bandwidth, extra processing and a lower overall level of performance delivered to the user. In terms of our mobility requirements it will impact performance transparency to the user.

When the Sony MHP routes packets between two MHs the packets are routed by an optimum route because each MH caches the current location of the other MH. The IBM MHP always uses close-to-optimum routes if both ends of the path implement LSSRs correctly. The Columbia MHP also uses close-to-optimum routes when operating in the local area. The IBM and Columbia MHPs may use a small number of extra hops depending on where the BASs and MSR are placed within the network topology.

There are situations, however, where each of the MHPs use very poor routes. An evaluation of the three MHPs depends on an assessment of how often these cases arise.

The worst case for the Sony MHP arises when a FH sends packets to a MH and there are very few Sony routers in the topology. In the example illustrated in Figure 4, the only Sony router is the MH's home gateway. All packets to the MH from the FH will thus be routed sub-optimally via the home gateway.

Configurations like this will be common in the short term. However, if the number of hosts and routers that understand the Sony MHP increases, the problem will gradually disappear. In the local area, it is practical, even in the short term, to optimise routing by placing an adequate number of Sony routers in the network topology.

Most existing host implementations do not process LSSRs correctly so the IBM MHP will, in the worst case, route all packets destined to a MH sub-optimally via its home MR. This problem will fade in the long term if hosts are updated to implement LSSRs correctly.

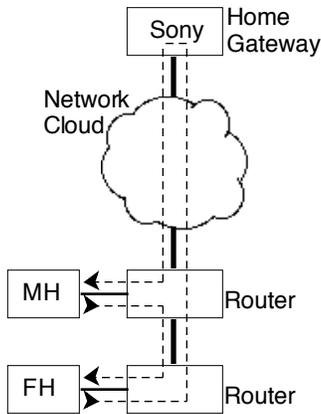


Figure 4. Non Optimum Sony MHP Routing

In the wide area, the Columbia MHP generally uses sub-optimum routes when a MH in *popup* mode communicates with a host not attached to the MH's home network. Consider the example illustrated in Figure 5 which shows two MHs from the same mobile subnet in *popup* mode connected to a foreign network. The route between the two MHs is through one or two MSRs connected at the home network.

If wide area mobile networking becomes widespread and MHs commonly want to access hosts not connected to the MH's home network, the Columbia MHP is inadequate. However as a short term solution, where local area mobility is likely to be most important, the Columbia MHP is ideal in terms of optimum routing.

In contrast, the IBM MHP is a longer term solution as it depends on modifications to a large number of hosts for optimum routing. Although the Sony MHP, with an appropriate investment in new routers, can provide close-to-optimum routing in the local area in the short term, it is also primarily a longer term solution.

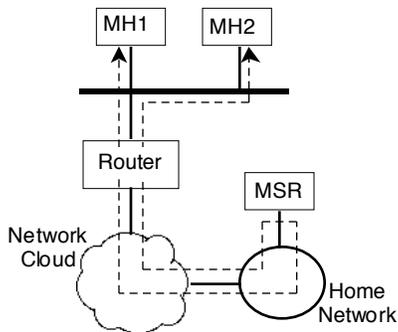


Figure 5. Wide Area Columbia MHP Routing

4.3 Intra-cell Communications

A topic closely related to optimum routing is that of intra-cell communications. Ideally two hosts that are connected to the same data link medium should be able to communicate directly without the assistance of a third party (eg BAS, MSR etc.). Depending on the lower layer protocols used, direct intra-cell communications may halve bandwidth requirements.

Only MHs can belong to a Columbia mobile subnet. This restricts so direct intra-cell communication to MHs when they are attached to the same network as an MSR that handles their mobile subnet. A difficulty arises when a MH migrates as the ARP cache of the remaining MHs must be updated to point at the MSR instead of the departed MH.

[Ioannidis, 92a] proposes that the MSR issue a gratuitous ARP to update the ARP caches of the remaining MHs when it discovers a MH has migrated. This method is unreliable as it depends on the error free reception of the gratuitous ARPs. An alternative has now been proposed [Ioannidis, 92c] that uses *Beacon* frames to flush the entire ARP cache of all MHs registered with the MSR when a MH migrates. Even though this technique causes a burst of ARP traffic just after a MH migrates, it is reliable and fast.

