IPv6 Operation for WAVE - Wireless Access in Vehicular Environments

Thomas Heide Clausen, Emmanuel Baccelli, Ryuji Wakikawa

INRIA

IPv6 Operation for WAVE - Wireless Access in Vehicular Environments

Thomas Clausen, Emmanuel Baccelli, Ryuji Wakikawa

N° 7383

September 2010

R apport de recherche

INRIA centre de recherche
SACLAY – ÎLE-DE-FRANCE

IPv6 Operation for WAVE - Wireless Access in Vehicular Environments

Thomas Clausen*, Emmanuel Baccelli†, Ryuji Wakikawa‡

Abstract:
   The IEEE WAVE protocol suite is providing communications services to applications in vehicular networks, by way of promising support for two protocol stacks: the Wave Short Message Protocol (WSMP) and IPv6. While WSMP is developed within the IEEE 1609 family of standards, the authors of this memorandum assert, that considerations for IPv6 operation for WAVE are less developed, and several issues are left unaddressed by the current IEEE 1609 specifications. This memorandum reviews these issues and analyzes the main challenges in providing proper IPv6 operation for WAVE networks.

Key-words: IEEE 1609, WAVE, DSRC, Vehicular Networks, IPv6

* LIX - Ecole Polytechnique, Thomas@ThomasClausen.org
† LIX - Ecole Polytechnique, Emmanuel.Baccelli@inria.fr
‡ Toyota-ITC, ryuji@us.toyota-itc.com

IPv6 Operation for WAVE - Wireless Access in Vehicular Environments

Résumé :

La suite de protocoles IEEE WAVE fournit des services de communications pour les réseaux véhiculaires, en promettant le support de deux piles de protocoles: le Wave Short Message Protocol (WSMP) d'une part, et IPv6 d'autre part. Si WSMP est spécifié dans la famille de standards IEEE 1609, les auteurs du présent papier affirment que les considérations concernant IPv6 dans WAVE ne sont pas assez dévelloppées, et que par conséquent, plusieurs problèmes ne sont pas résolus. Ce papier passe en revue ces problèmes, et analyse les principaux vérrous concernant le dévelloppement de solutions appropriées à ces problèmes.

Mots-clés : IEEE 1609, WAVE, DSRC, Rseaux Vhiculaire, IPv6

1 Introduction

The IEEE is currently undertaking standardization of a protocol suite for *Wireless Access in Vehicular Environments* (WAVE), with the objective of providing vehicles and pedestrians with the ability to communicate with each other and with road-side infrastructure. Possible applications hereof include emergency warning systems, cooperative cruise control and collision warning, as well as toll and parking fee collection. This protocol suite is developed in the IEEE 1609 working group, documented in [1, 2, 3, 4, 5, 6] and intended for operation over *Dedicated Short-Range Communications* (DSRC) – a set of wireless communications channels, dedicated for vehicular networking at 5.9GHz.

WAVE is providing communications services to applications, by way of promising support for two protocol stacks, the *Wave Short Message Protocol* (WSMP) and IPv6, as shown in figure 1. While WSMP is developed within the IEEE 1609 family of standards [4], considerations for operation of IPv6 for WAVE are less developed. The WAVE architecture specification [1] makes reference to the IETF[1] specification of IPv6 [7] and makes minimal observations regarding the use of IPv6 addresses, but no further specific recommendations as to IPv6 operation for WAVE are provided.

The authors of this memorandum assert that, while the IEEE 1609 family of specifications provides a set of necessary considerations for IPv6 operation over WAVE, these considerations are not sufficient for proper and correct IPv6 operation in this environment. This memorandum thus provides an analysis of IPv6 operation, as described in the IEEE 1609 family of specifications for WAVE networks, identifies where IPv6 operation for WAVE networks is underspecified, and presents a set of additional recommendations enabling proper IPv6 operation for WAVE networks.

1.1 Memorandum Overview

While IPv6, as defined in [7], principally concerns the data frame layout (header format, header extensibility, rules governing header construction and processing etc.), IPv6 operation implies operation of a set of basic protocols at the network layer, including NDP [8], stateless address autoconfiguration [9]. The IPv6 protocol stack provides additional protocols at other layers, such as the transport layer and the application layer. Most of these protocols make certain assumptions about properties of an underlying link model for their proper operation, and assume certain relationships between assigned IP addresses and communications ability across the underlying data link layer. This is discussed in details in section 2, elaborating on the link-model presented by a WAVE system, and presenting IPv6 network layer considerations for WAVE, resulting from the properties of that link-model. This memorandum , then, in section 3 presents additional issues for WAVE operation in an IP networking context, including "pseudonymity", transport and application layer challenges. This memorandum is concluded in section 4.

