This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the ICC 2008 workshop proceedings.

Using NEMO to Extend the Functionality of MANETs

Ben McCarthy, Dr. Christopher Edwards, Dr. Martin Dunmore
Computing Department, InfoLab 21
South Drive, Lancaster University
Lancaster, LA1 4WA
Lancashire, UK
Email: [b.mccarthy, ce, dunmore]@comp.lancs.ac.uk

***Abstract*—Network mobility is an established topic of research which has the potential capability to support many valuable scenarios. Specifically, the ability to support the mobility of entire networks of IP enabled devices that are oblivious to the changing network conditions beneath them is particularly useful to numerous scenarios such as vehicle based networks and Personal Area Networks (PAN). In this paper, we present an efficient and scalable approach that allows mobile networks to intercommunicate and be reachable via the Internet both directly (via their own Internet connection) and indirectly (via another mobile network with an Internet connection). The implemented approach is based on the concept of combining the beneficial features of Mobile Ad-hoc Networking (MANET) protocols and the Network Mobility Basic Support (NEMO BS) protocol to develop what is known as a MANEMO solution. In the paper we highlight the key performance characteristics of our protocol through analysis of our implementation in a testbed environment.**

I. Introduction

MANEMO is a relatively new and immature concept. The term MANEMO itself can be loosely defined as describing techniques which combine the properties of Mobile Ad-Hoc Networks (MANETs) and Network Mobility (NEMO). From its early inception, MANEMO has often been considered from the perspective of Nested NEMO networks. More specifically, MANEMO is often considered as a Route Optimisation solution to the pinball routing problem [1] that occurs in Nested NEMO networks if the NEMO Basic Support protocol (NEMO BS) [2] is used to support them. Considering MANEMO from this NEMO based perspective can be referred to as a NEMO-Centric MANEMO (NCM) approach, since we are considering scenarios that fundamentally revolve around the notion of typical NEMO connectivity (i.e. periodically moving between heterogeneous access networks, whilst predominantly communicating with nodes on the Internet). When utilised in this manner, the purpose of the MANET protocol that is run between the MRs is to optimise the path that packets traverse before they reach the Internet.

This approach to solving the Nested NEMO routing problem by combining MANET and NEMO is however just one facet off the overall MANEMO domain. The integration of MANET and NEMO technologies is a mutually beneficial process for both problem domains. NEMO-Centric MANEMO (NCM) refers to how the NEMO problem space can benefit from the introduction of MANET technologies. In this paper we introduce how the MANET problem space can benefit from the introduction of NEMO concepts, creating a so called MANET-Centric MANEMO solution (MCM) [3]. The rest of this paper is presented as follows: In Section II-A we introduce the concept of Network Mobility (NEMO), highlight its purpose and the complex Nested NEMO problem. In Section II-B we provide a brief overview of MANET technologies and discuss the considerations that influenced our protocol selection for our implementation. In Section III we outline the design of our Unified MANEMO Architecture (UMA) protocol and discuss the problems we encountered when implementing this protocol. In Section IV we present our preliminary lab based testing results. Finally in Section V we provide an analysis of the general performance capabilities of our protocol and conclusions regarding the feasibility of this approach.

II. Overview

A. Network Mobility (NEMO)

NEtwork MObility (NEMO) and, more specifically, the NEMO Basic Support Protocol (NEMO BS) offers a mobility solution based on the concepts used by Mobile IPv6 (MIPv6) [4] that is targeted at supporting entire networks of IP devices as opposed to just single hosts. Using NEMO BS, mobile networks can be provided with constant, uninterrupted Internet connectivity without any of the devices that are attached to the mobile network needing to be aware of their own mobility. In the NEMO BS model, the mobile entity is considered to be a Mobile Router (MR) that manages the mobility of the entire network over its Egress interface (i.e., its connection to the Internet) and presents its Ingress interface to IP devices as a normal, static IP connection. This is made possible through the use of a Home Agent (HA) situated on the Home Network of the MR; in the case of NEMO BS, the HA forwards packets destined for an entire prefix of addresses that are attached to the MR, known as the Mobile Network Prefix (MNP). This powerful solution supports many existing, real-life use case scenarios, such as vehicle based networks, Personal Area Networks (PAN) and Access networks on public transport where multiple IP devices wish to gain access to the Internet but cannot be expected to support any additional software or protocol stack.