The IBM MHP has a similar problem with direct intra-cell communications. [Perkins, 92a] suggests the use of short ARP time-outs in MHs. This approach has the drawback that it causes too many unnecessary ARP time-outs. Longer time-outs will result in noticeable breaks in network service to the user. However, there is no reason the technique suggested by [Ioannidis, 92c] could not be used.

The Sony MHP allows [Teraoka, 92c] a FH and a MH that belong to the same IP subnet to take part in direct intra-cell communications. When the MH migrates away from the home subnet the ARP cache in the FH can only be updated using existing protocols such as gratuitous ARPs or ARP time-outs. However, gratuitous ARPs are unreliable and many existing implementations do not have ARP cache time-outs.

A more serious problem associated with the Sony MHP relating to intra-cell communication between MHs is discussed in the next section.

4.4 Migration Procedures

It is important that the migration procedures in a MHP be efficient and reliable so that communications can restart as soon as possible after a MH migrates to a new location. Long migration procedures will result in unacceptable breaks in network connectivity.

Under normal circumstances both the IBM MHP and the Columbia MHP have efficient and reliable migration procedures. This is not the case with the Sony MHP.

Consider the situation where two Sony MHs are sending packets to each other. Now suppose one of the MHs migrates to a new location. The MH will notify its home gateway of its new location and the home gateway will then attempt to delete all cache entries throughout the network for the migrating MH.

Unfortunately the procedure for propagating the *VipDelAmt* packet relies on all routers understanding the Sony protocol. This is unlikely to be the case even in the very long term and so the cache deletion procedure will invariably fail.

Depending on subsequent updates from the migrating MH and on how many obsolete cache entries remain in the network fabric, communications with the migrating

MH might be impossible for up to the time-out period of the cache entries. If the time-out is too long, the user will notice a break in communication. If it is too short, packets will often be routed sub-optimally via the home gateway. The Sony MHP migration procedures are thus unreliable, which will degrade the quality of mobile service.

The Sony MHP migration procedure also potentially makes inefficient use of network transmission bandwidth. Every time a MH migrates to a new location an attempt is made to update all routers and hosts that contain information about the MH's location. The update process can use a large amount of bandwidth.

In contrast, the Columbia and IBM MHPs only update location information at one or two places. Other updates are made on an as-needed basis. This mechanism reduces the total amount of bandwidth required and spreads it over a period of time.

4.5 Packet Loss

Each of the three MHPs makes extensive use of management packets to update information around the network. It is important that there be efficient recovery mechanisms in place when a packet is lost, which is particularly likely in a wireless environment.

The IBM MHP and Columbia MHP use acknowledgments to ensure that the important management packets are received correctly. If an error does occur the recovery occurs after one or more retransmissions. The retransmission mechanism is designed so that it does not affect normal error-free operation of the MHP.

This is not the case for the Sony MHP. The situation described in the previous section can also arise from the loss of the *VipDelAmt* packet. A fast recovery depends on small cache time-out periods, which adversely affect the general operation of the MHP.

4.6 Component Loss

A MHP should be able to survive the loss of major components (MR and BAS in the IBM MHP, MSR in the Columbia MHP and Sony routers and home gateways in the Sony MHP). The recovery probably does not need to be very fast as major components should not fail very often.

The Columbia MHP will generally recover from the loss of a MSR after a cache time-out. Assuming that the MHs that were connected to the down MSR can attach to an alternative MSR most users will have almost no indication that there was any problem.

The Sony MHP concentrates a substantial portion of MHP functionality in a home gateway which makes it vulnerable to failure but there is no reason that the home gateway cannot be duplicated to increase resilience. The cost of home gateway duplication, besides dollars, is extra routing protocol traffic. Although not as robust as the Columbia MHP, home gateway duplication will provide sufficient reliability for most applications. The loss of a Sony router in the Sony MHP is no worse than the loss of a non-mobile router.

The IBM MHP can similarly use duplication of its MR to provide robustness against MR loss but the loss of a BAS is a more serious problem. If a BAS fails, any MHs registered with it will presumably register with a new BAS. However, hosts that were previously communicating with a MH attached to the down BAS will continue to unsuccessfully route packets through the down BAS.

