

Protocols to Efficiently Support Nested NEMO (NEMO+)

Ben McCarthy, Matthew Jakeman,
Dr Chris Edwards
Computing Department
Lancaster University
Lancaster, LA1 4WA, UK
[mccarhb, jakemanm, ce]
@comp.lancs.ac.uk

Pascal Thubert
Cisco Systems
Village d'Entreprises
400, Avenue de Roumanille
Sophia Antipolis 06410, France
pthubert@cisco.com

ABSTRACT

The NEMO Basic Support (NEMO BS) protocol provides a technique for enabling entire networks of IPv6 hosts to gain Internet access and remain reachable via constant, unaltered addresses whilst their underlying location in the Internet changes. In addition to individual hosts, this NEMO model also supports entire mobile networks connecting to other mobile networks, resulting in topologies known as Nested NEMO networks. In this paper we explain the inefficiencies that arise if NEMO BS is used to support this type of scenario and introduce our NEMO+ suite of protocols which are designed to optimise performance in Nested NEMO networks. We detail the Tree Discovery (TD), Network In Node Advertisement (NINA) and Reverse Routing Header (RRH) protocols that make up the NEMO+ suite and provide experimental evaluation results from a testbed comprising of our two distinct protocol implementation platforms (Linux and Cisco IOS). In addition we present simulation results based on scenarios of mass deployment of NEMO+ enabled mobile networks in order to determine the feasibility of our approach to efficiently support Nested NEMO networks.

Categories and Subject Descriptors

C.2.2 [Network Protocols]: Routing Protocols

General Terms

Design, Experimentation, Performance

Keywords

NEMO, Tree Discovery, NINA, RRH, Implementation

1. INTRODUCTION

Network Mobility is a powerful and important concept that is aimed at supporting the mobile operation of entire networks of IP enabled devices. As the Internet has progressively evolved and Internet access has become gradually

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiArch'08, August 22, 2008, Seattle, Washington, USA.
Copyright 2008 ACM 978-1-60558-178-1/08/08 ...\$5.00.

more ubiquitous, more and more mobile computing scenarios have arisen which could benefit from these types of network. Scenarios such as Internet access networks on modes of public transport like trains, buses and planes where commuters and travelers are beginning to seek out high quality Internet connectivity for the duration of their journey is one such example. Other examples include Personal Area Networks (PAN) and Vehicle Area Networks (VAN); in these scenarios numerous electronic devices (such as GPS units and Audio/Visual entertainment systems) which could benefit from Internet connectivity, can be expected to be physically collocated in a small scale network interconnected by either a short range wireless communication protocol like Bluetooth or a wired approach like Ethernet. The IETF defined NEMO Basic Support Protocol (NEMO BS) [1] is an approach to supporting network mobility scenarios that is based upon the Mobile IPv6 [2] routing model. One of the key advantages that using a Home Agent (HA) based approach ensures is that all of the IP devices connected to the mobile network need not be aware of their own mobility and can therefore connect to the mobile network as if it were a standard static network without running any mobility protocol themselves.

Unlike MIPv6 however, a NEMO Mobile Router (MR) presents an Ingress interface that other IPv6 devices can attach to, this creates the possibility that a NEMO MR could attach to the Internet via another NEMO MR. When NEMO MRs inter-connect in this manner they form a network topology known as a Nested NEMO and using NEMO BS to support this situation results in packets being transmitted via an extremely inefficient end-to-end path. In this paper we describe the NEMO+ trio of protocols (Tree Discovery, NINA and RRH) and highlight how these protocols can be used to efficiently route packets in Nested NEMO network topologies. We will introduce the two implementations we have of this protocol suite and present the results of our testing.

In Section 2 we explain how NEMO Basic Support operates, we provide an overview of the effects of Nested NEMO and we analyse the properties of Nested NEMO scenarios. In Section 3 we provide an overview of the NEMO+ trio of protocols. In Section 4 we detail the experimental evaluation and simulation work that we carried out and present the results of our testing. Finally in Section 5 we conclude the paper with a discussion of the work and the feasibility of this approach.