B. Mobile Ad-Hoc Networks (MANETs)

Mobile Ad-Hoc Networking protocols (MANET) support mobility (Mobile Host and Mobile Router) by deploying optimised routing protocols specifically designed to operate between mobile devices to predominantly support networking scenarios which have no prior infrastructure. MANET routing

978-1-4244-2052-0/08/$25.00 ©2008 IEEE

protocols can be classified as one of two main styles of protocol, Proactive or Reactive. Proactive MANET protocols such the Optimised Link State Routing protocol (OLSR) [5] periodically disseminate routing information between all of the mobile nodes in a MANET in the same way that a traditional routing protocol would, only in an optimised fashion. For example, OLSR assigns a subset of the mobile nodes in a MANET with the task of operating as Multi Point Relays (MPR); routing information is then disseminated only by this subset of nodes in order to reduce the amount of routing protocol overhead experienced in the MANET. On the other hand, Reactive MANET protocols such as Ad-Hoc On-Demand Distance Vector routing protocol (AODV) [6] do not disseminate routing information; instead mobile nodes utilising this kind of routing protocol only solicit for routing information as and when they need it.

C. Rationale

Mobile Ad-Hoc Networking is a maturing area of research and development that has seen many innovations; however the inherent complexity of the MANET problem domain has ensured that there still remains many challenges to overcome. In this paper we show how integrating the Home-Agent based approach of NEMO with a MANET protocol can introduce many benefits to MANET scenarios, such as Global reachability of MANET nodes, improved AAA and support for multihoming. MANET protocols were originally designed to support inter-communication between the nodes connected to a MANET, but over time they have been augmented to support communication with any node on the Internet from within the MANET via Internet Gateways [7]. Whilst different MANET protocols offer varying techniques for supporting this type of communication with nodes external to a MANET, there still remain many limitations. For instance, a Mobile Ad-Hoc Network cannot leak its routes directly into the Internet so therefore changing its point of attachment can be problematic. If a MANET of nodes wishes to roam across heterogeneous networks, each Access Router (AR) they connect to the Internet via must be equipped to support their connection. If a MANET of nodes roams onto a new access network and (in the most straightforward case) is connected to that access network via a single point of attachment (Internet Gateway) then something must be done in order to ensure packets can be routed out of and back in to the MANET. Typically this could involve one of two approaches. Either the MANET could use private addresses for intercommunication between the MANET nodes and then use a Network Address Translation (NAT) technique to communicate with nodes located externally in the Internet. This approach would rely on the AR performing NAT functionality on behalf of each of the MANET nodes and would also prevent the MANET node being reachable unless it specifically initiated a flow. The other approach would be to have each node configure an address that is topologically correct in relation to the access network that the gateway node has connected to. This approach would require every node to configure a new address, every time the gateway node changed its location. In addition, in this type of scenario a MANET of nodes will not be able to benefit from multiple simultaneous connections to different access networks (multiple Internet Gateways). This is because each node within the MANET will have initially configured a topologically acceptable IP address for use in its new location and therefore that address will only be topologically correct with respect to the access network that it was configured from.

Both of these approaches require the AR of a visited network to be 'MANET' aware (i.e. they must be augmented to support the attachment of MANETs). In addition, any change in the point of attachment by a Gateway node will result in the loss of sessions for all nodes within the MANET. Consider a scenario whereby a single Gateway node is providing access to the Internet for a cluster of 10 MANET nodes in total. If the Gateway node initially has a connection via a publicly available WiFi network, each of the MANET nodes will either configure a topologically correct address based on the ARs network prefix or a private address that will be registered in a NAT table in the AR. If the Gateway node then roamed away from the WiFi network and established a UMTS cellular connection, all of the MANET nodes would be required to carry out this initial address configuration process again, as they would every time the Gateway node changed its location. A change in the overall attachment point of the MANET to the Internet would also result in a change in the global IP address used by the MANET nodes to communicate externally with other nodes in the Internet. This change in addresses would subsequently break any TCP sessions that were in place; and therefore to prevent this an addition protocol to maintain session continuity would be needed.

