Henrik Petander, Eranga Perera, Kun-chan Lan, Aruna Seneviratne, Measuring and improving the performance of network mobility management in IPv6 networks, in IEEE Journal on Selected Areas in Communications (JSAC), vol. 24, number 9, pp. 1671-1681, September 2006.

Measuring and Improving the Performance of Network Mobility Management in IPv6 Networks

Henrik Petander, Eranga Perera, *Student Member, IEEE*, Kun-Chan Lan, *Member, IEEE*, and Aruna Seneviratne, *Member, IEEE*

***Abstract*—Measuring the performance of an implementation of a set of protocols and analyzing the results is crucial to understanding the performance and limitations of the protocols in a real network environment. Based on this information, the protocols and their interactions can be improved to enhance the performance of the whole system. To this end, we have developed a network mobility testbed and implemented the network mobility (NEMO) basic support protocol and have identified problems in the architecture which affect the handoff and routing performance. To address the identified handoff performance issues, we have proposed the use of make-before-break handoffs with two network interfaces for NEMO. We have carried out a comparison study of handoffs with NEMO and have shown that the proposed scheme provides near-optimal performance. Further, we have extended a previously proposed route optimization (RO) scheme, OptiNets. We have compared the routing and header overheads using experiments and analysis and shown that the use of the extended OptiNets scheme reduces these overheads of NEMO to a level comparable with Mobile IPv6 RO. Finally, this paper shows that the proposed handoff and RO schemes enable NEMO protocol to be used in applications sensitive to delay and packet loss.**

***Index Terms*—Handoffs, mobile router (MR), network mobility, route optimization (RO).**

I. Introduction

WITH THE ALMOST ubiquitous availability of computing and wireless communication capability in most electronic devices, the prediction that most devices will be connected to a network is fast becoming reality. An emerging form of this ubiquitous connectedness is vehicle networks, especially in public transport systems, which will enable groups of people to access network services, while on the move. In these environments, use of a dedicated device, a mobile router (MR), reduces the required complexity of the end devices, and provides numerous opportunities for optimizing the performance and operational costs. The IETF Network Mobility working group has standardized the network mobility (NEMO) basic support protocol [1] in which a MR manages the mobility of a moving network.

The performance of a moving network depends on the performance of the MR and the overhead of the network mobility management protocol. Hence, it is important to understand the impact of handoffs and protocol overhead in moving networks. Although numerous research activities have focused on the effects of Internet protocol (IP) extensions for providing support for host mobility, up to now there has been no systematic study of the performance of MRs and the network mobility management protocols. To address this, we have designed and implemented a testbed [2] and measured the performance of mobility management protocols in network mobility settings. In this paper, we use the testbed to measure and analyze the handoff performance with NEMO. We adapt two handoff performance enhancements, fast router advertisements (RAs) and optimistic duplicate address detection (DAD) to minimize handoff latency for mobile networks. Then, we show that the handoff performance of NEMO, even with these optimizations, is still not sufficient for performance-critical network applications, such as voice-over-IP. To overcome this, we propose a make-before-break (MBB) handoff scheme. We analyze its performance, including interference between the network interfaces and management of the NEMO protocol state. Through the analysis, it is shown that the proposed scheme enables lossless handoffs between networks with overlapping coverage area. In addition to studying the handoffs in NEMO, we perform extensive measurement and analysis of the protocol and routing overheads of NEMO in static and mobile scenarios. Again, the analysis is used to show that it is necessary to reduce the overheads of the protocol. We reduce these overheads by extending the OptiNets [3] protocol.

In summary, the contribution of this paper is threefold.

1) We design and implement a network mobility testbed for analyzing the performance of NEMO.
2) We propose a novel MBB handoff scheme which enables lossless handoffs between networks with overlapping coverage areas.
3) We study the overheads in NEMO signaling and routing via extensive measurement and analysis and show that these overheads can be minimized by using the extended OptiNets route optimization (RO) scheme.

The rest of this paper is organized as follows. In Section II, we discuss network mobility management with the NEMO protocol. In Section III, we present the design and implementation of the network mobility testbed. In Section IV, we propose and analyze MBB handoffs for NEMO. Reducing NEMO overheads using the extended OptiNets scheme is presented in Section V. This is followed by related work and the conclusions in Sections VI and VII, respectively.

II. Network Mobility Management

There are broadly two methods of providing mobility support, namely through redirection and indirection. A well-known

Manuscript received May 30, 2005; revised February 2, 2006.

The authors are with the National ICT Australia, Eveleigh, NSW 1430 Australia (e-mail: henrik.petander@iki.fi; eranga.perera@nicta.com.au; kun-chan.lan@nicta.com.au; aruna@nicta.com.au).

Digital Object Identifier 10.1109/JSAC.2006.875116

redirection scheme is the session initiation protocol (SIP) [4]. The indirection-based schemes use network agents to transparently reroute information. Mobile IPv6 (MIPv6) [5] and its variants are examples of schemes that use indirection. In this paper, we will focus on indirection schemes based on MIPv6 and NEMO. This section describes how NEMO manages the mobility of a moving network and presents a theoretical analysis of the handoff performance with NEMO and the communications overhead of using NEMO.

A. NEMO Operation

NEMO allows a MR to manage the mobility of the nodes inside a mobile network which are known as mobile network nodes (MNNs) with the help of a fixed mobility anchor point, home agent (HA). When an MR is in its home network, it is connected directly to its HA, so that all traffic to and from the mobile network is delivered via the HA and the MR. The mobile network is connected to the Internet via an IP-IP tunnel between the MR and the HA when the MR is away from home.

When a MR moves to a new network, it obtains a care-of-address (CoA) and sends a binding update (BU) to its HA. The BU binds the new CoA of the MR with its permanent address (home address). The HA sends a binding acknowledgement (BA) to inform the MR of the status of the update. A tunnel is then established between the CoA of the MR and the address of the HA. The MR and its HA then deliver all traffic between the mobile network and the Internet via this tunnel. This overlay routing hides the mobility of the MR from the CNs and also from the MNNs. Thus, the MNNs do not need any mobility management capabilities to take advantage of the mobile Internet access.

A MNN, which is not capable of managing its own mobility is known as a local fixed node (LFN). However, a mobile device managing its own mobility may enter a mobile network, treating it as a foreign network in which case the MNN is known as a VMN. An example of this is a passenger with a MIPv6 capable mobile device entering a train with a mobile network. In this case, the MIPv6 VMN will send a BU to its own HA $(HA_{\mathbf{VMN}})$ informing it to deliver all traffic to its new CoA using IPv6 tunneling. This results in two, nested levels of mobility management since a MR manages the mobility of the mobile network. However, the VMN can use MIPv6 RO to communicate more directly with CNs bypassing the $HA_{\mathbf{VMN}}$ using its CoA from the mobile network prefix.

B. Handoffs With NEMO

The handoff processs in an IPv6 network mobility setting can be divided into three main parts.

1) Link-layer handoff, in which the MR finds a new access point (AP) and associates with it. Thus, the link-layer handoff latency depends on the time it takes for the network interface to find a new AP and associate with it. This latency depends on the network technology.
2) IPv6 network attachment follows the link-layer handoff. Network attachment of the MR consists of router discovery and CoA configuration. In router discovery, the MR sends a router solicitation (RS) and receives a RA from a new access router. The access router waits a random delay before sending the RA message. This random delay is between 0–500 ms, so the average delay for receiving a RA is 0.25 s. The total delay of router discovery consists of the round-trip time (RTT) between the MR and the access router and the random delay. After discovering the access router, the MR acquires a new CoA from the foreign network, using either IPv6 stateless address autoconfiguration or a stateful mechanism, such as dynamic host configuration protocol (DHCP). The configuration of a new CoA requires the MR to ensure that the address is unique. In IPv6 this is done using the DAD procedure. If the MR uses the standard DAD procedure, it needs to wait for the procedure to finish before it can use the address and register its new CoA with its HA. The latency created by DAD is configuration dependant, and involves a random delay between 0–1 s. Minimum latency for the whole DAD procedure varies between 1–2 s, with an average of 1.5 s.
3) NEMO home registration latency, which represents the delay of the MR sending a BU to its HA and the HA replying with a BA. This consists of the propagation delays of the messages and the HA processing delays. The HA processing delay is dependent on the need for the HA to perform proxy DAD. Proxy DAD is performed only if the MR has a home address from a physical home link to guarantee that the home address is not used by another node on the link. Proxy DAD takes a minimum of 1 s.