This situation will only resolve itself if the MH happens to send a packet to the other host causing its LSSR for the MH to be updated. The root of the problem is that [Braden, 89] does not specify time-outs for the cached LSSRs. Higher layers will probably eventually timeout the connection and depending on the protocols and applications used may or may not recover the session.

The potentially infinite recovery time after the loss of a BAS will certainly result in the loss of performance transparency to the user. If network applications need to be restarted operational transparency will also be lost as user intervention is required.

One solution would be to change the specification for TCP and UDP handling of LSSRs to include a time-out. However, this is a long term solution as it requires changes to TCP and UDP to be specified and widely implemented.

Another alternative is to require a migrating host to send an update on all existing connections. This solution is also not backwardly compatible and it is not clear how UDP applications would be updated as UDP is not connection orientated.

It has also been suggested [Perkins, 92b] that a MH may act as its own BAS when one is not available, by acquiring a local address using methods similar to the Sony MHP and to the Columbia MHP in *popup* mode. This means that BASs will, in effect, go down every time a MH migrates, making an efficient solution to the problem even more important.

4.7 Network Breakage

An ideal MHP should allow a MH to operate normally, meaning continued communication with any hosts on the same segment, even if the network has been broken. Network breakage is similar to the loss of a large number of components and generally the same techniques can be used to increase robustness.

In the Columbia MHP, during wide area operation, continued operation for a MH after network breakage depends on whether the MH can contact an MSR. Placement of additional MSRs in the network away from the home network will increase robustness.

Recalling section 4.6 component duplication thus has a use in all three MHPs. If component duplication is not used then network breakage can easily result in a MH losing communication with hosts on both sides of the break.

The Sony MHP has a small amount of additional protection against breakage. Suppose a MH is separated from its home gateway. Hosts in the same segment as the MH will be able to send packets to the MH as long as they happen to pass through a Sony router that has the MH's current location cached. The mechanism is

not guaranteed and successful operation depends on traffic patterns, cache time-outs and the number of Sony routers in the topology.

4.8 Processing Bandwidth

One area in which the Columbia MHP has an advantage over the Sony MHP and IBM MHP is the processing associated with the MHP. A MHP that requires large amounts of processing will have lower network performance or higher cost. Lower network performance will affect performance transparency.

The Columbia MHP requires processing bandwidth to encapsulate and de-encapsulate packets destined for MHs. It also requires a small amount of protocol processing bandwidth (eg for sending *WhoHas*, *Redirect* packets, etc.)

The processing bandwidth requirement for the other two protocols is much larger. One source of extra processing results from the use of IP options. Existing routers generally have an optimised processing path for optionless packets and a much slower path for packets with options. In the long term, it might be possible to optimise the processing of specific options (such as the Sony option or the LSSR option) as well.

In addition, the Sony MHP requires every Sony router to process every packet passing through the router. The extra processing is substantial, involving table look-ups and sometimes modifying the packet. [Teraoka, 92b] showed that the total overhead is at least 29% more than the equivalent IP processing, resulting in a potential throughput bottleneck.

4.9 Impact on Existing Routing Protocols

The Columbia MHP has an impact on network processing bandwidth. To achieve optimum routing each MSR by advertising connectivity to the virtual mobile subnet will place an additional load on the routing protocols (such as RIP and OSPF).

In contrast, the Sony MHP and the IBM MHP only need to advertise connectivity to MHs at the home gateway and MR respectively.

4.10 Infrastructure

A MHP will have a better chance of widespread implementation if it requires a minimal additional infrastructure. This introduces a new requirement for MHPs; that they should be practical to implement in the real world.

The Columbia MHP, for local area mobility, requires a MSR to be connected to every subnet that is part of the virtual mobile subnet. For wide area mobility it needs facilities to acquire a local IP address in foreign subnets.

The Sony MHP always requires facilities to acquire a temporary IP address. It also requires a home gateway for each mobile subnet and for efficient routing it requires numerous Sony routers to be spread throughout the network. The main practical difficulty, in the short term, with a mixed community of FHs and MHs, is that optimum routing requires many existing routers be converted to understand the Sony MHP.