In this paper we present a mobile networking approach that has been designed to address these problems related to MANETs and global communication over the Internet. In addition the protocol we propose is primarily focused on supporting mobile networks of devices that are unaware of their own mobility. This is achieved by utilising a Mobile Router (MR) which in turn performs all IP mobility related functionality on the behalf of the attached devices in order to support the same key benefits as NEMO BS. However it is important to point out that the protocol implementation outlined in this paper also inherently supports the more simplistic case of single host mobility as well.

III. Approach and Implementation

The overall aim of our MANET-Centric MANEMO protocol is to ensure that mobile networks are consistently reachable on the Internet whenever they have access to an Internet connection irrespective of whether it is a direct connection or established via other mobile networks. In addition, the protocol has also been designed to introduce techniques for performing security and AAA for packets transmitted into the Internet from ad-hoc networks. The fundamental approach we have employed is based around combining the functionality of the Optimized Link State Routing (OLSR) protocol with a HA based approach to location registration (like the NEMO

Basic Support protocol uses); we call our approach the Unified MANEMO Architecture (UMA). This approach requires every UMA enabled Mobile Router (UMA-MR) to setup a MR-HA bi-directional tunnel whenever it has a direct connection to the Internet. This direct tunnel link between the UMA-MR and its respective UMA enabled HA (UMA-HA) ensures that the UMA-MR and its connected devices are permanently reachable (whenever the tunnel is present) via an address prefix that the HA advertises on the UMA-MRs behalf. In addition all UMA-MRs maintain a MANET interface which they are able to use to connect to other UMA-MRs (and standard MANET nodes if required). Our protocol then utilises this MR-HA tunnel to provide all other UMA-MRs that are unable to form their own direct Internet connections with a means to reach the Internet. Use of the MR-HA tunnel in this way ensures that the access network need not be aware of the protocol in any way. Finally, by using a technique of HA inter cooperation to ensure any UMA-MR can connect to the Internet via any other collection of UMA-MRs, we have been able to produce a simple, efficient solution to network mobility that pushes much of the associated complexity into the wired network (i.e. in the HAs).

With UMA, the UMA-MRs are only required to build a topologically correct Care-of-Address (CoA) if they establish a direct connection to the Internet via an access network (i.e. if they are a Gateway-MR). If their connection is indirectly established via other UMA-MRs, the newly attached UMA-MR can create a binding with whichever address it chooses to propagate in the MANET of UMA-MRs (i.e. its Home Address). In our solution we utilise the OLSR Host and Network Association (HNA) messages to provide the trigger for performing the BU process rather than movement detection based on IPv6 Neighbor Discovery messages. HNA messages are used within OLSR in order for a node to advertise its ability to reach other networks. The networks advertised by HNA messages can be both physically collocated networks (i.e. MNPs in the case of an MR) and also temporarily reachable networks (i.e. the Internet when a connection is available). When a UMA-MR successfully attains a direct connection to the Internet it becomes known as a Gateway-MR. The Gateway-MR then advertises its ability to reach the Internet in any subsequent HNA messages it sends to the other UMA-MRs that are connected to it (from here on we refer to any cluster of UMA enabled MRs such as this as a UMA-Stub). Therefore when a MR receives a HNA message containing this default route it knows it can begin its BU process with its HA via this Gateway-MR. In addition to advertising Internet reachability, the HNA record also carries the address of the Gateway-MRs HA.

