

Scalable Support for Globally Moving Networks¹

Marcelo Bagnulo, Alberto García-Martínez, Carlos J. Bernardos, Arturo Azcorra
Departamento de Ingeniería Telemática - Universidad Carlos III de Madrid
{marcelo, cjbc, alberto, azcorra}@it.uc3m.es

Summary

This paper proposes a scalable solution for the support for globally moving networks. It is basically oriented to airborne mobile networks built in commercial aircrafts in order to provide Internet access to the passengers. As opposed to currently used solutions, the proposed solution has no impact in the global routing tables while it provides optimized paths between the mobile network and the rest of the Internet. The proposed solution is an extension to the IETF standard network mobility support protocol and relies on the communication through multiple geographically distributed Home Agents in order to avoid panoramic routing imposed by single anchor points as in the case of a single Home Agent. The proposed solution includes a mechanism to select the best Home Agent to route new communications through.

1. Introduction

Internet access provision on board of long distance flights is a challenging problem in many aspects. Basically, the scenario under consideration is a network of several hundreds of devices that is moving at a speed close to the 1.000 km. per hour during several hours. The result is that the moving network may travel across half of the globe in a dozen of hours. The question is how to provide session continuity and an acceptable performance for the mobile network nodes under these conditions. According to Boeing, the aircraft network will access the fixed Internet using 3 different ground stations and 4 different satellite links during a typical flight between Europe and Asia [1]. The problem then is how to provide seamless handover during the changes of ground stations, especially considering that the different ground stations are very far apart both in the geographical and in the topological dimension. The problem is not only about session continuity but also about performance in the communications. Considering the large distances between two consecutive attachment points to the Internet, the performance of the Internet

access is also a critical issue. In other words, using an optimized route that is as much direct as possible between the node in the mobile network and the node in the Internet, without imposing panoramic routing due to fixed anchor points in the topology is an important requirement for any solution in this space.

As the day of today, there are several airlines that provide Internet access to the passenger of long distance flights. In the case of the Connexion Service by Boeing (see <http://www.connexionbyboeing.com/>), the selected approach to provide Internet connectivity is based on the injection through the Border Gateway Protocol of a route in the inter-domain routing system towards the prefix assigned to the moving network of the aircraft [1]. While this solution provides the required capabilities in terms of passenger service, it fails to meet the scalability requirements needed for a widespread adoption of such technique. This is so because each moving network contributes with an additional entry in the already oversized global routing table [2]. Currently, the in-flight Internet access service is limited to a reduced amount of aircrafts and selected flights, but as the service becomes a commodity, the impact in the global routing table in terms of churn will be considerable.

This paper proposes an alternative solution that provides Internet access on board to aircrafts while preserving the scalability of the inter-domain system. The proposed solution is an extension to the Network Mobility (NEMO) Basic Support defined in [3] [4]. In a nutshell, the idea is to use multiple Provider Aggregatable prefixes in the mobile network (one per geographically distributed Home Agent (HA) available) and then to provide enough information to the moving network so that it is able to select the closest Home Agent (HA) and its associated Mobile Network Prefix to initiate new communications.

The paper is structured as follows. The motivation for a scalable solution for globally moving networks is presented in Section 3. Section 4 provides an overview of the proposed mechanism, which is explained in more detail in Section 5. In Section 6, we explore different alternatives for the support of global moving networks

¹ This work has been supported by Optinet project TIC-2003-09042-C03-01 and IMPROVISA project

and compare them to our approach. Finally, Section 7 concludes the paper.

3. Background

3.1. Problem Characterization

The problem that is addressed in this paper is the provision of Internet access on board of flying aircrafts. The users of the Internet access are passengers and crew inside the aircraft that want to communicate with other nodes in the Internet. These are not critical communications in the sense that neither human life nor ship integrity depend on them. Air traffic control messages are not to be exchanged through these communication mechanisms, but alternative communication systems support such traffic. However, the following requirements are imposed to the solution for in-flight Internet access provision. First of all, session survivability through movement must be provided, meaning that ongoing communications must be preserved as the plane changes its attachment point to the Internet. In addition, the quality of the communication needs to be ensured, meaning the round-trip time must be bounded to acceptable limits, guaranteeing user satisfaction while there are communicating. This basically means that the solution needs to provide a somehow optimized path between the mobile network nodes and the peers in the Internet, especially considering the large distances involved in this particular scenario.

