A Thorough Investigation of Mobile IPv6 for the Aeronautical Environment

Christian Bauer and Serkan Ayaz
Institute of Communications and Navigation
German Aerospace Center (DLR)
Wessling, Germany
Email: {Christian.Bauer, Serkan.Ayaz}@dlr.de

Abstract—**Aviation is currently in the process of adopting IP as the network protocol for the Aeronautical Telecommunication Network (ATN), considering the fact that future communication services will be mainly based on data instead of analogue voice services. One of the crucial problems during this adoption is the mobility problem which is complicated by the fact that aircraft have a degree of world wide mobility and may make use of a variety of heterogeneous link technologies. Mobile IP, the standard IETF mobility protocol for the future, is investigated in this paper regarding its applicability to the aeronautical environment, based on realistic scenarios.**

I. INTRODUCTION

It has been identified that the communication infrastructure currently used for Air Traffic Management (ATM) can neither cope with the expected increase in air traffic nor can it support the envisaged paradigm shift in ATM. Related institutions in the area of civil aviation are currently developing the future communication systems that will make efficient data links available (when compared to today).

The International Civil Aviation Organization (ICAO) is currently working on guidance material for the Aeronautical Telecommunication Network (ATN) with IPv6 as the network procotol for air/ground communication.The global mobility problem is one of the key issues for which no solution has been decided on yet, although a comprehensive list containing potential solution candidates has been compiled [1].

The aeronautical use case is different from other domains (e.g. personal wireless communication) due to the existence of safety-related (i.e. Air Traffic Services) data which should be handled such that the required Quality of Service is fulfilled.

Mobile IPv6 (MIPv6) [2] is a protocol of the Internet Engineering Task Force (IETF) that supports node-based global mobility for a mobile node (MN) to allow communication with correspondent nodes (CNs) via the Home Agent (HA), which serves as a global mobility anchor. Both 3GPP and 3GPP2 have adopted the procotol in their most recent releases, and making use of it in the ATN would seem to be a reasonable choice. Two proposed extensions to MIPv6 – Network Mobility (NEMO) [3] to support a mobile router instead of just a mobile host and Mobile Nodes and Multiple Interfaces (MONAMI6) [4] to support multihoming – play an important role but rely both on the MIPv6 protocol.

An investigation of Mobile IPv6 for an aeronautical scenario has already been performed in [5], however this paper was based on several simplifications and did not take into account the service requirements published in the ICAO approved Communications Operating Concept & Requirements document (COCR) [6]. Our simulations are based on realistic flight data, simplified terrestrial and satellite link technologies and an accurate MIPv6 model.

The paper is structured as follows: we first introduce the link technologies that are foreseen to be used in the future, followed by a short insight into the communication services in aviation. In section IV an analytical investigation of Mobile IPv6 overhead is conducted, taking into account the bandwidth available on aeronautical data links. We then explain the assumptions and structure of our test scenarios and the results we have obtained from the simulations. We then analyze the obtained data and conclude with several suggestions.

II. LINK TECHNOLOGIES

The link/access technologies that are envisaged for use in aviation are as follows:

- The short range data link 802.16eAV, based on the WiMAX standard, used for communication at the airport.
- A long range data link, also called L-band Digital Aeronautical Communication System (LDACS), for the so-called terminal manoevering area and en-route flight domains, during which an aircraft is flying over continental areas.
- One or several satellite systems especially for use in non-continental areas (i.e. Oceanic, Remote and Polar regions) where no other link technology is available.

The LDACS, with limited bandwidth and high delay when compared to terrestial technologies s.a. 3GPP LTE, will become the standard IP based wireless link in the future ATN. Our investigations will therefore be based mainly on this link technology. Furthermore, Table I gives a list of reasonable values for all three link categories as described above. It is important to keep in mind that the available bandwidth is shared among all aircraft within one cell (radius usually 120

This work is partially funded by the European Commission through the NEWSKY project under contract no. 37160.

TABLE I
TYPICAL VALUES FOR LINK TECHNOLOGIES FORESEEN TO BE USED IN AVIATION. 62 USERS, FROM [7]

Technology	Throughput	One-way delay (in ms)
IEEE 802.16eAV	14 Mbps	<20-70
LDL[1]	37.5-100 kbps	N/A
P-34[1]	76.8-691.2 kbps	<1240
B-AMC[1]	300 - 1600 kbps	<1000
DVB-S2[2]	FL[3] 80 Mbps	400
	RL[3] 8 Mbps	
Inmarsat BGAN[4]	22 Mbps	<400

nautical miles for LDACS), which can be a number in the magnitude of several hundreds.

