NAT-Based Internet Connectivity for On-Demand Ad Hoc Networks

Paal Engelstad and Geir Egeland

Telenor R&D,
1331 Fornebu, Norway
{Paal.Engelstad, Geir Egeland}@telenor.com

Abstract. A prerequisite for widespread and successful deployment of on-demand ad-hoc networking technology is its ability to provide easy access to the Internet. Existing solutions for Internet access are mainly based on modifying Mobile IPv4 (MIPv4). An easier approach, yet poorly documented in published material, is to implement Network Address Translation (NAT) on Internet Gateway nodes in the ad hoc network. In this paper we describe problems experienced in our lab test-bed with NAT-based solutions under the condition of site multi-homing. Based on this experience, we propose a working solution for multi-homed ad hoc networks.

1 Introduction

IP-based applications, such as web browsing, e-mail, telnet and ftp, mainly communicate with servers or peers over the Internet. A mobile ad-hoc network (MANET [1]) has no fixed infrastructure and services on the Internet might not be available in these networks. A likely scenario is that nodes on an ad-hoc network in some cases also want to connect to nodes on the Internet, using services available here. For a widespread and successful deployment of MANETs, the ability to provide easy access to the Internet is therefore a prerequisite.

A possible solution is to let a node that is participating in a MANET operate as an Internet gateway and provide other nodes on the MANET with Internet access. One approach is to implement a Mobile IPv4 Foreign Agent (MIP-FA) on the gateway ([2], [3]). MANET nodes that require Internet access, implement a Mobile IPv4 (MIPv4 [4]) client, and register the globally routable IPv4 address of the gateway as a care-of-address with their MIPv4 Home Agents (HA). However, as we describe in this paper, a Mobile IP-based solution has a number of drawbacks and introduces high complexity to implementations.

An easier way to provide IPv4 Internet access to MANET nodes is to implement Network Address Translation (NAT) on the gateways [5]. Although NAT solutions have emerged in different test-bed implementations (see for example [6]) little has been documented, neither in scientific papers nor in IETF Internet Drafts.

A challenging issue with the use of NAT in general relates to site multi-homing. Since a NAT-device translates both outgoing and incoming packets, all

R. Battiti, M. Conti, and R. Lo Cigno (Eds.): WONS 2004, LNCS 2928, pp. 344–358, 2004.
© IFIP International Federation for Information Processing 2004

traffic belonging to one communication session (e.g. TCP session) must traverse through the same NAT-device, otherwise the session will break. When the site is multi-homed, it can be difficult to control that all packets from one session are consistently routed through the same NAT-device.

This issue is also highly relevant to MANETs. A MANET is without infrastructure, and it is difficult to eliminate the possibility of multi-homing. It is not possible to control the network behavior in such a way that there is only one node on the MANET operating as a MANET Internet gateway. Even if there existed a simple solution to suppress a second node to operate as a gateway, the problem would re-emerge at the moment when two MANETs, each with a NAT-based gateway, merge into one network. Shutting down one of the gateways would mean to break all on-going communication sessions over that gateway.

In this paper we address the lack of a good mechanism for IPv4 Internet access in on-demand MANETs, i.e. a MANET that is routed with a reactive routing protocol, such as the Ad hoc On-demand Distance Vector (AODV [7]) routing protocol or the Dynamic Source Routing (DSR [8]) protocol. The paper examines the use of NAT for this purpose and point out potential problems with NAT-based solutions and multi-homing as experienced in our test-bed. Based on this, a working solution for multi-homing is proposed.

In Section 2 we present proposed solutions for providing Internet connectivity on on-demand MANETs. In Section 3 we analyze how a NAT-based solution work in scenarios of multi-homing. A proposed solution for NAT-based gateways based on our findings is presented in Section 4. The applicability of our findings to MIP-FA-based solutions is also discussed in Section 5.

The main focus is set on on-demand networks that use AODV as a routing protocol. However, as the analyses and proposals presented in this paper are of a general nature and not strictly dependent of the reactive routing protocol, we will deal with applicability of our findings to DSR in a separate section by the end of the paper.

2 Background

2.1 Reactive ("On-Demand") Routing Protocols

A number of reactive routing protocols have been proposed. The most widely studied and popular proposals include the AODV and the DSR protocols.

Reactive protocols allow source nodes to discover routes to a destination IP-address on demand. Most proposals, including AODV and DSR, work as follows: When a source router needs a route to a destination for which it does not already have a route, it issues a Route Request (RREQ) packet. The packet is broadcasted by controlled flooding throughout the network, and sets up a return route to the source.

