

Mobility Support for Nimrod : Challenges and Solution Approaches

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

We discuss the issue of mobility in Nimrod. While a mobility solution is not part of the Nimrod architecture, Nimrod does require that the solution have certain characteristics. We identify the requirements that Nimrod has of any solution for mobility support. We also classify and compare existing approaches for supporting mobility within an internetwork and discuss their advantages and disadvantages. Finally, as an example, we outline the mechanisms to support mobility in Nimrod using the scheme currently being developed within the IETF - namely, the Mobile-IP protocol.

Table of Contents

1	Introduction.....	1
2	Mobility : A Modular Perspective.....	2
3	Effects of Mobility.....	4
4	Approaches.....	6
5	Solution using IETF Mobile-IP.....	10
5.1	Overview	10
5.2	Protocol Details.....	11
6	Security Considerations.....	15
7	Summary.....	16
8	Acknowledgements.....	16
9	Author's Address.....	17

1 Introduction

The nature of emerging applications makes the support for mobility essential for any future routing architecture. It is the intent of Nimrod to allow physical devices as well as networks to be mobile.

Nimrod, as a routing and addressing architecture, does not directly concern itself with mobility. That is, Nimrod does not propose a solution for the mobility problem. There are two chief reasons for

this. First, mobility is a non-trivial problem whose implications and requirements are still not well understood and will perhaps be understood only when a mobile internetwork is deployed on a large scale. Second, a number of groups (for instance the Mobile-IP working group of the IETF) are studying the problem by itself and it is not our intention to duplicate those efforts.

This attitude towards mobility is consistent with Nimrod's general philosophy of flexibility, adaptability and incremental change.

While a mobility solution is not part of the "core" Nimrod architecture, Nimrod does require that the solution have certain characteristics. It is the purpose of this document to discuss some of these requirements and evaluate approaches towards meeting them.

We begin by identifying the precise nature of the functionality needed to accommodate mobile entities (section 2). Following that, we discuss the effects of mobility on Nimrod (section 3). Next, we classify current and possible approaches to a solution for mobility (section 4) and finally (in section 5) we describe how mobility can be implemented using the IETF's Mobile-IP protocol.

This document uses many terms and concepts from the Nimrod Architecture document [CCS96] and some terms and concepts (in section 5) from the Nimrod Functionality document [RS96]. Much of the discussion assumes that you have read at least the Nimrod Architecture document [CCS96].

2 Mobility : A Modular Perspective

Nimrod has a basic feature that helps accommodate mobility in a graceful and natural manner, namely, the separation of the endpoint naming space from the locator space. The Nimrod architecture [CCS96] associates an endpoint with a globally unique endpoint identifier (EID) and an endpoint label (EL). The location of the endpoint within the Internetwork topology is given by its locator. When an endpoint moves, its EID and EL remain the same, but its locator might change. Nimrod can route a packet to the endpoint after the move, provided it is able to obtain its new locator.

Thus, providing a solution to mobility in the context of Nimrod may be perceived as one of maintaining a dynamic association between the endpoints and the locators. Extending this viewpoint further, one can think of mobility-capable Nimrod as essentially consisting of two "modules": the Nimrod routing module and the dynamic association module (DAM). The DAM is an abstraction, embodying the functionality pertinent to maintaining the dynamic association. This is a valuable paradigm because it facilitates the comparison of various mobility schemes from a common viewpoint. Our discussion will be structured based on the DAM abstraction and will be in two parts, the themes of which are :

- o What constitutes mobility for the DAM and Nimrod? Is the realization of mobility as a mobility "module" that interacts with Nimrod viable? What then are the interactions between Nimrod and such a module? These points will be discussed in section 3.
- o What are some of the approaches one can take in engineering the DAM functionality? We classify some approaches and compare them in section 4.

