

Using NEMO to Support the Global Reachability of MANET Nodes

Ben McCarthy, Christopher Edwards, Martin Dunmore

Computing Department, InfoLab 21

South Drive, Lancaster University

Lancaster, LA1 4WA

Lancashire, UK

Email: [b.mccarthy, ce, m.dunmore]@comp.lancs.ac.uk

Abstract—Mobile Ad hoc Network (MANET) routing protocols have been the focus of an accomplished research effort for many years within the networking community and now the results of this effort are beginning to show. With protocol development maturing (and now typically concentrating on a smaller number of standardised routing protocols), increasing numbers of deployment successes are materialising. However, despite these successes and the relative stability of the protocol implementations, seamlessly incorporating MANETs into the Internet still presents many challenges that have hindered their deployment in important mobile scenarios. In this paper we discuss the inherent properties that have affected the adoption of MANET solutions and present an innovative new protocol which has been designed to comprehensively address these challenges. Using performance results acquired from our experimental testbed, we demonstrate how our approach can be used to produce MANET solutions that are highly suited to use in synergy with the current Internet architecture. Our protocol is based on the concept of integrating MANET routing protocols with Network Mobility (NEMO) technologies to produce what is termed a MANEMO solution. This has meant that by utilising the properties of both of these technologies we have been able to realise a solution that provides mobile networks with the efficient localised communication and robustness of MANETs, as well as the global reachability and the ability to provide structured AAA that a NEMO approach can support.

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 the NEMO Basic Support protocol (NEMO BS) [1] to produce solutions which benefit from the positive characteristics of both approaches. In this paper we introduce how the MANET problem space can benefit from the introduction of NEMO concepts. By incorporating an innovative Home Agent and bi-directional tunneling approach similar to the one utilised by NEMO BS into a MANET routing protocol we have been able to produce a routing technique that provides a comprehensive solution to some of the main problems affecting the uptake of MANETs today. The key aim of this technique is to ensure that MANET nodes are able to maintain persistent global reachability in the Internet whenever any one node in a MANET has Internet connectivity, and maintain that reachability irrespective of any subsequent roaming across heterogeneous Internet access

networks that takes place. Supporting this functionality is a well documented aim of the MANET community and in this paper we are presenting the design and evaluation of the first formal implementation of a solution that solves this goal using a MANEMO approach. We call our solution the Unified MANEMO Architecture (UMA) and in addition to supporting this primary aim of persistent global reachability, the UMA protocol also incorporates a number of other positive aspects including the ability to provide structured AAA solutions to unstructured MANET topologies and the potential to support a Multihoming approach that does not effect the size of the global routing tables. Finally, the UMA protocol has been designed to be a near term solution to the scenarios it supports and as a result can be deployed for use across any IPv6 enabled access network without the need to update any of the infrastructure or the access network itself.

The rest of this paper is presented as follows: In Section II we describe the problem of supporting global addressing for MANET nodes and highlight existing solutions that have been proposed. In addition we introduce our proposal of how NEMO technologies can be employed to create an entirely new approach to solving this problem domain. In Section III we present the design of our proposed solution, the Unified MANEMO Architecture (UMA) protocol and highlight its benefits and capabilities. In Section IV we present our testing results and an analysis of the general performance capabilities of our experimental implementation of UMA. Finally in Section VI we provide our conclusions of this work and discuss the feasibility of the UMA approach.

II. BACKGROUND

MANEMO protocols incorporate the use of both MANET and NEMO. In this section we analyse the current outstanding problems that arise when MANET nodes are introduced into the Internet. We describe why this problem of supporting global communication for MANET nodes occurs and highlight existing solution approaches to the problem. Finally, we provide a brief overview of NEMO BS and explain how it can be used to solve the MANET global reachability problem by creating a MANEMO solution.

A. Mobile Ad Hoc Networks (MANETs)

Mobile Ad Hoc Networking (MANET) protocols support node mobility (Mobile Host and Mobile Router) by performing routing information exchanges that have been specifically optimised to support mobile networking between nodes which predominantly have no prior infrastructure in place. MANET routing protocols can be classified as one of two main styles of protocol, Proactive or Reactive. Proactive MANET protocols such as the Optimised Link State Routing protocol (OLSR) [2] 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 the Ad Hoc On-Demand Distance Vector routing protocol (AODV) [3] 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.

B. Current Limitations of Global Communication with MANETs

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. 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 [4]. 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 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, discussed in the following subsections.

1) *NAT Based Approach:* A potential approach for supporting communication in these scenarios is to ensure that the addresses used by nodes within a MANET are always considered private and never used externally to the MANET itself. The Access Router in the access network could then be augmented to use a Network Address Translation (NAT) [5] technique specifically designed to support the connection of MANETs [6]. In this approach the Access Router that the

MANET used to transmit packets into the Internet would be required to map the private addresses of MANET nodes to a topologically correct global IP address. Therefore in order for a MANET that is reliant on this approach to roam onto an access network, the Access Router would need to have been upgraded to support this technique prior to the MANET's roaming attempt. This means that basic access networks that have only been designed for single host connections would not be able to support the introduction of MANETs. This approach could also potentially require the Access Router to install large amounts of state at once if a single Gateway with a large cluster of MANET nodes attached behind it roamed onto the Access Router's network.