If a router receiving the RREQ is either the destination or has a valid route to the destination, it unicasts a Route Reply (RREP) packet back to the source along the reverse route. The RREP normally sets up a forward route. Thus, the

RREQ and RREP messages set up two uni-directional unicast routes in opposite directions between source and destination. Once the source node receives the RREP, it may begin to forward data packets to the destination. (The acronyms RREQ and RREP are borrowed from AODV.)

Different reactive routing protocols have different strategies to deal with route maintenance and route repair. Most protocols let routes that are inactive eventually time out. If a link break occurs while the route is active, the routing protocol normally implements an algorithm to repair the route.

Different protocols have different ways to manage routing state information. With AODV, for example, correct forwarding of packets between source and destination relies on stored state information in routing tables on intermediate nodes in the network. With DSR, on the other hand, the sender of the packet encodes the entire route explicitly into the packet, and the packet is source-routed from source to destination.

2.2 Basic Functionality for Internet Connectivity

Providing Internet access to a MANET node – hereafter referred to as a "Source Node" (SN) - via another MANET node that has direct access to the Internet – hereafter referred to as a "gateway" (GW) - involves a number of steps.

Before packets from the SN are routed to a node outside of the MANET, one has to determine somehow that the destined node is not present on the MANET. If the destination IP address is not found on the MANET, available gateways must be detected, and one of these must be selected.

The selected gateway must somehow provide a source IP address to SN's outgoing traffic to make sure that the incoming traffic returning from the Internet is correctly routed back to the gateway. Due to the shortage of globally routable IPv4 addresses, an Internet gateway is likely to be assigned only one IP-address on the external network to which it is connected. Hence, all MANET nodes that use this gateway to access the Internet must share the external IP address of the gateway for traffic returned from the Internet.

Some explicit signaling between the SN and the gateway might be required (e.g. for a MIP-FA-based solution), or implicit translation at the gateway is required (e.g. for a NAT-based solution). Finally, as an integrated part of the previous steps, or as a separate step, there must be established forward and return routes between the SN and the gateway.

In summary, the basic functionality required for Internet connectivity includes:

- Determining that an address is not present on the MANET
- Detection of different available gateways
- Selection of one gateway
- Signaling with selected gateway (if required)
- Providing a globally routable IP address for incoming traffic
- Routing of outgoing traffic via gateway
- Routing of incoming traffic via gateway

In the following, a MIP-FA-based solution and a NAT-based solution are presented and described in relation to this list of required functionality.

2.3 MIP-FA for AODV

There exist at least two proposals that attempt to implement a MIP-FA-module on a gateway for AODV. The MIPMANET proposal [3] takes an academic approach and proposes a number of schemes to optimize the use of Mobile IPv4 for this purpose. The scheme proposed by Belding-Royer *et. al.* [2] represents a similar, but considerably simpler engineering approach. We will discuss the latter approach, since this proposal is more relevant to the analyses undertaken in this paper.

A MIP-FA-based gateway will periodically flood Agent Advertisements throughout the MANET. From this other nodes in the MANET will learn the IP address of the gateway and the care-of-address that it offers. The SN can select a gateway and use AODV mechanisms to discover a route to it. Once the route is known, the SN unicasts a Registration Request message to the gateway, thereby registering with the FA-module on the gateway and SN's own HA.

When the SN searches for a node, it sends a RREQ, whether the node is present locally on the MANET or not. Upon reception of a RREQ message for a destination that the gateway believes is present outside the MANET, the gateway transmits a RREP to the requesting node with an 'F'-bit set to indicate that the destination node can be reached through the gateway.

The destination IP-address in the RREP is set to the IP-address of the node the SN is searching for. This means that the gateway sends out "Proxy RREPs" on behalf of nodes that might be present on the Internet. The advantage of this approach is that the SN can send packets to external nodes in the same way as it would to nodes that are present on the MANET.

A mechanism is required to ensure that communication will not escape through a gateway when the destined node is present on the MANET. This can be solved by ensuring that the destination sequence number in an RREP from a MANET node present on the network is always larger than an that of an RREP from a MANET gateway, when the two RREPs are triggered by the same RREQ. Since the highest sequence number gives preference during route discovery, routes to MANET nodes present on the MANET will always have preference over routes established by Proxy RREPs sent by gateways.