2. 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) that is targeted at supporting entire networks of IPv6 devices as opposed to just single hosts. Using NEMO BS, mobile networks can provide hosts with Internet connectivity via a constant unchanging IPv6 address regardless of the actual location of the network and without any of the hosts 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 an entire prefix of addresses that are attached to the MR, known as the Mobile Network Prefix (MNP). An individual NEMO MR can have multiple MNPs registered to it, to ensure that complex configurations can be supported on the mobile network. 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, the HA then intercepts any packets destined for the MR or for addresses that match the MNPs that are attached to the MR and forwards them to the MRs current location via the bi-directional tunnel.

2.1 Nested NEMO Overview

NEMO BS can effectively support straightforward scenarios involving single, distinct mobile networks that sporadically change their point of attachment to the Internet. However, scenarios exist where mobile networks would benefit from the ability to directly inter-connect and send packets toward the Internet via one another or communicate directly with each other. NEMO BS can fundamentally be used to support scenarios where NEMO MRs connect to other NEMO MRs (which is known as Nested NEMO), however the resulting routing that occurs in these Nested NEMO scenarios quickly becomes extremely inefficient and is not suitable for real life deployment solutions.

To illustrate how inefficient these Nested NEMO Networks can become consider Figure 1. This illustration shows how even in a simple Nested NEMO scenario where only two mobile networks are connected together, packet transfer will follow a highly suboptimal path. For a packet to be sent from Node A that is connected to the mobile network NEMO 1 to Node B on NEMO 2, the packet must first be tunneled through the destination NEMO (NEMO 2) and its Home Network (Network 2) before it is finally decapsulated by its Home Agent (HA1). The packet is then forwarded (without tunnelling) to the destination NEMO's Home Network and is then finally delivered to NEMO 2 over its MR-HA tunnel. The inefficient routing model that occurs in Nested NEMO Networks is commonly referred to as "Pinball Routing". This example also illustrates how using NEMO BS to support these scenarios requires the HAs to be permanently reachable, even if a direct connection between the two mobile networks exists. Again this is another fundamental flaw that many use case scenarios cannot accommodate.

In more complicated scenarios, such as vehicle commu-

nication where the number of inter-connecting networks is potentially much higher, the resulting network topologies that would arise if NEMO BS were relied upon would ultimately be massively more complex. Consider a vehicle inter-communication scenario that allows vehicles to access the Internet via the mobile networks of other vehicles. This type of scenario would proliferate the availability of the Internet by allowing vehicles to mesh together and access the Internet indirectly via other vehicles when they are unable to form their own direct Internet connection. This kind of scenario would result in many mobile networks connecting to each other and could therefore be expected to generate Nested NEMOs that were many layers deep, further exacerbating the inefficiencies of the NEMO BS model.

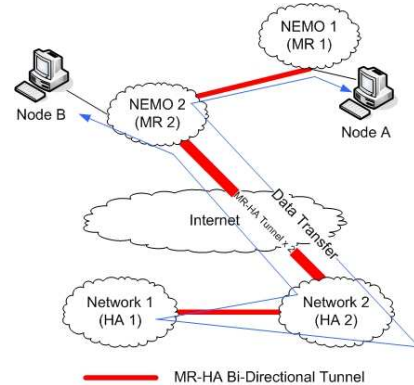


Figure 1: Nested NEMO supported by NEMO BS

2.2 Nested NEMO Solution Properties

A Nested NEMO network forms under specific circumstances. Unlike a typical Mobile Adhoc Network (MANET) where communication is predominantly supported in order to facilitate data transfers amongst the members of the MANET, a Nested NEMO network configuration principally arises when mobile networks are specifically trying to attain Internet connectivity. Therefore configuring an efficient default route to the Internet is the main characteristic that should be achieved by a solution for supporting Nested NEMO. Whilst supporting the efficient routing of packets to and from the Internet is of key importance, consideration must be applied for the possibility that packet transfers may occur between two nodes in the same Nested NEMO. In this case, it is undesirable that any such packet transfer should first be routed out into the Internet, via any respective HAs since a far more efficient, local route through the Nested NEMO should be available. In addition to ensuring that packets traverse a direct, optimised route when they are transmitted across a Nested NEMO, the route that packets follow in order to reach the entry point of the Nested NEMO (i.e. the Gateway-MR) must also be taken into consideration. Principally this means addressing the problem of Pinball Routing and thus aiming to ensure that packets are only ever transmitted via one HA.