2) *Auto-Address Configuration Based Approach:* To support MANET networks connecting to arbitrary access networks, an approach called *Auto-Address Configuration* [7] could be employed. With this approach each node configures an address that is topologically correct in relation to the access network that the Gateway node has connected to. Therefore in order to work, this approach requires the Gateway node to disseminate the topologically correct access network prefix to all of the nodes within the MANET. Once a node in the MANET recognises the availability of a new access network prefix it would construct a new address from this prefix. Then each MANET node would need to configure a new address from this prefix, every time the Gateway node changed its location. 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 configure a topologically correct address based on the Access Router's network prefix. 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. This problem would potentially become increasingly worse as the numbers of nodes within a MANET increased, as information about the topologically correct prefix for each new Access Network would need to travel further (more hops) away from the Access Router as the MANET grew in size.

C. How NEMO Can Help

By default, if an IP enabled node changes its point of attachment to the Internet it must configure a new address that is topologically correct for use in its new location. This resulting alteration in the address that the node utilises to communicate with other nodes in the Internet ultimately breaks any existing sessions that the node initiated using its previous address. In addition if a mobile node has a constantly changing address this also prevents other nodes in the Internet from actually initiating a communication session with it because it does not have a permanent, static address it is always contactable via. Mobile IPv6 (MIPv6) [8] is a technique that is designed to address this problem and Network MOBility (NEMO) is the augmentation of MIPv6 that supports the

movement of entire networks of moving nodes. Using the NEMO Basic Support protocol (NEMO BS), mobile networks can be provided with a persistent set of globally reachable Internet address prefixes that can be routed to a mobile network's ever changing location, 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 IPv6 devices as a normal, static IPv6 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 entire prefixes of addresses that are attached to the MR. As the MR moves around and changes its point of attachment to the Internet, it configures a new Care-of-Address (CoA) based on the prefix of the access network it is currently connected to. The MR then subsequently registers this CoA with its HA and they build a bi-directional tunnel between one another, this registration process is called the Binding Update (BU). When a node attached to the MR then sends packets into the Internet, they are first tunneled via the HA and conversely, for any packets sent from a node in the Internet towards the MR, the HA intercepts them and forwards them to the MRs current location via the bi-directional tunnel.

Our previous research carried out in this area led us to identify that it would be possible to leverage the approach adopted by the NEMO BS protocol to allow entire MANETs of nodes and mobile networks to legitimately transmit packets into the Internet without any cooperation from the access network. Essentially, the bi-directional tunnel approach imposed by NEMO BS prevents packets transmitted by nodes in a mobile network from being Ingress Filtered [9] by the access network and ensures that packets destined for those nodes are delivered to the correct location in the Internet. If this functionality is therefore incorporated into the gateway node of a MANET it can then feasibly perform these operations on behalf of all of the MANET nodes that it provides an Internet connection for. This is the basic principal that has driven the design of our Unified MANEMO Architecture (UMA) solution.

III. UNIFIED MANEMO ARCHITECTURE

In this section we present the design of our MANEMO protocol, the Unified MANEMO Architecture (UMA). With the UMA protocol every MANET node sets up a bi-directional tunnel with its HA whenever it is able to establish its own direct connection to the Internet. At the same time, all of these nodes also maintain an ad hoc communication interface over which they participate in Optimized Link State Routing (OLSR) protocol exchange to establish a MANET with other MANET nodes around them. When a MANET node is able to establish a direct connection to the Internet, it assumes the role of a Gateway. In this role, the MANET node provides the opportunity for other nodes connected to the MANET to utilise its pre-existing connection with its HA to forward their

packets into the Internet. Once packets leave the MANET, UMA then employs a technique of HA inter cooperation that ensures that any MANET nodes can connect to the Internet via any available MANET. It is important to note that for our implementation we have specifically designed the UMA protocol to support entire mobile networks of hosts and therefore whenever we refer to a MANET node we considered this node to be a mobile router. However UMA does also support the more simplistic case whereby each of the MANET nodes operates as an individual host.

A. Roaming

With UMA, a MANET node can access the Internet via one of two methods. Either it can establish its own direct connection to the Internet, in which case the node performs the role of a Gateway. Or it can establish a connection to the Internet indirectly, via an existing Gateway. MANET nodes are only required to build a topologically correct Care-of-Address (CoA) if they fall into the former category and establish their own direct Internet connection via an access network. If their connection is indirectly established via other MANET nodes then packets are routed into and out of the MANET based on the tunnel that the Gateway already has in place with its HA. To ensure this is possible, when a MANET node performs a Binding Update (BU) via a Gateway the HA must therefore record which tunnel the BU request arrived via and install routes to the MANET nodes and its associated prefixes via this tunnel. Once incoming packets reach the Gateway via its bi-directional tunnel, they are then routed onward to the correct node in the MANET using the routes installed by the OLSR MANET routing protocol. This approach, as described so far, relies on the newly connecting MANET node sharing the same HA as that of the Gateway. In order for the packets to reach the HA in the first place, they must first be routed normally through the Internet. This would therefore only happen for MANET nodes that the HA already maintains prefixes for, as these would be advertised by the HA into the Internet infrastructure.