TABLE I
THEORETICAL MINIMA FOR NETWORK-LAYER HANDOFF LATENCIES WITH NEMO

Handoff type	**NEMO without optim.**	**With Fast RAs**	**With ODAD**	**With Fast RAs and ODAD**
Home-Foreign	2.75s	2.5s	1.25s	1s
Foreign-Foreign	1.75s	1.5s	0.25s	0s
Foreign-Home	0.25s	N/A	0.25s	0s

Of the above factors, the network attachment latencies are independent of the access technology and network topology, and we use two techniques to minimize these latencies. The random delay associated with router discovery can be removed by using the fast RA mechanism proposed in [6]. The DAD delay can be mitigated by using optimistic DAD (ODAD) [7]. The theoretical handoff latencies given in Table I are derived and explained in Appendix A. The table does not include the link-layer and NEMO signaling latencies.

It is evident from the above analysis, that the use of protocol optimizations, such as fast RAs and ODAD, reduces the network attachment latency substantially. However, we still need to address the link-layer handoff and the NEMO signaling latencies which have a significant impact on the handoff performance. This will be discussed further in Section IV-A.

C. Overhead of Using NEMO

In NEMO, the MR uses an overlay route via a fixed anchor point to hide the mobility from the nodes in the mobile network. This overlay routing leads to less than optimal routing and adds a protocol header overhead to every packet. In addition to the protocol header overhead for data packets, NEMO also incurs a

TABLE II
OVERHEADS OF NEMO LFN, VMN, AND MIPV6 MN IN BYTES

Mob. man scheme	Per packet OH	Signaling OH/s	Signaling OH per handoff
NEMO LFN	40	N/A	208
NEMO VMN w/o MIPv6 RO	80	N/A	208
NEMO VMN with MIPv6 RO	64	1.66	208
MIPV6 MN with RO	24	1.33	424

signaling overhead between the MR and its HA every time the MR performs a handoff.

Use of NEMO introduces an overhead to each packet which a MNN and a CN exchange when the MR is in a foreign network. The overhead is caused by the IPv6 tunneling and it is 40 bytes for every packet. The signaling overhead of NEMO with LFNs is caused by the BU-BA exchange between the MR and its HA. The size of these messages depends on how IPsec is used to protect them [8]. The total mobility management protocol header overhead will be larger, if the MNN is a MIPv6-capable VMN which uses MIPv6 to guarantee session continuity and reachability. This leads to higher protocol overheads and also potentially inefficient routing. If the VMN uses RO with the CN, the per packet overhead will be reduced. However, RO requires extra signaling between VMNs and CNs. The per packet overhead, per MNN signaling overhead, and handoff related signaling overhead are presented in the Table II. The values in the table are derived in Appendix B.

In addition to reducing the payload size available to applications, NEMO also introduces an extra routing leg to the routing path between MNNs and CNs. The effect of routing packets via a HA depends largely on the network topology. If foreign networks are topologically close to the home network and the HA, then the effect may be negligible, but in the case of intercontinental mobility the effects may be large on applications sensitive to the RTT, even if the long routing legs have no other effects, such as packet reordering, packet loss, or packet duplication.

III. NETWORK MOBILITY TESTBED

A. Testbed Architecture

Our testbed consists of three logical parts: 1) the wide area network connecting the access routers, the $HA_{\mathbf{VMN}}$ and $HA_{\mathbf{MR}}$, and the CN; 2) the wireless access network consisting of two foreign access networks and a home access network; and 3) the mobile network which consists of the MR and a MNN connected to the MR via a local area network (LAN). The MNN can act either as a LFN or as a VMN. This logical topology of the test network is shown in Fig. 1. Three IEEE 802.11 b APs are used for wireless access. We use NISTNet [9] to emulate the Internet by introducing network latency between the nodes.

B. NEMO Implementation

The NEMO implementation is based on the MIPL Mobile IPv6 implementation by the Helsinki University of Technology [10]. It consists of NEMO-based MR and NEMO capable

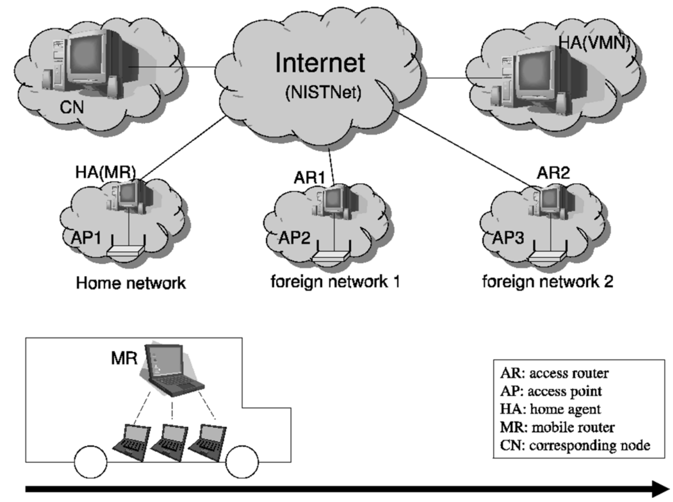

Fig. 1. Logical network topology of the testbed.

HA prototypes for testing and measuring the performance of NEMO and its extensions. The MR uses the information from the link-layer to trigger handoffs when it moves to a new wireless network.

C. Hardware and Software Configuration

Our hardware consisted of six desktop computers and five laptops with processor speeds between 350 MHz and 3 GHz and memory sizes between 128 MB and 512 MB. We used Cisco 1200 series WLAN APs and an integrated Intel IPW2100 card, a PCMCIA Lucent silver card and a Demarctech Prism 2.5-based PCMCIA card for our IEEE 802.11 b wireless access network.

The use of 802.11 b access networks for experiments in this paper affects the results of the experiments to some degree. However, since our analysis consists of comparison studies between different schemes for handoffs and routing, this isolates the effects of the access technology.

Our testbed used three modified software components in addition to standard IPv6 capable Linux operating system and a NEMO MR and a HA. First, we used a modified *radvd* daemon [11] developed at Monash University to send fast RAs [6] in all the experiments. Second, a modified DHCPv6 client and server were used in the MR and the access router to achieve the extended OptiNets RO described in Section V. Finally, the CN, the VMN, and the $HA_{\mathbf{VMN}}$ were equipped with MIPv6 capabilities to test nested mobility management. We used MIPL 2.0 for MIPv6 CN and $HA_{\mathbf{VMN}}$ and a modified VMN based on MIPL 1.1.

