

# Support of Anonymity in VANETs – Putting Pseudonymity into Practice

Emanuel Fonseca, Andreas Festag  
NEC Deutschland GmbH  
Heidelberg, Germany  
Email: {fonseca|festag}@netlab.nec.de

Roberto Baldessari  
NEC Network Laboratories  
Heidelberg, Germany  
Email: baldessari@netlab.nec.de

Rui Aguiar  
Universidade de Aveiro  
Instituto Telecomunicaes  
Email: ruilaa@det.ua.pt

**Abstract**—Despite the great advantages offered by vehicular ad hoc networks (VANETs), they also introduce challenges with respect to security and privacy. Using unique identifiers for communication, a vehicle can easily be located and tracked. Today, people are more and more concerned about their privacy. As a technical solution to protect drivers' privacy, the use of changing pseudonyms has been proposed. Existing work mainly focus on algorithms for pseudonym change and neglect the practical implications. For deployment and integration of pseudonymity into a VANET communication system, open issues need to be solved. This paper analyzes the practical challenges and proposes solutions for protocol- and implementation-related issues necessary to turn pseudonymity support into practice. Finally, the paper concludes by means of analysis and measurements that the affects of pseudonymity can be alleviated at reasonable costs and compromises in anonymity support.

## I. INTRODUCTION

Vehicular ad hoc networks (VANETs) enable vehicle-to-vehicle and vehicle-to-roadside communication. Vehicles equipped with on-board units (OBUs) and fixed communication units along the road (road-side units, RSUs) can share information. Applying short range wireless technology based on IEEE 802.11 [1], [2], multi-hop communication utilizing geographic positions facilitate information exchange among network nodes that are not in direct communication range [3]. The technology offers a wide range of safety applications, for example, vehicles can warn others from hazardous events to avoid accidents. While safety is the main focus, VANETs also enable applications for improving traffic efficiency, and infotainment applications, such as location-based services along the drive. Despite all these benefits, a VANET system design should respect the drivers' desire for privacy [4].

In principle, the communication protocols of a VANET node publicly disclose data, such as node address, position, speed, heading, and time. An attacker can potentially link these data to the user's identity and invade the privacy of the user. Cryptography and encryption are not appropriate means to protect these data since the exchange of these information is mandatory for network operation. In order to preserve the privacy of users, the use of randomly chosen and changing identifiers – referred to pseudonyms – has been considered [5]–[10]. Pseudonymity considerably aggravates the association of pseudonym and user identifier, but cannot prevent that an attacker collects personal data.

Few existing papers have investigated pseudonymity, mainly

from the perspective of how often a node should change a pseudonym and with whom it should communicate. [6] proposes to use a silent period in order to hamper linkability between pseudonyms, or alternatively to create groups of vehicles and restrict that vehicles in one group can hear messages of other groups. In [9] a user can cloak information before sending it, by providing location information at a coarse granularity in terms of time and space. [10] studies *mix zones* to protect location privacy of location-based services from service providers.

In comparison to existing work this paper investigates the integration of pseudonymity into a real VANET communication system. Challenges include an addressing concept across layers of the protocol stack, issues in geographical routing (location service, forwarding), and cross layer information exchange, as well as problems related to implementation design and performance. As contributions we present 1) a multi-layer addressing scheme with support for pseudonymity enabling user privacy at the different layers, 2) enhanced packet forwarding schemes based on pseudonym caching, and 3) a location service capable to work with periodic changing pseudonyms, enabling unicast communications. Finally, we discuss implementation and performance issues related to pseudonymity on the basis of our prototype implementation of a VANET communication system.

The paper is organized as follows: Sec. II describes the system architecture and lists possible threats against privacy. Sec. III analyzes challenges of applying changing pseudonyms in the architecture. Sec. IV presents the proposed addressing concept, pseudonymity-enhanced location service and forwarding schemes, and link-layer callbacks. Sec. V discusses implementation and performance aspects, Sec. VI presents our conclusions.