As a result a process must also be in place that ensures that when a MANET node that is registered with a different HA attempts to utilise the Internet connection of a Gateway, it is also supported and its registration request is also carried out with the its HA as well as the HA of the Gateway. To give an example of where this problem may occur, consider an emergency situation such as the breakout of a fire in a building. In this scenario the emergency services crews that attend the fire would all employ systems based on the use of UMA (to support PAN networks for the individual emergency servicemen and vehicle networks for their fire engines and ambulances). 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 the MANET mobile routers carried by the firefighters will have originated from the 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 mobile routers into the existing network should also be supported. However since their mobile routers 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 MANET.

In protocol terms, if a MANET node can only obtain an indirect connection to the Internet via a Gateway, then the principal concern is whether that node is registered with the same HA as the Gateway or a different one. So referring to the fire scenario, if a paramedics mobile router seeks to establish a connection to the Internet indirectly via a Gateway then it first ascertains whether the Gateway is registered with the paramedics HA or the fire brigade HA. If the newly connecting MANET node detects that the Gateway is registered with the same HA as itself (i.e. the paramedics HA in this case) we refer to this as an Aggregated Roaming Scenario. If however the newly connecting MANET node detects that the Gateway is registered with any other HA (i.e. the fire brigade HA) we refer to this as a Non-Aggregated Roaming Scenario. Here we outline how the operation of the MANET node and the HAs will differ in these two distinct scenarios:

1) *Aggregated Roaming Scenario*: In this scenario (illustrated here in Figure 1) the newly connecting MANET node will send its registration request message directly to the Gateway's HA (which is also its own HA). Receiving this type of request message signals to the HA that a MANET node is trying to indirectly bind to it via one of its existing tunnel connections that it already has in place with another one of its MANET nodes. Subsequently, after ensuring that the newly connecting MANET node is permitted to register with itself, the HA will record the details of the tunnel that the registration request was received via and then install routes to the newly binded MANET node and its connected prefixes, via that tunnel.




Fig. 1. Aggregated Roaming Scenario

2) *Non-Aggregated Roaming Scenario*: The Non-Aggregated Roaming Scenario (illustrated here in Figure 2) refers to the situation where a MANET node tries to establish a connection via a Gateway that is registered with a different HA. In this situation the Gateway's HA will behave as a Proxy-HA, forming an indirect link between any legitimate

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

Forming this direct HA-HA tunnel link between the Proxy-HA and the Target-HA ensures that the Target-HA is able to transmit packets to the MANET nodes Home Address via an address that would otherwise be unreachable (the address the MANET node propagates throughout the MANET). Since the addresses used within the MANET are not leaked into the Internet, the combination of the HA-HA tunnel and the pre-existing Gateway tunnel provide a direct link for the Target-HA and its associated MANET node to transmit packets via.

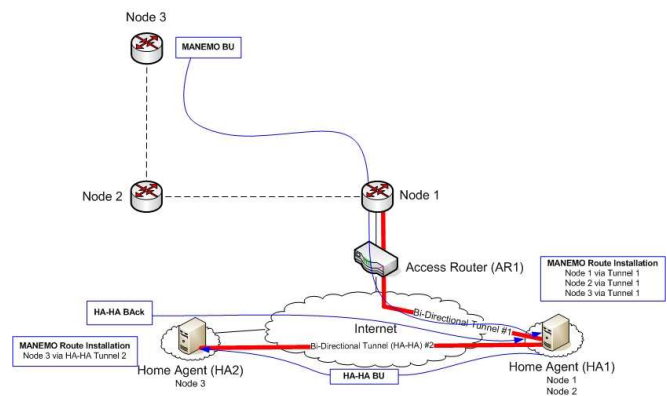


Fig. 2. Non-Aggregated Roaming Scenario

B. Inter-MANET Mobility

Once the connections between Gateways and their HAs have been established, the routing of packets from HAs to all of the nodes in a MANET is based upon tunnel ID. This simple and efficient concept means that mobility between different MANETs only translates into the process of a single route modification on the HA. To demonstrate this, again consider the fire in a building scenario; in a situation where all of the MANET nodes in the scenario have formed into a topology consisting of two distinct MANETs (as illustrated in Figure 3). In this illustration, the movement of the paramedic Node 3 from MANET 2 to MANET 1 would result in the need to only update the paramedic HA with the new tunnel ID number