The basic IBM MHP requires a BAS to be located anywhere a MH may be connected as well as a MR for every mobile subnet. This has the practical difficulty that it is unlikely a BAS will be available everywhere a user would like to connect to the network. The universal requirement for a BAS could be relaxed if a *popup* mode similar to that employed in the Columbia MHP was added. It has also been suggested [Perkins, 92b] that a MH could act as its own BAS but this will run into the problem of disappearing BAS's discussed in Section 4.6.

It is difficult to come to a definitive conclusion about the relative merits of the MHPs with regard to infrastructure as each requires substantial infrastructure for effective mobility.

4.11 Migration Detection

The Sony MHP makes extensive use of local facilities to acquire an IP address every time a MH attaches to a network. The Columbia MHP makes use of similar facilities when a MH is in *popup* mode. Neither proposal defines the exact protocols used although the use of DHCP [Droms, 92] is suggested [Ioannidis, 92a].

A MH needs to know when it has migrated to a new subnet so that it can carry out the appropriate migration procedures. The IBM MHP and the Columbia MHP (when not in *popup* mode) use *Beacons* transmitted by the local BAS and MSR respectively to detect migration. MHs in the Sony MHP and in the Columbia MHP in *popup* mode must use some other method.

One choice is the MHP use an indication from the data link layer every time the host connects to the network. This approach is unsatisfactory as connection to a data link layer does not guarantee subnet migration and many data link layers are incapable of providing an appropriate indication. A second possibility is the provision of network layer facilities to detect migration. Neither the Sony MHP nor the Columbia MHP specifies such facilities and DHCP is not suitable.

Unless appropriate facilities are specified, the Sony MHP and the Columbia MHP in *popup* will require user intervention to indicate network migration has occurred. This is unsatisfactory as it does not provide transparent migration to the user.

There is no reason that these facilities cannot be defined and implemented widely in the longer term. However the facilities are likely to be specific to MHPs. Therefore it can be argued that whereas wide area mobility in the IBM MHP is limited by the spread of BASs, wide area mobility for the Columbia MHP and Sony MHP is limited by the spread of these new protocols.

4.12 Address Usage

A limitation of IPv4 is that it is rapidly running out of address space. It is thus desirable that a MHP does not put additional pressure on this scarce resource.

The IBM MHP is the best in this regard as each MH only ever has one address assigned to it. The only additional addresses needed are for each BAS in the topology.

The Sony MHP, in contrast, has a huge requirement for extra addresses. Each subnet must reserve a block of addresses equal in size to the maximum number of MHs expected to be connected to that subnet at any one time. Network wide, the number of addresses that need to be reserved is many times the number of actual MHs.

The Columbia MHP is address efficient in the local mode of operation needing only one address for each MH and MSR but has a similar problem to the Sony MHP in the wide area mode of operation.

4.13 Security

Security requirements are not well understood even in the context of non-mobile hosts. The use of MHs simply adds complexity to the problem. The three proposed MHPs recognise the importance of security by providing appropriate hooks for the addition of security functions but none of them define specific procedures.

The two most common approaches to network security today are to ignore it altogether or to make access very inconvenient through the use of facilities such as fire walls. Wide area mobile networking makes it vital that mechanisms are designed to allow users access to facilities from anywhere in the network in a manner that does not beg security violations.

4.14 A Partial Conclusion

This section has reviewed three proposed MHPs that provide host mobility in a TCP/IP environment by adding functionality to IP. The goal of each MHP is to provide optimum, robust routing to and from MHs in a way that is practical to implement and compatible with the existing network infrastructure.

Under the proposed criteria, the question as to which MHP is the best is still open to debate. Table 1. summarises the criteria discussed. Each has a range of advantages and disadvantages which more or less balance each other. Moreover, none of the problems highlighted for each of the MHPs are serious enough to discount it from consideration as a contender for the best MHP proposal.

Consideration of their appropriateness for short term and long term application, however, serves as a useful basis of comparison. In the short term, local area mobility and practical considerations will be more important, while in the longer term, wide area mobility will become paramount and it will be feasible to implement more of the fundamental changes required in network operation.