A word of caution: the DAM should not be thought of as something equivalent to the current day Domain Name Service (DNS) - the DAM is a more general concept than that. For instance, consider a mobility solution for Nimrod similar to the scheme described in [Sim94]. Very roughly, this approach is as follows: Every endpoint is associated with a "home" locator. If the endpoint moves, it tells a "home representative" about its new locator. Packets destined for the endpoint sent to the old locator are picked up by the home representative and sent to the new locator. In this scheme, the DAM embodies the functionalities implemented by all of the home representatives in regard to tracking the mobile hosts. The point is that the association maintenance, while required in some form or other, may not be an explicitly distinct part, but implicit in the way mobility is handled.

Thus, the DAM is merely an abstraction useful to our discussion and should not be construed as dictating a design.

In summary, we view the Nimrod architecture as carrying a functional "stub" for mobility, the details of the stub being deferred for later. The stub will be elaborated when a solution that meets the requirements of Nimrod becomes available (for instance from the IETF Mobile-IP research). We do not, however, preclude the modification of any such solutions to meet the Nimrod requirements or preclude the development of an independent solution within Nimrod.

3 Effects of Mobility

One consequence of mobility is the change in the locator of an endpoint. However, not all instances of mobility result in a locator change (for instance, there is no locator change if a host moves within a LAN) and a change in the locator is not the only possible effect of mobility. Mobility might also cause a change in the topology map. This typically happens when entire networks move (e.g., an organization relocates, a wireless network in a train or plane moves between cells, etc.). If the network is a Nimrod network, we might have a change in the connectivity of the node representing the network and hence a change in the map.

In this section, we consider the effects of mobility on the two "modules" identified above: Nimrod, which provides routing to a locator, and a hypothetical instantiation of the DAM, which provides a dynamic endpoint-locator association, for use by Nimrod. We consider four scenarios based on whether or not the topology and an endpoint's locator changes and comment on the effect of the scenarios on Nimrod and the DAM.

Scenario 1. Neither the locator nor the topology changes. This is the trivial case and affects neither the DAM nor Nimrod. An example of this scenario is when a workstation is moved to a new interface on the same local area network (This is not true for all LANs, only those in which all interfaces are part of the same Nimrod node) or when mobility is handled transparently (by lower layers).

Scenario 2. The locator changes but the topology remains the same. This is the case when an endpoint moves from one node to another, thereby changing its locator. The DAM is affected in this case, since it has to note the new endpoint-locator association and indicate this to Nimrod if necessary. The effect on Nimrod is related to obtaining this change from the DAM. For instance, Nimrod may be informed of this change or ask for the association if and when it finds out that the mobile host cannot be reached.

Scenario 3. The locator does not change but the topology changes. One way this could happen is if a network node moves and changes its neighbors (topology change) but remains within the same enclosing node. The DAM is not affected because the endpoint-locator association has not changed. Nimrod is affected in the sense that the topology map would now have to be updated.

Scenario 4. Both the locator and the topology change. If a network node moves out of its enclosing node, we have a change both in the map and in the locators of the devices in the network. In this case, both Nimrod and the DAM are affected.

In scenarios 3 and 4, it may not be sufficient to simply let Nimrod handle the topological change using the update mechanisms described in [RS96]. These mechanisms are likely to be optimized for relatively slow changes.

Mobile wireless networks (in trains and cars for instance) are likely to produce more frequent changes in topology. Therefore, it might be necessary that topological updates caused by mobility be handled using additional mechanisms. For instance, one might send specific updates to appropriate node representatives, so that packets entering that node can be routed using the new topology. We observe that accommodating mobility of networks, especially the fast moving ones, might require a closer interaction between Nimrod and the DAM than required for endpoint mobility. It is beyond the scope of this document to specify the nature of this interaction; however, we note that a solution to mobility should handle the case when a network as a whole moves. Current trends [WJ92] indicate that such situations are likely to be common in future when wireless networks will be present in trains, airplanes, cars, ships, etc.

In summary, if we discount the movement of networks, i.e., assume no topology changes, it appears that the mobility solution can be kept fairly independent of Nimrod and in fact can be accommodated by an implementation of the DAM. However, to accommodate network mobility (scenarios 3 and 4), it might be necessary for Nimrod routing/routers to get involved with mobility.