that packets destined for Node 3 should be forwarded via. This is because both of the Gateways are also paramedics and therefore Node 3's binding with its HA will remain valid so just the path it can be reached via must be updated. Similarly, in the case also illustrated where the fire brigade Node 4 moves from MANET 1 to MANET 2, again a simple route modification will be all that is required. When registering with its HA via MANET 1, the fire brigade Node 4 will have formed a proxy connection with its HA via the paramedic HA. If any subsequent movement made by Node 4 results in it registering via the same Proxy-HA (in this case the paramedic HA) the change in location need not be reported beyond that Proxy-HA. Since the fire brigade HA and the paramedic HA already have a HA-HA tunnel in place and a valid binding for Node 4, the Paramedic HA can simply update the tunnel ID number that Node 4 can be reached via without needing to inform the fire brigade HA of any change. It is also important to highlight that the HA-HA tunnel connection need only be instantiated once, irrespective of the number of proxy connections which then subsequently utilise it. Again, considering Figure 3, if two of the MANET nodes in MANET 1 were registered with the Fire Brigade HA and two of the MANET nodes in MANET 2 were also registered with this HA, there would still only ever be a requirement for one HA-HA tunnel. Once the first fire brigade MANET node has successfully initiated its connection with the fire brigade HA, any subsequent connections made via the same Proxy-HA (the paramedic HA in this case) simply utilise the existing HA-HA tunnel which is already in place.

In addition, this routing approach of forwarding packets to the appropriate MANET based on the tunnel ID they are reachable via has the additional benefit that any subsequent movement the Gateway performs can be carried out transparently to the rest of the MANET that is connected behind it. This means that the movement of an entire MANET of nodes across heterogeneous access networks caused by the changing point of attachment of a Gateway, only ever requires a single Binding Update exchange to take place between the Gateway and its HA, irrespective of whether the MANET contains 10, 50 or 1000 MANET nodes.

IV. TESTING AND RESULTS ANALYSIS

In this section we present and analyse the results of our experimental evaluation of our UMA implementation which we developed on the 2.6.22 version of the Linux kernel. In order to perform the testing we configured two distinct testbeds. Testbed 1 (illustrated in Figure 4) refers to the standalone setup we devised where all of the associated entities of the testbed are located in the Computing Department at Lancaster University. This setup consisted of five UMA-enabled laptop PCs (configured to operate as MANET nodes), each consisted of a 2Ghz AMD Athlon Processor, 512MB RAM, an onboard Atheros Chipset 802.11b/g wireless interface and a Cardbus Atheros Chipset 802.11a/b/g wireless interface. This testbed also included two UMA-enabled Linux desktop PCs with 2Ghz CPU, 512MB RAM & 80GB hard drives (configured to operate as HAs), three static IPv6 enabled Cisco

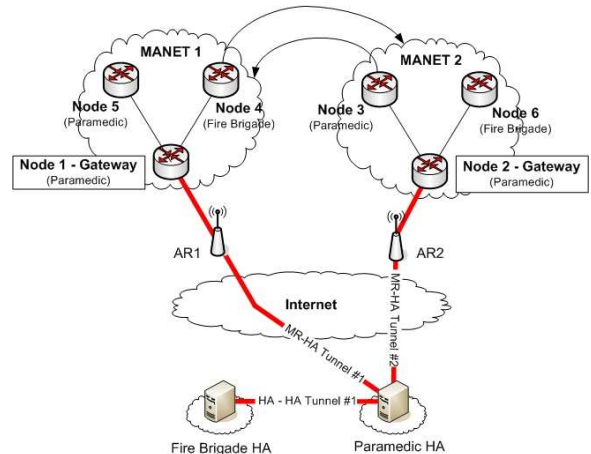


Fig. 3. Inter-Stub Mobility

routers (labelled Access Router 1 - 3) and three IPv6 enabled Cisco Aironet WiFi Access Points. In all of the experiments, separate interfaces on two of the static routers were used to provide one Home Network and one Access Network per router. In addition, all three of the static routers were also interconnected together using a further interface to provide an Ethernet backbone between all of the networks. Connected to each Home Network interface via Ethernet was an individual PC configured to operate as a HA (labelled HA1 and HA2 in the diagram). Connected to the Access Network interface of each static router was an IPv6 enabled Aironet Access Point configured to operate in 802.11g mode. Finally, the five UMA-enabled Linux laptop PCs were configured to operate as MANET nodes and therefore form one or more MANETs during testing. Testbed 2 (illustrated in Figure 5) on the other hand was designed to illustrate UMA's potential to be deployed for use over the Internet at present, so we therefore incorporated the use of geographically dispersed UMA-enabled HAs (which we located at the University College London's computing department and Lancaster University's main campus network) and Wide Area Internet access technologies (such as a HSDPA link via Vodafone's cellular data network and a satellite communication link via SES Astra's satellite network).