D. Experiment Setup

We used *iperf* [12] to generate and measure user datagram protocol (UDP) and transmission control protocol (TCP) traffic and *ethereal* [13] to capture packet traces. To understand the effect of NEMO handoffs on real-time traffic, we generated UDP traffic with a small packet size. Additionally, we measured TCP traffic to analyze the effect of handoffs, protocol header, and routing overhead on bulk and interactive TCP traffic. We used the default maximum window size in Linux for TCP, 16 Kbytes, for all measurements. We used downstream traffic in all handoff

TABLE III
NETWORK LATENCY BETWEEN NODE PAIRS SET BY NISTNET

Node Pair	Network Latency
MR to AR	2 ms
AR to $HA_{\mathbf{MR}}$	10 ms
AR to $HA_{\mathbf{VMN}}$	50 ms
AR to CN	40 ms
$HA_{\mathbf{MR}}$ to CN	40 ms
$HA_{\mathbf{VMN}}$ to CN	40 ms
$HA_{\mathbf{MR}}$ to $HA_{\mathbf{VMN}}$	40 ms

experiments, i.e., CN acted as the source of traffic and MNN as the sink which had an effect on packet loss.[1]

Table III contains the one-way latencies between the nodes in our testbed. The wireless links between the MR and the access routers in the testbed showed a latency of 2 ms consistently in all experiments. We experimented with multiple values for access router—$HA_{\mathbf{MR}}$ latency, and chose a relatively small value for the access router—$HA_{\mathbf{MR}}$ latency to emulate the case where the HA is in the network of the same ISP. The effect of this latency on the overall handoff time was as described in Section II-B. We also experimented with other emulated latencies and noticed that these did not have a measurable effect on the handoff performance. However, the end-to-end network latency does have an effect on TCP performance of the MNNs. Therefore, we have measured the effect of different end-to-end network latencies on NEMO routing performance in Section V.

IV. MBB HANDOFF MECHANISM FOR LOSSLESS NEMO HANDOFFS

There have been proposals for reducing MIPv6 and NEMO signaling latencies [14], [15], but these proposals have been based on the assumption of the MR being connected to only one AP at a time. When the MR can connect to only one AP, it is forced to break the connection to its current network before reattaching itself to a new network. With this type of handoff, referred to as a BBM handoff, packet loss is hard to eliminate completely. However, if it is possible to simultaneously listen to multiple APs, the MR could establish a connection to the new network before breaking its current connection, thus mitigating or reducing the impact of handoff latency. This could be done by equipping the MR with multiple interfaces. We propose the use of two interfaces to enable MBB handoffs for reducing packet loss due to handoff latency.

In the proposed scheme, one interface is used for data communication, and the other is used for scanning for networks which can provide better connectivity. Once a network with better connectivity is found, the scanning interface takes over the data transmission, and the other reverts to a scanning role. This, as well as being access technology independent, allows lossless handoffs with uninterrupted connectivity for data communications since the MR maintains its connection to the old network using one interface, while performing a handoff to a new network using the other interface.

[1]The use of upstream traffic would have resulted in packet loss ending with the MR sending a BU in foreign-to-foreign and foreign-to-home handoffs, instead of the HA receiving the BU.

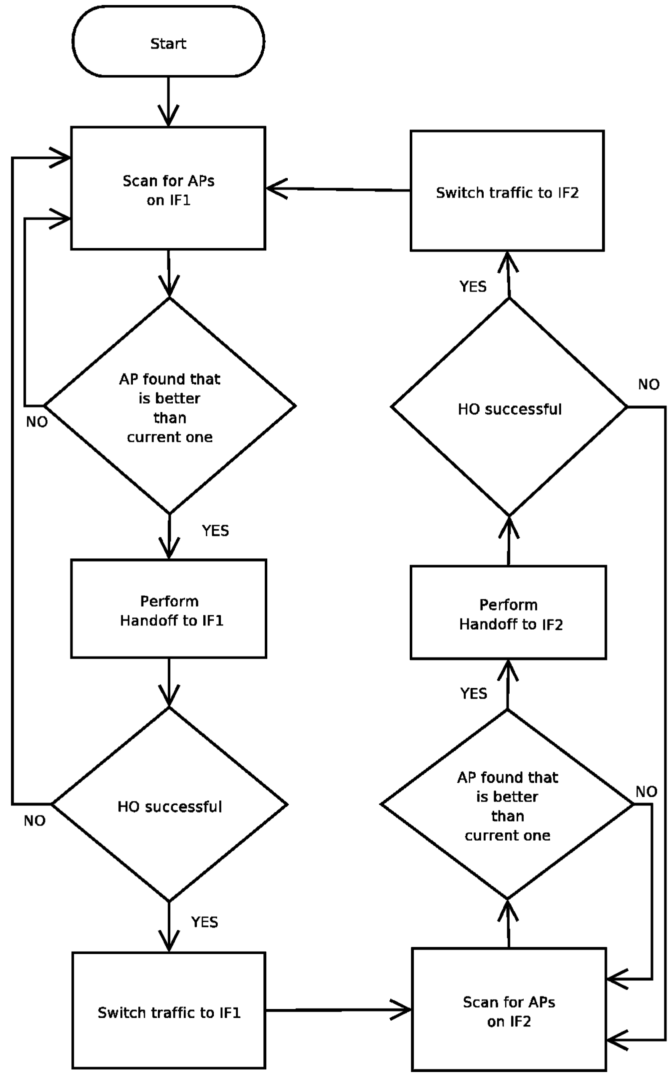

Fig. 2. MBB handoff algorithm with two interfaces.

A. MBB Handoff Algorithm Using Two Network Interfaces

The proposed MBB handoff scheme uses the algorithm in Fig. 2. The handoff decision can be made using techniques such as signal-to-noise ratio comparisons [16] combined with movement prediction algorithms [17]. The proposal is to compare the signal strength of the candidate AP with the current one. If the difference is greater than a threshold value, MR performs a handoff to the new network. An ideal threshold value would be high enough to prevent ping-pong movement, but still trigger handoffs early enough to prevent packet loss. A dynamically adaptive algorithm for choosing and adjusting the threshold value would allow a MR to make more optimal handoff decisions and avoid fluctuations between APs.

Using the algorithm, it is possible to perform completely lossless handoffs, provided that the coverage of the old access network and the new access network overlap sufficiently and the handoff decision is done at the correct time. The required overlap (l_{overlap}) depends on the speed of movement (v_{mr}) and latency of the handoff (t_{ho}) : $l_{\text{overlap}} = v_{mr} * t_{ho}$. Thus, even with two interfaces, it is worthwhile to minimize the handoff time.

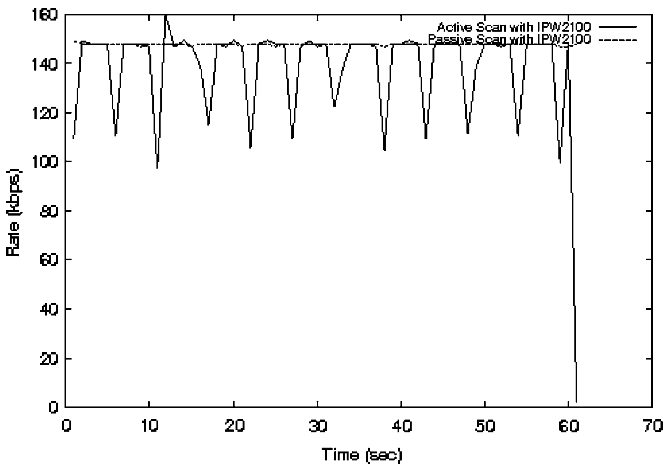

Fig. 3. Effects of active scanning on UDP downstream traffic.