---

[1]http://www.ietf.org

Layer	Name	Data Plane w. IPv6 stack	Data Plane w. WAVE stack
5	Application	HTTP & other Applications	WAVE Applications
4	Transport	IPv6 Protocol Stack (TCP, UDP, ND...)	WAVE Protocol Stack (WSMP...)
3	Network		
2	Link	802.2 LLC	
		WAVE MAC	
1	Physical	WAVE physical Layer	
		WAVE Physical Medium	

Figure 1: Dual stack, IPv6 and WAVE.

2 WAVE Network Layer Challenges with the IPv6 Protocol Stack

IPv6 operation is, beyond the use of the IPv6 frame format [7] on the network layer, generally understood to also imply assumptions of a specific and well-defined link-model reflected in a well-defined addressing model [10], and operation of a set of supporting protocols [8], [9].

The IPv6 addressing model defines different address families (*e.g.*, Link Local or Global addresses), with associated properties. This enables applications or protocol to have certain expectations of communication abilities, corresponding to the addresses they use. For example, an application using a Global address as destination address expects the network to be able to ensure multi-hop communication to that destination address. The network, then, expects such addresses to be assigned in a way such that by inspection of the address, it can be determined if the destination is reachable directly, or reachable only via a (and, in that case, also via which) router.

In an IPv6 network, the supporting *Neighbor Discovery Protocol* for stateless autoconfiguration (of addresses, default routers etc) and duplicate address detection [8], [9], is assumed to be running – and that protocol expects specific *link model* and *addressing model*.

Thus, IPv6 operation entails (i) using the IPv6 frame format, (ii) certain assumptions of a well-defined *link-model*, reflected in an (iii) *address model*, and (iv) proper operation of NDP. This is detailed further in the following sections.

2.1 IPv6 Link Model

[11] points out that network protocols and applications are designed with specific assumptions on the nature of an IP link, illustrated in figure 2 and summarized as follows:

- all hosts (H) with network interfaces configured with addresses from within the same prefix `p::`, and with the same prefix `p::` assigned to the interfaces, can communicate directly with one another; *i.e.*:

  - IPv6 datagrams are not forwarded at the network layer when communicating between interfaces which are configured with addresses from within the same prefix; hence

  - hop-limit in IPv6 datagrams are not decremented when communicating between interfaces which are configured with addresses from within the same prefix, and;

  - multicast/broadcast IPv6 datagrams with a hop-limit of 1 are (modulo data loss) delivered to all interfaces within the same subnet.

- link-local multicasts and broadcasts are received by all interfaces configured with addresses from within the same prefix without forwarding.

The IPv6 Link Model, in figure 2, axiomatically assumes that neighbor relationships are symmetric: if communication from air interface A to air interface B is possible in one hop, then communication in the reverse direction is also possible – in other words, connectivity between neighbor interfaces is assumed symmetric.

An even shorter summary of the IPv6 link model is to say that *an IPv6 link looks like an Ethernet.*

2.2 IPv6 Addressing Model, Address Scopes and Uniqueness

As described in section 2.1, the notion of an "IPv6 link" is tied with that of an IPv6 subnet prefix: all interfaces which are configured with the same subnet prefix are considered to be on the same IP link and, thus, for communication between nodes on the same subnet, no forwarding is required and no decrement of TTL/hop-limit is performed. In addition to this relationship between link and prefix, IPv6 introduces address scopes – *Link-Local* and *Global* – and mechanisms by which addresses are constructed using *Interface IDs*.

A Link-Local address is valid for communication with a device on the same link: an IPv6 datagram with a Link-Local source or destination address is not to be forwarded on the network layer, but is to be received by a destination on the same link – or not received at all. The only requirement for an unicast Link-Local address to be useful is, thus, that it is unique on the local link; the
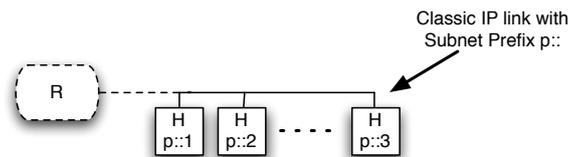

Figure 2: IP Link Model: hosts (H) connected to the same link have assigned IP addresses from a common prefix, possibly assigned by a router (R).

same Link-Local address may well be in use on another, disjoint, link, however as IPv6 datagrams with Link-Local addresses are never to be forwarded, no ambiguities exist.