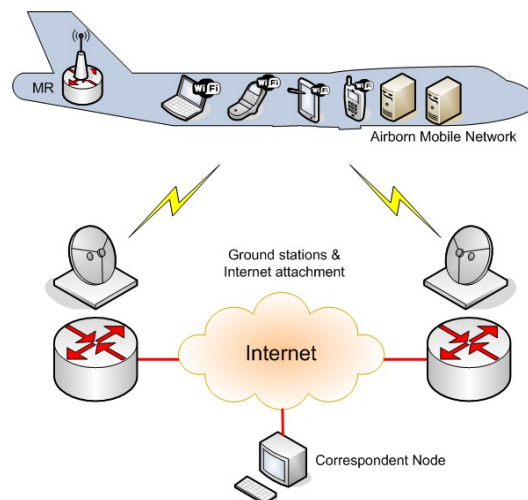
So, the considered scenario is depicted in figure 1. We have a mobile network located inside the aircraft that is moving around the globe. It is changing its attachment point to the Internet when it changes the ground station that it is using to communicate. While the changes in the attachment point to the Internet are not very frequent (one every few hours) the distance within the Internet topology between two consecutive attachment points can be considerably big (e.g., they can be located in different ISPs in different continents).

3.2. Limitations of NEMO protocol

The Network Mobility protocol as defined in [3] [4] has been designed to provide support for moving networks in the Internet. It is a general tool that is well fitted for most common scenarios. It has been defined as an extension to Mobile IP [5] [6] in order to perform bindings between prefixes assigned to the mobile network (called Mobile Network Prefixes, MNPs) and the current location of the Mobile Router (MR) that provides access to the mobile network. In the general NEMO scenario, mobile network nodes configure an address from the mobile network prefix and they obtain Internet connectivity from the MR. The mobile network prefix is part of the home network prefix, which is

announced to the Internet through the ISP of the Home Network. This means that packets addressed to the mobile network will be routed to the Home Network. In the case that the mobile network is roaming, packets addressed to the mobile network that reach the home network are forwarded to the MR by the HA.

Figure 1: Aircraft Mobile Network



This general setup is designed to provide mobile network support in a general scenario. However, in this particular case in hand, it fails to provide an optimized support. This is so, because the distance between the Home Network/HA and the MR can be considerably large. In this case, routing all packets through the Home Network imposes additional penalties in terms of delay that render the approach described above unacceptable.

4. Overview of the solution

The proposed approach is based on the usage of multiple HAs distributed geographically in multiple points along the trajectory of the plane. This is possible because, first the plane flight plan is known before-hand (as opposed to the general mobility case) and the airline company has commercial relationships with the local providers in the different locations, making possible to assume that it is acceptable for the local ISP to host a Host Agent for the air company.

Each of these HAs suppose a local Home Network in each of these locations. The result is that the mobile network will have multiple home networks and multiple HAs distributed along the globe. Each of these home networks will have a home network prefix that belong to the Provider Aggregatable (PA) block of the local ISP, implying that they do not suppose an additional contribution to the global routing table, since the home network prefix is contained in the local ISP aggregate. This implies that the mobile network will have multiple mobile network prefixes, one per Home network and

that mobile network nodes will have multiple addresses, one per available mobile network prefix. Note that the selection of the address of the mobile network node to use in a communication determines the HA through which the communication will be routed, which implies selecting the attachment point of the Internet through which the communication will be coursed. This means that if the prefix associated to a closer HA is selected for a new communication, the communication performance will be probably acceptable, while if the prefix that is associated to a far away HA is selected the communication performance may be really poor. However, current NEMO protocol does not provide any support for determining which HA is closer. In order to fill this gap, a new extension to NEMO is proposed in this paper, to assess the distance between the HA and the MR and enable an optimal address selection on the mobile network nodes. The basic idea is to enable to HA to measure the distance with the current location of the MR. The proposed metric for this measurement is the AS path length. The proposal is to provide the BGP information to the HA, so that given the CoA of the MR, the HA can determine the distance with the MR measured in AS hops. Once the HA has obtained such information, it can be passed it to the MR, using an extension to the NEMO protocol defined in this paper. Through this mechanism, the MR will obtain information about the distance with all available HAs, so it is capable of determining which HA is closer. Once that the MR has identified the closest HA, it informs the mobile networks nodes about which is the preferred prefix to be used for new communications, so that an optimized path is used.

5. Scalable Support for Globally Moving Networks

In this section, we describe in detail the components of the proposed solution for a scalable support for globally moving networks.

By deploying multiple globally distributed Home Networks, a globally moving network is allowed to choose the closest Home Agent to forward its traffic. In this way, an improvement in the performance of the communications of a mobile network – when compared to the case of a single HA scenario - is enabled, since the delay introduced by the MR-HA bidirectional tunneling can be reduced. The limit on the delay depends on the number of distributed HAs deployed and on their particular locations.

The solution basically consists of three different parts: Home Agent Selection, Address Management within the moving network and Route Optimization of long-lasting communications established before performing a handover.