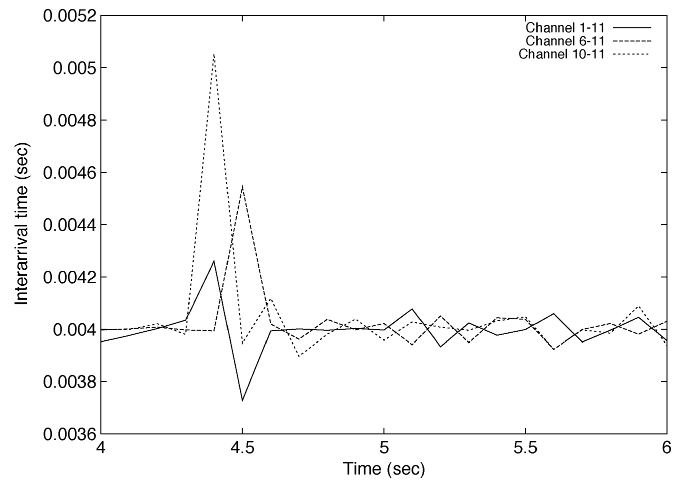

Fig. 4. Effects of channel separation on UDP downstream traffic during handoff.

B. Analysis of Factors Affecting MBB Handoff Performance

MBB handoffs are in theory lossless. However, in our experiments, we found two major causes for packet loss. First, the MR and its HA have inconsistent protocol state during NEMO handoffs due to NEMO signaling and binding management being designed for BBM handoffs. Second, the use of two co-located wireless interfaces in the MR results in intercard interference. Addressing these two causes completely eliminated any packet loss which is evident from the results in Section IV-C. However, there are more general issues pertaining to the wireless network environment, such as fading, which could affect performance also during handoffs. In this paper, we consider only the physical-layer effects which are specific to the proposed scheme, namely, the interference between the two co-located network interfaces in the MR.

We experimented with several 802.11 b/g card pairs and found that the Intel-Prism card pair performed most consistently. Therefore, we used this pair of cards in our experiments for MBB handoffs.

802.11 b/g has several channels, most of which overlap to some extent. Even the nonoverlapping channels can cause interference in the case of co-located wireless interfaces due to the limited adjacent channel (1 and 6, 6 and 1, 6 and 11, or 11 and 6) rejection rate in most 802.11 hardware [18]. Therefore, the transmissions on the scanning interface will cause interference when the active interface is receiving data from an AP on a different channel. In the case of MBB handoffs, this interference occurs in two cases: 1) active scanning of candidate APs using the scanning interface and 2) transmissions of handoff signaling and outgoing traffic on the scanning interface during the handoff.

With active scanning, the scanning card sends a probe on each channel and waits for a response from APs for a certain period before moving on to the next channel. This decreases the scanning time when compared with passive scanning. However, the active scanning resulted in significant interference, as shown in Fig. 3.

The effects of the intercard interference during the handoff depend on the channels that the old and new AP use as can be seen in Fig. 4, which compares the channel pairs 1-11 for nonadjacent channels, 6-11 for adjacent channels, and 10-11 for partially overlapping channels. The handoff between channels 1-11 was lossless, although the interarrival time fluctuated during the handoff, whereas the handoffs with the pairs 6-11 and 10-11 showed some packet loss.

The network latency between the MR and its HA causes them to have a different state for the CoA of the MR during the time it takes for the BU to be delivered from the MR to its HA. During this time, the HA will deliver packets to the old CoA, but the MR will send packets using the new CoA. NEMO and MIPv6 binding management would lead to the MR dropping the incoming packets due to the packets containing an incorrect CoA, as described in [19]. The packet loss is the product of MR-HA delay and bandwidth, and thus the impact of this inconsistency could be significant for fat and long pipes, e.g., a fast satellite connection. To overcome this, we modified the binding management in the MR to accept packets on the old CoA. This removed the packet loss completely during handoffs between foreign networks on channels 1 and 11. However, when performing a handoff between home and foreign networks, we observed another cause for packet loss: HA did not accept tunneled packets after getting the BU from the MR before proxy DAD processing had finished. This resulted in a 1 s period during which the HA dropped incoming packets from MR. We resolved this issue by allowing the HA to process incoming tunneled packets from the MR during the Proxy DAD process.

C. Comparison of NEMO Handoff Performance With MBB and Break-Before-Make Handoffs

We measured UDP packet loss for NEMO without ODAD, NEMO with ODAD, and NEMO with MBB handoffs using two interfaces. We used the Lucent card for the BBM handoffs in the comparison study in this section, since it showed the lowest link-layer handoff latency when compared with the values seen with different cards. The results for UDP packet loss during a handoff for a 100 kB stream from the CN to the LFN are shown in Fig. 5. There is no packet loss for the MBB handoffs

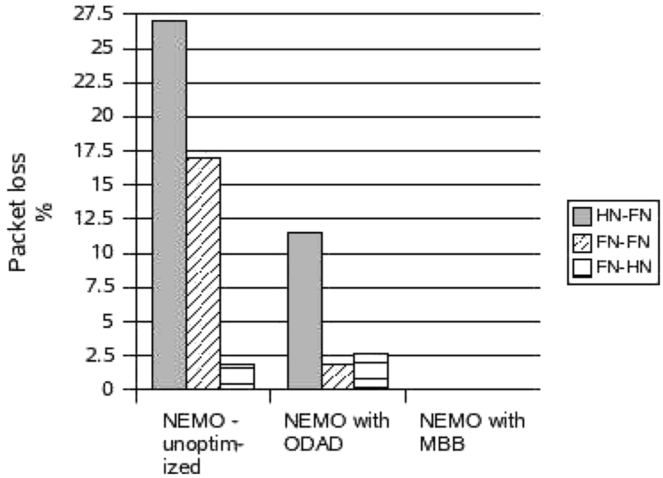

Fig. 5. UDP packet loss comparison during handoffs.

performed using two interfaces due to the simultaneous connectivity to both the old and the new network, whereas in the BBM handoffs (NEMO unoptimized and NEMO with ODAD) the handoff latency reflects directly on the packet loss.

The measurement results for TCP during home-to-foreign handoffs are presented in Fig. 6(a). The negative effects of packet loss in BBM handoffs are amplified by the congestion control mechanisms, whereas the TCP traffic is not affected when using MBB handoffs. The foreign to foreign and foreign to home network handoff results in Fig. 6(b) and (c) are as expected for the BBM handoffs. In Fig. 6(b), it is visible that the TCP throughput increases temporarily during the MBB foreign to foreign handoff due to the use of the new access network for sending acknowledgements, while still receiving data via the old access network. This increase is not visible in the handoffs to and from the home network.

The BBM handoff results presented in this section are dependent on the link-layer technology. However, the results for MBB handoffs are independent of the link-layer handoff latency. Therefore, the analysis of MBB handoff performance presented in this section is valid also for other link-layer technologies.

In summary, we showed here that a MR can reduce the impact of handoffs by optimizing the IPv6 network attachment procedures with fast RAs and ODAD. In the case of BBM handoffs, packet loss is hard to avoid. MBB handoffs with two interfaces can achieve fully lossless handoffs. However, binding management and interference between the interfaces are potential limiting factors to the performance of the handoffs. The impact of interference depends on the network design, i.e., channel separation between adjacent cells, and the hardware design of the wireless interfaces in the MR. The interference can be minimized in the MR by separating the antennas as proposed in [18], or by modifying the wireless cards to avoid interference from the co-located transmitters, for example, by using bandpass filters or polarized antennas.