One solution is to mandate that a MANET node that is present on the MANET (while not operating as a gateway) must update its destination sequence number before issuing the RREP in response to an RREQ with the F-flag set. The MIP-FA-based gateway, on the other hand, will always copy the AODV-specific destination sequence number of the RREQ into the corresponding field of the RREP. Another solution is to mandate that SN always set the Unknown Sequence Number flag to 1 and the destination sequence number to zero in RREQs with the F-flag set. Then the RREPs from MANET gateways will always have a zero destination sequence number.

2.4 AODV-UU NAT Solution

Uppsala University's implementation of AODV [6] includes a NAT-solution. Mobile hosts are unaware of its presence, hence there is no NAT discovery and NAT uses Proxy RREP to reply to RREPs destined for hosts on the Internet.

Unlike the MIP-FA solution, the gateway uses its own AODV destination sequence number when replying with Proxy RREPs. This may potentially introduce a problem since a route will always be set up to the node that sends the RREP with the highest sequence number. Hence, in cases where the destination sequence number of the gateway is higher than that of the destination node, packets will be routed out through a gateway even if the destination node is present on the MANET.

To prevent this, the AODV-UU NAT solution mandates that all addresses on the MANET belong to the same subnet prefix. Hence a NAT will reply with a Proxy RREP only if the destination IP address is not under this prefix.

This prefix limitation conflicts with the widely accepted notion that a node should be able to bring any address into the MANET, and use this while present on the network. The prefix limitation might also introduce problems when other gateway technologies are present on the network. One example is that the NAT-based gateway and a MIP-FA-based gateway are mutually exclusive, since Mobile IP mandates that mobile nodes use their home address on visited links, i.e. an address that not necessarily belongs to a prefix assigned to the MANET.

We did not find the AODV-UU NAT solution satisfactory, so we implemented an alternative approach to the prefix limitation. The NAT was allowed to answer to all RREQs, similar to the MIP-FA-based solution of Belding-Royer et al. [2]. A targeted node would have to increase its destination sequence number before replying with an RREP, while a MANET gateway would copy it from the destination sequence number field of the RREQ. Hence, if the targeted destination address were present on the MANET, a direct route to that node would have preference, since the RREP from the targeted node would have a larger destination sequence number.

In our further references to NAT-based solutions, we refer to solutions without the aforementioned prefix limitation.

2.5 Comparison of MIP-FA-Based and NAT-Based Solutions

A drawback with a MIP-FA solution was revealed after implementing it in our research lab. The scheme requires changes to the Mobile IP implementation on both the Mobile Host (MH) side (i.e. on the SN requiring Internet access) and on the Foreign Agent (FA) side (i.e. on the gateway). Since the MH and the FA are no longer on-link, both sides will have to deal with Agent Solicitations and Agent Advertisements in a different way; TTL values and IP destination addresses must be set differently; ARP must be used differently and MAC-addresses are no longer relevant for communication between MH and FA. Moreover, independent co-existing implementations of MIPv4 and a MANET routing protocol

are not trivially managed, since both implementations will make unsynchronized modifications to the routing table.

Another drawback that limits the applicability of a MIP-FA based solution is that it assumes that the care-of-address of the gateway is globally routable. However, the IPv4 address space is a scarce resource, and many MANET gateways might only be able to acquire a private IPv4 address on the external network to which they are connected.

A NAT-based mechanism, on the other hand, appears as a considerably easier solution. The NAT functionality may be in the form of Basic NAT, however NAPT (i.e. NAT with port translation) is a more applicable solution, since many MANET Internet Gateways might only be able to acquire a single IP-address on the external network to which they are connected [5]. NAT-devices can be nested, and this solution will work even when the MANET Internet Gateway acquires a private IP address from the external network.

In the next section we take a closer look at problems arising with NAT-based solutions, especially when the MANET is multi-homed.

3 Problems with NAT and Multi-homing

3.1 NAT-Based Multi-homing with One Source Node

We tested site multi-homing in an on-demand MANET with two NAT-based gateways present. The source node (SN) communicated with an external host (XH), and both gateways (GW1 and GW2) were reachable through the same Intermediate Node (IN). The AODV-UU implementation for Linux [6] was used as routing modules on all MANET nodes (i.e. SN, IN, GW1 and GW2), and WLAN 802.11b was used for the wireless communication. All nodes were located in the same room (9-by-9 meters), and MACKILL [6] was used to emulate that SN was not in direct radio-range with the gateways. The test configuration is illustrated in Figure 1.

Initially the SN established a route to XH over one of the gateways, i.e. GW1. The SN then established a TCP connection with XH, and periodically sent over short messages on 1 seconds intervals. Using the parameters proposed by the AODV specification, the route between the SN and GW1 would time out after 3 seconds. However, since AODV uses the data traffic to update routes, the route would never time out before the TCP connection sent a new packet. Hence, the route remained active, stable and unaltered without experiencing any significant problems.