A Global address is valid for communication beyond the local link: an IPv6 datagram with a Global source and destination address can be forwarded on the network layer and, thus, be received by a destination on the same or on a different link – or not received at all. For an unicast Global address to be useful, it must, thus, be unique across the entire network.

It is important that these address uniqueness requirements are universally satisfied. This is ensured in IPv6 by having an interface detect when it connects to a link (typically, by way of a discrete link-layer trigger), upon which it constructs a Link-Local IPv6 address by concatenating the Link-Local Prefix (`FE80::/10`) with an *Interface ID*, typically derived from the MAC address of that interface. Duplicate Address Detection (DAD) [9] is then performed, to verify that this address is not already in use on the link. DAD employs Link-Local Multicast, interrogating all other interfaces on the link as to if they are already using that address, by way of Neighbor Solicitation (NS) messages. Absent a reply to this interrogation – by way of Neighbor Advertisement (NA) messages – the address is assumed unique on the link and henceforth used. As all Link-Local addresses share the same Prefix (`FE80::/10`), this DAD procedure in reality verifies that the chosen *Interface ID* is unique across the link.

Global addresses are constructed by concatenating the *Global prefix* of a link with the *Interface ID* of an interface, verified to be unique during the configuration procedure for Link-Local addresses. The *Global Prefix* is obtained from a router on the link[2], by way of Router Solicitation / Router Advertisement messages [8]. It follows that uniqueness of a Global address for an interface relies on (i) the *Interface ID* being unique on the link to which the interface is connected, and (ii) unique prefixes being delegated to routers. It follows, then, that a Global address is valid only as long as that interface is connected to the link on which the router providing the *Global prefix* is present.

An interface must also detect when it disconnects from a link (typically also by way of a discrete link-layer trigger), upon which it must cease to use the previously configured addresses. Thus, in IPv6, a link describes a well-determined set of network interfaces, all able to communicate directly with each other without forwarding, and with all interfaces in a single (link-local multicast) transmission be able to reach all other interfaces on the same link. This set of network interfaces is maintained by way of explicit and discrete signals, allowing an interface to detect its connection to a given link.

A summary description of the *IP addressing model* is therefore, that (i) addresses are of a specific *validity scopes*, global or local, where (ii) within validity scope of an address, it must be used by no more than one interface, and (iii) an address of global validity scope assigned to an interface must be *topologically correct, i.e.,* it must match the Global prefix provided by the router on the IP link to which the interface connects.

---

[2]Global addresses are only relevant in case the network can provide multi-hop communication, *i.e.* a router *is* present on the link.

2.3   IPv6 Network Layer Considerations Regarding WAVE

As indicated in section 1, the IEEE 1609 family of specifications present a minimal set of considerations for IPv6 operation for vehicular networks. Devices in vehicular networks are separated into On-Board Units (OBUs) and Road-Side Units (RSUs), with the latter providing, as needed, infrastructure and configuration support for the former. With respect to IPv6 operation in such networks, the IEEE 1609 family of specifications simply state that:

- IPv6 is provided as a data plane protocol, and that the "standard IPv6 protocol" is used;

- IP configuration parameters (global prefixes, ...) are provided in the WAVE Routing Advertisement (WRA) messages;

- OBUs advertising services to other OBUs do so using Link-Local addresses: OBUs provide services to direct (1-hop) neighbors only, and therefore acquiring and maintaining *topologically correct* Global addresses is wasteful;

- RSUs are identified by either Link-Local or Global addresses;

- Link-Local addresses are derived by the device, are not globally unique and are not usable for routing;

- NDP [8], otherwise used for populating the neighbor cache, is asserted to generate a substantial and unacceptable amount of traffic, and thus other means for populating the neighbor cache (ICMPv6, IPv6 PDUs) are employed;

- NDP is, however, not excluded for "cases where it might be needed".

2.4   WAVE air Interface "Link Model"

The air interfaces of a WAVE system, and the "links" to which they attach, have different characteristics from those described in section 2.1, and are therefore detailed in this section. The resulting *"WAVE link model"* does not provide for a direct mapping to the IPv6 link model, thus considerations for operating IPv6 over this "WAVE link model" are detailed in section 2.5.