Beyond the constraints imposed by the interaction with Nimrod, it is desirable that the mobility solution have some general features. By general, we mean that these are not Nimrod specific. However, their paramount importance in future applications makes them worth mentioning in this document. The desirable features are :

- o Support of both off-line and on-line mobility. Off-line mobility (or portability) refers to the situation in which a session is torn down during the move, while on-line mobility refers to the situation in which the session stays up during the move. While currently much of the mobility is off-line, trends indicate that a large part of mobility in the future is likely to be on-line. A solution that only supports off-line mobility would probably have limited applications in future.

- o Scalability. One of the primary goals of Nimrod is scalability, and it would be contrary to our design goals if the mobility solution does not scale. The Internet is rapidly growing and with the advent of Personal Communication Systems (PCS) [WJ92], the number and rapidity of mobile components in the Internet is also likely to increase. Thus, there are three directions in which scalability is important : size of the network, number of mobile entities and the frequency of movement of the mobile entities.

Note that for any given system with minimum response time (to a move) of α seconds, if the mobile entity changes attachment points faster than $1/\alpha$ changes per second, the system will fail to track the entity. Augmenting traditional location tracking mechanisms with special techniques such as predictive routing might be necessary in this case. Hooks in the mobility solution for such augmentation is a desirable feature.

- o Security. It is likely that in the future, there will be increased demand for secure communications. Apart from the non-mobility specific security mechanisms, the solution should address the following :

- Authentication. The information sent by a mobile host about its location should be authenticated to prevent impersonation. Additionally, there should be mechanisms to decide if a mobile user who wishes to join a network has the privileges to do so or not.
- Denial of service. The schemes employed for handling mobility in general could be a drain on the resources if not controlled carefully. Specifically, the resource intensive portions of the protocol should be guarded so that inappropriate use of them does not cause excessive load on the network.

4 Approaches

As discussed in section 2, the problem of mobility in the context of Nimrod may be viewed as one of maintaining a dynamic association (DAM) and communicating this association and changes therein to Nimrod. Approaches to mobility may be classified based on how different aspects of the DAM are addressed.

Our classification identifies two aspects to the mobility solution :

1. How and where to maintain the dynamic association between endpoints and locators? This may be perceived as a problem of database maintenance. The database may be maintained in a centralized fashion, wherein a single entity maintains the association and updates are sent to it by the mobile host or in a distributed fashion, wherein there are a number of entities that store the associations.

A (distributed) database that stores the endpoint-locator mapping is required by Nimrod even in the absence of mobility. If this service can accommodate dynamic update and retrieval requests at the rate produced by mobility, this service is a candidate for a solution. However, we note that the availability of such a system should not be a requirement for the mobility solution.

2. Where to do the remapping between the endpoint and locator, in case of a change in association? By remapping, we mean associate a new locator with the endpoint. Some candidates are : the source, the "home" location of the host that has moved and any router (say, between the source and the destination) in the network.

Many of the existing approaches and perhaps some new approaches to the problem of mobile internetworking may be seen to be instantiations of a combination of a dynamic association method and a remapping method. We

(Re-mapping location)				
v				

		Source	Home	Routers

(Assoc.	Centralized	A1	X	X
maint)->	-----			
	Distributed	X	A2	A3

Table 1 : Classification of approaches based on how the association is maintained and where the remapping is done.

consider some combinations as illustrated in Table 1. We discuss three combinations (marked A1 - A3 in the table) and examine their advantages and disadvantages in the context of our requirements. The other combinations (marked X in the table) are possible, but do not represent a substantially different class of solutions from the ones discussed and hence are not considered here.

Note that this is but one classification of mobility schemes and that the remapping and endpoint-locator maintenance strategies mentioned in the table are not exhaustive. The main intention is to help understand better the kinds of approaches that would be most suitable for Nimrod.