As we increased the transmission interval of data packets to a period of 4 seconds, however, we experience serious problems. When a new packet was to be sent, the route had already timed out and SN had to discover a new route to the XH over GW1. Due to dynamics in the system, sometimes the RREP from the other gateway (GW2) would be the first to arrive at the IN. In these cases, the IN established a route to the XH through GW2.

This happens since the RREPs from both gateways carry the same destination sequence number, because both copy it from the RREQ sent by the SN.

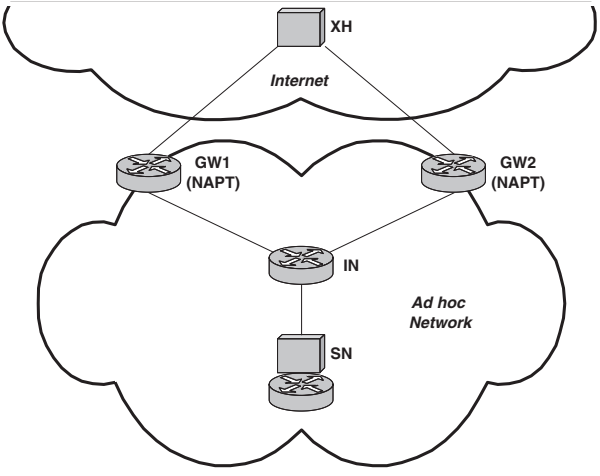

Fig. 1. Test-bed implementation with one Source Node (SN) communicating with an external host (XH) over the Internet. There are two NAT-based gateways (GW1 and GW2) present on the MANET, and all communication passes through an Intermediate Node (IN).

Furthermore, both RREPs carry the same hop-count since both gateways are one hop away from IN. Hence, when the RREP from GW1 finally arrived approximately a millisecond later, the IN did not update its route for the XH via GW1. According to the AODV specification, the IN has to discard and not process RREPs for a valid route unless its destination sequence number is higher or its hop-count is lower than those of the RREP that established the route. (Figure 2.)

When the outgoing TCP packets passed out through the NAT-module of GW2, the module naturally translated the source IP address of the TCP packets to a different address than the one used by the NAT-module at GW1. The packets were not recognized when they finally arrived at XH, and the TCP session therefore broke.

We experienced that 11% of the time the RREPs from GW2 were the first to arrive at the IN and thus established the route through GW2, whereby the TCP session would break due to the route change. Sessions were able to send only approximately 10 packets on average before they would break.

When the route changes, the XH will normally respond by a "TCP Reset" message, to which different applications will react differently. When we tested Telnet [9], for example, the application would shut down the Telnet connection immediately upon reception of the "TCP Reset" message.

There will always be a certain amount of non-deterministic dynamics in wireless ad-hoc networks, due to factors such as radio fading, node mobility, packet collisions and so forth. The two gateways will also have different performance due to the internal states of the operating systems. To study these effects un-

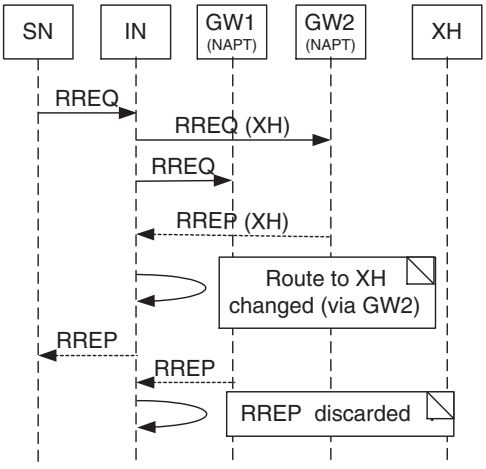

Fig. 2. A sequence diagram describing how the race conditions occur in the test-bed configuration described in Figure 1.

der more controllable conditions we emulated varying transmission times over the links by allowing the gateways to delay the transmission of outgoing Proxy RREP packets in response to RREQ requests for the XH. Each Proxy RREP was delayed for an arbitrary time between zero and T_{max} ms. (The possible varying state of the link between SN and IN was not considered significant in this trial.)

For each chosen value of T_{max}, we performed a series of 40 trials, each consisting of 400 RREQs sent with a packet transmission interval of 4 seconds. For each series of trials we incremented T_{max} from 0 to 5 ms. Results are presented in Figure 3. It shows that $T_{max} = 0$ corresponds to the case where only 11% of the RREPs from GW2 were the first to arrive at the IN. As T_{max} increases the introduced variation begins to dominate over the test-bed-specific variations, and the ratio of RREPs from GW2 arriving first at IN approaches 50% as expected, where a session breaks after having transmitted only approximately 3 packets on average.