In scenarios supported by this approach, one fundamental consideration is the makeup of the UMA-Stub. Mobile Routers within the UMA-Stub can either all belong to the same Home Network or can originate from a number of different Home Networks. To give an example of how these scenarios can arise, consider an emergency situation such as the breakout of a fire in a building. Depending on the nature of the emergency the fire brigade are likely to be the first of the emergency services to attend the scene. In the period in which only the members of the fire brigade are present, all MRs carried by the firefighters will have originated from same Home Network (i.e. the fire brigade HQ) and therefore will all be registered with the same HA. If it is possible that people have been or may be injured because of the fire then the scene will also be attended by paramedics. When the paramedics arrive they would undoubtedly benefit from the ability to communicate on the same network as the fire brigade and therefore incorporating their MRs should also be supported. However since their MRs will originate from a different Home Network (the hospital network) and therefore be registered with a different HA, the UMA protocol must behave differently to support these new additions to the UMA-Stub.

In protocol terms, if a UMA-MR can only obtain an indirect connection to the Internet via a Gateway-MR, then the principal concern is whether that UMA-MR is registered with the same HA as the Gateway-MR or a different one. When the MR receives a HNA record from a Gateway-MR advertising Internet reachability, the MR will first extract the address of the Gateway-MRs HA and compare this address with the address of its own HA to determine which type of Binding Update (BU) to perform. We refer to the case where the MR sending the BU is registered with the same HA as the Gateway-MR as the Aggregated Scenario. In this situation (illustrated here in Figure 1) the MR will first mark its BU message with the appropriate flags (detailed later) and send the message directly to the Gateway-MRs HA. Receiving a BU message with both these flags set signals to the HA that an MR is trying to indirectly bind to the HA via one of its existing MR-HA Tunnel connections. Subsequently, after ensuring that the MR is registered with itself, the HA will record the MR-HA tunnel that the UMA BU request was received via and then install routes to the newly binded MR and its MNPs via that Tunnel. Once the correct routing entries are in place, the HA returns a Binding Acknowledgment (BA) to the newly registered UMA-MR to signal whether the bind was a success or failure. Upon receiving a successful BA, the UMA-MR will then also install the appropriate routes in its routing table and communication with Correspondent Nodes (CN) in the Internet can continue. In this scenario, our approach ensures that the overhead imposed by tunneling is kept to a minimum. Any UMA-MRs that are connected to the Internet indirectly will not impose any further tunneling, therefore beyond the Gateway-MR packets can be transmitted without any additional tunnel headers.

Supplementary to this Aggregated Scenario many real world examples exist where a UMA-MR attaching to an UMA-Stub will not be registered with the same HA as the Gateway-MR. We refer to this as the Non-Aggregated Scenario (illustrated here in Figure 2). In this situation the Gateway-MR's HA will behave as a Proxy-HA, forming an indirect link between any legitimate MRs and their actual HA. In this case the MR again determines which HA to initially contact via the information carried in the Gateway-MRs HNA messages. However when

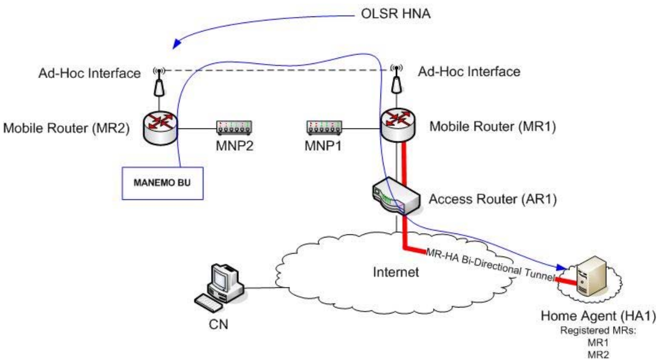

Fig. 1. UMA BU - Aggregated Scenario

the MR recognises that the advertised HA is not located on its Home Network it switches its operational mode to perform a UMA Proxy-Bind. To perform a Proxy-Bind the MR sends its BU to the Proxy-HA (The Gateway-MR's HA) and also inserts the address of its own HA (the Target-HA) into the BU message. When the Gateway-MR's HA receives this BU it assumes the role of a Proxy-HA and begins the Home Agent to Home Agent (HA-HA) binding process. To do this the Proxy-HA extracts the address of the Target-HA from the BU message and sends a separate HA-HA BU to the Target-HA requesting simultaneously, the setup of HA-HA bidirectional tunnel to carry packets directly to the MR and a binding registration for the MR itself. The Target-HA then registers the CoA of the MR as the MANET address and sets up a route to that address as being reachable via the newly created HA-HA tunnel.