3. NEMO+ PROTOCOLS

The NEMO+ suite is a trio of protocols that were designed to interoperate in order to optimise the delivery of packets in Nested NEMO network configurations. The interoperation

of these protocols results in a three phase procedure that is illustrated in Figure 2. Phase 1 depicts the flow of messages generated by the first protocol, Tree Discovery (TD) out of the MRs' Ingress interfaces. Phase 2 illustrates the subsequent flow of messages generated by the second NEMO+ protocol, Network In Node Advertisement (NINA) out of the MRs' Egress interfaces. Finally phase 3 illustrates the direct establishment of MR 4's MR-HA tunnel using the third protocol, Reverse Routing Header (RRH). Each of the NEMO+ protocols operate independently, and rely upon each other to varying degrees in order to produce the most optimised outcome.

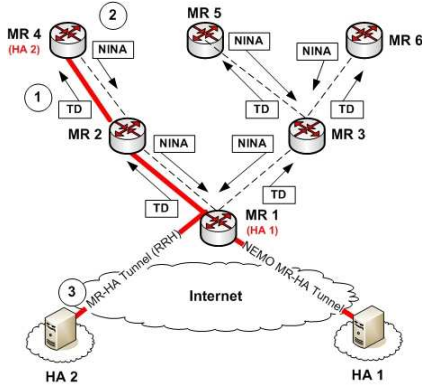


Figure 2: TD, NINA and RRH Interaction

3.1 Tree Discovery (TD)

The Tree Discovery (TD) protocol [3] forms the foundation which the other NEMO+ protocols operate on top of. TD constitutes the first phase of the operation of NEMO+. By augmenting the IPv6 Neighbor Discovery (ND) [4] Router Advertisement (RA) messages that a MR transmits, TD allows NEMO MRs to distribute information to any other NEMO MR that connects to its Ingress interface about any Nested NEMO networks they are connected to. TD achieves this by carrying the relevant information in an ICMPv6 Option known as a Tree Information option (TIO), which once included in an RA is referred to as a RA+TIO. The information carried in these messages ensures that NEMO MRs can make an informed decision when determining whether to connect to another NEMO MR, and it also helps prevent loops from forming in the Nested NEMO. Loop avoidance is a very important concern when forming Nested NEMO networks; when several MRs attach to each other to form a Nested NEMO, loops can be created if they are not explicitly avoided. In the simplest case, when the Egress and Ingress interfaces of a NEMO MR are all wireless, a NEMO MR may be listening to RAs from its own ingress interface, creating a conflict problem as a NEMO MR will essentially attempt to access the Internet via itself. Generally, the arbitrary attachment of NEMO MRs will ultimately form graphs that are not exempt of loops. For example consider a situation where MR1 has a direct connection to the Internet (it is a Gateway-MR) and separately, MR3 is attached to MR2. If MR2 is in a position whereby it can communicate with both MR3 and MR1 over its wireless egress interface then it must make a decision as to which access network to connect to. If MR2 selects MR1, then connectivity to the Internet will be provided for all. However, if MR2 selects MR3, then MR2

and MR3 will end up forming a loop and will therefore be disconnected from their Home Agents. With NEMO BS, a NEMO MR has no means of making an informed decision as to which access router it should select in order to avoid a loop from forming, TD provides MRs with this information. In addition to loop avoidance, TD also carries other information to help MRs make an informed connection decision based on other criteria such as whether a Nested NEMO currently has Internet connectivity, or how deep an individual MR is in a Nested NEMO (i.e. its distance in hops to the Internet) or the bandwidth capabilities of the MRs Internet link.