It could be possible to implement a simple heuristic rule to eliminate the detected route race condition:

When a source node sends an RREQ to re-establish an external route, only the NAT-based gateways that have established NAT-state for that node are allowed to respond with a Proxy RREP.

Such a scheme could be implemented as follows: When a SN wants external access, it sends a RREQ with a "NAT-initiate" bit set, indicating that it has not yet established any NAT-state at any gateway. All NAT-based gateways reply with Proxy RREPs, and the first RREP returned from a gateway closest to the SN will form the outgoing route. When the SN sends out the first packet, NAT-state will be established at the gateway through which the packet is routed. If the route eventually times out, the SN sets a "NAT-reestablish" bit in the

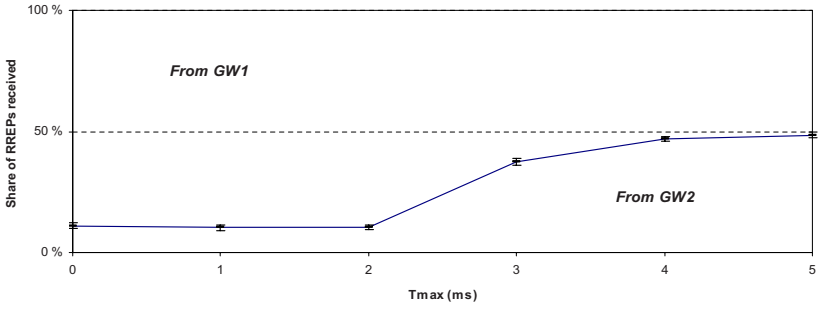

Fig. 3. The share of Proxy RREPs reaching the SN in Figure 1 (99% confidence interval).

RREQ in order to reestablish the route. Only the gateway that has established NAT-state for that SN is allowed to answer with a Proxy RREP.

Another way to eliminate the route race condition is to prohibit the use of Proxy RREPs. Instead, all NAT-based gateways would answer with RREPs establishing an outgoing route to the IP-address of the gateway. The SN can now tunnel the outgoing packets to the gateway, using for example IP-in-IP encapsulation [10] (Figure 4). The IP-address of the interface on the MANET-side of the NAT-based gateway should be used as destination IP-address of the encapsulating (outer) IP-header. The encapsulated (inner) IP-header is destined for the IP-address of XH and uses SN's IP-address as source address. The gateway must decapsulate the packets from SN *before* sending the inner IP-packets to the NAT-module.

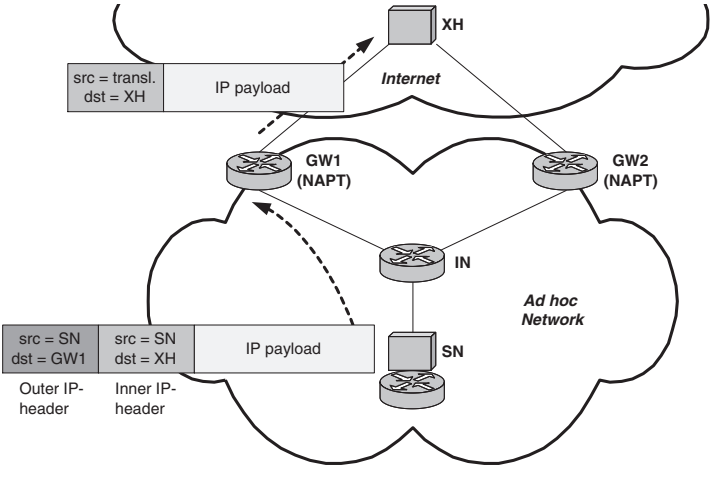

Fig. 4. Tunneling of packets to external host via gateway.

In Sub-section 3.2, it is discussed whether the Proxy RREP or the tunnelling approach is preferable.

Both these solutions allow the SN to be aware of the presence of gateways. Hence, instead of letting the gateways or the routing protocol determine whether a node is present on the MANET or not, it is possible to introduce an alternative 2-step SN-aware approach:

1. A source node first floods the MANET with RREQs to find a route to a node that is present. Gateways are not allowed to answer on behalf of external nodes.
2. If the source node does not succeed to locate the MANET locally, it issues a new RREQ with a "gateway"-bit set. When the bit is set, gateways are allowed to answer with RREPs on behalf of external nodes.