In the short term, the Columbia MHP seems to offer the best overall solution. It provides close to optimum routing to and from MHs in the local area with a minimum of infrastructure in a manner that is robust and makes minimal demands on the existing network. In the wide area, it offers an unsatisfactory but workable solution that uses sub-optimum routes and relies on protocols that do not currently exist.

Although the Sony MHP can also provide optimum routing, to do so requires at least some existing routers to be modified to understand the MHP. This makes its use less attractive from a practical point of view.

Additionally, it is not as robust as the Columbia MHP, requiring facilities (local address acquisition) that have not yet been fully specified. It will also reduce the performance of existing routers due to its use of options and is incompatible with many existing PC implementations that discard any packets containing options. For these reasons the Sony MHP can only be a second choice for the short term.

The IBM MHP cannot be considered as a possibility in the short term as very few hosts implement LSSRs correctly resulting in most packets being routed sub-optimally. Even those hosts that do implement LSSRs as interpreted by [Perkins, 92a] will be unable to route UDP optimally. Additionally, the use of an IP option results in incompatibility with most PCs and reduced performance from existing routers.

In other respects, the IBM MHP is an effective protocol and if the appropriate modifications can be made to the TCP and UDP specifications, in the longer term it could provide a good homogenous solution to the problem of MHs.

In contrast, the lack of efficient wide area functionality makes the Columbia MHP unsatisfactory in the long term. As users make use of wide area network facilities it is unacceptable for packets to and from a MH to be routed through the home network. Address space usage also makes the Columbia MHP less desirable in the wide area in the longer term.

The Sony MHP is better in the long term than in the short term as more hosts and routers can be expected to understand the protocol. The main concerns with the Sony MHP are the amount of processing that is needed at intermediate nodes, a general lack of robustness and the requirement for a huge pool of reserved addresses for MHs.

The number of proposals for MHPs has not diminished in recent times. Already improvements to the three reviewed MHPs have been suggested.

The authors of the IBM MHP, for example, have recently proposed [Perkins, 93] a variation of the IBM MHP that uses encapsulation rather than LSSRs to carry the mapping between the identification information and the location information. The result is a very similar MHP without the problems associated with using and modifying LSSRs. It cannot optimally route packets from a FH to a MH, however it is likely to be the basis of a good long term solution as more hosts understand the MHP.

A similar variation has been by [Wada, 92]. However, this has two main disadvantages compared to [Perkins, 93]. Firstly, it wastes scarce address space by requiring MHs to acquire a temporary address. In contrast, the IBM MHPs always use the BAS address as an MH's location contact point.

Although the use of an individual temporary address allows hosts to send packets directly to a MH without a PFS (equivalent to a BAS) hop, it also means that the PFS must promiscuously process every packet on the network just in case a host sends a packet to a MH that has subsequently migrated. This is unsatisfactory. Overall [Perkins, 93] is a better choice than [Wada, 92].

5 Mobility Considerations At Other Layers

The three MHP proposals discussed here generally confine themselves to the IP layer. A key question which remains is this; is it sufficient to tackle the mobility problem at this layer alone?

This section considers a number of areas where support for mobility is needed at layers other than the network layer. Part of the discussion is based on results from an implementation of the Columbia MHP written at Macquarie University using modified HP27285A Router ERs as MSRs and PCs running a modified NCSA Telnet as MHs.