3.2 Network In Node Advertisement (NINA)

The Network in Node Advertisement (NINA) protocol [5] works in direct partnership with TD. Its main role is to ensure that routes to all of the mobile networks throughout a Nested NEMO are available within the tree structure that TD creates. To achieve this, upon receiving a RA+TIO, a MR running the NINA protocol responds to the source of the RA+TIO with a NINA message which details all of the prefixes that the MR currently maintains. When the source of the RA+TIO receives the NINA response it too transmits a NINA message (if it is not a Gateway-MR) to the MR it is currently attached to, containing both the prefixes it maintains and the prefixes of the MR that just attached to it. This process is then repeated upwards by all MRs until the top of the Nested NEMO tree (the Gateway-MR) is reached. As with TD, NINA augments the functionality of IPv6 ND. The NINA message itself is an IPv6 Neighbor Advertisement (NA) message with an additional ICMPv6 Option in it known as a Network In Node Option (NINO). NINA extends NA messages in this manner because of the inherent behavior of NEMO MRs. A NEMO MR behaves as a typical IPv6 router in relation to its Ingress interfaces (i.e. the interfaces that other nodes connect to in order to attain Internet access). Hence this is the reason TD augments RA messages. However, a NEMO MR presents itself as an individual host over its Egress interface (i.e. the interface that the NEMO MR itself uses to obtain Internet connectivity) and therefore only advertises itself using NA messages over this interface. By carrying additional information about network prefixes it can reach in these NA messages, the NEMO MR is advertising reachable networks in a single node advertisement (hence the name NINA).

The route propagation model employed by NINA produces an efficient update technique which hides any topology changes in the Nested NEMO tree that happen beneath a MR (i.e. the mobility of a MR's sub tree). A movement happening within a sub-tree is hidden from the parent of the sub-tree root as long as the route simply remains reachable via the MR's Ingress interface. As a result, MRs that are located deeper in a tree (further away from a Gateway-MR) maintain fewer routes to prefixes but observe more MR mobility, whereas MRs closer to a Gateway-MR maintain routes to more prefixes but observe less MR mobility.

3.3 Reverse Routing Header (RRH)

Once TD and NINA have configured an optimised tree-based layer 3 routing model on to the Nested NEMO, and the appropriate routes have been disseminated throughout, the Reverse Routing Header (RRH) protocol [6] is introduced to ensure packets are routed efficiently beyond the

Gateway-MR and into the Internet in general. When a MR attaches to a Nested NEMO it configures its CoA from the MNP of the MR it attaches to, even with TD and NINA in place MRs will still inherently build their MR-HA tunnels via the HAs of the other MRs in the Nested NEMO as described in Section 2.1. To prevent this inefficient Pinball Routing from occurring, RRH allows MRs within a Nested NEMO to update their HAs with the actual location of the Gateway-MR, to ensure packets can be more directly delivered to their current location in the Internet. To achieve this, RRH applies a technique which combines a source routing approach within the Nested NEMO network with traditional IP routing used beyond the Gateway-MR in the Internet. With RRH enabled, any MR within a Nested NEMO that isn't a Gateway-MR will insert the Reverse Routing Header into any packets that either it or its attached hosts generate. The RRH is placed into the packet after the outer IPv6 header and behaves like a multihop version of the MIPv6 Type 2 routing header, because it holds multiple IPv6 addresses which are stored or utilised at different stages of the Nested NEMO. Essentially, the RRH is used to record the route that packets traverse (i.e. the CoAs of every MR) as they travel towards the Internet through a Nested NEMO. To do this, each MR that the packet travels through overwrites the source field of the outer IPv6 header and records the existing source address that it overwrites in the RRH. By following this procedure at every MR in the Nested NEMO, the packet leaves the Gateway-MR with the topologically correct CoA of the Gateway-MR as its source address and with an RRH that describes the route that can be taken back through the Nested NEMO to reach the actual source MR directly. The HA then stores this information and subsequently sets the destination of the outer IPv6 header to the CoA of the Gateway-MR and inserts the same RRH back into packets it transmits from a CN to the MR in the Nested NEMO. This way packets are first delivered to the Gateway-MR's CoA and then routed to their ultimate destination based on the CoAs recorded in the RRH.