## II. SYSTEM ARCHITECTURE AND PRIVACY THREATS

The assumed architecture is split into three domains (Fig. 1): the in-vehicle domain is a sub-network with an *On Board Unit* (OBU) and potentially several *Application Units* (AUs). AUs are user devices (PDA, mobile phone) attached to the OBU that execute dedicated applications. The ad hoc domain is composed of OBUs and stationary units along the road, termed *Road-Side Units* (RSUs). OBUs and RSUs are equipped with wireless technology based on IEEE 802.11 technology. They

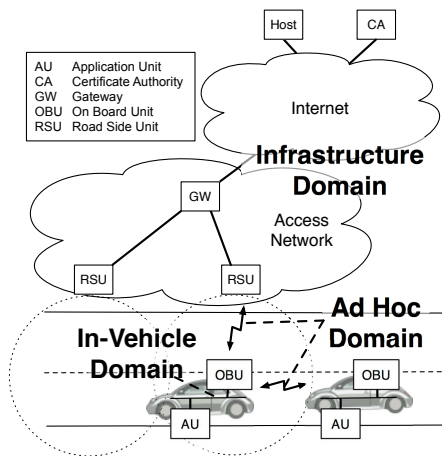


Fig. 1. High-level system architecture view

can directly communicate when they have direct wireless connectivity. Otherwise, multi-hop communication is used where data packets are forwarded from one node to another until they reach the destination. The infrastructure domain provides connectivity to Internet nodes and servers and access to Internet services. The infrastructure access is provided via RSU, if available.

For efficient and scalable communication in the ad hoc domain, position-based routing (PBR) is applied [11]. It provides wireless multi-hop communication between two nodes (unicast using so called ‘greedy forwarding’) and efficient broadcasting of data packets in geographical areas (geocast). Integral parts of PBR are beaconing and location service: Using beacons, a node periodically advertises its identifier and position to the direct neighbor nodes. With the location service a destination node is searched for its current position. Every OBU maintains a soft-state *Location Table* (LT) with (at least) identifier, position, speed, heading, and timestamp of the node the OBU communicates with. The LT comprises data from both, neighbors and non-neighbors, The data are acquired from beacons, data packets, etc. that are processed by a node. The LT is used for routing of unicast packets in order to make decisions on the next hop of a packet.

There are several potential attacks to VANETs that can threaten users privacy. An adversary can potentially *hack* into the OBU of a vehicle being then able to access information by reading communication logs. Also, an adversary might eavesdrop the wireless communication, such as beacons (passive attack), or even misuse the location service to track nodes (active attack). The point of attack is the fact that every node carries a unique identifier that is utilized for both, to establish a session between nodes, and for forwarding along the routing path. When the identifier represents an *electronic license plate* [12], an adversary can even link the identifier to the personal data of the registered vehicle owner.

In order to provide secure communication and enable anonymity we assume that every node has a unique and static *Global ID* (GID) and a set of pseudonyms (PID) used as network addresses that can change during communication. Nodes would have to exchange their GID and the GID public

key prior to unicast communications between them.<sup>1</sup> Despite we can have unicast communications using PIDs we cannot know the other persons identity. Applying asymmetric cryptography, a pseudonym is associated with a public/private key pair. A certificate signed by a *trusted certificate authority* (CA) binds public key and GID. A node can change its pseudonym and associated signature and certificate in order to prevent identity and location tracking. Key distribution, key revocation, procedures and services for secure communication – authentication, authorization, integrity, non-repudiation – are important features of the system architecture. The VSC project [13] provides a framework for safety-oriented communication and can principally adopted for a security architecture of future VANETs.

### III. CHALLENGES

This section lists conceptual problems encountered when applying pseudonymity to a VANET communication system.

**Cross-layer addressing and information exchange.** Assuming a layered protocol architecture for VANET nodes [14], a node uses multiple addresses simultaneously on the different layers. Changing the pseudonym only at one layer implies the risk that an adversary can link two pseudonyms by the unchanged address at another protocol layer. Consequently, supporting pseudonymity requires to change all addresses across a node’s protocol stack at the same time. In order to achieve non-repudiation, the pseudonyms on the different layers should be selected from a preloaded set of addresses. An assignment scheme that allows to derive addresses for the different protocol layers from a single pseudonym would ease pseudonym management. Since information exchange among protocol layers (e.g. link state information) is typically based on network addresses and designed under the assumption that the addresses do not change during operation, mechanisms for cross-layer information fail when pseudonyms change.