In the following, we use the term source to refer to the endpoint that is attempting to communicate with or sending packets to a mobile endpoint. The source could be static or mobile. We use the term mobile destination to refer to the endpoint that is the intended destination of the source's packets.

- A1. In this approach, all endpoint-locator mappings are maintained at a centralized location. The source queries the database to get the locator of the mobile destination. Alternatively, the database can send updates to the source when the mobile destination moves. The main advantage of this scheme is its simplicity. Also, no modification to routers is required, and the route from the source to a mobile destination is direct.

The main disadvantage of this scheme is vulnerability - if the centralized location goes down, all information is lost. While this scheme may be sufficient for small networks with low mobility, it does not scale adequately to be a long term solution for Nimrod.

- A2. This approach uses distributed association maintenance with remapping done at the home. This is the approach that is being used by the Mobile-IP working group of the IETF for the draft proposal and by the Cellular Digital Packet Data (CDPD) consortium. In this approach, every mobile endpoint is associated with a "home" and a "home representative" keeps track of the location of every mobile endpoint associated with it. A protocol between a mobile endpoint and the home representative is used to keep the information up-to-date. The source sends the packet using the home locator of the mobile destination, and the home representative forwards the packet to the mobile destination. The advantage of this scheme is that it is fairly simple and does not involve either the source or the routers in the network. Furthermore, the mobile destination can keep its location secret (known only to the home representative) - this is likely to be a

desirable feature for mobile hosts in some applications. Finally, most of the control information is confined to the node containing the home representative and the mobile host and this is a plus for scalability. The main disadvantage is a problem often referred to as triangular routing. That is, the packets have to go from the source to the home representative first before going to the mobile destination. This is especially inefficient if, for instance, both the source and mobile destination are in, say, England and the home representative is in, say, Australia. Also, there is still some vulnerability, since if the home representative becomes unreachable, the location of all of the mobile hosts it tracks is lost and communication from most sources to the mobile host is cut-off. It is also not clear how well this scheme will scale to mobile internetworks of the future.

Nevertheless, we feel that this approach or a modification thereof might be a viable first-cut mobility solution for Nimrod.

- A3. In each of the previous cases, the routers in the network were not involved in tracking the location of the mobile host. In this approach, state is maintained in the routers. An example is the approach proposed in [TYT91] wherein the packets sent by a mobile host are snooped and state is created. The packets contain the mobile host's home location and its new location. This mapping is maintained at some routers in the network. When a packet intended for the mobile host addressed to its home location enters such a router, a translation is made and the packet is redirected to the new location.

An alternate mechanism is to maintain the mapping in all of the border routers (e.g., forwarding agents) in the node within which the movement took place. A packet from outside the node intended for a destination within the node would typically enter the node through one of the border routers. Using the mapping, the border router could figure out the most recent locator of the mobile destination and send the packet directly to that locator. If most of the movements are within low level nodes, this would scale to large numbers of movements. Furthermore, the packet takes an optimal path (or as optimal as one can get with a hierarchical network) to the new location within the time it takes for the node representative to get the new information, which is typically quite small for low-level nodes.

The main disadvantage of this scheme is that routers have to be involved. However, future requirements in regard to scalability and response time might necessitate such an approach. Furthermore, this solution has closer ties with Nimrod routing and is better suited to handling scenarios 3 and 4 where the topology changes as a result of mobility.

All of these approaches seem potentially capable of handling scenarios 1 and 2 of the previous section. Scenarios 3 and 4 are best handled by an approach similar to A3. However, approaches like A3 are more complex and involve more Nimrod entities (e.g., routers) than may be desirable.

We have tried to bring out the various issues governing mobility in Nimrod. In the final analysis, the tradeoffs between the various options will have to be examined vis-a-vis our particular requirements (for instance, the need to support network mobility) in adopting a solution. It is likely that general requirements such as scalability and security will also influence the direction of the approach to mobility in Nimrod.