However, RRH's source routing approach within the Nested NEMO is only required if an appropriate NINA route is not already in place. If NINA is supported in the Nested NEMO then its routes can be relied upon and the CoA of each MR does not need to be recorded. Therefore, in a Nested NEMO where NINA is supported by all of the MRs that a packet traverses out to the Internet, the principal role of the RRH protocol is to record the actual CoA of the originating MR and the topologically correct CoA of the Gateway-MR and report it to the MR's HA. Again, the return flow of communication from the HA to the MR will be transmitted directly to the Gateway-MR, at which point the originating MR's actual CoA will be inserted as the destination address and the packet will be routing using the NINA routes as opposed to being source routed.

4. EXPERIMENTAL EVALUATION

We have developed experimental implementations of the NEMO+ trio of protocols for both Linux and for Cisco IOS. In order to test the capabilities of our implementations and ensure their ability to interoperate we carried out a number of different roaming procedures in a testbed, consisting of both Linux PCs running Ubuntu 7.10 (Kernel version 2.6.22) and Cisco 3200 Series Mobile Access Routers (MARs). The testbed we configured and the roaming procedure we fol-

lowed are depicted in Figure 3. As illustrated the testbed consisted of 2 HAs and 4 MRs, where MR 2 and MR 3 were Linux PCs and MR 1 and MR 4 were Cisco MARs; each MR's relative HA is highlighted below its title. All of the connections used in the testbed were formed using Ethernet cabling except when handover testing was performed, in which case the final connection between the roaming network and its destination network was altered to use 802.11g. We chose to use a wired approach to ensure that a reliable, consistent connection existed and therefore any comparative differences that we observed during our performance tests could be attributed to the mobility protocols and the resulting network configuration rather than an unexpected external event/interference. The testing was carried out in a number of stages, in Figure 3 the numbered arrows depict the movement performed at each of these stages:

- Stage 1: MR 4 roams away from a direct connection to the Internet and obtains connectivity via MR 2, which is part of a Nested NEMO network.
- Stage 2: MR 4 again roams, this time by detaching from the Ingress interface of MR 2 and reattaching via MR 3.
- Stage 3: MR 1 (the Gateway-MR) roams from Access Network 1 to Access Network 2, which in turn will consequently effect the connectivity of the MRs connected in the Nested NEMO behind it.
- Stage 4: MR 3 (with MR 4 still connected to its Ingress interface), connects via MR 2. This results in a single branch topology of MRs which are connected as follows: MR 1->MR 2->MR 3->MR 4.

For each stage in our experimental evaluation we first recorded the latency experienced in the network, to do this we used the Ping6 utility to collect multiple sets of 5000 Round Trip Time (RTT) measurements, to obtain our average latency value. Once the latency tests were completed the available bandwidth was determined by recording the TCP throughput rate that was achieved by the iPerf bandwidth measurement tool. Then finally we would roam the appropriate mobile network according to the aforementioned testing stages by performing an 802.11g handover, this roaming procedure was then repeated 20 times to produce an average handover value. For all of the experiments, the RA output interval on all routers (static and mobile) was set to 500ms. All of these tests were carried out between two Linux laptop PCs (MNN 1 and CN 1), this communication represents end-to-end packet transfer between a node connected to a mobile network (Mobile Network Node (MNN)) and a node located in the Internet (Correspondent Node (CN)).