Criteria	Section	Sony MHP	Columbia MHP	IBM MHP
Fully Backward Compatible?	4.1	NO <ul style="list-style-type: none"> No multicast support. IP Option dropped by some hosts. 	NO <ul style="list-style-type: none"> No multicast support. 	NO <ul style="list-style-type: none"> No multicast support. IP Option dropped by some hosts. LSSR unimplemented by existing hosts.
Uses Optimal Routing?	4.2	LONG TERM or LOCAL AREA <ul style="list-style-type: none"> Requires majority of Sony gateways or Sony hosts to be deployed. 	LOCAL AREA ONLY <ul style="list-style-type: none"> Always uses sub-optimal route in <i>popup</i> mode. 	LONG TERM or LOCAL AREA <ul style="list-style-type: none"> Requires LSSR to be implemented correctly.
Optimum Intra - Cell Comms.?	4.3	NO <ul style="list-style-type: none"> Cache deletion after migration unreliable. 	LIMITED <ul style="list-style-type: none"> MHs in same virtual mobile subnet only. 	POSSIBLE <ul style="list-style-type: none"> Needs protocols to be defined.
Migration Detection Procedures Defined?	4.11	NO <ul style="list-style-type: none"> Assumes data link support. 	LOCAL AREA ONLY <ul style="list-style-type: none"> Assumes data link support in <i>popup</i> mode. 	YES
Reliable Migration Procedures?	4.4	NO <ul style="list-style-type: none"> Cache deletion after migration unreliable. 	YES	YES
Robust against Component Loss?	4.6	YES <ul style="list-style-type: none"> Requires <i>home gateway</i> duplication. 	YES <ul style="list-style-type: none"> Requires MSR dup. in wide area. 	NO <ul style="list-style-type: none"> Loss of BAS critical. Requires MR dup.
Use of Network Resources?	4.8 4.12 4.9	INEFFICIENT <ul style="list-style-type: none"> Req. large pool of temporary addresses. IP Options req. extra processing in routers. Extra processing req. in Sony gateways. 	INEFFICIENT <ul style="list-style-type: none"> Req. large pool of temporary addresses in <i>popup</i> mode. Each MSR must advertise connectivity. 	INEFFICIENT <ul style="list-style-type: none"> IP Options req. extra processing in routers.
Security Defined	4.13	NO <ul style="list-style-type: none"> Hooks only. 	NO <ul style="list-style-type: none"> Hook only. 	NO <ul style="list-style-type: none"> Hooks only.
Existing Infrastructure Adequate?	4.10	NO <ul style="list-style-type: none"> Home gateway. Many Sony gateways. Address acquisition facilities. 	NO <ul style="list-style-type: none"> MSR wherever a MH may connect. Address acquisition facilities for <i>popup</i> operation. 	NO <ul style="list-style-type: none"> MR. BAS wherever a MH may connect.

Table 1. Summary of Evaluation Criteria

5.1 Data Link Layer Support

Regardless of the network layer MHP chosen, data link layer support is required to provide mobility across bridges. To understand the reason for this consider the example illustrated in Figure 6.

Suppose MH2 is connected to the right hand interface of bridge B and is communicating with FH1. Both bridges will learn that MH2 is accessible through their right hand interfaces and FH1 is accessible through their left hand interface. Now suppose MH2 migrates to the left interface of bridge A.

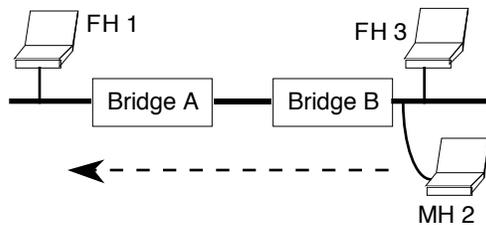


Figure 6. A bridged mobile architecture.

MH2 and FH1 will continue communicating successfully. However FH3 will not be able to contact MH2 unless MH2 transmits a packet that passes through bridge B forcing it to modify its forwarding table, or the forwarding table entry times out. Time-outs may not always work as many commercial bridge products have been found to not time-out forwarding table entries.

A solution to this problem, which has been implemented at Macquarie University on a bridged Ethernet, is that when MH2 reconnects to the network on the left hand interface of bridge A it transmits a broadcast packet forcing all bridges to update their forwarding tables.

In the Macquarie University implementation, the broadcast packet is initiated by the user (violating the requirement for operational transparency) as there was no reliable way of detecting within the MH that it has connected to a different Ethernet segment. This might not be a problem when using wireless media as it is more likely that they will be capable of providing a suitable indication.

The transmission of a broadcast packet is potentially unreliable as the technique relies on the broadcast packet not being corrupted or dropped as it propagates through the bridged network and on the bridges learning on every packet. One commercial bridge at Macquarie University has been found to only update the forwarding table on every fourth packet. The MH could transmit multiple broadcasts to reduce the possibility that not all the forwarding tables are updated, however this will only decrease the odds of a problem occurring and does not provide a guarantee.