This approach allows source nodes to have better control with whether the packets are routed out of the MANET or not, which might also be a useful feature for name resolution in on-demand MANETs [11].

This solution also allows the SN to implement an integrated 1-step solution, i.e. the SN sets the "Gateway"-bit in every RREQ that it floods, and prioritizes an RREP returned directly from a targeted node that is present on the MANET.

3.2 NAT-Based Multi-homing with Two Competing Source Nodes

A site multi-homing test was performed, where the two source nodes (SN1 and SN2) both communicated with the same external host (XH). Two gateways (GW1 and GW2) were still present and both were reachable through the same Intermediate Node (IN). This is illustrated in Figure 5.

We first tested the heuristic rule using the Proxy RREP solution and no tunneling. This configuration experienced a problem caused by the fact that the IN can only have one active unicast route to the IP address of XH. This is a general limitation of unicast routing based entirely on destination IP-addresses. Due to the route race condition presented in the previous sub-section, SN1 may easily establish NAT-state over GW1 and SN2 over GW2. In this situation, SN1 and SN2 cannot communicate with XH simultaneously. The IN will either establish the outgoing route over GW1 in which the communication session of SN2 will break or over GW2 in which the communication session of SN1 breaks.

We did not experience such problems with the tunneling solution proposed in Sub-section 3.1. By tunneling packets out via the gateway, the IP-address of the XH will not appear in the routing protocol internally on the MANET. The routing protocol of the MANET will not be polluted with IP-addresses not really belonging to the MANET.

When we first introduced tunneling of outgoing packets we did not make any changes to the processing of incoming packets returned from the Internet. The NAT-based gateway would translate the incoming packets and inject them into the MANET. However, AODV uses the source and destination IP addresses of regular data-packets to reset the timers of the forwarding and return routes.

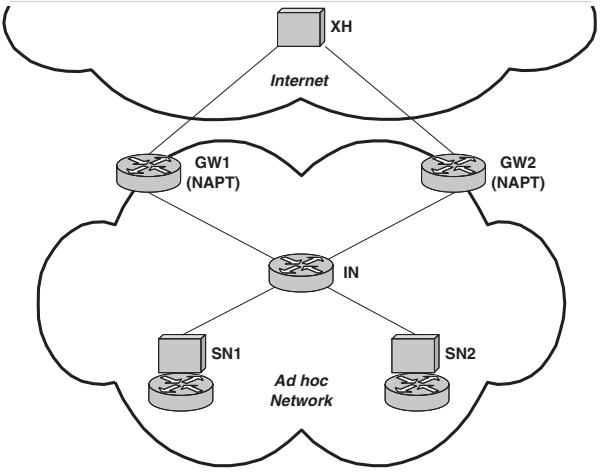

Fig. 5. A similar test-bed implementation as in Figure 1. Here there are two Source Nodes (SN1 and SN2) communicating with the same external host (XH) over the Internet.

Since an incoming packet destined for the SN carries XH's IP-address as the source address, the intermediate nodes were not able to update the timers of the outgoing route (since this is now a route to the IP-address of the NAT-based gateway).

The easiest way to deal with this problem is to let the NAT-based gateway tunnel the incoming packets to the destined source nodes, as well. Then, the gateway's IP-address appears as the source IP-address of the packets forwarded to the SN.

4 Summary and Proposed Solution

Our experiments show that the Proxy RREP solution without tunneling is not appropriate for routing between the source node and the gateways used for communicating with external hosts on the Internet. Instead a tunneled solution is required. The source node should tunnel traffic destined for the Internet to a selected gateway, and the gateway should tunnel return traffic from the Internet to the source nodes.

The following steps should be implemented to provide Internet access to MANETs:

– A SN should use the aforementioned 2-step solution when it searches for a destination. For detection of gateways the SN set a "Gateway"-bit in the RREQ when it believes that the targeted IP-address might be present on the Internet.

- Upon reception of a RREQ with this bit set, a gateway will return an RREP that establishes a route to its own IP-address. There will be no route race conditions between different gateways, which all have different IP addresses. Furthermore, there will not be race conditions between an RREP from a targeted node present on the MANET and an RREP from a gateway, independent of how the destination sequence numbers are set in the RREPs. The reason is that both the targeted node and the gateway will return an RREP that establishes a route to its own IP-address.
- A RREP returned from a gateway should contain an extension containing information about the capabilities of the gateway, e.g. whether it is a NAT-based gateway or a MIP-FA-based gateway. The source node might use this information when selecting an appropriate gateway that best matches its preferences. The extension might also include the IP-address that the source node searched for in the RREQ, so that the source node will be able to match a received RREP with a previously sent RREQ.
- To eliminate the need for additional signaling between SN and the NAT-based gateway, the gateway must set up a tunneling interface for the SN's IP-address as it returns an RREP to the SN. The tunneling interface enables the gateway to receive tunneled packets from the source node. Furthermore, the NAT will use it to tunnel packets returned from the Internet to the source node.
- The first packet that the SN to an external host establishes state in the NAT-module of the gateway after having been decapsulated.