4.1 NEMO Basic Support Testing

To act as a comparative baseline, we began by configuring all of the MRs in our testbed to operate using the NEMO Basic Support protocol. This network configuration therefore results in a Nested NEMO topology that is supported using multiple layers of tunneling (MR-HA tunnels within MR-HA tunnels) and Pinball Routing via numerous HAs. We present all of the results from our experimental evaluation in Table 1. Before describing the outcome of the performance tests of each individual phase it was possible during the testing to identify a general trend with regards

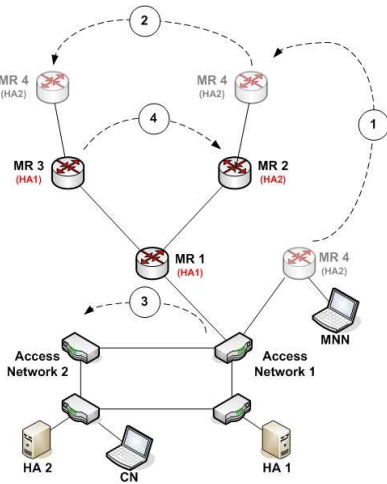


Figure 3: Testbed Configuration

to all of the handover tests performed. What was evident during testing and is also evident from our results was that the resulting complexity of the configuration of our testbed in any of the testing stages had no distinguishable effect on the average handover times we recorded. The latency times experienced by an MR when performing its Binding Update (BU) process have the ability to effect the overall time it takes for an MR to complete its registration process. However in our testbed the latency times were so low relatively that they were insignificant in comparison to the combined time it took the layer 2 handoff and the MR to subsequently configure its CoA to complete.

The latency and throughput tests carried out in stage 1 (before MR 4 roams into the Nested NEMO) represent the performance of the mobile network in a standard NEMO configuration (i.e. where the MR has a direct connection to the Internet). Consequently the performance results recorded in this state highlight the maximum performance bounds achievable in our testbed setup. The roaming procedure carried out in Stage 1 puts MR 4 into the Nested NEMO for the first time. In stage 2, MR 4 will register with its HA via HA1 and every packet transmitted out of its network will have a further three IPv6 headers attached to it because of the multiple levels of tunneling employed. These inefficiencies are immediately evident in the performance results, where we recorded more than a doubling in the latency time experienced and a halving of the rate of throughput achievable. In stage 3, packets transmitted out of MR 4's network will travel via MR 3 first (which is registered with a different HA to MR 2) instead of MR 2 (which is registered with HA 1). Whilst this results in a change in the overall process performed, the end-to-end path actually remains the same, with the only difference being that HA 2 rather than HA 1 has to perform double decapsulation/encapsulation of packets. This fact is reflected in the results where we recorded almost identical performance characteristics as in the previous stage. Finally in stage 4, the roaming of MR 3 and MR 4 to behind MR 2 results in a branch that is 4 MRs long. This means that packets sent from MNN 1 are encapsulated in 4 tunnels in total and transmitted via a pinball route that results in the packet travelling first via HA 1, then HA 2 and then back to HA 1 and then finally back to HA 2 before it is fully decapsulated out of all of the nested

tunnels. As would be expected, this increase in encapsulation and the distance of the overall end-to-end path has dire consequences for the recorded performance, with throughput dropping to less than a quarter of that attainable in a non Nested NEMO scenario and the latency increasing by almost a factor of 4.

4.2 NEMO+ Testing

We then performed the same testing procedure with the same testbed but configured to support our NEMO+ protocols. Since stage 1 represents the operation of a MR in a non Nested NEMO scenario, and the NEMO+ protocols only engage in Nested NEMO environments, these results remain unchanged from the NEMO BS testing phase. In stage 2 a stark improvement is already noticeable, with latency increasing by less than 1 second on the non Nested NEMO scenario and throughput only dropping by 300kbps. Again, in stage 3 the resulting network topology imposes exactly the same end-to-end path as in stage 2 (as was the case in the NEMO BS phase of testing), the results reflected this. Finally, the efficiency gains achievable using NEMO+ protocols in contrast to NEMO BS become most apparent in stage 4. Whereas with NEMO BS the network performance of MR 4 suffered a degradation of service by around a factor of 4, using the NEMO+ protocols only introduces a single additional hop to the end-to-end path, since packets will travel via the most direct route. As the results demonstrate, this culminates in throughput and latency readings that are only fractionally worse than those recorded in stages 2 and 3.