5.1 Home Agent Selection

The Mobile Router, when moving around the globe, updates its current location on all the distributed HAs, but selects a primary one to be used to route the new communications of the moving network. The Mobile Router needs to select the most appropriate Home Agent (i.e., the closest one) among the possible ones distributed geographically. In order to do that, the MR is provided with a metric that enables it to choose the closest HA. The proposed metric is the BGP AS path length.

When a HA receives a Binding Update (BU) from a MR, the HA calculates the distance with the MR's CoA measured in AS hops and conveys this information back to the MR into a new mobility option in the Binding Acknowledgement (BA) message. Once the MR has received all BA messages it can compare the received AS path lengths and choose the closest HA in terms of BGP AS hops for routing its traffic.

There are other possible metrics that could be considered besides AS path length. A metric based on the Round Trip Time (that could be obtained by the HA sending ping probes to the MR's CoA) could provide a more precise estimation of the distance between HA and MR in some scenarios, but this kind of metric is more complex and requires additional signaling to be exchanged. Besides, the path estimation resulting from measuring a few probe packets may not be good enough to forward all the traffic associated with all the communications of the mobile network. A manually pre-established configuration (e.g., the MR is configured beforehand with the sequence of HAs that should be used for a given flight) could also be possible, but it would require to configure the MR for each particular flight.

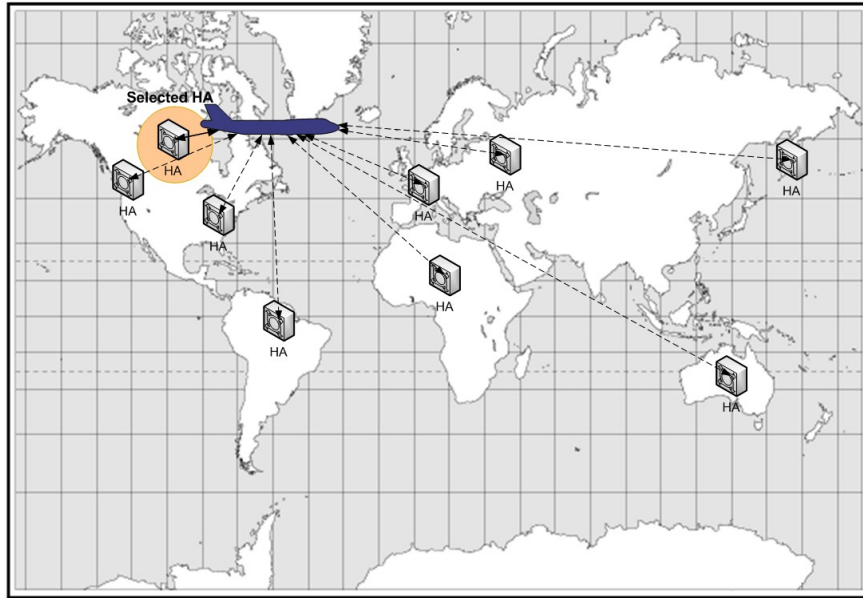


Figure 2: Home Agent selection

5.2 Address management

When the MR detects that there is a HA that is closer than the one that it is using, it switches to the closest one for new communications, so the additional delay due to the MR-HA path is decreased.

The HA used to forward the traffic of a node of the mobile network is determined by the address that this node is using, e.g., if a node is using as its source address for the packets it sends an IPv6 address that has been configured from the Mobile Network Prefix of a HA X (MNP X), then these packets will be sent by the MR through HA X. Similarly, packets destined to that node will be received by the MR from HA X. Therefore, if a particular HA is preferred to be used for new communications, the MR has to inform the nodes of the moving network of the address they should use in upcoming attempts of communication. To do that, the MR deprecates the prefix that was being used in the network and starts advertising the prefix belonging to the chosen HA (that is, the MNP of the closest HA) as the preferred one [9] [10]. In this way, nodes of the mobile network will keep using their old addresses for established communications (during the lifetime of these communications) and will use the currently preferred address for new communications.

5.3 Route Optimization of long-lasting sessions

The optimization proposed in this paper is not suited for long-lasting traffic, since it does not support to change the HA used to route traffic of a particular communication during its lifetime. We consider

additional mechanisms to deal with long lasting sessions, which remain established after changes in the closest HA. In this case, the address used for initiating the communication is no longer the one that provides optimal routing. In order to deal with these cases, additional route optimization mechanisms are required to improve the performance.