III. AIR TRAFFIC SERVICES

There are four different service classes in aviation, and we are focussing on the most important one for our investigation: Air Traffic Services (ATS), which covers all communication between the cockpit and the controller on the ground in order to perform Air Traffic Management. The ATS services are specified in the COCR [6], which defines the amount of data to be exchanged and delay requirements that have to be fulfilled. The communication end point for ATS is an Air Traffic Services Unit (ATSU), which changes from time to time depending on the geographical position of the aircraft.

IV. DELAY ANALYSIS

The delay during a Mobile IPv6 handover (HO) is usually in the order of seconds [8] and consists of several steps (we are assuming Stateless Autoconfiguration [9] for our purpose as it is the fasted method):

- Perform layer 2 handover
- Movement Detection: receiving a Router Advertisement (RA) from a new access router (AR) and potentially detecting the unavailability of the old AR
- Configuring a care-of address: Duplicate Address Detection (DAD) involves sending a Multicast Listener Discovery (MLD) report message to join the solicited-node multicast group and sending a Neighbor Solicitation (NS) to this group to check whether any other node is already using the tentative IP address.
- Sending Binding Update (BU) to and receiving a Binding Acknowledgement (BA) from the Home Agent
- Performing route optimization (RO) with the CN that involves Return Routability with a Care-of Test (CoT) and Care-Of Test Init (CoTI) message exchange between MN and CN directly as well as Home Test (HoT) and Home Test Init (HoTI) between MN and CN via the HA, followed by an BU/BA exchange between MN and CN.

We use the variables listed in Table II to denote the time necessary for the individual signalling messages.

TABLE II
STAGES IN THE HO PROCESS AND VARIABLES DENOTING THE IMPOSED DELAY. RTT DENOTES ONE ROUND TRIP TIME

Signalling/action	Variable	Typical Value
Layer 2 Handover	t_{L2HO}	L2 specific
Receiving new RA	t_{RA}	$30 - 70$ms
Movement Detection	$t_{movement}$	$2 * t_{RA}$
MLD report message	t_{MLD}	$0 - 1$sec
DAD	t_{DAD}	1sec
BU/BA to/from HA	t_{B-HA}	1 RTT MN–HA
CoTI/CoT	$t_{CoT(I)}$	1 RTT MN–CN
HoTI/HoT	$t_{HoT(I)}$	1 RTT MN–HA–CN
Return Routability	t_{RR}	MAX($t_{CoT(I)}$, $t_{HoT(I)}$)
BU/BA to/from CN	t_{B-CN}	1 RTT MN–CN

TABLE III
SIGNALLING OVERHEAD. MESSAGE SIZES IN BYTES, W/O IPSEC.

Message	IPv6 Header	HoA Option	Mobility Headers	Sum
BU to HA	40	20	6 + 6	72
BA from HA	40	0	6 + 6	52
HoTI	40+40	0	6 + 10	96
HoT	40+40	0	6 + 18	104
CoTI	40	0	6 + 10	56
CoT	40	0	6 + 18	64
BU to CN	40	20	6+6+14+6	92
BA from CN	40	0	6+6+14+2	68

Especially the MIPv6 signalling can be dominated by the wireless link due to the significant delays on the aeronautical links, as discussed in Section II.

V. OVERHEAD INVESTIGATION

Our next step in the investigation of the suitability of MIPv6 for the aeronautical environment is to analyze the signalling overhead of the route optimization procedure and compare it to the bandwidth available from the LDACS, which is the most bandwidth critical link (cf Section II).

The mobility signalling overhead can be seen in Table III.

The lifetime of a correspondent binding is limited to 7 minutes (MAX_RR_BINDING_LIFETIME in [2]) for security reasons and therefore requires constant Return Routability (RR) and BU/BA signalling. We are estimating the resulting overhead as follows:

- Both RR (CoTI/CoT and HoTI/HoT) and BU/BA takes place every seven minutes.
- The accumulated value of the signalling is 480 bytes / 7 minutes (cf. Table III).
- The number of aircraft is based on the Peak Instantaneous Aircraft Count (PIAC) per cell [10].
- We are assuming high density regions s.a. in Europe, where the PIAC in a radio cell could indeed be constant over a long period of time.
- The increase in payload due to IPSec, which is applicable to all signalling messages that are routed over the MN-HA tunnel, is not taken into account.