Over each testbed we performed a 4 stage roaming procedure that tested each of the different potential UMA Binding Update processes that can take place. For each stage we recorded the handover times experienced, the overall throughput achievable once the handover had taken place and also the effect that the UMA approach had on the overall end-to-end latency between a host on the Test Node network and the Correspondent Node. Our testing for each stage of UMA mobility we configured was based on the following three step procedure:

- 1) For each stage we first determined the handover time experienced by using the Ping6 utility in collaboration with the network packet analyser Wireshark. This involved setting the ping request interval to a high value (1

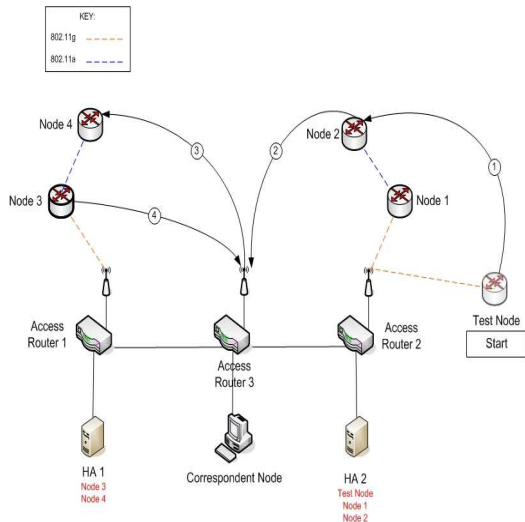


Fig. 4. “Standalone” Experimental Testbed Setup

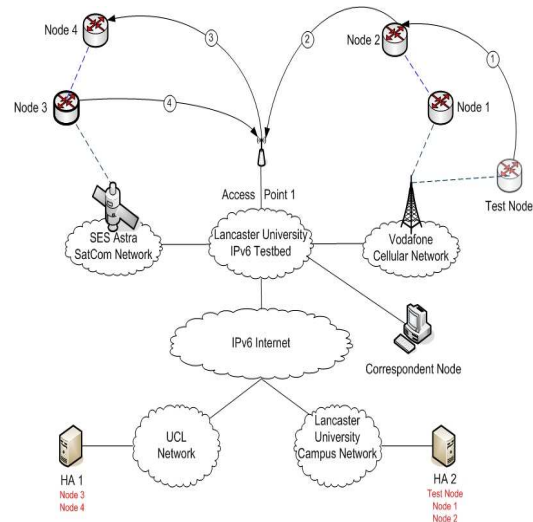


Fig. 5. “Global” Experimental Testbed Setup

request every 0.01 seconds) and then recording the time difference between the time the last ping reply was received (i.e. the beginning of the roaming procedure) and the time the next reply was successfully received (i.e. the point at which the connection was reestablished).

- 2) Once the connection was established, the Ping6 utility was then used to collect 1000 Round Trip Time (RTT) measurements to obtain an average latency value.
- 3) Finally, once the latency test was completed, TCP throughput was determined using the NetPerf bandwidth measurement tool.

For each step in the testing procedure, this regime was repeated 20 times to ensure the results were consistent. We present all of the results from our experimental evaluation over both testbeds in their respective sections below and provide a summary of the results for Testbed 1 in Table I and Testbed 1 in Table II.

A. Stage 1: UMA Aggregated Roam

Stage 1 of our testing process was designed to emulate an Aggregated Roaming scenario. In this stage Node 1 is connected to its respective HA (HA 2) and therefore acts as a Gateway providing an in-direct connection to the Internet to Node 2 over its ad hoc interface. The Test Node then loses its own direct connection to the Internet but is presented with the opportunity to reestablish its connection via Node 2. Therefore because the Test Node was configured to originate from the same HA as the Gateway (Node 1) in this situation, no HA-HA communication would be required to take place as binding requests from the Test Node would immediately reach its own HA (HA 2) after being tunneled out of the MANET. In each of the roaming procedures where the Test Node establishes a connection via an existing MANET, the node is able to perform a ‘Make-before-break’ handover, whereby it establishes a layer-2 connection with a MANET node first which it can utilise as soon as it loses its direct connection to the Internet. In addition the node is also able to register its own

ad hoc interface address as its Care-of-Address (CoA) with its HA as this address is already distributed within the MANET, which means the node is able to avoid the costly process of configuring a topologically correct address as it must do if it establishes a direct connection to the Internet. This therefore results in a relatively quick handover time of under 1 second in Testbed 1. In Testbed 2 this figure unavoidably increases not only because the round trip time between the Test Node and its HA is much greater, but also because of the lossy nature of the link. In many cases we observed the loss of the initial BU message that the Test Node transmitted, which ultimately causes a longer handover as the node waits to retransmit a second (and in some cases third) BU message. In Testbed 1 in this scenario we also saw slight increases in the overall latency experienced and a slight decrease in the throughput measured. This is expected since additional hops via the ad hoc wireless connections between Node 1, Node 2 and the Test Node are introduced into the end-to-end path. However in the case of Testbed 2 these increases were undetectable because the fluctuation in latency and throughput caused by the HSDPA network link were so large that they effectively masked any performance degradation experienced within the MANET itself.