Symmetric to the IPv6 Link Model in figure 2, figure 3 illustrates the relationship between WAVE air interfaces, and serve for elaborating the "WAVE link model" in the discussions in this section.

Each air interface is a (radio) broadcast interface, able to establish a direct link layer communication with air interfaces which are within its coverage area. In figure 3, this coverage area is approximated by a simple disc of fixed radius (light gray discs) – in the real world, both the shape and size of the coverage area is variable as a function of the interface, interference from the environment etc. Referring to figure 3 if, *e.g.,* if N3 transmits, then this transmission may be received by N2 and N4, but not by N1 and N5. This implies that, *e.g.*, N3 and N4 – despite being neighbors and on the same "link" – do not share the same view of which other nodes are neighbors and on the same "link": N3 considers that it is on the same "link" as N2 and N4, whereas N4 considers itself to be on the same "link" as N3 and N5.
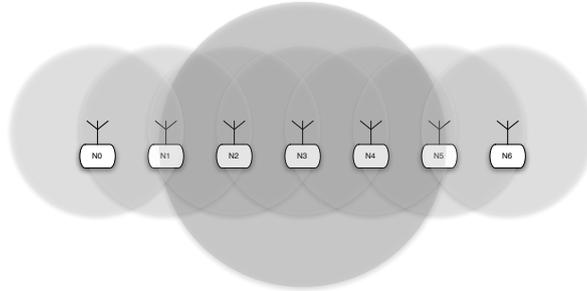
Figure 3: Nodes (N) with air interfaces. The light grey area indicates the **coverage area** of each air interface. The dark grey circle indicates the **interference area** of the air interface of N3.

Thus, a set of air interfaces within a region – even if using the same channels and modulation – may not all be able to communicate to all other air interfaces, without intermediate relaying. A link-local multicast transmission from one air interface may not (even disregarding losses) be able to be received by all other air interfaces; indeed, a multicast transmission from one air interface may not be able to reach the same set of air interfaces as would a multicast transmission from its closest neighbor air interface. This is the case in figure 3, where no two air interfaces can directly transmit to the same set of other air interfaces.

An air interface has an *"interference area"* which may be greater than its coverage area, *i.e.* a transmission by N3 in figure 3 will, as indicated above, be correctly received by the interfaces N2 and N4. At the same time, however, this transmission may be propagating to interfaces of N1 and N5 where, while the transmission can not be correctly decoded, it can be detected, and cause interference with other transmissions which could otherwise be correctly received over the air interfaces of N1 and N5 (such as transmissions from N0 and N6).
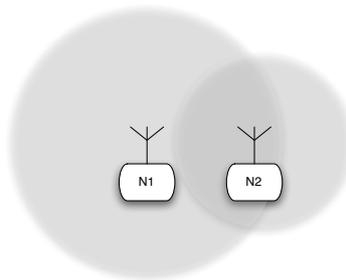

Figure 4: Neighbor asymmetry.

Figure 4 illustrates a situation where, for some reason (powerful transmitter, environmental interference, large antenna, ...), the air interface of N1 has a large enough coverage area for its transmissions to be received and correctly decoded by the air interface N2. The air interface of N2, on the other hand, has a much smaller coverage radius, such that transmissions from the air interface of N2 can not be received and correctly decoded at the air interface of N1. Thus

asymmetric – or more precisely, unidirectional – connectivity between the air interface of N1 and the air interface of N2 exists: N2 sees N1 as a neighbor (since the air interface N2 can receive transmissions from the air interface of N1), whereas N1 does not see N2 as a neighbor (since the air interface of N1 can not receive transmissions from the air interface of N2).

A vehicular network, naturally, represents a dynamic topology: OBUs move relatively to each other and to RSUs. The resulting network is a highly dynamic graph, where the neighborhood of an air interface is also dynamic and varies over time – due to mobility, and due to changing environmental factors: two air interfaces which were not in communications range a moment ago may become neighbors, and vice-versa.

Thus, neighboring air interfaces may experience distinctly different neighborhoods, may not even agree on if they are or are not neighbors, and may at any time become, or cease to be, neighbors.

Finally, as the set of air interfaces "on a link" are communicating via radio waves rather than electrical wires, there are no implicit physical signals, allowing an air interface to detect its association or disassociation with a given set of other air interfaces "on the same link". And the set of air interfaces "on a link" may be subject to constant and rapid change. In a certain way, it is tempting to add "this is just as well", as the other "on a link" properties expected in the IPv6 Link Model do not hold, as described above.