**Resolution service for pseudonyms.** In order to establish a communication session among nodes, the communication peers need to identify each other. Imagine your friends were constantly changing their phone number without letting you know in advance. How would you know which number to call if you wanted to talk to one of them? The same happens in VANETs if vehicles keep changing the addresses. Therefore, a system must offer a service that resolves the true identifier to the current pseudonym, more specifically GID to PID. In order to preserve the users’ privacy, the resolution service should disclose the link between GID and PID only to trusted nodes and communication peers.

**Affect on routing.** Since VANET nodes typically move at high speeds, it happens frequently that nodes have outdated forwarding entries. From the perspective of a certain node, a pseudonym change of another node can not be distinguished from the case that a new neighbor leaves its transmission range. In both situations, the node would have an outdated entry for forwarding of unicast packets, which might result

<sup>1</sup>This process is similar to what happens with phone numbers. We have to give our phone number to someone else if we want them to call us.

in a wrong forwarding decision and packet loss. Arguing that the change of addresses aggravates the problem of stale entries in the location tables, pseudonymity can significantly degrade the system performance by increasing unicast packet loss. The forwarding of broadcast packets using simple broadcast strategies (referred to as naive flooding) is not impaired, since – in contrast to unicast forwarding – it is based on broadcast addresses on link and network layer. However, enhanced broadcast algorithms attempt to alleviate the ‘broadcast storm’ problem [15] based on knowledge of neighbors. In such schemes, a node with a changed pseudonym would be regarded as a new node and a packet would unnecessarily be forwarded leading to redundant packet transmission.

Though the above list summarizes the practical issues to be solved, it is not all encompassing. Two issues appear particularly important for pseudonym support, but are beyond the scope of this paper. First, pseudonym change causes an interruption of established communication sessions. Session continuity could be achieved by session re-establishment or a VANET-specific Mobile IPv6 solution (e.g. proposed in [16]). Second, pseudonymity affects security services, such as non-repudiation, and should be considered in a security protocol architecture.

#### IV. FRAMEWORK FOR PSEUDONYMITY SUPPORT

Finding solutions to the challenges listed in the previous section is essential for applying pseudonymity. This section presents a framework for pseudonymity support in a VANET communication system. The combination of the individual solutions in the framework allows to turn pseudonymity into reality.

##### A. Cross-Layer Addressing Concept

The proposed addressing concept presumes that a node simultaneously uses multiple addresses at different protocol layers:

**MAC address.** Unique identifier attached to the wireless interface with fixed length of 48 bits.

**PBR address.** Is used for routing among OBUs and RSUs in the ad hoc domain by the PBR protocol and has a fixed length of 64 bits.

**IP version 6 address.** Standard IP address as defined by RFC 3513 with a fixed length of 128 bits. Each node has multiple IPv6 addresses simultaneously. An IPv6 link local default address is used for compatibility and test. An IPv6 unique local unicast address (RFC 4193) is used for communication inside of the ad hoc domain, whereas this address must not be propagated in the Internet. A node builds a globally routable IPv6 address by means of auto-configuration when RSU access is available.

In order to create the addresses specific procedures are used (Fig. 2): Given a MAC address (EUI-48), an EUI-64 is built by encapsulating the EUI-48 address. The resulting interface identifier is then combined with a 64-bit IPv6 prefix. For communication in the ad hoc domain, a Unique Local IPv6 Unicast Address is adopted with an arbitrary prefix (FD::/16). If




Fig. 2. Address assignment scheme

access to a RSU that provides automatic address configuration is available, the IPv6 address is configured by means of the automatically assigned prefix.

Applying the addressing concept allows a software instance to control the use of pseudonyms by a simple trigger carrying a new pseudonym. The VANET communication system then automatically determines the needed addresses and configures the network interfaces.

##### B. Extended Location Service

As briefly described in Sec. II, PBR (in its basic version without support for pseudonymity) already makes use of a location service that resolves a given identifier for its current location. A simple instantiation of such a location service is based on scoped flooding (Reactive Location Service, RLS): A requester issues a *location query* message that carries a certain look-up identifier. Every node re-broadcasts the query within a scope of wireless hops defined by the initiator. On reception of a query, the searched node sends a *location reply* to the requester with its current position. Then, the requester can forward data packets towards the destination by means of ‘greedy forwarding’.