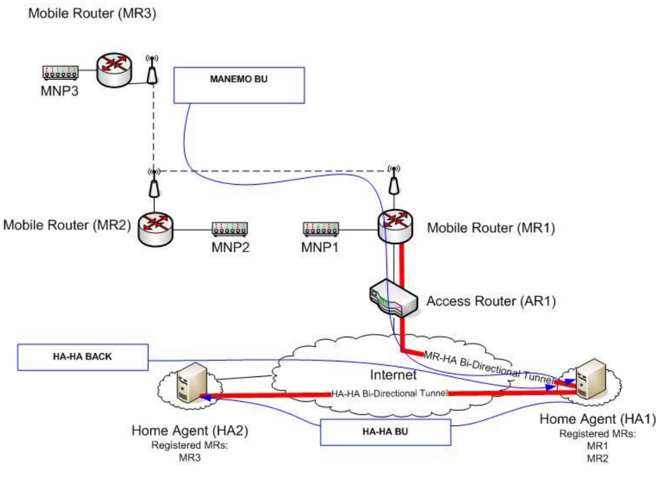

Fig. 2. UMA BU - Non-Aggregated Scenario

Figure 3 illustrates the Binding Update (BU) message format that the UMA protocol utilises. The layout of the UMA messages has intentionally been designed to have a similar overall format as those used by NEMO BS. In addition to the NEMO related fields the UMA BU message introduces 2 new flags, the (U) flag and the (P) flag. The U flag (Unified MANEMO Architecture flag) is used to indicate to HAs (both Proxy and Target) that the BU is a request for a UMA binding and therefore the MR does not have a direct attachment to the Internet. The (P) flag (Proxy Registration flag) is always used in conjunction with the U flag by both MRs and Proxy-HAs and incorporates the use of the (H) flag (Home Registration flag) from MIPv6 in order to differentiate its use. If a HA receives a BU with just the (U) and (H) flag set, this signifies that the MR is directly registered with that HA but is performing a UMA bind via an existing Gateway-MR's tunnel connection. Whereas if a HA receives a BU with just the (U) and (P) flags set (H flag not set), then this signifies that the MR is not directly registered with the HA and is therefore requesting that the HA performs a proxy bind with its Target-HA on behalf of the MR. Finally, if a HA receives a BU with all 3 flags set (U), (P) and (H) this signifies that the BU message is from a Proxy-HA that is requesting to establish a HA-HA tunnel on behalf of the MRs HoA that is contained within the BU message. As with the BU message, the UMA Binding Acknowledgment (BA) message is designed to have the same fundamental format as the NEMO BA message however it introduces new status numbers to support the new proxy binding functionality. These new status numbers are set by the Target-HA in response to receiving a Proxy-BU request from a Proxy-HA and they instruct the Proxy-HA as to whether the Proxy-BU was accepted or rejected (and if it was rejected, for what reason).

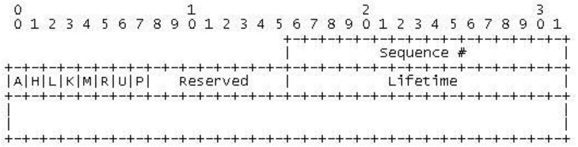

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |          Sequence #           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|A|H|L|K|M|R|U|P|   Reserved    |           Lifetime            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Fig. 3. UMA Binding Update Message Format