2.5 Considerations for IPv6-over-WAVE

Considering the differences between the IPv6 link model, described in section 2.1, and the WAVE link model, described in section 2.4, verbatim use of the standard IPv6 protocol stack, as the IEEE 1609 family of specifications stipulate, is not sufficient.

2.5.1 Air Interface Addresses

addresses for OBUs are specified to be link-local in [1]; it is further stated that configuring air interfaces of OBUs with Global Addresses is undesired due to the need to maintain topological correctness of such Global Addresses. This is necessary, but not sufficient, for these interfaces to be configured with "valid" addresses. Link Local addresses assigned to interfaces should in addition be globally unique, *i.e.,* must be derived from some globally unique token. The reason for this is, as vehicles – and so their OBUs – may move, any two OBU air interfaces may at some point in the future become direct neighbors. On the other hand, since there is no guarantee that an arbitrary pair of air interfaces of OBUs will always remain neighbors, no IPv6 subnet prefix can be configured on an interface: vehicle movement may render such two air interfaces unable to communicate, requiring reconfiguration of the on-link IPv6 subnet prefix to respect the IPv6 Link Model assumption that "two interfaces with the same subnet prefix can communicate directly", as described in section 2.1. The recently published RFC5889 [12] describes an IP addressing model for ad-hoc networks; the considerations described herein apply equally to air interfaces in WAVE networks, notably:

- An IP address configured on an air interface should be unique, at least within the routing domain (in this case, the vehicular network at large), and

- No on-link subnet prefix is configured on this air interface.

.

2.5.2 Supporting Protocols Employing Link-Local-Multicast

protocols such as DHCP, NDP and Stateless Address Autoconfiguration, assume the multicast characteristics of the IPv6 Link Model; as stated, these do not hold for the WAVE Link Model. NDP basic mechanisms such as Neighbor Solicitation (NS) do not operate as expected: the set of air interfaces which will receive such a NS is the set of air interfaces which, at the time of emission of the NS, happen to be within radio range. Thus, *e.g.,* Duplicate Address Detection (DAD) will not ensure the desired uniqueness properties of IPv6 Link Local Addresses.

2.5.3 Discrete Link Layer Association Triggers

such triggers, otherwise used for initiating IPv6 interface configuration, are absent on the WAVE Link Model. Thus, protocols, including [8], [9], and the address configuration assumption that an interface can detect when it "disconnects" and thus should cease using previously used addresses, can not rely on such. Information can thus not be "solicited when events happen" but must be beaconed, and protocols adopted accordingly.

2.5.4 Communications Bidirectionality

link bidirectionality cannot be assumed. Experience with air interfaces show that if a device A hears service advertisements from another device B (OBU or RSU), this does not guarantee that device B can hear any service request sent by device A – *i.e.,* it is not uncommon that links are unidirectional. Thus any alternative mechanism to be developed should at least verify link bidirectionality before relying on it.

3 Other Challenges with WAVE using the IPv6 Protocol Stack

The IPv6 protocol stack includes various additional protocols, above the network layer protocols described in section 2. Regardless of how the described network layer issues are resolved, attention must be paid to operation of the transport and the application layers. This section briefly overviews some of these additional considerations.

3.1 Transport

At the transport layer, the IPv6 protocol stack proposes two types of protocols: TCP, a reliable, rate-adapting mechanism enabling end-to-end transport of application data across several IP hops and requiring bi-directional communication

between the peers for acknowledgements etc. The second IPv6 transport protocol is UDP, a much simpler protocol providing no rate-adapting or reliability mechanism and so no signaling from the destination to the sender in a traffic flow.

It is worth noting that TCP is often very inefficient in wireless ad hoc environments [13], especially when faced with mobility: TCP was designed to interpret packet loss as traffic congestion and to diminish sending rates in this case, whereas in wireless networks, packet loss may have causes that are other than traffic congestion, such as interfaces moving out of reach, collisions or interference. Also, if a TCP connection is established between two air-interfaces, subsequently moving out of range before the connection is terminated, connection-state remains for timing out (and possibly causing extraneous transmissions), not cleared up by the usual end-of-connection signaling. Therefore, TCP is usually not employed in VANETs, which leaves UDP as the only viable alternative within the standard IPv6 stack. [1] recommends the use of UDP, however the reasons given (in section 6.4.3 of [1]) relate to the matching of *"the connection-less nature of WAVE transmissions"* only. The authors submit that there are also technical reasons for why TCP might be a lesser appropriate choice for this environment.