We propose to integrate the services for location resolution and pseudonym resolution. Suppose the exemplary scenario in Fig. 3 with an initiator *A*, forwarder *B* and responder *C*. Node *A* creates a *location query* message with its own GID and the look-up GID. The initiator’s GID and the look-up GID is encrypted with *C*’s public key using a probabilistic encryption to ensure some random value in the cipher. When *B* receives the *location query* it attempts to decrypt the lookup GID. This operation fails, since *B* does not have the corresponding private key. As a result, the decrypted look-up GID does not match *B*’s GID. *B* derives that the packet is not destined to itself and rebroadcasts the query. When *C* receives the query, *C* can successfully decrypt the look-up GID identifying itself as the destination node and *A* as the initiator. Then, the responder *C* decides to reveal its GID to the initiator *A* (or to ignore it as explained below). *C* creates a *location reply* message with its own GID encrypted with *A*’s public key, *C*’s pseudonym, current position and timestamp, and sends the reply to initiator *A*’s pseudonym. Again, a forwarder *B* cannot decrypt *C*’s pseudonym and forwards it to *A*. When *A* receives the *location reply*, it caches *C*’s PID, time stamp and position, and proceeds with the unicast communication.

The *extended location service* assumes to be working on

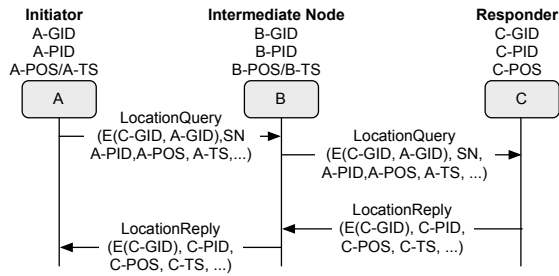


Fig. 3. Extended location service

a network with a secure routing protocol which provides message authentication through digital signatures using asymmetric cryptography.<sup>2</sup> This assumption will allow mutual authentication of the initiator and the responder. The need for authentication is required in order to provide a scheme for identity management where nodes control the exposure of identities and pseudonyms with respect to the peer identity, location, and time. Identity management can be seen as a mechanism to provide authorization to the *extended location service* since the receiver of a location query request would be able to choose whether to authorize the location reply or not. The use of encryption using asymmetric cryptography allows the identities involved on the *extended location service* to keep secrecy of their identity (GID) from any nodes that do not possess the specific private key.

### C. Pseudonymity-Enhanced Packet Forwarding Schemes

In order to leverage the problem of forwarding errors, we propose to enhance forwarding schemes where a node caches previously used pseudonyms of its own and apply these for forwarding (Fig. 4).

For pseudonymity support it is important to understand that with position-based routing (PBR) every data packet carries address, position, timestamp for the destination and additionally the same data fields for source and sender. When a node forwards a data packet it selects the next hop by adding the next hop's MAC address as the destination address of the MAC frame.

The simplest forwarding scheme is based on the assumption that a node uses only a single pseudonym simultaneously. More precisely, a node listens to packets that carry its current pseudonym  $PID_n$  in the MAC header and then adds its current pseudonym  $PID_n$  as source address in the header of the data packet forward (Fig. 4(a)). When the pseudonym changes, the node accepts and potentially forwards data packets only for the current pseudonym and drops all packets that carry the previous pseudonym(s). Since for reception and send only a single pseudonym is used simultaneously, we refer to this simple scheme as the  $'1:1'$  scheme for pseudonymity.

In the  $'m:1'$  scheme (Fig. 4(b)), a node caches the last  $m-1$  pseudonyms in addition to the current pseudonym. On reception of a data packet, the node compares the MAC address in the packet header with the cached pseudonyms. If one of the pseudonym matches, the packet handler forwards




Fig. 4. Schemes for pseudonymity-enhanced forwarding schemes

the data packet to the next hop with the current pseudonym  $n$  as source address.

With the  $'m:1'$  forwarding scheme (Fig. 4(c)), a node handles a received packet in the same way as with the  $'m:1'$  scheme. When a packet is forwarded, the packet header carries the same pseudonym as when it was received, i.e. when a packet with the  $(m-k)$ th pseudonym as next hop is received, the source address of the forwarded packet also carries the  $(m-k)$ th pseudonym.