B. Stage 2: NEMO

The movement in Stage 2 represents the Test Node roaming away from the MANET it joined in Stage 1 and establishing its own direct connection to the Internet via Access Router 3. In this situation the Test Node detects that it should act as a Gateway because it becomes involved in the IPv6 Neighbor Discovery process over the interface connected to Access Router 3 and therefore configures a topologically correct address that is valid for use in that network. Again, in this situation it is possible for the MANET node to simultaneously establish an alternative (direct) connection to the Internet at layer 2 whilst it continues to communicate with Internet nodes via its existing connection. This therefore results in

the MANET node again being able to quickly perform a handover once it chooses to switch interfaces as it will already have configured a topologically correct address with which to communicate over the Access Network as well. The resulting configuration that remains in place once the handover in this Stage has been performed offers the best overall latency and throughput performance capabilities because the Test Node is directly connected to the Internet and therefore doesn't transmit its packets over any additional wireless hops. When carried out over Testbed 2, the resulting network configuration from this roaming stage again provided the best performance results as the Test Node was ultimately connected to the highest quality link and registered with the closest HA (HA 2).

C. Stage 3: UMA Non-Aggregated Roam

In addition to testing the Aggregated Roaming Scenario it is then important to understand the implications that the additional overhead imposed by the Non-Aggregated Roaming Scenario has on the performance of UMA. Therefore in this stage of the testing we caused the Test Node to perform a similar handover to an existing MANET by roaming it from Access Router 3 to Node 4. This movement subsequently causes the Test Node to initiate a Non-Aggregated Binding Update because Node 3 (the Gateway) is registered with HA 1 whilst the Test Node is registered with HA 2. This situation therefore highlights the performance implications of the proxy bind request and of the HA-HA communication that is associated with it. What we witnessed in this testing stage was a slight but acceptable increase in the overall handover time required in comparison to the Aggregated Roaming Scenario and similarly acceptable degradation in the latency and throughput performance. This overall performance hit could obviously be expected since the binding process in this scenario involves an additional party and an increase in the overall amount of processing that must be performed. In addition, the introduction of the Proxy-HA into the network configuration also impacted on the overall latency and the achievable throughput. Packets in this scenario were transmitted via an additional hop via the Proxy-HA before reaching the Test Node's own HA, but also incurred the processing overhead of a further IPv6-in-IPv6 encapsulation phase between the two HAs. It is important to note that this procedure represents the most complicated roaming event that can occur with UMA, and therefore no other UMA roaming scenario results in an operation with any greater level of processing overhead. Whilst variable factors such as the density of the MANET (and thus the number of ad hoc wireless hops packets must travel before they are delivered) or the distance between two inter-communicating HAs will affect the overall service received by a node in a MANET using UMA, the amount of processing in the HAs imposed by the protocol never increases. Regardless of the number of nodes in a MANET or the possible configuration of the HAs, for any individual MANET node the UMA protocol will only result in one level of IPv6-in-IPv6 tunnel (no nested tunneling

is performed) and a potential connection with one Proxy-HA. With this in mind, these preliminary evaluation results are an encouraging display of the overall capabilities of the UMA protocol. In addition to the performance observations we made when this stage was carried out over Testbed 1, this roaming stage also involved communication over the satellite link when we performed it over Testbed 2. Utilising this link therefore imposed harsh limits on the level of throughput achievable and also increased the Latency experienced significantly. We also observed the highest level of loss over this link and this contributed negatively to the average handover time we recorded.

D. Stage 4: UMA Gateway Roam

Finally, we wanted to test the implications of roaming the Gateway node when it has a Non-Aggregated connection in place via a Proxy-HA. To achieve this, we used the resulting MANET network configuration that remained in place after the testing performed in Stage 3 and instead of causing the Test Node to perform a roaming procedure, we roamed the Gateway (Node 3) from Access Router 1 to Access Router 3. Since all of the packets that are transmitted between the Test Node and its HA in the Non-Aggregated Roaming Scenario are routed based on the appropriate tunnel ID numbers (i.e. both the HA-HA tunnel ID and the Gateway's tunnel ID), the extent of the packet loss experienced by the Test Node is only determined by the loss of availability of the Gateway connection. For this reason, the roaming of a Gateway from one Access Router to another is the same procedure that a NEMO mobile network performs when it changes its point of attachment to the Internet. This is because the Gateway node must first break its connection to an Access Network in order to subsequently reestablish it with another, different Access Network. As with the NEMO BS protocol, this layer 2 handover time imposes significant performance implications on the overall network layer handover time experienced in these scenarios. This stage in our testing highlights that the performance experienced when using the UMA protocol is only ever at worst equal to the performance that is supported by the NEMO BS protocol.

In this stage of the testing, the resulting network configuration that is in place after the Gateway's roaming procedure has completed is exactly the same as in testing Stage 4 (i.e. Node 3 performing the role of Gateway node with Node 4 and the Test Node attached behind it). Therefore the latency and throughput results from the testing performed over Testbed 1 were observed to be very similar. In contrast however, the resulting configuration in this stage when we performed this testing over Testbed 2 culminates with the Gateway node (Node 3) accessing the Internet via an 802.11g Access Point connected to the IPv6 Testbed at Lancaster University. This therefore resulted in much improved performance to that experienced in Stage 3 where the Gateway was attached to the satellite network of SES Astra. However, in comparison to the results attained over Testbed 2 during Stage 2 of our evaluation (when again the Gateway node has a connection to the Internet via

an 802.11g Access Point connected to a relatively high speed network) the throughput performance was considerably lower. In this situation we ascertained through additional analysis that the bottleneck in this situation was in fact imposed by the path available to the HA located at UCL, we were able to determine that even a direct transfer between these two sites was constrained to similar levels of throughput as those we recorded with UMA.