The user intervention problem and the lack of robustness can both be solved if the MH regularly transmits a broadcast packet. The disadvantage of this approach is that either it wastes significant transmission bandwidth or if the time between broadcasts is more than about a second, performance transparency is

compromised. While many other data link layer MHPs are possible, this solution illustrates the trade-offs that have to be made in design.

5.2 Transport Support

If a TCP entity is transmitting data when the network layer temporarily disconnects it is possible that packets will be lost. In this situation the TCP protocol will back off and retransmit the packets after a timeout. The back-off time will increase, to a maximum, every time a packet is not acknowledged.

In a wireless mobile environment it is likely that short network layer breaks will occur quite often as a MH migrates between wireless cells. In a wired mobile environment the breaks will not occur as often but are likely to be much longer.

In experiments at Macquarie University with the TCP protocol running over a Columbia MHP implementation the transport layer delay from the user's perspective was noticeable even when the disconnection time was short. A network layer break of about ten seconds often resulted in a transport layer delay of thirty to forty seconds³. This behaviour is unacceptable in terms of performance transparency. The Sony MHP and the IBM MHP have similar transport layer performance problems.

The obvious conclusion is that mobility functions cannot be confined simply to the network layer. The experiments described above indicate a clear need for modifications at the transport layer to avoid amplified breaks.

[Teraoka, 91] suggested modifications that halted transport connection during a break to provide for what it called *on-line* migration. Unfortunately, they assumed that a MH can send messages before it migrated which is unlikely to be the case in many situations. However it is not difficult to design alternative modifications that perform the required functions. The main problem is that the new protocols are not likely to be backward compatible with existing protocols, so their introduction can only be a long term objective.

5.3 Application Support & Soft Hand-Over

In a wired mobile environment or in a wireless environment where there are holes in the wireless coverage, breaks in network access are inevitable. Potentially the breaks in access can be long. For example, consider a user that wants to carry a wireless notebook computer down a hall, with no access points, to another room. Many minutes could elapse before network access is again established.

If the break in connectivity causes applications to crash then operational transparency will be lost. A MHP that effectively requires a host to be restarted every time it migrates is not mobile but merely portable.

Therefore support is even needed from applications to ensure true mobility. Network Applications must be designed so that they can continue to operate after a

³ In more recent versions of Unix it was found that sending three packets to the transport entity short circuited the break as the estimated round trip delay was reduced.

temporary break in network access. The hurdle is that applications must somehow distinguish between temporary and permanent loss of connectivity to the network. There is no obvious way to do this.

In many applications it is very important that network connectivity is maintained at all times during migration. The cellular telephone system attempts to do this by using soft hand-over techniques. In soft hand-over a MH does not deregister from its current wireless cell until it has registered with the next cell.

Soft hand-over can require support at the physical layer (multiple transmitters and receivers), the data link layer (registration within two cells) and the network layer (routing packets to two locations while handover occurs). Soft hand over is yet another example of why network layer mobility support alone is insufficient in some cases.

5.4 Facility Management

Consider the example of a user connected to a foreign network who would like to use local print facilities. Many existing networks require that the user know the name and address of the local printer. However protocols need to be defined so that a mobile user can determine what local printers are available.

Local facility management does not restrict itself to printers. Directories of databases (such as a local telephone directory), disks (for temporary storage) and processors (to off-load computations) will all be relevant to the mobile user. True operational transparency and thus mobility depends on the appropriate protocols and applications being available.

5.5 File Transparency

Users connected to a home network commonly access a local file server and become accustomed to a certain level of service when accessing information. Now suppose the user migrates to another location. Ideally the user will be able access the information with a similar perceived performance.

This can be achieved by defining protocols for caching the information on a file server on the foreign network. In this way the user can achieve performance transparency. If the user now migrates to another network the cached information should somehow follow the user. Issues of file locking and synchronisation must also be addressed. Some work has been done in this area [Badrinath, 92], [Tait, 91].

6 Conclusion

True mobility requires MHPs that allow transparent access to network services in the local and wide area. A MHP, in addition to having operational and performance transparency, must satisfy some practical criteria that allow them to be implemented in the real world.