Stage	NEMO BS			NEMO+		
	HO	Lat	E/W	HO	Lat	E/W
1	1.69 S	2.806 Ms	13.1 Mbps	1.29 S	2.806 Ms	13.1 Mbps
2	1.44 S	5.984 Ms	6.44 Mbps	1.43 S	3.787 Ms	12.8 Mbps
3	1.52 S	5.985 Ms	6.44 Mbps	1.48 S	3.787 Ms	12.8 Mbps
4	1.38 S	9.693 Ms	3.24 Mbps	1.71 S	3.903 Ms	12.7 Mbps

Table 1: Testing Results Summary

4.3 Convergence Simulations

In addition to the experimental evaluation we performed using our multi-platform testbed, we also implemented the TD and NINA protocols in a simulation environment in order to gather simulated results of the performance of these protocols. Where in the experimental evaluation we used our testbed to compare handover, throughput and latency characteristics in NEMO+ enabled networks, our simulation work was designed to highlight the convergence times that could be expected in networks using these protocols. Our simulations were carried out on the OMNeT++ Discrete Event Simulation Platform. Once we had augmented the base functionality of the OMNeT++ code to support the TD and NINA protocols, we facilitated the convergence simulations by also incorporating Gateway-MRs. Using varying combinations of MRs and Gateway-MRs, a number of simulation runs were performed which illustrate the convergence properties of these protocols in large scale networks.

In the first set of simulation runs we carried out, we wanted to determine the effect that the number of MRs had on the propagation of MNP routes in a Nested NEMO network. To carry out these tests we recorded the length of time it took for either 1, 2 or 3 Gateway-MRs to establish routes to every mobile network in simulation tests consisting of from 25 up to 400 MRs. The results of these simulations are illustrated

in Figure 4. These results demonstrate that the stabilisation time of NINA scales almost linearly with the number of MRs present.

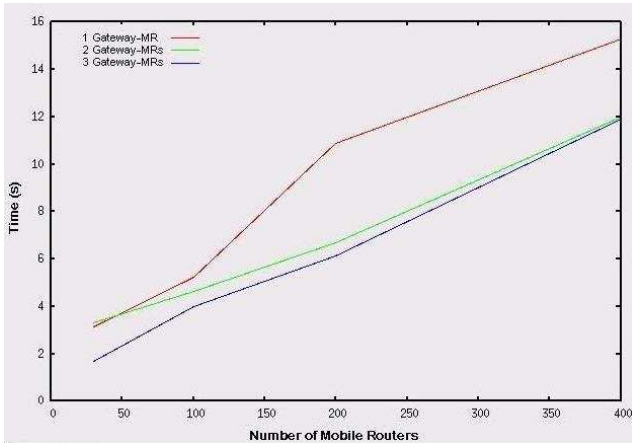


Figure 4: Nested NEMO Stabilisation Time

In the next set of simulation runs we performed five separate tests. In each test 400 MRs converged and propagated routing information (each was assigned a single MNP) with reachability to a varying number of Gateway-MRs (1,2,3,4 and 5 Gateway-MRs in total). For each of the tests, again the number of routes to the MNPs reachable from the collective of Gateway-MRs was recorded over time to determine the point at which NINA had stabilised and could provide routes to every mobile network. Figure 5 illustrates how in the best case scenario where 5 Gateway-MRs were reachable, routes were available to all MRs within 4 seconds, whilst in the worst case scenario where only 1 Gateway-MR was available all 400 MRs converged in just under 15 seconds.