A disadvantage is that tunneling comes at a cost, since every packet sent out of or into the MANET will carry an overhead of 20 bytes if IP-in-IP encapsulation is being used. Furthermore, the gateway will have to set up a tunneling-interface for each SN that wants to use it for Internet access, which will require it to store state information for every external connection. However, this can be managed by letting the gateway establish tunneling interfaces as soft state and remove interfaces that are not used within a certain time by setting an appropriate time out value.

The tunneling functionality can easily be integrated into the routing modules of MANET nodes. The fact that the tunneling solution is SN-aware (unlike a traditional NAT-solution) does not introduce any problem since it does not affect protocols and applications above the IP networking layer.

An IETF Internet Draft that describes explicit gateway discovery and tunneling in more detail has been published as a result our work. [12].

5 Applicability to MIP-FA-Based Multi-homing

Belding-Royer *et. al.* [2] do not explicitly consider multi-homing, and we therefore anticipate problems with race conditions also for this scheme. If there are several MIP-FA-based gateways at the same number of hops away from the SN, the outgoing route will be determined by a race condition and it might also alternate between different gateways. Since Mobile IP does not depend on outgoing

traffic to pass by the FA to which the mobile node is registered, this might not cause any significant problems (unless the outgoing traffic is subject to ingress filtering on the external network).

In a multi-homed scenario one must assume that NAT-based and MIP-FA-based gateways might be present simultaneously on a MANET (Figure 6). In this case, however, the race conditions may easily cause problems. If the SN uses NAT for external communication, outgoing packets escaping through the MIP-FA gateway will not be correctly translated. Similarly, if the SN uses the MIP-FA gateway for external communication, outgoing packets escaping through the NAT will be translated and will therefore not be recognized by the external host. However, by using explicit tunneling to the NAT-based gateway, the latter problem is eliminated.

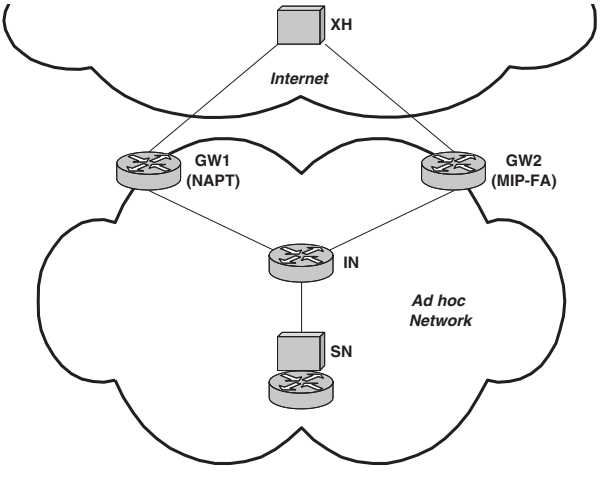

Fig. 6. A similar test set-up as in Figure 1. Here there are a NAT-based and a MIP-FA based gateways present on the MANET. The SN is capable of using both gateways. (The MIPv4 home agent of the SN on the Internet is not shown.)

To eliminate the problem of packets intended to pass through the NAT-gateway escaping through the MIP-FA gateway, we propose the following heuristic rule for MIP-FA gateways:

A gateway must not answer to an RREQ with a Proxy RREP unless the node originating the RREQ is registered with the MIP-FA on the gateway.

This rule also eliminates the problem of outgoing routes fluctuating between different MIP-FA-based gateways. Instead the outgoing traffic is forced to pass consistently out through the MIP-FA to which the SN is registered.

More work is required to determine whether this heuristic rule is sufficient for the MIP-FA-based solution, or whether such solutions also should be based on tunneling (as is proposed by the MIPMANET effort [3]). A source node might

for example register with two different MIP-FA-based gateways simultaneously (by setting the 'S'-bit of Mobile IPv4) in order to do a "make-before-break" change of gateways. In this situation, the race conditions between the different MIP-FA-based gateways will probably reoccur despite our proposed measures. Hence, tunneling might be required not only for NAT-based gateways, but also for MIP-FA-based solutions.