In comparison to the  $'1:1'$  scheme, with the  $'m:1'$  and  $'m:m'$  packet forwarding scheme a node can still use a neighbor node as forwarder though this node has changed its pseudonym. Eventually, the enhanced forwarding schemes eliminate stale entries caused by pseudonymity and reduces the packet loss. It is worth noting that the proposed enhanced schemes also allow a node to receive data packets when this node is the destination.

Despite the benefit, we identify the following issues related to anonymity: **In the  $'m:1'$  scheme**, an adversary might link a previous and the current pseudonyms of a node by comparing source and destination addresses before and after packet forwarding. In case the addresses are different, the attacker can presume a pseudonym change at the forwarder and associates both PIDs. Eventually, the attacker can track the node, which would render the pseudonym change as useless. However, we argue that the attacker model is restricted: Since with PBR the next forwarder is addressed by its MAC address, an attacker is required to be located in direct wireless range. This type of attack is reasonable compared to conventional methods (e.g. police car with a radar pursues another car). **With the  $'m:m'$  scheme** a node accepts data packets with a previous pseudonym as source address and also forwards the data packets with the same pseudonym. The problem is that other nodes along the path still update their location table for the forwarder's outdated pseudonym. As a result, the forwarder's entry does not time out and can be selected by subsequent forwarding decisions. An adversary may potentially misuse this feature for tracking a node by its previous pseudonym though a node has already changed the pseudonym. As a countermeasure we propose that the forwarder cloaks its position and time information in the packet header such that the following nodes do not update the corresponding entry, e.g. zero the position information. A next hop node would then check the plausibility of the values, drop the position information, but would process the remaining part of the data packet.

<sup>2</sup>The use of a Message Authentication Code (MAC) using symmetric cryptography is also possible.

#### D. Link Layer Callback

We propose the usage of link layer callback as a mechanism to alleviate the impact caused by changing pseudonyms.

Link layer callback exploits the capability of IEEE 802.11's positive acknowledgment and MAC level retransmissions. In order to remove stale entries in the location table, the link layer informs the upper layer about an unsuccessful data frame delivery. The upper layer PBR interprets the trigger either i) as a forwarder's pseudonym change or ii) a loss in wireless connectivity due to movement or bad wireless conditions. In both cases the node removes the stale entry from the location table and another node is selected as forwarder for subsequent packets. The use of link layer callback shortens the duration of time a node selects a stale location table entry as a forwarder. Without the callback the entry would remain in the table until it expires, typically in the range of seconds and subsequent packets waiting in the link layer queue would be lost. With link layer callback, a stale entry is removed after the first transmission attempt – including link layer retransmissions – in a time frame much smaller than a table entry lifetime. Passing a link layer frame back to the network layer allows to re-forward the data packet using an alternative table entry – potentially the same node, but with a new pseudonym. As a results, link layer callback reduces packet loss and saves bandwidth for packets sent via stale forwarding entries.

It is worth noting that the link layer callback requires acknowledgements for data frames as trigger for the callback. In IEEE 802.11 broadcast frames are not acknowledged and hence, link layer callbacks are not applicable for broadcast communication.

### V. ANALYSIS AND EXPERIMENTS

As proof of concept we extended an existing implementation of an OBU's communication system by support for pseudonymity, and conducted tests and performance measurements in an experimental testbed. The testbed is currently based on conventional notebooks each equipped with GPS device, IEEE 802.11a WLAN network interface card using Atheros chipset, and an external antenna to be fixed on a car roof. The notebooks run the Linux operating system with kernel version 2.6.10, the MADWIFI-NG driver for the WLAN card, and a prototype implementation of PBR in user-space.

In order to assess the performance of pseudonym change we define the following two metrics: *Local Pseudonym Configuration Delay*  $T_{LPC}$  and *Forwarding Delay*  $T_{FORW}$ .

#### A. Local Pseudonym Configuration Delay $T_{LPC}$

$T_{LPC}$  is referred to the duration of time between reception of the trigger for pseudonym change until the new pseudonym locally comes into effect. When  $T_{LPC}$  is small, applications – in particular for road safety – can send and receive data seamlessly and instantly inform the driver of hazardous situations.