V. FUTURE WORK

In addition to the direct benefits we have outlined, the UMA approach also introduces potential advantages in a number of other areas, including Multihoming and Authentication, Authorisation and Accounting (AAA) for MANETs, these will be explored more as part of our future work.

A. Multihoming

Multihoming in mobility scenarios is a highly useful concept. The use of multiple available connections to the Internet can help improve the resilience and reliability of a node's Internet connectivity as well as provide an opportunity to perform more seamless handovers. Figure 6 illustrates a scenario in which a newly attaching MANET node (Node 3) has three available Internet access options. Two indirect connections to the Internet via existing MANETs and one direct connection via a UMTS interface. This type of communication situation could be feasibly expected to arise in many typical MANET scenarios. By leveraging the concept of Multiple Care-of-Address Registration (MCoA) [10] within UMA, the newly attaching MANET node could make use of both of the available in-direct connections to the Internet as well as establish its own direct connection via its UMTS interface to register three simultaneous bindings with its HA (HA3). Using this approach the MANET node is able to register simultaneous locations that it is reachable at with its HA. Once registered, the MANET node and the HA can then choose which connection to transmit packets via based on policy or connectivity quality. In addition, this approach also would also enable MANET nodes to perform near instantaneous handovers since parallel layer 3 connections can be established and then switched between as and when required, resulting in almost no disruption.

B. Authentication, Authorisation and Accounting (AAA)

The ability to efficiently and accurately perform Authentication, Authorisation and Accounting (AAA) is a fundamental component of most successful networking solutions. Performing effective AAA in Mobile Ad hoc networks is inherently difficult because of the infrastructureless nature of ad hoc networks. UMA has been designed in a manner which attempts to provide a potential solution to these AAA considerations by leveraging the structured approach of the Inter-HA communication process imposed by the UMA protocol. This process ensures that there is always static entity available that is directly associated with any MANET node (i.e. the HA). As the HA is always involved in the communication

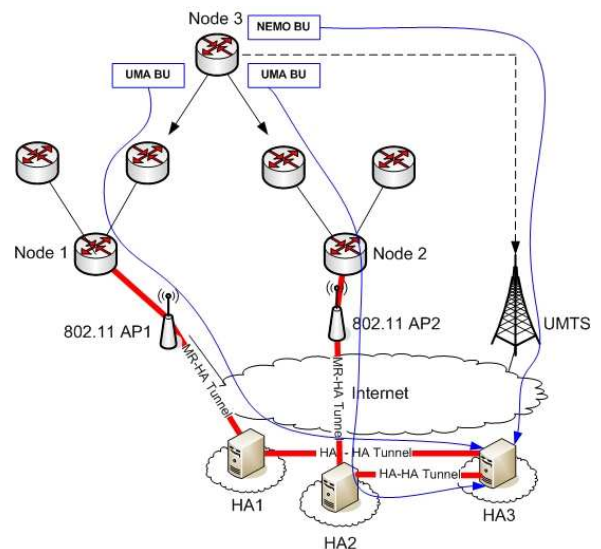


Fig. 6. UMA Multihoming

process it is therefore constantly available to authenticate the MANET Node and can be subsequently billed for the nodes service usage if necessary. Accountability is important because if we consider a typical Mobile Ad hoc Networking scenario whereby nodes in the MANET wish to communicate externally with nodes in the Internet, the Gateway nodes are required to perform an unfair role in the overall communication model. This is because the Gateway nodes will be required to carry the traffic of other nodes in the MANET as well as its own. Arbitrarily requiring nodes to perform this function may be infeasible in certain scenarios, especially if the Gateway node accesses the Internet via a potentially resource constrained or financially expensive access medium. In these scenarios, the Gateway node will suffer a degradation in their own service and possibly incur additional costs. The Inter-HA communication system employed by UMA ensures that the HAs of Gateways are potentially able to maintain accurate records of which other MANET nodes have utilised the Gateways Internet connection and how much traffic they have transmitted in total. At present our implementation only performs basic Access and Authentication checks, but it is our intention to integrate a comprehensive AAA solution into the UMA model in the future, in order to demonstrate the potential benefits available through using this approach.