This paper reviewed three proposed MHPs that are compatible with the IP and evaluated them against suggested requirements of mobility. It has been shown that all three MHPs need further work to provide true mobility.

The Columbia MHP is the most suitable for short term application as it is most effective in local area operation. Its lack of efficient wide area capability restricts its long term use. However the IBM MHP (and its variation specified in [Perkins, 93]) could well be the best basis for a longer term solution.

Regrettably, mobility impacts all layers in the protocol suite and cannot be confined to the network layer. In this sense these proposals are only a first step towards providing true mobility to the user.

Within the next few years a new version of IP will almost certainly be introduced. It seems that this is a good opportunity to take the best ideas from the Sony, IBM and Columbia MHPs and combine them into the new IP. Mobility is so important that it should be fully integrated into the architecture at all levels rather than simply being bolted on as an after thought.

References

- [Myles, 92] Andrew Myles, Macquarie University, *IP Options*, mobile-ip mailing list, 29□July 1992.
- [Badrinath, 92] B.R.Badrinath, T.Imielinski, Rutgers University, *Replication and Mobility*, 1992.
- [Perkins, 92a] Charles Perkins, Yakov Rekhter, T.J.Watson Research Centre, IBM, *Short-cut Routing For Mobile Hosts*, draft RFC, 17 July 1992.
- [Perkins, 92b] Charles Perkins, T.J.Watson Research Centre, IBM, personal communication, 26□August□1992.
- [Perkins, 93] Charles Perkins, Yakov Rekhter, *Support for Mobility with Connectionless Network Layer Protocols (Transport layer Transparency)*, draft RFC, January 1993.
- [Tait, 91] C.Tait, D.Duchamp, Columbia University, *Service Interface and Replica Consistency Algorithm for Mobile File System Clients*, Conf. Parallel and Distributed Information Systems, Dec. 1991.
- [Teraoka, 91] Fumio Teraoka, Yasuhiko Yokote, Mario Tokoro, Sony, *A Network Architecture Providing Host Migration Transparency*, Proceedings of SIGCOM'91, ACM, pp 209-220, September 1991.
- [Teraoka, 92a] Fumio Teraoka, Sony, *VIP: IP Extensions for Host Migration Transparency*, Draft RFC, July□1992.
- [Teraoka, 92b] Fumio Teraoka, Kim Claffy, Mario Tokoro, Sony, *Design Implementation and Evaluation of Virtual Internet Protocol*, Proceeding of the 12th International Conf. on Distributed Computing Systems, pp 170-177, June 1992.
- [Teraoka, 92c] Fumio Teraoka, personal communication, 20□August 1992.
- [Ioannidis, 92a] John Ioannidis, Dan Duchamp, Gerald Q. Maquire Jr., Columbia University, Steve Deering, Xerox PARC, *Protocols for Mobile Internetworking*, Draft RFC, June□1992.
- [Ioannidis, 92b] John Ioannidis, Columbia University, *Source Routing Experiments*, mobile-ip mailing list, 27□July□1992.

- [Ioannidis, 92c] John Ioannidis, Columbia University, personal communication, 24 Aug. 1992.
- [Ioannidis, 92d] John Ioannidis, Dan Duchamp, Gerald Q. Maquire Jr., Columbia University, *IP-based Protocols for Mobile Internetworking*, Proceedings of SIGCOM'91, ACM, Sept. 1991, pp 235-245.
- [Postel, 81] Jon Postel, *Internet Protocol*, RFC 791, September 1981.
- [Carlberg, 92] Kenneth G. Carlberg, SAIC, *A Routing Architecture That Supports Mobile End Systems*, MILCOM'91, vol 1, Oct.1992..
- [Deering, 92] Steve Deering, Xerox PARC, personal communication, 20 August 1992.
- [Wada, 92] Hiromi Wada, Tatsuya Ohnishi, Brian Marsh, *Packet Forwarding for Mobile Hosts*, Draft RFC, November 1992.
- [Braden, 89] R. Braden, *Requirements for Internet Hosts -- Communication Layers*, RFC 1122, October 1989.
- [Droms, 92] R. Droms, *Dynamic Host Configuration Protocol*, Draft RFC, August 1992.