6 Applicability to DSR

Although our lab trials were based on AODV, the same arguments apply to DSR. However, the basic functionality that is part of the DSR routing protocol will inhibit that the aforementioned race conditions with multi-homing will occur.

A source node will use source routing - which can be considered as a sort of tunneling - to get the packet out through the selected gateway. The gateway will use source routing to get incoming packets forwarded to the source node.

The reserved 'L'-bit of DSR can be used to distinguish RREPs from gateways from a RREP returned directly from a node present on the MANET. However, an additional option containing the capabilities of the gateway might be required to make it easier for source nodes to select the right gateway with the right capabilities (e.g. describing to which extent the gateway supports MIP-FA, NAT or other kinds of gateway technologies).

7 Concluding Remarks

A number of proposed solutions for providing external connectivity to AODV-based on-demand networks allow a source node to send packets to an external host in the same way as to MANET-local hosts (i.e. without tunneling). By using Proxy RREPs, gateways respond to RREPs on behalf of external nodes, and all traffic destined for that node will be received by the gateway.

By simple replicable lab experiments with multi-homing, however, we have shown that this easily leads to routing race conditions. This again makes it difficult for a source node to control which gateway that will be used for external communication. Traffic might be directed over different gateways in a non-deterministic way.

With a NAT-based gateway, all outgoing packets belonging to the same communication session must pass through the same gateway, and the aforementioned race conditions must be avoided. The most obvious solution is to mandate that a source node tunnel packets to a selected gateway. This means that the source node must explicitly discover available gateways, and select one based on their different capabilities.

With a MIP-FA-based solution, the race conditions are less critical, since packets are allowed out through any gateway (unless the external network is subject to ingress filtering). We have proposed a simple heuristic rule to avoid routing race conditions, even when the Proxy RREP solution without tunneling

is being used. However, whether these measures are sufficient, or whether MIP-FA based solutions also should be based on tunneling, is a subject for further work.

Since test-beds are prone to non-deterministic changes of a number of physical parameter that are difficult to control, the analyses and results presented in this paper should be confirmed by simulations using an appropriate simulation tool. Simulations can provide numerical results that describe the multi-homing problems experienced in a test-bed in greater detail.

References

1. MANET Working Group of the *Internet Engineering Task Force (IETF),* homepage, *http://www.ietf.org/html.charters/manet-charter.html.*
2. Belding-Royer *et al.*, "Global connectivity for IPv4 Mobile Ad Hoc Networks", *IETF Internet Draft*, draft-royer-manet-globalv4-00.txt, November 2001 (Work in Progress).
3. Alriksson, F., and Jönsson, U., "MIPMANET - Mobile IP for Mobile Ad Hoc Networks", Master of Science Thesis, Royal Institute of Technology (KTH), http://www.e.kth.se/∼e94_fal/mipmanet.pdf, August 1999.
4. Perkins C. (ed.), "IP Mobility Support for IPv4", *RFC 3344, Internet Engineering Task Force (IETF),* August 2002.
5. Srisuresh, P. and Holdrege, M., "IP Network Address Translator (NAT) Terminology and Considerations", *RFC 2663, Internet Engineering Task Force (IETF),* August 1999.
6. Uppsala University's implementation (version 6) of AODV. http://user.it.uu.se/∼henrikl/aodv/
7. Perkins, C.E., Royer, E.M., and Das, S.R., "Ad-hoc On Demand Distance Vector (AODV) Routing", *RFC 3561, Internet Engineering Task Force (IETF),* July 2003.
8. Johnson, D.B., Maltz, D.A., Hu, Y.-C. and Jetcheva, J.G., "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR)", *IETF Internet draft.* draft-ietf-manet-dsr-09.txt, April 2003 (Work in Progress).
9. Postel, J. and Reynolds, J., "Telnet Protocol Specification", *RFC 854, Internet Engineering Task Force (IETF),* May 1983.
10. Perkins, C.E., "IP Encapsulation within IP", *RFC 2003, Internet Engineering Task Force (IETF),* October 1996.
11. Engelstad, P.E. and Egeland, G., "Name Resolution in On Demand MANETs", *IETF Internet draft*, draft-engelstad-manet-name-resolution-00.txt, February 2003 (Work in Progress).
12. Engelstad, P.E. and Egeland, G., "Explicit Gateway Discovery in IPv4 On Demand MANETs", *IETF Internet draft*, draft-engelstad-manet-gateway-discovery-v4-00.txt, October 2003 (Work in Progress).