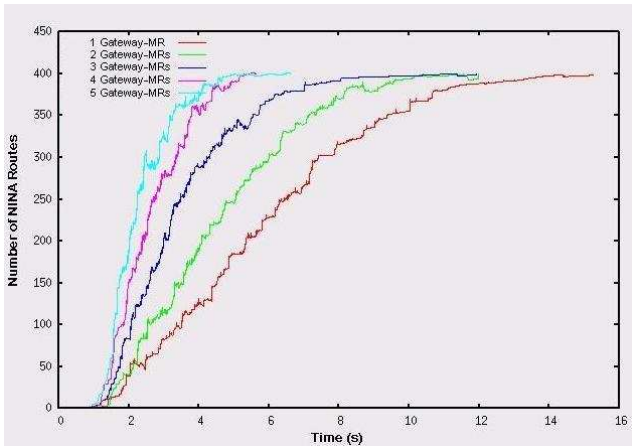


Figure 5: Route Propagation (400 MRs)

5. CONCLUSION

In this paper we have provided an introduction to the notion of Nested NEMO networks and the inefficiencies that are incurred if the NEMO Basic Support protocol is relied upon to support these network topologies. In addition we present an overview of the NEMO+ suite of protocols which have been designed specifically to provide improved support for Nested NEMO networks, to ensure communication can

occur in a much more efficient and direct manner in these scenarios. One of the key design principals adhered to for two of these protocols (TD and NINA) is that they were both developed to be extensions of the existing Neighbor Discovery (ND) process that NEMO MRs perform as part of their Basic Support procedure. By augmenting the Router Advertisement (RA) and Neighbor Advertisement (NA) message transfers that a NEMO MR performs, TD and NINA are able to operate efficiently without generating any additional distinct message transfer and therefore benefiting from reuse of the existing communication that already occurs. This ND based approach also facilitates a simplified and therefore quicker development time, since working implementations need only augment existing protocol code rather than be developed from the ground up. Using a testbed consisting of NEMO+ enabled Cisco IOS and Linux Mobile Routers we carried out an experimental evaluation of our implementations that demonstrated the interoperability of our implementations and the performance improvements that can be provided over NEMO BS. We also presented simulation work we have carried out in order to determine the convergence times of large scale deployments of the NEMO+ protocols in Nested NEMO scenarios. In our experimental evaluation we were able to show the stark performance improvements that NEMO+ can provide over NEMO BS, but more importantly we were also able to show how the introduction of increasingly complex Nested NEMO configurations that resulted in severe service degradation using NEMO BS only marginally affected performance when the NEMO+ protocols were enabled. In addition to the positive results we received from our experimental evaluation we were also encouraged by the potential scalability of the TD/NINA approach that was highlighted by our simulation runs. The simulations provided results for scenarios involving up to four hundred inter-connected Mobile Routers, which would be considered as an extremely large Nested NEMO. Therefore total convergence times of four seconds in these environments should certainly be satisfactory. To summarise, we believe that NEMO+ is an efficient, lightweight approach to supporting Nested NEMO networks that is specifically tailored to the inherent operation of a NEMO MR. By augmenting existing message transfers and intelligently utilising any available information, NEMO+ is able to provide clear benefits with comparatively little additional overhead.

6. REFERENCES

- [1] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert. "NEMO Basic Support Protocol". IETF RFC 3963, January 2005.
- [2] J. Arkko D. Johnson, C. Perkins. "Mobility Support for IPv6". IETF Request For Comments 3775, June 2004.
- [3] P. Thubert, C. Bontoux, and N. Montavont. "Nested Nemo Tree Discovery". IETF Draft (Work In Progress), July 2007.
- [4] T. Narten, E. Nordmark, W. Simpson, and H. Soliman. "Neighbor Discovery for IP version 6 (IPv6)". IETF RFC 4861, September 2007.
- [5] P. Thubert, R. Wakikawa, C. Bernardos, R. Baldessari, and J. Lorchat. "Network In Node Advertisement". IETF Draft (Work In Progress), February 2007.
- [6] P. Thubert and M. Molteni. "IPv6 Reverse Routing Header and its application to Mobile Networks". IETF Draft (Work In Progress), February 2007.