A. Access, Authentication and Accounting

The MANET-Centric MANEMO protocol outlined in this paper permits a two phase Access, Authentication and Accounting (AAA)model to Mobile Ad-Hoc Network scenarios which are notoriously difficult to provide meaningful AAA for. Firstly a newly connecting MR must gain access to the local wireless communication provided by the UMA-Stub. This initial access would most likely be protected by encryption techniques such as WEP and WPA. Once the MR has gained access to the wireless portion of the network it can begin communicating directly with other MRs in the UMA-Stub. Once the MR has established a connection, any further communication with nodes external to the UMA-Stub must pass a second phase of AAA which is performed by the Gateway-MR's HA. Whether the Gateway-MR's HA will act as a Proxy-HA or as the ultimate Target-HA, it will first decide whether to accept any incoming request for connection based on either a locally stored policy or a more dynamic process such as the response from a remote AAA server. If the connection cannot be authorised then the unauthorised UMA-MR cannot transmit any packets into the Internet, which prevents any illegitimate and potentially damaging traffic from originating within the UMA-Stub. If a connection is then accepted and established, the Gateway-MR's HA can then record all further throughput and potentially bill the appropriate customer once the session has completed. This functionality provides network administrators with a level of control that is often unachievable in ad-hoc networks.

IV. Testing and Results Analysis

We have developed an implementation of the protocol we outline in this paper on a Linux Platform (Ubuntu 7.10 distribution) using the 2.6.22 Linux kernel. At present we have carried out our preliminary tests on a testbed consisting of 4 Linux PCs. Each machine has multiple Ethernet network cards to provide wired local access to devices. All of the tests detailed in this section were performed between two laptop PCs that were directly connected to one of the MNP interfaces of the UMA-MRs using Ethernet cables. Also, both the HAs and the Gateway-MR were all attached to separate access networks, each interconnected by Cisco 3200 Series routers. Wired Ethernet links were used throughout our testing to reduce the possibility that the results were effected by external influences (as wireless links can be effected by interference). Using wired links in the testing phase results in throughput figures that are far greater than would be expected using wireless links, however what we are most interested in is the comparative performance of the protocol in each of the different scenario stages. One of the most important measurements when considering the performance of a mobility protocol is the handover times. In addition to handover times one of the other primary factors that we wish to analyse is the effect that this approach has on the overall end-to-end path once it is established. To analyse these two factors we developed a testing regime whereby for each appropriate mobility scenario the handover time taken to form the new network layout was first measured. Then once the handover had been performed and the connection was setup we analysed how the overall performance along the end-to-end path was effected for each new configuration we tested. For each of the tests we used ICMP Echo request/replies (pings) to determine the round trip times between the test machines and the iperf bandwidth measurement tool to determine the throughput. To determine the handover times we again used the Ping6 utility, in collaboration with the network packet analyser Wireshark. By setting the ping request interval to a high value (1 request every 0.001 seconds) and then analysing the time difference between the time of the first request not to receive a corresponding reply and the time the next reply was received.

A. Scenario 1 - NEMO

Firstly, to act as a baseline for our results we began by measuring the handover time and end-to-end performance of the UMA MR when it is directly connected to an access network. In this scenario the UMA-MR behaves exactly as a NEMO-MR except that it will also run the MANET protocol over its MANET interface as well.

- Handover: 6.65s
- End-to-end Throughput: 10.7Mbps
- End-to-End Latency: 2.52ms

B. Scenario 2 - UMA (Aggregated)

In the Aggregated scenario an MR moves and subsequently changes the Gateway-MR that it is connected to the Internet via. In this scenario, the Gateway-MR and the adjoining UMA-MR are both registered to the same HA.

- Handover: 4.2s
- End-to-end Throughput: 10.6Mbps
- End-to-End Latency: 2.85ms

C. Scenario 3 - UMA (Non-Aggregated)

Finally, in the Non-Aggregated scenario the same movement occurs as in the previous test, however the Gateway-MR is connected to the same HA that the BU is ultimately destined for. Therefore a Proxy-HA connection will be established.