$T_{LPC}$  affects both, unicast and broadcast packets. It is mainly determined by the execution of consecutive steps to set the wireless interface down, set the MAC address, and set the interface up, and finally to restore the network routes. It




Fig. 5. Unicast forwarding delay with table entry renewal

is independent of the applied forwarding scheme (Sec. IV-C) and can simply be calculated as follows:

$$T_{LPC} = T_{\text{face down}} + T_{\text{Set MAC Addr}} + T_{\text{face up}} + T_{\text{Restore IP Routes}} \quad (1)$$

#### B. Forwarding Delay $T_{FORW}$

We define  $T_{FORW}$  as the duration of time the pseudonym change causes for forwarding of a packet in addition to the normal delay for processing it without pseudonym change. A small  $T_{FORW}$  ensures that a node is continuously aware of neighbor node's presence though pseudonym change and can use this node as a forwarder to distribute information to other nodes in the region.

$T_{FORW}$  is different for geocast and unicast communication. Since geocast packets do not carry a destination address, but rather a geographical region as destination, these packets can immediately be forwarded independently of the next hop's address. After a pseudonym change the next hop is able to process the packet after  $T_{LPC}$ . The impact on unicast packets, however, can be worse and we distinguish between 3 different cases according to our proposals in Sec. IV. For simplification we assume a scenario with four nodes, denoted as *predecessor*, *best forwarder*, *backup forwarder* and *successor*, whereas the terms *best* and *backup* refer to the greedy routing metric of PBR. The *best forwarder* changes its pseudonym and causes the appearance of a forwarding delay.

**'1 : 1' forwarding scheme with table entry renewal.** After the pseudonym change, the predecessor keeps selecting the old identifier of the *best forwarder* as next hop, until this node sends the first beacon with the new pseudonym (Fig. 5). Assuming that  $I_B$  is the beacon interval,  $1/T_{data}$  is the constant rate of the data traffic such as  $T_{Data} < I_B$  and that  $t_{last}$  is the time when the *best forwarder* forwarded the last data packet using the old pseudonym, the first beacon with the new pseudonym will be sent at  $t_{last} + I_B$ . Indeed, in order to reduce the consumption of wireless resources, each node restarts the timer for sending beacons after each data transmission. Therefore, for  $I_B \gg T_{LPC}$ , the forwarding delay is expected to be  $T_{FORW} = I_B$ .

**'1 : 1' forwarding scheme with link layer callback.** Right after the *best forwarder* has changed its pseudonym, the predecessor sends the next data packet using the outdated pseudonym. By means of missing link layer acknowledgements, the predecessor detects that the *best forwarder* is not




Fig. 6. Experimental setup

available anymore at the previous pseudonym. A possible strategy for the predecessor to deliver the packet is to select an alternative next hop (i.e. the *backup forwarder*) until the *best forwarder* will announce its new pseudonym. Using this strategy, the forwarding delay is drastically reduced. Assuming that  $T_{First}$  is the time interval between the pseudonym change trigger and the first packet sent by the predecessor to the *best forwarder*'s old pseudonym, for values of  $T_{First}$  such as  $T_{First} > T_{LPC}$ , the forwarding delay is expected to be  $T_{FORW} = T_{First} + T_{Retrans}$ , where  $T_{Retrans}$  is the time needed to perform the link layer retransmissions. In case that  $T_{First} < T_{LPC}$  and data packets are sent at a constant rate of  $1/T_{Data}$  packets/second, we expect that  $T_{FORW} \leq T_{LPC} + T_{Data} + T_{Retrans}$ .

**'m : 1' and 'm : m' forwarding schemes.** The forwarding delay is minimal because the predecessor can still use the old pseudonym entry until this entry times out.

### C. Experimental Results

In order to quantitatively evaluate the impact of pseudonymity we set up an experimental testbed of nodes (Fig. 6), equipped as described above.

In a first test, we instrumented the PBR daemon implemented in user space to log time stamps at begin and end of the execution of pseudonym operations and measured  $T_{LPC}$ . As a result we measured a mean value of  $31.51 \mu s$  and a standard deviation of  $1.36 \mu s$  (number of tests  $n = 100$ ).