Table IV shows the resulting overhead, taking the PIAC and multiplying it with 244 bytes for the return link (HoTI+CoTI+BU to CN) and with 236 bytes for the forward link (HoT+CoT+BA from CN).

---

[1] Candidate for becoming LDACS

[2] Digital Video Broadcasting - Satellite Second Generation (S2); assuming GEO satellite

[3] FL = Forward Link, RL = Return Link

[4] Broadband Global Area Network; GEO satellite system

TABLE IV

MIPv6 RR SIGNALLING OVERHEAD DEPENDING ON PIAC AND CAPACITY.

PIAC		45	62	204	522
Capacity FL[5]	Kbps	150	150	300	500
Overhead FL	Kbytes/7 min	82.97	114.31	376.13	962.44
	Kbits/sec	1.58	2.18	7.16	18.33
Ratio	in %	1.05	1.45	2.39	3.66
Capacity RL[5]	Kbps	30	30	40	50
Overhead RL	Kbytes/7 min	85.78	118.19	388.89	995.06
	Kbits/sec	1.63	2.25	7.41	18.95
Ratio	in %	5.45	7.50	18.52	37.91

TABLE V

VALUES USED FOR LINK BUDGET AND DELAY. LDACS VALUES FROM [14].

Technology / Property	LDACS		Satellite
	medium	high	
Forward Link Delay (ms)	72	271	250
Return Link Delay (ms)	169	552	250
Thermal noise (dBm)	−117		−162
TX power (W)	3.2364		105.6896
RX sensitivity (dBm)	−107		−155
Bandwidth (Mhz)	0.5		36
Bitrate (Mbps)	0.355		14.71

The results reveal that the signalling overhead in the worst case (RL for 522 aircraft) amounts to nearly 38% of the (minimum) available capacity, while it is 3.7% for the forward link, also for the worst case. The binding between MN and HA has to be renewed every one hour and the corresponding overhead is therefore negligible.

VI. SIMULATION

Our simulation environment is based on the discrete event simulator Omnet++ [11] and its IPv6 framework INET. Our focus is on simulating inter-access network handovers [12] which trigger the MIPv6 protocol; the measurement of the handover delay and transmission time over the final end-to-end path to the correspondent nodes is of special importance for us and will be compared with the service requirements from COCR.

A. Links and Mobile IPv6

For the sake of computational efficieny (and due to the fact that we do not want to perform an analysis of the link technologies) we have implemented a *generic network interface card* (generic NIC) that allows specifying link budget and delay per NIC as a flexible input to the simulation, as shown in Table V. The NIC also provides *Link_Down* signalling in case the association to a base station is lost - this allows to perform a relatively fast switch from one network card to another one, although it is not as powerful as the primitives provided by e.g. the IEEE 802.21 MIHF [13].

Physical channel and Medium Access Control (MAC) related errors (i.e. bit errors, collisions) have been omitted in the simulations. However the specified values that were used to simulate the delay of the MAC have been taken from appropriate simulation results which already take into account retransmissions on the link layer and are also relative to the number of aircraft in the radio cell (PIAC) – we have focussed on results for medium and highly loaded cells; see Table V.

An aircraft can also be seen from the NEMO [3] perspective, as a mobile network with several on-board hosts for which the mobile router handles mobility. Assuming pre-configured mobile network prefixes there is no additional signalling necessary in comparison to MIPv6. Hence our investigation is also relevant for the case when a mobile router is on the aircraft, although it is important not to forget that NEMO does not yet support RO.

Movement detection was simplified in our simulation as to not wait for any missed RAs from the old AR - as soon as a RA with a new prefix is received, the MN immediately starts to configure a CoA based on it. Hence, the variable $t_{movement}$ (cf. Section IV) is ignored for the subsequent analysis.

B. Ground Network

Another important factor is the structure of the ground network, due to the additional delays for packet forwarding. Besides standard IPv6 routers we have added additional entities within our network:

- Home Agent, located in Cologne/Germany (e.g. Airline HQ).
- ATSU correspondent node Gander OCC, located in Canada, is the responsible ATSU for the western half of the north atlantic airspace.
- ATSU node Madrid ACC covers the airspace around Madrid up to the Spanish–French border.