5 A Solution using IETF Mobile-IP

The Mobile-IP Working Group of the IETF is in the process of standardizing a protocol that allows an IPv4 capable network to support mobile hosts. In this section, we outline how mobility can be implemented within Nimrod using the same mechanism and indeed, the same protocol headers defined in [Sim94]. Not all functionality described in [Sim94] are covered - only those that form the "core" of mobility support.

In order to follow this section, the reader is required to have some familiarity with the IETF Mobile-IP protocol (see [Sim94]).

5.1 Overview

The general scheme employed by the IETF Mobile-IP protocol is as follows. A Mobile Host (MH) has a predefined Home Agent (HA) that is responsible for the MH's whereabouts. Typically, the MH spends most of its time in the network containing the HA. Let us assume that the MH wanders to a new network. The MH then contacts a Foreign Agent (FA) at the new network that will act on its behalf and sends a registration request to the HA via the FA. This serves the purpose of informing the HA of the MH's new whereabouts and also is a means of verification of the MH's authenticity. It also contains the address of the FA as the new Care-of-Address. A correspondent host (CH) wishing to send a message to the MH uses the MH's Home IP address. This message is captured by the HA and tunnelled using encapsulation

to the FA whereupon the FA decapsulates and sends the original message to the MH.

If the MH can get itself a new transient address then there is no need for a Foreign Agent. The transient address will be sent as the Care-of-Address. The packets will be tunnelled directly to this address by the Home Agent. Note, however, that some networks may require that a mobile host go through a Foreign Agent.

A fundamental difference between IP and Nimrod is that in the latter an endpoint has both a (topologically sensitive) locator and a (topologically insensitive) endpoint-id (EID). In IP, the IP address serves as both the EID and the locator. Thus, it should be possible to use the Mobile-IP protocol for providing mobility support in Nimrod by simply using the EID of the MH wherever its Home IP Address was being used and by appropriately using the EID and locator of the FA and HA in place of their IP addresses (An issue is the format and length compatibility between EIDs and IP addresses. For the discussion here, we assume that an EID can fit into an IP (v4 or v6) address given in Figure 1). We give below the details of the protocol fields and the actions taken by the MH, FA and HA to show that this is possible and that it is quite simple.

5.2 Protocol Details

There are two kinds of protocol headers relevant to our discussion - the Mobile-IP Protocol (MIPP headers) and the headers for data packets transported by Nimrod (NP headers). It is our intent that Nimrod use, as much as possible, the next generation IP (IPv6) header. The NP header contains as a subset fields that would eventually be present in the IPv6 header.

In the scheme given below, the MIPP header is enclosed within the NP packet (i.e., MIPP operates over NP). The details of the fields constituting the NP header are beyond the scope of this document. However, without venturing into bit lengths, etc., we identify below a few fields that are relevant to our discussion:

- o Source EID (S-EID) : The endpoint ID of the source entity originating the packet.
- o Destination EID (D-EID) : The endpoint ID of the destination.
- o Source locator (S-LOC) : Locator of the entity originating the packet.
- o Destination locator (D-LOC) : Locator of the destination.

The MIPP header fields are described in [Sim94].

In what follows, we describe the values that must be assigned to the relevant NP and MIPP fields in order for Nimrod to work with Mobile-IP. There are three phases we must consider : agent discovery, registration and forwarding [Sim94]. A pictorial summary of the control and data packets is given in Figure 1.

Agent Discovery: In this phase, the MH discovers the foreign agent, if any, that will act on its behalf. In MIPP, this is done using the ICMP Router Discovery messages.

When an MH attaches to a Nimrod network (node), foreign agent discovery is done as follows. We assume that a link-level connection is established between the MH and a node N belonging to the network. For instance, this node could be a wireless equipped base station that establishes a signalling channel for communication with the MH.

If the MH is itself a node then N and the MH execute an arc formation procedure between themselves as described in [RS96]. This results in a locator being assigned to the MH and to the arcs between N and MH.