Applications requiring rate-adapting or end to end transport reliability services may not be satisfied with what the standard IPv6 protocol stack has to offer.

3.2 Pseudonimity

The IEEE 1609 family of specifications also promise to support MAC address changes to provide pseudonymity, *i.e.,* to ensure that a device's non-temporary identity, and its long-term patterns of behavior, cannot be deduced from its network traffic and are only available to authorized parties.

However, IEEE 1609 specifications do not provide a way to generate or assign pools of globally unique network addresses, aside of basic "random" local generation which is likely to provide duplicate network addresses if devices change their network addresses too frequently. Note that the specifications do not define how frequently MAC address changes should take place in order to provide pseudonymity. Nevertheless, it is doubtful that pseudonymity can be achieved without a significant probability of duplicate network addresses using the current IEEE 1609 specifications.

In the same vein, in order to provide pseudonymity at layers above the network layer (which may be necessary to provide user pseudonymity in the end), similar issues are bound to arise if devices must change their IP addresses frequently. More generally, providing pools of globally unique IDs, dynamically and in a distributed manner, becomes a hard problem if the number of possible IDs is not extremely large.

Therefore, applications' requirements concerning pseudonymity may not be entirely satisfied with what the standard WAVE and IPv6 specifications have to offer so far, and thus, it is presently left to these applications to provide such services (if at all possible). Moreover, one should note that, as quickly mentioned in the IEEE 1609 specifications, frequent MAC or IP address changes will disrupt most applications.

4   Conclusion

This memorandum has provided an overview of the issues concerning IPv6 use over WAVE, the protocol suite for wireless access in vehicular environments currently developed by the IEEE. While WAVE promises communications services to applications via IPv6 stack support – defined in the IEEE 1609 family of standards – this memorandum has shown that this support underspecified, and described the issues that are so far left unaddressed concerning IPv6 operation for WAVE. This memorandum has also analyzed the challenges in designing solutions to overcome these issues, and provided guidelines regarding the design of appropriate solutions.

References

[1] The Wireless Access in Vehicular Environments (WAVE) Working Group of the Intelligent Transport Systems (ITS) Committee, "IEEE P1609.0/D0.1: Draft Standard for Wireless Access in Vehicular Environments (WAVE) - Architecture," April 2010.

[2] —-, "IEEE P1609.1/D1.3: Draft Standard for Wireless Access in Vehicular Environments (WAVE) - Remote Management Services," May 2010.

[3] —-, "IEEE P1609.2/D5: Draft Standard for Wireless Access in Vehicular Environments (WAVE) - Security Services for Applications and Management Messages," June 2010.

[4] —-, "IEEE P1609.3/D7.0: Draft Standard for Wireless Access in Vehicular Environments (WAVE) - Networking Services," June 2010.

[5] —-, "IEEE P1609.4: Draft Standard for Wireless Access in Vehicular Environments (WAVE) - Multi-Channel Operation," June 2010.

[6] The 802.11 Working Group, "IEEE 802.11: Wireless Lan Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Amendment 6: Wireless Access in Vehicular Environments," June 2010.

[7] R. Hinden and S. Deering, "RFC2460: Internet Protocol, Version 6 (IPv6) Specification", IETF RFC, http://www.ietf.org/rfc/rfc2460.txt, 1998.

[8] T. Narten, E. Nordmark, and W. Simpson, "RFC2461: Neigh- bor Discovery for IP Version 6 (IPv6)", IETF RFC, http://www.ietf.org/rfc/rfc2461.txt, 1998.

[9] S. Thomson, T. Narten, and T. Jinmei, "RFC2462: IPv6 state- less address autoconfiguration", IETF RFC, http://www.ietf.org/rfc/rfc2462.txt, 1998.

[10] R. Hinden and S. Deering, "RFC3513: Internet Protocol Version 6 (IPv6) Addressing Architecture", IETF RFC, http://www.ietf.org/rfc/rfc3513.txt, 2003.

[11] D. Thaler, "RFC4903: Multilink Subnet Issues", IETF RFC, http://www.ietf.org/rfc/rfc4903.txt, 2007.

[12] E. Baccelli, M. Townsley, "RFC5889: IP Addressing Model in Ad Hoc Networks", IETF RFC, http://www.ietf.org/rfc/rfc5889.txt, 2010.

[13] G Holland, N Vaidya, "Analysis of TCP performance over mobile ad hoc networks," Wireless Networks, 2002.

Contents