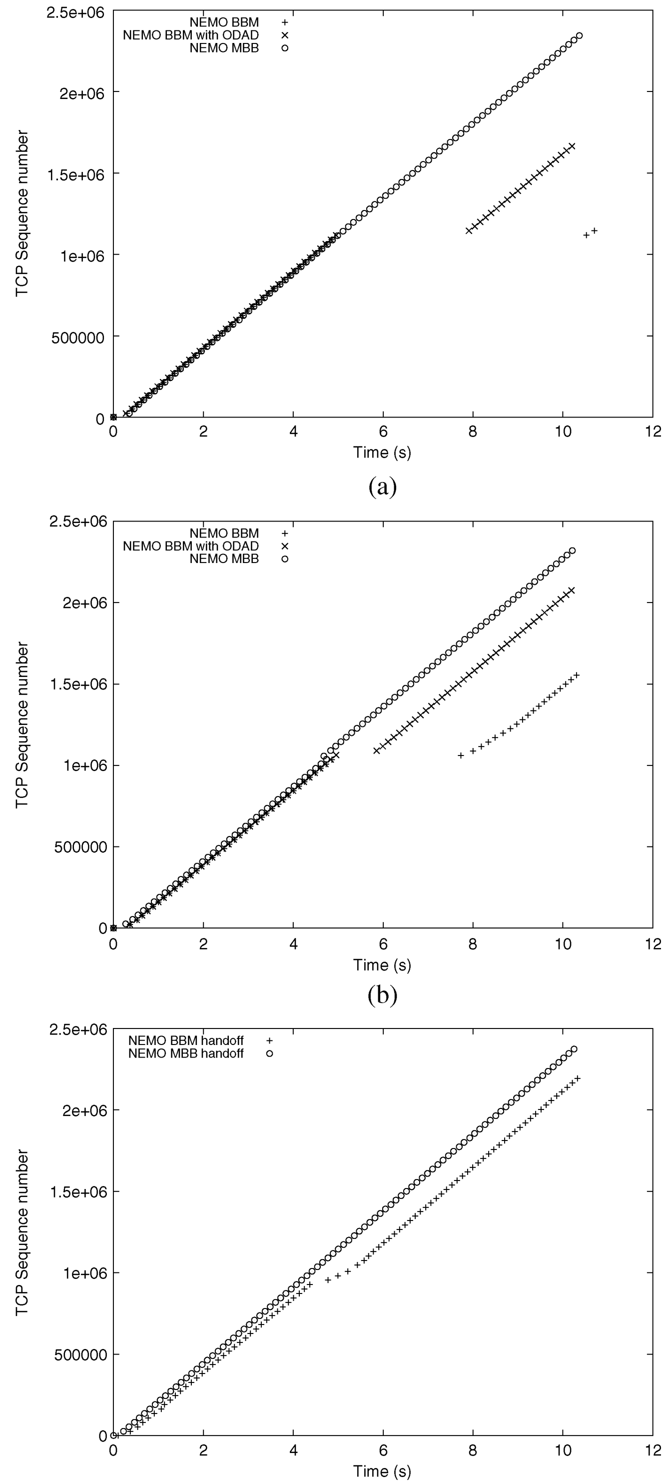

Fig. 6. TCP sequence number diagrams for handoffs. (a) Home-foreign network handoff. (b) Foreign-foreign network handoff. (c) Foreign-home network handoff.

V. Reducing NEMO Overhead

As discussed in Section II-C, employing the NEMO protocol gives rise to nonoptimal routing and protocol header overheads. In this section, we extend the OptiNets RO scheme [3] to address these overheads.

A. The Extended OptiNets RO Scheme for VMNs

In order to cater to the nodes present in the network that have no mobility capabilities, the NEMO basic support protocol assumes that all nodes present in the mobile network have no MIPv6 capabilities. It is evident that this assumption restricts the MIPv6 enabled nodes from achieving better performance.

If the VMNs within the mobile network were aware of the current location, then these nodes would be able to perform standard MIPv6 RO and avoid indirect routing via both HAs, i.e., $HA_{\mathbf{MR}}$ and $HA_{\mathbf{VMN}}$. In the OptiNets scheme, this is achieved by having the MR advertise a topologically correct network prefix on its ingress interface. This enables the MIPv6 capable nodes within the mobile network to autoconfigure a location specific CoA. The MR acquires the topologically correct network prefix from the foreign network using DHCPv6 prefix delegation. As a part of the delegation process, the access router updates its routing table to deliver packets to the prefix via the CoA of the MR. This ensures that MNNs will receive packets to their topologically correct CoAs as long as the MR is connected to the same access router.

In our implementation of OptiNets, the MR runs a DHCPv6 client on its egress interface and obtains a prefix from an access router running a DHCP server. The MR then advertises this prefix on its ingress interface using a special RO prefix option in the RA message. Using this prefix, the VMNs would autoconfigure a CoA for route optimization (RO-CoA). The active VMNs would then send a CoA test init (CoTi) to the CN with the source address being the new CoA. Upon receiving a CoA test (CoT) message from the CN, the VMN sends a BU to the CN by generating a key by combining the new token from the received CoT and the token from a home test (HoT). VMN receives the HoT message similarly as in MIPv6 RO.

In this work, we improve the OptiNets technique by restricting the use of the location specific CoA only for the purpose of RO with CNs. This ensures that only the VMNs which are actively communicating would perform a handoff when the MR changes its point of attachment to the Internet. Further, we use a special ICMPv6 option in the RA for the foreign network prefix advertised by the MR in order to ensure that nonmobility capable nodes do not use the prefix to configure addresses.

The extended OptiNets scheme reduces the per-packet overhead considerably. However, it creates a certain amount of extra signaling when compared with NEMO with LFNs. Prefix delegation is performed every time MR moves and results in a total of 180 bytes being sent over the air interface in addition to the BU-BA exchange between the MR and the $HA_{\mathbf{MR}}$. The remaining part of the signaling overhead results from MIPv6 RO. Every time a VMN switches to a new CoA, it performs a return routability test for the CoA, and sends a BU to the CN. The return routability test for the home address is performed every 210 s. Based on the message sizes in Appendix B, the OptiNets handoff dependent signaling overhead is 388 bytes per handoff for MR and 216 per handoff for each VMN. In addition to this, the VMNs perform MIPv6 return routability for their home address with the CN resulting in a signaling overhead of 360 bytes every 210 s.

B. Results and Discussion

We compared the TCP performance of a LFN, a VMN with no RO, a VMN with MIPv6 RO (i.e., avoiding the $HA_{\mathbf{VMN}}$), and a VMN with OptiNets (i.e., avoiding both $HA_{\mathbf{VMN}}$ and $HA_{\mathbf{MR}}$). We measured the performance in a static case, in which the MR was located in a foreign network and a dynamic case in which the MR moved between two foreign networks. The results for the static case, as shown in Fig. 7, indicate that the performance of the other schemes decreases as the latency between the MR and the $HA_{\mathbf{MR}}$ increases, whereas the performance of the OptiNets scheme is not affected. The results for the dynamic case in Fig. 8 show that the performance of the OptiNets scheme is comparable with the static case. The TCP performance of the other schemes was the same as the static case, and therefore we only show the NEMO LFN performance for comparison.

The per packet header overhead did not have an effect in the previous two measurements since the TCP performance was limited by the end-to-end latency and not by the available bandwidth (2 Mbits/s) due to the use of the default TCP window size. In Fig. 9, we analyze the relative overhead of the different schemes. We used a 64 Kbits/s constant bit rate stream with 220 Byte packets as traffic and calculated the amount of signaling and per packet protocol overhead relative to the total

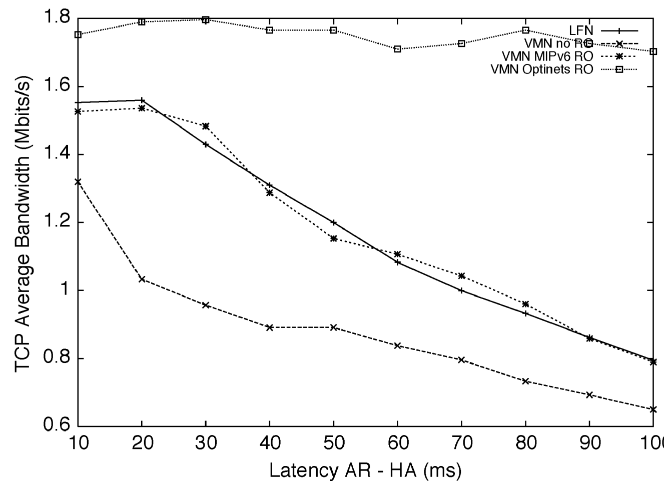

Fig. 7. TCP performance comparison in static case.