VI. CONCLUSION

In this paper we have presented our approach to extending the functionality of MANET routing protocols through the use of a MANEMO based solution. By integrating NEMO techniques with existing MANET technology our UMA protocol is able to provide a comprehensive solution to providing global connectivity to MANET scenarios. The UMA approach has been designed to support entire mobile networks of hosts. In doing so it ensures that all hosts connected to any UMA enabled mobile networks are not required to take part in any

Stage	HandOff	Latency	Throughput
Stage 1	0.89 Secs	3.789 MSecs	9.06 Mbps
Stage 2	0.84 Secs	3.482 MSecs	11.27 Mbps
Stage 3	1.28 Secs	6.368 MSecs	8.84 Mbps
Stage 4	1.47 Secs	6.413 MSecs	8.81 Mbps

TABLE I
TESTING RESULTS - "STANDALONE" TESTBED

Stage	HandOff	Latency	Throughput
Stage 1	3.16 Secs	461 MSecs	140.23 kbps
Stage 2	1.04 Secs	11.321 MSecs	11.04 Mbps
Stage 3	5.48 Secs	637 MSecs	100.9 kbps
Stage 4	1.89 Secs	22.432 MSecs	1.76 Mbps

TABLE II
TESTING RESULTS - "GLOBAL" TESTBED

form of mobility signalling themselves as the UMA enabled mobile router will perform this functionality on their behalf. Supporting this capability ensures that any nodes connected to UMA enabled mobile networks (such as Personal/Vehicle Area Network nodes) need only be standard IPv6 hosts. This in turn ensures that nodes connected to UMA enabled mobile networks can communicate constantly across the Internet using the same address regardless of their physical location, without their TCP sessions being dropped whenever a roam takes place. It also means that nodes in the MANETs can be directly contacted from anywhere in the Internet, without having to establish a prior flow of packets. In addition to these benefits, we also strived to ensure that the UMA approach did not affect the current Internet architecture by requiring any augmentation of the core infrastructure or in any access networks. By achieving this aim we are able to present UMA as a mobility solution that is immediately suitable for use across IPv6 enabled networks. This is an important consideration since the number of different providers offering Internet access is already significantly large. Therefore a solution which relies on Access Routers in these provider networks being augmented to support its functionality could be excessively difficult to deploy. As well as a large number of Internet access providers, there also exists numerous different technologies that can be used to connect to the Internet that each possess differing network characteristics. The ability to simultaneously utilise as many of these connectivity options as possible through the use of a Multihoming approach can provide significant improvements to the robustness of mobile networking scenarios. For this reason UMA was designed and implemented in a manner which inherently supports this capability through the use of multiple simultaneous network location bindings, and this will be explored further as part of our future work.

Using a testbed configured to replicate a realistic UMA communication scenario we also carried out a performance evaluation of the experimental implementation of our protocol. The results of this experimentation were considered to be very encouraging. Our intention for the UMA protocol was to design an approach which could provide global reachability for MANET nodes with a handover performance that was as good as or better than shown by the NEMO Basic Support protocol with individual mobile networks. Through the results of our testing with the UMA protocol we have shown that in every configuration that arises when using UMA we achieve that goal and in certain cases, notably improve on the performance of NEMO BS. In addition to highlighting the actual performance of the UMA protocol we also configured

a second testbed using wide area Internet access technologies and distributed Home Agents that was intended to demonstrate UMAs suitability for immediate deployment over the existing Internet infrastructure. Using the UMA protocol we were able to provide MANET nodes with the benefits of global reachability via access networks including a satellite communication link and a HSDPA cellular connection. This capability would simply not be possible using any other existing proposed solutions to this problem as it would require permission to install experimental software on the Access Routers of the respective networks.

Therefore by combining the properties of both MANET and NEMO techniques we feel we have been able to produce and demonstrate a versatile and efficient approach to extending the functionality of MANETs that is immediately deployable without any alterations required to the existing Internet architecture.

VII. 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] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert. "NEMO Basic Support Protocol". IETF Request For Comments 2693, January 2005.
- [2] T. Clausen and P. Jacquet. "Optimized Link State Routing Protocol (OLSR)". IETF Request For Comments 3626, October 2003.
- [3] C. Perkins, E. Belding-Royer, and S. Das. "Ad hoc On-Demand Distance Vector (AODV) Routing". IETF Request For Comments 3561, July 2003.
- [4] 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.
- [5] G. Tsirtsis and P. Srisuresh. "Network Address Translation - Protocol Translation (NAT-PT)". IETF Request For Comments 2766, February 2000.
- [6] H. Zhou, M.W. Mutka, and L.M. Ni. "Ip address handoff in the manet". In *Proceedings of Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2004)*, 7-11 March, 2004, Hong Kong, China, volume 4, pages 2425–2433. IEEE, IEEE, March 2004.
- [7] IETF Ad-Hoc Network Autoconfiguration Working Group Charter Homepage. <http://www.ietf.org/html.charters/autoconf-charter.html>.
- [8] D. Johnson, C. Perkins, and J. Arkko. "Mobility Support for IPv6". IETF Request For Comments 3775, June 2004.
- [9] P. Ferguson and D. Senie. "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing". IETF RFC 2827, May 2000.
- [10] R. Kuntz. "Deploying Reliable IPv6 Temporary Networks Thanks to NEMO Basic Support and Multiple Care-of Addresses Registration". Second International Workshop on Network Mobility (WONEMO), January 2007.