The ground network itself is based on realistic data:

- Structure of the Autonomous Systems (AS): We have investigated interconnections/peerings between various ASs with the online RIPE NCC [15] and ARIN [16] databases to define a reasonable ground network structure.
- Delays between the different ASs are modeled based on information from the DIMES project [17] and on the guaranteed maximum delays from the service level agreement of a tier one backbone operator.

This network can be seen in Fig.1. Two backbone routers provide transatlantic connectivity for the access networks of the Aeronautical Communication Service Providers (ACSPs): ARINC in America and Europe, SITA in Europe and the satellite operator Inmarsat (assuming IP connectivity via the BGAN service) with a satellite covering the area from the US east coast up to Europe and a corresponding gateway (GW) in the Netherlands. The end host Gander OCC is connected via the American ARINC network, whereas Madrid ACC is connected via SITA. The connection of the HA to the EU backbone router is simplified as that it is modeled as a single ethernet line without any intermediate hops, although taking into account the delay from a regional service provider to the backbone as it is in reality. Figure 1 also contains the precise delay values used – the connections between the base stations and the access routers do not produce any delay.

---

[5] Capacity requirement for the LDACS as defined in [6]

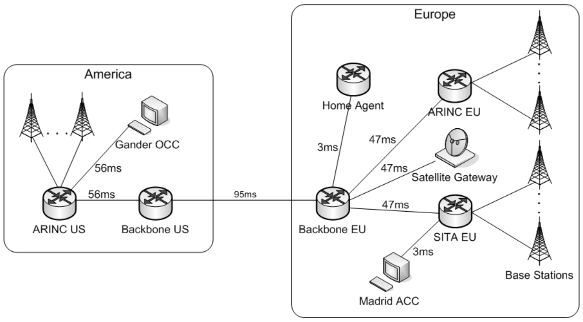

Fig. 1. Global network structure with reasonable delay values (geographically simplified).

C. Scenarios

We are using realistic flight data that covers all arrivals at Frankfurt Airport for a certain day in a 30min time window, from which we have extracted those flights which are interesting for our handover simulation. The scenarios we have decided on are as follows:

1) Flight number DLH443, departing from Detroit. This scenario only covers the starting period where the aircraft attaches to the LDACS system in the US which provides us with values on the mobility signalling between the aircraft and the Europe based HA.
2) As Scenario 1, but the aircraft now leaves LDACS coverage over Canada at a position of about 50°N where a handover to the satellite system becomes necessary. CN is the ATSU Gander OCC.
3) Flight number DLH4417, departing from Madrid, starts with an attachement to SITA and performs an inter-access network HO to ARINC, which could be due to a network-initiated HO triggered by traffic overload. LDACS performance delays are based on medium delay (cf. Table V). CN is the ATSU Madrid ACC.
4) As Scenario 3, but involves different delay values for the LDACS (high delay).

D. Results

30 simulation runs were performed for each scenario; the results for the handover delays are presented in Table VI. $t_{CoT(I)}$ and $t_{HoT(I)}$ provide us with the RTT of the end-to-end path (MN–CN) that can be used to determine the delay for communication over the optimized path and the MN–HA–CN path. Base stations or GWs send beacons in random intervals, therefore also randomizing the point in time when a HO is performed.

For all scenarios, except the third one, t_{L2HO} takes a relatively large amount of time. The layer 2 HO with its necessary request-response messaging already consumes one RTT over the wireless link between the MN and the base station/gateway, which is in the order of 300+ milliseconds; the rest is due to the random reception delay of the base station beacon. It should be noted that a real satellite HO usually takes even more time due to antenna repointing, gateway log-off/-on, etc.

TABLE VI
HANDOVER RESULTS (IN MS).

Variable	Scenario 1	2	3	4
t_{L2HO}	-	697	365	782
t_{RA}	-	26	27	27
$t_{MLD} + t_{DAD}$	-	1521	1553	1563
t_{B-HA}	2461	1621	903	2052
$t_{CoT(I)}$	497	1029	452	894
$t_{HoT(I)}$	1217	1042	596	1071
t_{B-CN}	498	2029	486	861
Complete HO	-	6936	3941	6347

The accumulated value of $t_{MLD} + t_{DAD}$ is within the expected bounds: The expectation value of the MLD report message equals to 0.5s whereas DAD is always a constant of 1s.