It should be noted that the 98% of Internet communications last less than 15 minutes [11], so they are likely to obtain a reasonable performance by simply using the previously described mechanisms, since the same HA will be used during the lifetime of the communication. However, for the remaining 2% of the communications we propose to use a Route Optimisation solution to avoid the sub-optimal routing of this long-lasting traffic

MIRON [12] is a proposal of a solution for Route Optimization that does not require upgrades in CNs, MNNs, or HAs. MIRON has two modes of operation: the MR performing all the Route Optimisation tasks on behalf of those nodes that are not mobility capable - thus working as Proxy MR - and an additional mechanism, enabling mobility-capable nodes (i.e. Mobile Nodes attached to a NEMO) that actually have mobility and Route Optimization capabilities to manage their own Route Optimization. Given the characteristics of MIRON, it seems to be a perfect candidate for optimizing the traffic of long-lasting communications in globally moving networks.

6. Comparison with previous work

This section compares the proposed mechanism with alternative approaches proposed in the available literature.

The approach followed in the Connexion by Boeing service [1] basically consists in using BGP as a mobility solution, by means of the use of the global routing table and selective route announcements and withdrawals as planes move. Each aircraft is assigned a public IP /24 address block, to allow for the mobility of the aircraft to be seen through the global default free zone. The aircraft's route is announced from the ground station that is currently serving the plane. When the aircraft moves to another ground station, the public IP address block is withdrawn from the BGP table in the "old" ground station and is then announced by the next ground station that will provide service to the plane.

This solution has several problems. First it introduces instability in the global routing tables (the ASes used by the Connexion by Boeing service are within the top 20 in terms of number of added and withdrawn routes [13]). Besides, the solution is not scalable, since a /24 route is injected in the BGP global routing table per aircraft. Last, but not least, the convergence time required by BGP has a noticeable impact in the handover time.

Another proposed approach for supporting globally moving networks is the Global HAHA protocol [7]. The solution basically assumes that several Home Agents serve the same home network together at the same time. To do so, the same route of the aggregated home network is advertised from separated home links, therefore potentially suffering from similar scalability concerns that the approach of Connexion by Boeing.

There are several generic Route Optimization mechanisms for NEMO already proposed [8]. Most of them could be also used to optimize the traffic of globally moving networks. However, since statistical studies of the distribution of session lifetimes show that over 98% of the sessions last less than 15 minutes [11], our proposed approach enables to likely obtain an optimal performance without needing to use generic route optimization mechanisms - that would increase the complexity of the solution - for all the traffic, but only for long-lasting communications.

7. Conclusion

The NEMO Basic Support protocol enables whole networks to move and change their point of attachment, transparently to the nodes of the network. However, this solution introduces some limitations and problems in terms of performance (increased delay in packet delivery). The impact of these problems is exacerbated in global moving networks, since panoramic routing is imposed due to the use of single fixed Home Agents.

To overcome these limitations, we propose an extension to the NEMO Basic Support protocol that

relies on the communication through multiple geographically distributed Home Agents. A simple metric based on the BGP AS path length between the HA and the MR enables to select the best Home Agent to route new communications through.

The proposed approach provides a simple yet efficient solution for the problem in hand. The proposed solution enables an almost optimal path for the communication without imposing much complexity and preserves the scalability of the global routing system.

8. References

- [1] A. Dul, "Global IP Network Mobility", in Proceedings IETF 62, March 2005.
- [2] G. Huston, "Commentary on Inter-Domain Routing in the Internet", RFC 3221, December 2001.
- [3] K. Leung, G. Dommety, V. Narayanan, A. Petrescu, "IPv4 Network Mobility (NEMO) Basic Support Protocol", Internet-Draft (Work in progress), February 2006.
- [4] V. Devarapalli, R. Wakikawa, A. Petrescu, P. Thubert, Network Mobility (NEMO) Basic Support Protocol, RFC 3963, January 2005.
- [5] D. Johnson, C. Perkins, J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [6] C. Perkins, "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [7] P. Thubert, R. Wakikawa, V. Devarapalli, "Global HA to HA protocol", Internet-Draft (Work in progress), October 2005.
- [8] C. Ng, F. Zhao, M. Watari, P. Thubert, "Network Mobility Route Optimization Solution Space Analysis", Internet-Draft (work in progress), February 2006.
- [9] T. Narten, E. Nordmark, W. A. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [10] S. Thomson, T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [11] N. Brownlee, K.C. Claffy, "Understanding Internet traffic streams: dragonflies and tortoises", IEEE Communications Magazine, Vol. 40, no. 10, pp. 110-117, October 2002
- [12] M. Calderón, C. J. Bernardos, M. Bagnulo, I. Soto, A. de la Oliva, "Design and Experimental Evaluation of a Route Optimisation Solution for NEMO"; to appear in IEEE Journal on Selected Areas in Communications (J-SAC), issue on "Mobile Routers and Network Mobility";
- [13] "CIDR REPORT for 26 Jun 06", <http://bgp.potaroo.net/cidr/>