- Handover: 5.4s
- End-to-end Throughput: 9.7Mbps
- End-to-End Latency: 6.3ms

V. Analysis and Conclusion

The overall outcome of the results from our initial testing phase have been quite positive. The results from scenario 1 represents the amount of time it takes to perform a normal NEMO Binding Update. From our results it is possible to see that a simple handover in this configuration can take a considerable length of time. A lot of this time is spent waiting for the interface and the CoA to be configured rather than on the process of actually communicating the CoA back to the HA. In scenario 2 we witnessed an improved handover time since the MANET interface is already configured with an address and must only wait to receive the appropriate HNA message. This removes the reliance on waiting for the Neighbor Discovery process to complete. In this scenario we see an increase in the overall latency experienced as is expected since the an additional hop has been introduced to the end-to-end path. Since no additional tunnel need be instantiated, this scenario also results in a reduced level of processing on the HA once it receives the BU. Scenario 3 highlights the implications of carrying out the proxy bind request that is carried out by HA2 on behalf of MR1. As the results show this approach does introduce implications on the overall latency and the achievable throughput. Obviously, since packets in this scenario must be transmitted via a secondary Proxy-HA this additional step will incur an increase in the length of the end-to-end path. In addition, the process of encapsulation and decapsulation for both the MR-HA tunnel and the proxy HA-HA tunnel impose some restrictions on the amount of data that it is possible to force through the network. Handover times also increase since the BU process now involves an additional party, however since the additional steps are only required on the HAs, the overall process does not introduce too considerable a delay.

These results show that the UMA approach outlined in this paper introduces much advantageous functionality without introducing too much overhead in terms of processing. One further factor that must be considered (as with any Home-Agent based approach) is the network distance between the HAs in the Non-Aggregated scenario. In our testbed setup, the networks that interconnect the Proxy-HA and the Target-HA directly connected together and therefore the latency of packet delivery between them is small. On the other hand, if the HAs were geographically far apart (i.e. one HA was in London, the other in Tokyo) the Proxy-HA connection would obviously suffer from increased latencies and handover times. It is also important to consider that this UMA approach ensures that any change in the network structure is only communicated out to the appropriate HAs when it is absolutely necessary. The first example of this is that any changes to the structure of the UMA-Stub that don't effect whichever MRs are acting as Gateways (i.e. any movement within the MANET) will not be reported beyond the UMA-Stub. This efficient technique is also applied to the Non-Aggregated Scenario; if an MR is connected back to its HA via Proxy any changes to the Gateway-MRs location need not be reported all the way back to the Target-HA. The Target-HA need only be updated when the Gateway-MR itself is changed and therefore a new HA-HA tunnel must be established.

We feel that the UMA approach outlined in this paper offers many useful benefits that could potentially be applied to numerous different scenarios and application areas. By combining the mutually beneficial properties of MANET and NEMO techniques we have been able to produce an extremely strong solution to network mobility that is both capable and also immediately deployable without any alterations required to the existing Internet architecture.

VI. Acknowledgement

The authors wish to thank Cisco Systems for their support of the Mobile Networks project at Lancaster University under which this work was completed.

References

[1] M. Watari C. Ng, P. Thubert and F. Zhao. “Network Mobility Route Optimization Problem Statement”. IETF Request For Comments 4888, July 2007.

[2] P. Thubert et al. “NEMO Basic Support Protocol”. IETF Request For Comments 2693, January 2005.

[3] C. Edwards B. McCarthy and M. Dunmore. Network Transparency in a Mountain Rescue Domain. *Internet Research Journal (Emerald Publications)*, Volume 17 Number 5, November 2007.

[4] J. Arkko D. Johnson, C. Perkins. “Mobility Support for IPv6”. IETF Request For Comments 3775, June 2004.

[5] T. Clausen and P. Jacquet. “Optimized Link State Routing Protocol (OLSR)”. IETF Request For Comments 3626, October 2003.

[6] E. Belding-Royer C. Perkins and S. Das. “Ad hoc On-Demand Distance Vector (AODV) Routing”. IETF Request For Comments 3561, July 2003.

[7] Anders Nilsson, Charles E. Perkins, Antti J. Tuominen, Ryuji Wakikawa, and Jari T. Malinen. AODV and IPv6 Internet Access for Ad Hoc Networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 6(3):102–103, 2002.