More interesting is the time it takes to receive the valid BA from the HA, denoted by t_{B-HA}, which – except for Scenario 3 – takes over 1s. The explanation is that both the satellite GW and the ARs in the LDACS subnet do not yet have a neighbor cache entry for the MN: Upon reception of a neighbor advertisement (NA) – as happened previously during the MNs DAD procedure – the message is discarded as no entry for the target address exists (Section 7.2.5 of [18]). When the GW/AR receives the BA with a destination address that is not in the neighbor cache it will send a Neighbor Solicitation and wait for the MN's NA. This adds an additional RTT that increases the accumulated number of message exchanges over the wireless link to 2 RTTs (BU+NS+NA+BA). Yet another problem is the retransmission timer for sending the BU: If a BU is not acknowledged with a BA by the HA after 1 second (INITIAL_BINDACK_TIMEOUT in [2]), another BU is sent. The sequence number in the binding update list (BUL) at the MN is incremented and the validity check for the BA (refering to the first BU) therefore fails as the sequence number does not match the one in the BUL. As the satellite GW or the AR do possess a neighbor cache entry now the second BA is forwarded to the MN within a reasonable time span and finally accepted.

$t_{HoT(I)}$ and $t_{CoT(I)}$ are nearly the same in Scenario 2 (±7ms) and similiar in Scenarios 3+4 (roughly ±150ms), when compared to the overall RTT. This can be explained by the very similiar routing path - in all three cases the HA is located very closely to the AR of the aircraft and therefore the necessary additional hops for the HoTI/HoT messages are negligible.

The retransmission problem that already affected the BU is valid for HoTI/CoTI transmission as well, as the 1 second timeout is missed in three out of the four scenarios. In scenario 1 HoTI is retransmitted due to the delay imposed by the long routing path from the US to Europe, in 4 it is due to the high delay of the LDACS and in 2 both HoTI and CoTI are retransmitted because of the delay of the satellite link and the ground network, as the test-init messages get routed from Europe to America and back. Yet this issue does not pose any problem as the BU can be sent to the CN nevertheless, as valid care-of and home tokens are available with the arrival

TABLE VII

REQUIREMENTS ON LATENCY AND MESSAGE EXCHANGES FOR CERTAIN SERVICES (FOR EN-ROUTE DOMAIN). LATENCY IN SECONDS AND UPLINK/DOWNLINK MESSAGE SIZES IN BYTES.

Service	Latency	Uplink	Downlink	RTT(s)
ACL	3.0	2×93	2×93	2
ACM	3.0	1×126	1×8	1
COTRAC (Interactive)	5.0	3×1969	4×1380	3.5
FLTPLAN	30	9×968	9×92	9
WXGRAPH	30	4×21077	6×93	5.5

of the first HoT/CoT batch, but adds additional unnecessary signalling overhead.

As expected, in Scenario 2 the same holds for t_{B-CN}: a BA is not received before the retransmission timeout, a new BU therefore sent, the BA to the first BU discarded, making it necessary to wait for the second BA.

The overall handover process, from establishing the layer 2 association up to having completed route optimization with the CN, takes between 4 and 7s for the different scenarios.

VII. COMPARISON TO SERVICE REQUIREMENTS

We have taken the most stringent services from [6] and listed them in Table VII, together with the maximum acceptable delay and number of message exchanges (in RTT).

Using $t_{HoT(I)}$ (denotes the RTT MN–HA–CN) from the simulation results seem to indicate that without RO it would still be possible to meet the requirements in all scenarios. Even COTRAC in Scenario 1 (3.5×1217ms) could be met, where the aircraft is communicating with its CN in the US via the Europe-based HA. However this will probably not be the case in reality due to link loss or using a protocol like TCP, which would add additional delays, especially due to the handshake phase that introduces additional 1.5 RTTs during connection establishment. At least for the COTRAC service, but also for ACL, this would violate the latency requirement ($(3.5+1.5) \times 1217 > 5000$ms). The same holds for Scenario 4, where MN, CN and HA are on the same continent but the delay of the LDACS is high ($((3.5+1.5) \times 1071 > 5000$ms)).