If the MH is not a node but only an endpoint, then MH initiates locator acquisition procedure as described in [RS96]. This results in a locator being assigned to the MH.

The MH then sends a Foreign Agent Request message to N. This message contains, amongst other information, the EID and locator of the MH. If N is not itself the foreign agent, then we assume that it knows of and has the ability to reach a foreign agent.

The foreign agent (FA) notes the EID of the MH in its Visitor List and sends a Foreign Agent Reply to the MH. This contains the EID and the locator of the FA and will be used as the "Care-of-Address" (COA) of the MH for a prespecified period.

Registration: In the registration phase, information is exchanged between the MH and the Home Agent (HA). The HA could, for instance, be the endpoint representative of the endpoint in its home location. The registration procedure is used to create a mobility binding which the HA uses to forward data packets intended for the MH. Another purpose of registration is to verify the authenticity of the MH.

There are four parts to the registration. We describe the values assigned to the relevant fields. Recall that there are two headers we must create - the Nimrod Protocol (NP) header and the Mobile-IP Protocol (MIPP) header. The NP fields are as described above and the MIPP fields are as in [Sim94]. The fields mh-eid(mh-loc), fa-

eid(fa-loc), ha-eid(ha-loc) are used to refer to the EID (locator) of the mobile host, foreign agent and home agent respectively.

1. The MH sends a Registration Request to the prospective Foreign Agent to begin the registration process.

- o NP fields : S-EID = mh-eid; D-EID = fa-eid; S-LOC = mh-loc ; D-LOC = fa-loc.

- o MIPP fields : Home Agent = ha-eid; Home Address = mh-eid; Care-of-Address = fa-eid.

Note that the mh-loc is known to the MH by virtue of the locator acquisition (see paragraph on "Agent Discovery") and that the fa-eid is known to the MH from the Foreign Agent Reply. The FA caches the mh-eid for future reference.

2. The Foreign Agent relays the request by sending a Registration Request to the Home Agent, to ask the Home Agent to provide the requested service.

- o NP fields : S-EID = fa-eid; D-EID = ha-eid; S-LOC = fa-loc; D-LOC = ha-loc.

- o MIPP fields : Same as in (copied from) (1) above.

The HA caches the (Home Address, Care-of-Address) as a mobility binding. Optionally, for efficiency, it may also cache fa-loc.

3. The Home Agent sends a Registration Reply to the Foreign Agent to grant or deny service.

- o NP fields : S-EID = ha-eid; D-EID = fa-eid; S-LOC = ha-loc; D-LOC = fa-loc.

- o MIPP fields : Home Address = mh-eid; code = as in [Sim94].

The S-EID and D-EID fields are taken from the Request and swapped, as are the S-LOC and D-LOC fields. The Home Address in the MIPP is the same as the Home Address in the Request. The code indicates whether or not permission was granted by the Home Agent.

4. The Foreign Agent sends a copy of the Registration Reply to the MH to inform it of the disposition of its request.

- o NP fields : S-EID = fa-eid; D-EID = mh-eid; S-LOC = fa-loc; D-LOC = mh-loc.

o MIPP fields : Same as (copied from) (3) above.

At this point the MH is registered with the HA (provided the registration request is approved by the HA) and packets can be forwarded to the MH.




Figure 1 : The control and data packets for mobility handling using the Mobile-IP protocol. The packets bordered as # denote data packets and those bordered * denote control packets. Only the crucial information conveyed in each message is shown (i.e., locators and EIDs in packet headers are not shown). The associations maintained at HA and FA are shown.

Forwarding Data: We describe the manner in which a packet from the correspondent host (CH) intended for the MH is encapsulated and forwarded by the HA.

- o At HA : Suppose that a packet P intended for MH arrives at HA. For instance, P first comes to the router for the local network and the router finds that MH is unreachable. The router then forwards P to the HA for possible redirection.