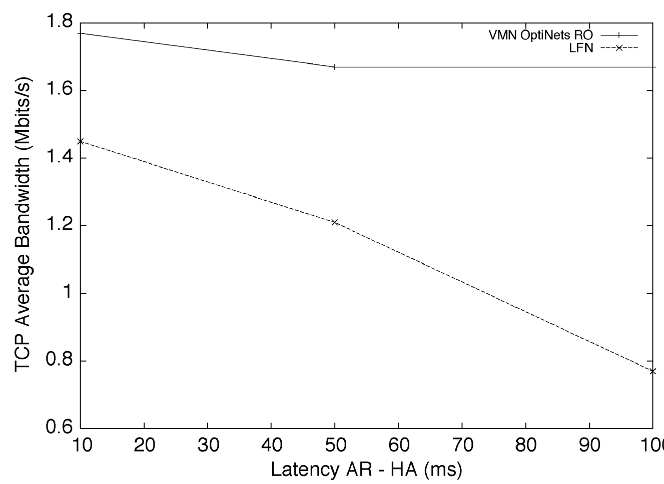

Fig. 8. TCP handoff performance comparison for LFN, MIPv6 MN OptiNets RO.

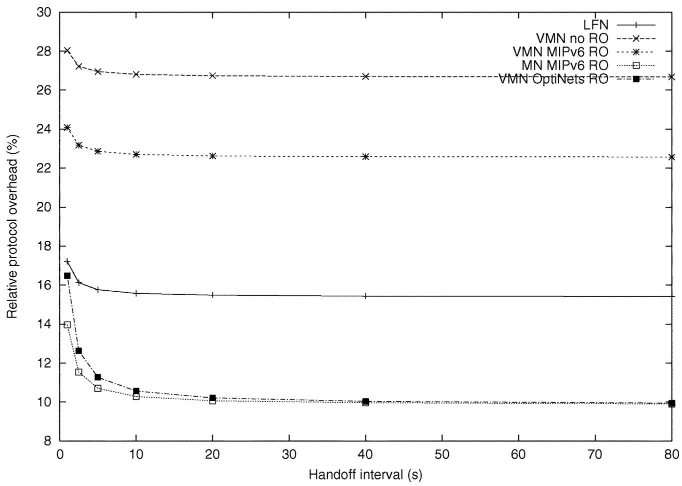

Fig. 9. Overhead comparison for 1 MNN running CBR traffic with varying handoff interval.

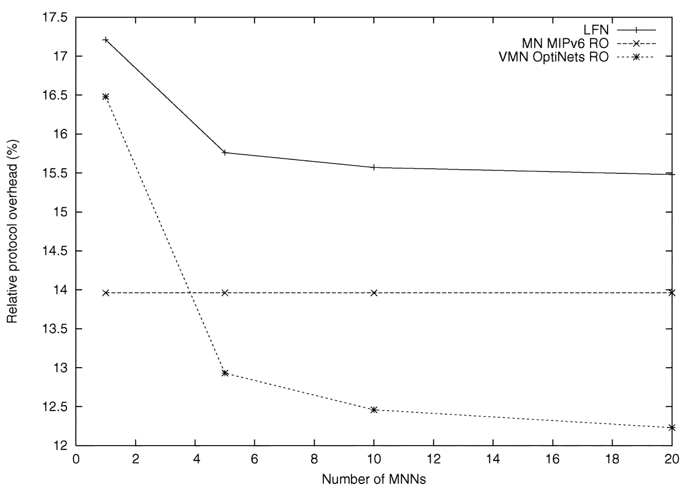

Fig. 10. Overhead comparison for LFN, OptiNets RO and mobile IPv6 MN, with varying number of MNNs.

amount of data sent over the air interface between the MR and the access router. It can be seen that the use of OptiNets incurs the smallest total overhead of the NEMO variants regardless of the handoff frequency, when 1 MNN is communicating up to 1 handoff per second which is the maximum frequency specified in [5]. With OptiNets it can be seen that we are able to reduce the per packet overhead to a level comparable to that of a route optimized MIPv6 mobile node connecting directly to the access router, bypassing the MR.

We also analyzed the effect of multiple MNNs with the same traffic type as in Fig. 9 and the results in Fig. 10 indicate that the relative overheads of NEMO and OptiNets decrease as the number of MNNs increases. This is due to the aggregation of the mobility signaling.

VI. Related Work

A. Wireless Network Testbed

There are several testbeds related to network mobility. The OverDRiVE project [20] focused on UMTS enhancements and coordination of existing radio networks into a hybrid network to enable the delivery of spectrum-efficient multicast and unicast services to vehicles. The iCar [21] testbed utilized multiple-access technologies for connecting a car network to the Internet. The mobile access router (MAR) testbed [22] focused on evaluating the performance of vertical handoff and load balancing in a vehicular mobile network environment. The Nautilus project testbeds have been designed to verify the applicability of the NEMO protocol implementations in different scenarios, such as the E-wheelchair [23]. Our network mobility testbed is geared towards evaluating the performance of NEMO, related IPv6 protocols and the proposed optimizations.

B. Handoff Performance Improvement

Previous research on handoff performance improvement on the network-layer has mostly focused on improving the performance of BBM handoffs since most mobile devices can only connect to a single-access network at a time.

Hierarchical MIPv6 [15] reduces the packet loss by introducing additional functionality to the foreign network infrastrcuture for localizing the handoffs. Fast MIPv6 [14] emulates MBB handoffs by allowing a mobile device to connect virtually to its new and old access router at the same time. These approaches are well suited to networks with a large number of mobile devices since they allow for simple mobile devices by moving complexity to the edge of the network. However, in NEMO, the MR acts as an aggregation point for mobility management and routing, and thus the benefits of reduced complexity in the MR do not necessarily outweight the costs of additional complexity in the infrastructure. Further, previous research [24] suggests that handoff prediction may be successful on the average only 50% of the time, thus reducing the performance of Fast MIPv6 significantly.

MBB handoffs have been utilized in cellular networks at the link-layer. However, as a part of the IP level mobility management, MBB handoffs could be used independently of the underlying link-layer technology.

C. Route Optimization (RO)

Several RO techniques have been proposed in the context of single-level and nested mobile networks. The schemes for nested mobile networks, such as Kang *et al.*'s [25] proposal, Thubert's reverse routing header (RRH) protocol [26] and Ohnishi *et al.*'s [27] Hierarchical MIPv6-based approach, reduce the overheads of multiple levels of nested mobile networks.

There are several schemes for RO for unnested mobile networks. The optimized route cache (ORC) protocol[28] reduces the overhead of tunneling by introducing correspondent routers that can be configured anywhere in the Internet to be an anchor router for the mobile network. The performance gained from using ORC scheme depends on the vicinity of correspondent routers to CNs. However, this scheme requires significant support from the network infrastructure. Jeong *et al.* [29] proposed an optimization mechanism for MIPv6 enabled nodes in which the MR acts as a bridge and a neighbor discovery proxy between VMNs and the foreign link. Although this technique does not require any support from the infrastructure, it increases the signaling load on the wireless link since each MR performs neighbor discovery on the link for each MNN. Thus, it is more applicable to mobile networks with relatively few nodes. In this work, we extended and implemented our previously proposed OptiNets RO technique [3] which requires support from the access routers for prefix delegation, but scales to a larger number of MNN.