VIII. CONCLUSION

With the MIPv6 RO procedure consuming up to one third of the available capacity, Enhanced Route Optimization [19] is an attractive option to reduce the signalling overhead.

For continental flights as in Scenarios 3 and 4, performing RO does not significantly decrease the overall communication delay, given the HA is on the same continent, unless the wireless link is highly loaded. For inter-continental flights RO becomes important and improves performance significantly as can be seen from Scenario 1 (MN and CN in US, HA in Europe) where delay is reduced by a factor of 1.44 when comparing the HoTI/CoTI RTTs. For the satellite as in Scenario 2, the gain from RO is eliminated due to the satellite GW being close to the HA but far distant from the CN.

This observation however is based on a UDP like message exchange without any link losses – using a protocol like TCP in an inter-continental flight scenario would, for some services, violate the communication requirements. RO should therefore be considered as a must in this case and as desirable for continental flights if the wireless link is highly loaded.

The retransmission timer INITIAL_BINDACK_TIMEOUT of 1s proved to be problematic, and an increase to 1.75 or even 2s, in case the value ought to be an integer, would eliminate the associated retransmission problem.

Neighbor Discovery over a wireless link with a high delay proved to be problematic - the NS/NA message exchange to resolve the IP to the MAC address consumes an additional RTT at the AR. It would be reasonable to generate neighbor cache entries for nodes from the first datagram sent by the MN or when providing an IP address through DHCPv6 [20].

A HO while performing or commencing communications would result in such a huge delay that meeting the service requirements would be impossible. Multihoming and link signalling as in [13] to allow proactive HOs are therefore important additions for an aeronautical IP mobility solution.

REFERENCES

[1] ICAO ACP WG I-01/WP-06, "Analysis of candidate ATN IPS Mobility Solutions," June 2007.

[2] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," RFC 3775 (Proposed Standard), June 2004.

[3] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol," RFC 3963 (Proposed Standard), Jan. 2005.

[4] R. Wakikawa (Ed.), "Multiple Care-of Addresses Registration," Internet Draft (work in progress), Jan. 2008.

[5] W. Ivancic and D. Tran, "Mobile-IP Aeronautical Network Simulation Study," NASA TM-2001/210751, June 2001.

[6] Eurocontrol/FAA Future Communication Study, "Communications Operating Concept and Requirements for the Future Radio System," May 2007, COCR version 2.0.

[7] ICAO ACP WG T, "First meeting of the working group," 2007, Montreal, Canada.

[8] Christian Vogt, "A Comprehensive and Efficient Handoff Procedure for IPv6 Mobility Support," in *Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, Niagara Falls, NY, USA*, June 2006.

[9] S. Thomson, T. Narten, and T. Jinmei, "IPv6 Stateless Address Autoconfiguration," RFC 4862 (Draft Standard), Sept. 2007.

[10] Eurocontrol/FAA Future Communication Study Evaluation Scenarios, "Future Communications Infrastructure – Technology Investigations," May 2007, Appendix of ACP/1-WP/21.

[11] András Varga et al., "The OMNeT++ discrete event simulation system," http://www.omnetpp.org, 2005.

[12] J. Manner and M. Kojo, "Mobility Related Terminology," RFC 3753 (Informational), June 2004.

[13] IEEE P802.21/D8.0, "Draft Standard for Local and Metropolitan Area Networks: Media Independent Handover Services," December 2007.

[14] DLR, "Expected B-AMC System Performance," Report D5, September 2007.

[15] Réseaux IP Européens Network Coordination Centre, "RIPE Database," http://www.ripe.net/db/, January 2008.

[16] American Registry for Internet Numbers, "ARIN WHOIS Database Search," http://ws.arin.net/whois, January 2008.

[17] Yuval Shavitt and Eran Shir, "DIMES – Letting the Internet Measure Itself," http://www.arxiv.org/abs/cs.NI/0506099, June 2005.

[18] T. Narten, E. Nordmark, W. Simpson, and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)," RFC 4861 (Draft Standard), Sept. 2007.

[19] J. Arkko, C. Vogt, and W. Haddad, "Enhanced Route Optimization for Mobile IPv6," RFC 4866 (Proposed Standard), May 2007.

[20] R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," RFC 3315 (Proposed Standard), July 2003, Updated by RFC 4361.