The HA extracts the destination EID from the NP header for P. If no match is found in its mobility binding, then the MH is deemed as unreachable. If a match is found, the corresponding fa-eid is extracted. A new header is prepended to P. For this header, S-EID = ha-eid, D-EID = fa-eid, S-LOC = ha-loc and D-LOC = fa-loc. The fa-loc may be obtained from the Association Database [CCS96]. Alternatively, if it was cached in (2) above, it could be obtained from the cache.

- o At FA: By looking at the next header field in the Nimrod Protocol packet header, the FA knows that the packet is an encapsulated one. It removes the wrapping and looks at the EID in P. If that EID is found in the Visitor List then the FA knows the locator of the MH and can deliver the packet to the MH. Otherwise, the packet is discarded and an error message is returned to HA.

Other Issues: We have not addressed a number of issues such as deregistration, authentication, etc. The mobility specific portion of authentication can be adapted from the specification in [Sim94]; deregistration can be done in a manner similar to registration.

The protocol in [Sim94] describes a registration scheme without the involvement of the Foreign Agent. This is done when the MH obtains a transient IP address using some link-level protocol (e.g. PPP). A similar scheme can be given in the context of Nimrod. In this case, the MH obtains its locator (typically inherited from the node to which it attaches) and sends this locator as its Care-of-Address in the Registration Request. The HA, while forwarding, uses this as the locator in the outer NP header and thus the encapsulated packet is delivered directly to the MH which then decapsulates it. No Foreign Agent Discovery is needed. Apart from this, the fields used are as described for the scheme with the FA.

We note however that many networks may require that the registration be through a Foreign Agent, for purposes of security, billing etc.

6 Security Considerations

The registration protocol between a mobile host and the network (for instance, in the mobile-ip protocol, the MH and the HA) contains security mechanisms to validate access, prevent impersonation etc.

This document is not a protocol specification and therefore does not contain a description of security mechanisms for Nimrod.

7 Summary

- o Nimrod permits physical devices to be mobile, but does not specify a particular solution for routing in the face of mobility.
- o The fact that the endpoint naming (EID) space and the locator space are separated in Nimrod helps in accommodating mobility in a graceful and natural manner. Mobility may be perceived, essentially, as dynamism in the endpoint - locator association.
- o Nimrod allows two kinds of mobility:
 - Endpoint mobility. For example, when a host in a network moves. This might cause a change in the locator associated with the host, but does not cause a change in the topology map for Nimrod.
 - Network mobility. For example, when a router or an entire network moves. This might cause a change in the topology in addition to the locator.
- o Endpoint mobility may be handled by maintaining a dynamic association between endpoints and locators. However, network mobility requires addressing the topology change problem as well.
- o Apart from the ability to handle network mobility, it is desirable that the mobility solution be scalable to large networks and large numbers of mobile devices and provide security mechanisms.
- o There are a number of existing and emerging solutions to mobility. In particular, adaptation of solutions developed by the IETF is a first cut possibility for Nimrod. As the description given in section 5 shows, it is relatively easy to implement the scheme being designed by the Mobile-IP working group in the context of Nimrod.

8 Acknowledgements

We thank Isidro Castineyra (BBN), Charles Lynn (BBN), Martha Steenstrup (BBN) and other members of the Nimrod Working Group for their comments and suggestions on this memo.

9 Author's Address

Ram Ramanathan
BBN Systems and Technologies
10 Moulton Street
Cambridge, MA 02138

Phone : (617) 873-2736
Email : ramanath@bbn.com

References

- [CCS96] Castineyra, I., Chiappa, N., and M. Steenstrup, "The Nimrod Routing Architecture", RFC 1992, August 1996.
- [RS96] Ramanathan, S., and M. Steenstrup. Nimrod functional and protocol specifications, Work in Progress.
- [Sim94] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.
- [TYT91] F. Teraoka, Y. Yokote, and M. Tokoro. A network architecture providing host migration transparency. In Proceedings of ACM SIGCOMM, 1991.
- [WJ92] K. A. Wimmer and J. B. Jones. Global development of pcs. IEEE Communications Magazine, pages 22--27, Jun 1992.