VII. Conclusion

In this paper, we measured the performance of NEMO in a network mobility testbed. We analyzed the handoff performance and protocol and routing overheads of a NEMO-based network mobility system. The analysis showed that unoptimized handoff performance of NEMO would be unsuitable for most applications due to handoff latencies of up to 2.75 s. Even with protocol optimizations the handoff latencies would still limit the suitability for performance sensitive applications, such as voice-over-IP. Further, from the analysis it was evident that the protocol and routing overheads of NEMO would lead to inefficient use of scarce wireless network resources. To address these shortcomings we proposed the use of multiple interfaces for MBB handoffs and extended our previously proposed OptiNets RO scheme.

The MBB handoffs make it possible for a fast moving MR to take advantage of high-speed but short-range radio technologies without compromising the service it offers to MNN. However, there are a number of potential drawbacks to using multiple interfaces in mobile devices, such as an increase in power consumption, interference caused by the usage of multiple interfaces, and increased size and cost. These drawbacks apply mostly to mobile hosts and do not limit the use of multiple interfaces on MRs to the same extent for the following reasons.

1) An on-board MR is not limited by power constraints in the same way as battery powered mobile devices since it will be powered by the vehicle.
2) The ability to use physically separated external antennas on a MR will reduce the effects of interference.
3) A small increase in the size and cost of a MR can be easily justified by the fact this increase benefits a large number of nodes due to aggregation of mobility management at the MR.

The extended OptiNets RO scheme enables a VMN to bypass NEMO tunneling, and use its own mobility capabilities, such as Mobile IPv6. Although we have used only the MIPv6 protocol in our measurements for OptiNets, it is important to note that the scheme could be used with any other host mobility protocol, possibly with even greater performance increases. For example, SIP [4] does not require extra headers for RO. Thus, it would be possible to remove the per packet overhead of NEMO completely by using OptiNets with SIP.

In summary, we showed that the MBB handoff scheme and the extended OptiNets RO scheme alleviate the performance issues of the NEMO protocol. With these optimizations NEMO could be used even with applications highly sensitive to delay and packet loss.

Appendix A
NEMO Handoff Latency Calculations

The total handoff latency equals to the link-layer handoff latency T_{L2} + the IPv6 network attachment latency, consisting of RS advertisement exchange $T_{\mathrm{RS-RA}}$ and DAD (T_{DAD}) + NEMO protocol latency consisting of the RTT between the MR and its HA (T_{RTT}) and possibly Proxy DAD (T_{PDAD}).

When the MR moves from its home network to a foreign network the latency is

$$T_{h2f} = T_{L2} + T_{\mathrm{RS-RA}} + T_{\mathrm{DAD}} + T_{\mathrm{RTT}} + T_{\mathrm{PDAD}}$$

Average value for $T_{\mathrm{RS-RA}}$ is 0.25 s and the average value for T_{DAD} is 1.5 s and the value for T_{PDAD} is 1 s. The average value for the handoff then equals to $T_{h2f} = T_{L2} + T_{\mathrm{RTT}} + 2.75$ s.

When the MR moves between two foreign networks proxy DAD is not performed, thus the handoff latency becomes

$$T_{f2f} = T_{L2} + T_{\mathrm{RS-RA}} + T_{\mathrm{DAD}} + T_{\mathrm{RTT}}$$

This equals to a median of $T_{f2f} = T_{L2} + T_{\mathrm{RTT}} + 1.75$ s.

Finally, when the MR returns home no DAD is performed since the HA acts as a proxy for the home and the link local addresses of the MR. Then the latency consists of only the L2 handoff latency, the RS-RA delay and the RTT between the MR and its HA

$$T_{f2h} = T_{L2} + T_{\mathrm{RS-RA}} + T_{\mathrm{RTT}}$$

Thus, the average minimum for handoff latency T_{f2h} is 0.25 s + link-layer dependent delays when returning home.

TABLE IV
Message Sizes With IPsec and Tunneling in Bytes

Message	Base Size	With IPsec	With NEMO Tunnel
BU MR-$HA_{\mathbf{MR}}$	80	104	N/A
BA $HA_{\mathbf{MR}}$- MR	80	104	N/A
BU MN-HA	80	104	N/A
BA HA-MN	80	104	N/A
DHCPv6 solicit MR-AR	96	N/A	N/A
DHCPv6 reply AR-MR	184	N/A	N/A
VMN MIPv6 BU VMN-CN	96	N/A	136
HoTi VMN-$HA_{\mathbf{VMN}}$-CN	96	136	176
HoT CN-$HA_{\mathbf{VMN}}$-VMN	104	144	184
HoTi MN-HA-CN	96	136	N/A
HoT CN-HA-MN	104	144	N/A
VMN MIPv6 CoTi VMN-CN	56	N/A	96
VMN MIPv6 CoT CN-VMN	64	N/A	104
OptiNets VMN-CN CoTi	56	N/A	N/A
OptiNets VMN-CN CoT	64	N/A	N/A
MIPv6 MN-CN CoTi	56	N/A	N/A
MIPv6 MN-CN CoT	64	N/A	N/A
OptiNets BU VMN-CN	96	N/A	N/A
MIPv6 BU MN-CN	96	N/A	N/A

Appendix B
Message Sizes and Overheads for NEMO, MIPv6, and OptiNets RO

The Table IV shows the message sizes in bytes for all signaling messages sent over the air interface, excluding IPv6 neighbor discovery and link-layer signaling messages. IPsec encapsulating security payload (ESP) [30] is used for protecting the HoT and HoTi messages and IPsec AH [31] for protecting BUs between the MR and the $HA_{\mathbf{MR}}$. The NEMO protocol overhead consists of signaling and per packet overheads. For a

NEMO LFN this is always 40 bytes per packet due to use of the MR-HA IPv6 tunnel. For a VMN which uses MIPv6 with tunneling with the $HA_{\mathbf{VMN}}$ the overhead becomes 80 bytes due to the double tunnel. For a VMN which uses MIPv6 RO with a 24 byte extension header the overhead is 64 bytes per packet and for a MIPv6 MN it is 24 bytes.

Nested mobility may result in per MNN time-dependent signaling overhead which does not exist for LFNs or is negligible for VMNs not using RO since they register with the $HA_{\mathbf{VMN}}$ only when entering the mobile network. MNs and VMNs performing MIPv6 RO send BUs and return routability messages to CNs according to MIPv6 specification. This signaling consists of a CoA return routability test between the MN and the CN, a BU to the CN, and a BU to the HA every time the MN moves and a HoA return routability every 210 s, or every time before MN sends a BU to the CN, if this is more seldom than every 210 s. Thus, the per MN time-dependent overhead for a MIPv6 MN becomes the size of HoTi + HoT protected with IPsec every 210 s. For a VMN CoT, CoTi and BU are exchanged with a CN every 420 s with HoTi and HoT packets being tunneled between the MR and the $HA_{\mathbf{MR}}$.

Handoff signaling is handled solely by the MR and the $HA_{\mathbf{MR}}$ for a NEMO LFN and consists of a BU and a BA protected by IPsec. For a MIPv6 MN, the handoff signaling consists of a BU to its HA, and a CoTi-CoT-BU exchange with each CN.

References

[1] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, "Network mobility (NEMO) basic support protocol," Internet RFC: RFC3963, Jan. 2005.

[2] K.-C. Lan, E. Perera, H. Petander, L. Libman, C. Dwertmann, and M. Hassan, "MOBNET: The design and implementation of a network mobility testbed," in *Proc. 14th IEEE Workshop on Local and Metropolitan Area Netw.*, Crete, Sep. 2005, pp. 1–6.

[3] E. Perera, A. Seneviratne, and V. Sivaraman, "OptiNets: An architecture to enable optimal routing for network mobility," in *Proc. Int. Workshop on Wireless Ad-Hoc Networks*, May 2004, pp. 68–72.

[4] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session initiation protocol," Internet RFC: RFC3261, Jun. 2002.

[5] D. Johnson, C. Perkins, and J. Arkko, "Mobility support in IPv6," Internet RFC: RFC3775, Jun. 2004.

[6] J. Kempf, M. Khalid, and B. Pentland, "IPv6 fast router advertisement," Jul. 2004, draft-mkhalil-ipv6-fastra-05.txt.

[7] N. Moore, "Optimistic duplicate address detection for IPv6," Feb. 2005, draft-ietf-ipv6-optimistic-dad-05.txt.

[8] J. Arkko, V. Devarapalli, and F. Dupont, "Using IPsec to protect mobile IPv6 signaling between mobile nodes and home agents," Internet RFC: RFC3776, Jun. 2004.

[9] M. Carson and D. Santay, "Nist net: A linux-based network emulation tool," *ACM Comput. Commun. Rev.*, pp. 111–126, 2005.

[10] MIPL mobile IPv6 for Linux. [Online]. Available: http://www.mobileipv6.org

[11] Fast router advertisements implementation. [Online]. Available: http://www.ctie.monash.edu.au/ipv6/fastra.htm

[12] Iperf network performance measurement tool. [Online]. Available: http://dast.nlanr.net/Projects/Iperf/

[13] Ethereal network analyzer. [Online]. Available: http://www.ethereal.com/

[14] R. Koodli, Ed., "Fast handovers for mobile IPv6," Internet RFC: RFC4068, Jul. 2005.

[15] H. Soliman, C. Castelluccia, K. El-Malki, and L. Bellier, "Hierarchical MIPv6 mobility management," Jul. 2002, draft-ietf-mipshop-hmipv6-04.txt.

[16] S. Aust, D. Proetela, N. Fikouras, C. Paupu, and C. Gorg, "Policy based mobile IP handoff decision (POLIMAND) using generic link layer information," in *Proc. 5th IEEE Int. Conf. Mobile and Wireless Commun. Netw.*, Oct. 2003.

[17] G. Y. Liu and G. Maguire, Jr., "A class of mobile motion prediction algorithms for wireless mobile computing and communications," *MONET*, vol. 1, no. 2, pp. 113–121, 1996.

[18] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy, "Experimenting with a multi-radio mesh networking testbed," in *Proc. 1st Workshop on Wireless Netw. Measurements (Winmee)*, Garda, Italy, Apr. 2005.

[19] H. Petander and E. Perera, "Improved binding management for make before break handoffs in mobile IPv6," Oct. 2005, draft-petander-mip6-mbb-00.txt.

[20] R. Tönjes, K. Mößsner, T. Lohmar, and M. Wolf, "OverDRiVE—Spectrum efficient multicast services to vehicles," in *Proc. IST Mobile Summit*, Thessaloniki, Greece, Jun. 2002. [Online]. Available: http://www.comnets.rwth-aachen.de/~o_drive/index.html

[21] T. Ernst, K. Uehara, and K. Mitsuya, "Network mobility from the Internet CAR perspective," in *Proc. 17th Int. Conf. Advanced Inf. Netw. Appl.*, Xi'an, China, Mar. 2003, pp. 19–25.

[22] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee, "MAR: A commuter router infrastructure for the mobile Internet," in *Proc. 2nd Int. Conf. Mobile Syst., Appl. Serv.*, 2004, pp. 217–230.

[23] T. Ernst, "E-Wheelchair: A communication system based on IPv6 and NEMO," in *Proc. 2nd Int. Conf. Smart Homes and Health Telematic*, Singapore, Sep. 2004.

[24] T. C. Schmidt and M. Wählisch, "Predictive versus reactive—Analysis of handover performance and its implications on IPv6 and multicast mobility," *Telecommun. Syst.*, vol. 30, no. 1–3, pp. 123–142, Nov. 2005. [Online]. Available: http://dx.doi.org/10.1007/s11235-005-4321-4

[25] H. Kang, K. Kim, S. Han, K.-J. Lee, and J.-S. Park, "Route optimization for mobile network by using bi-directional between home agent and top level mobile router," Jun. 2003, draft-hkang-nemo-ro-tlmr-00.txt.

[26] P. Thubert and M. Molteni, "IPv6 reverse routing header and its application to mobile networks," Jun. 2004, draft-thubert-nemo-reverse-routing-header-05.txt.

[27] H. Ohnishi, K. Sakitani, and Y. Takagi, "HMIP based route optimization method in a mobile network," Oct. 2003, draft-ohnishi-nemo-ro-hmip-00.txt.

[28] R. Wakikawa, S. Koshiba, and K. Uehara, "ORC: Optimized route cache management protocol for network mobility," in *Proc. 10th Int. Conf. Telecomm. (Protocol for Network Mobility)*, Papeete, Tahiti, French Polynesia, Feb. 2003.

[29] J. P. Jeong, K. Lee, J. Park, and H. Kim, "ND-proxy based route and DNS optimizations for mobile nodes in mobile network," Feb. 2004, draft-jeong-nemo-ro-ndproxy-02.txt.

[30] S. Kent and R. Atkinson, "IP encapsulating security payload (ESP)," Internet RFC: RFC2406, Nov. 1998.

[31] ——, "IP authentication header," Internet RFC: RFC2402, Nov. 1998.

Henrik Petander received the M.Sc. degree in electrical engineering from the Helsinki University of Technology, Helsinki, Finland, in 2002. Currently, he is working towards the Ph.D. degree in computer science at the Helsinki University of Technology.

He worked at the Helsinki University of Technology on MIPL Mobile IPv6 implementation as a Researcher from 2002 until joining the Network and Pervasive Computing Program, National ICT Australia, in August 2005. His current research interests are mobility management in IP networks and security of mobile networking.

Eranga Perera (S'03) received the B.Sc. degree in mathematics and computing (First Class Hons) from the University of London, London, U.K. the M.S. degree in information science from the University of New South Wales, Sydney, Australia. Currently, she is working towards the Ph.D. degree in electrical and telecommunications engineering at the University of New South Wales.

She is working in the Ambient Networks Project which is part of the European Union's Sixth Framework Program. Her research interests are in mobile and wireless systems.

Kun-Chan Lan (M'06) received the B.A. degree from the National Cheng Kung University, Tainan, Taiwan, the M.S. degree from the State University of New York (SUNY) at Stony Brook, Stony Brook, NY, and the Ph.D. degree in computer science from the University of Southern California, Los Angeles, in 2004. During his Ph.D. study, he has maintained and contributed codes to the main simulation software ns-2.

He joined the Network and Pervasive Computing Program, National ICT Australia (NICTA), Eveleigh, as a Researcher in May 2004. His research interests include traffic modeling, simulation, network measurement, and wireless network. His past research work focused on rapidly generating realistic application-level simulation models from distributed network measurements. His current research interests are mainly in wireless mesh networks, vehicular ad hoc networks, realistic wireless simulations, and VoIP over wireless.

Aruna Seneviratne (M'94) received the B.Sc. (Hons) degree from Middlesex Polytechnic, London, U.K., and the Ph.D. degree from the University of Bath, Bath, U.K.

Since graduating, he has held academic appointments at the University Bradford (U.K.), Curtin University, Australian Defense Force Academy (UNSW), and the University of Technology, Sydney. Outside academia he has worked at the Standard Telecommunication Laboratories (U.K.), Muirhead and Company (U.K.), and Telecom Australia (Telstra). He has also spent time at the MASI Laboratory, University of Pierre Marie Curie (Paris), and the Rodeo Group at INRIA (Nice) as a Visiting Professor. He is the Director of the Australian Technology Park Laboratory, National ICT Australia. He also holds the Mahanakorn Chair of Telecommunication in the School of Electrical Engineering and Telecommunication, University of New South Wales. His current research interests are in mobile data communication systems.