In a second test we set up four nodes A, B, C and D such that they are statically positioned as described in Sec. V-B: Node B and C are neighbor of A and D, and the sender A can choose among B and C as next hops to reach D (Fig. 6). Being closer to the destination D, node C is the *best forwarder* and changes its pseudonym at a random time during the multi-hop communication. Node A sends *ICMPv6 echo requests* to node D with a periodic interval of  $T_{Data}$ . By encapsulating the ICMP packets in PBR packets, they are forwarded from A via C to D. *ICMPv6 echo replies* are in turn sent from node D, via C, to A. We approximated the forwarding delay as  $T_{FORW} = T \cdot N_{lost}$ , where  $N_{lost}$  is the number of *ICMPv6 echo requests* lost due to the pseudonym change. For the '1 : 1' scheme with table entry renewal we yielded a mean value  $\mu T_{FORW} = 2.042 s$  and a standard deviation of  $\sigma T_{FORW} = 0.106 s$ , and with link layer callback  $\mu T_{FORW} = 0.102 s$  and  $\sigma T_{FORW} = 0.020 s$  (number of tests  $n = 100$ ). The duration between two consecutive pseudonym changes were  $T = 0.1 s$  and the beacon interval was set to  $I_B = 2 s$ .

In summary, the test results show that the delay in the forwarding process due to a pseudonym change can be ef-

fectively decreased, such that it becomes tolerable also for unicast applications with stringent requirements on end-to-end delay. This can be achieved with reasonable compromises with the location privacy (as for link-layer callbacks and enhanced forwarding schemes).

## VI. CONCLUSIONS

Changing pseudonyms is regarded as a technical means to protect the driver's privacy in future VANETs offering services for road safety, traffic efficiency, and infotainment. We have investigated the challenges that arise from the perspective of deploying and implementing pseudonymity in a VANET communication system. As solutions we presented four different aspects *i)* cross-layer addressing concept, *ii)* extended location service, *iv)* pseudonymity-enhanced packet forwarding schemes, and *v)* link layer callbacks. The solutions represent a framework to put existing algorithms for pseudonymity into practice. An analysis substantiated with experimental results shows that the costs of pseudonymity in terms of delay can be decreased. Applying the framework, the impact on the overall performance can be neglected and a reasonable tradeoff between the driver's privacy and deployability of changing pseudonyms can be made.

## REFERENCES

- [1] "Dedicated Short Range Communications (DSRC) Working Group," <http://grouper.ieee.org/groups/scc32/dsrc/>.
- [2] IEEE, "IEEE 802.11p/D0.21 Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Wireless Access in Vehicular Environments (WAVE)," January 2006.
- [3] "NoW - Network on Wheels," <http://www.network-on-wheels.de>.
- [4] E. Fonseca and A. Festag, "A Survey of Existing Approaches for Secure Ad Hoc Routing and Their Applicability to VANETS," NEC Network Laboratories, TR NLE-PR-2006-19, March 2006, 28 pages.
- [5] P. Golle, D. Greene, and J. Staddon, "Detecting and Correcting Malicious Data in VANETS," in *Proc. of VANET 2004*, Oct. 2004, pp. 29-37.
- [6] K. Sampigethaya et al., "CARAVAN: Providing Location Privacy for VANET," in *Proc. of ESCAR 2005*, Nov. 2005, 15 pages.
- [7] A. Aijaz et al., "Attacks on Inter Vehicle Communication Systems - an Analysis," in *Proc. of WIT 2006*, March 2006, pp. 189-194.
- [8] M. Raya and J.-P. Hubaux, "The Security of Vehicular Ad Hoc Networks," in *Proc. of SASN 2005*, Nov. 2005, pp. 11-21.
- [9] M. Gruteser and D. Grunwald, "Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking," in *Proc. of MobiSys'03*, May 2003, pp. 31-42.
- [10] A. Beresford and F. Stajano, "Mix Zones: User Privacy in Location-aware Services," in *Proc. of PerSec 2004*, March 2004, pp. 127-131.
- [11] H. Füllner et al., "Position-Based Routing in Ad-Hoc Wireless Networks," in *Inter-Vehicle Communications Based on Ad Hoc Networking Principles - The FleetNet Project*. Publisher University Karlsruhe, Germany, 2005, pp. 117-143.
- [12] J. Hubaux, S. Capkun, and J. Luo, "The Security and Privacy of Smart Vehicles," *IEEE Security and Privacy Magazin*, vol. 2, no. 3, pp. 49-55, May/June 2004.
- [13] VSC Project, "Final Report," May 2005, <http://www-nrd.nhtsa.dot.gov/pdf/nrd-12/060419-0843/>.
- [14] H. Füllner et al., "Thoughts on a Protocol Architecture for Vehicular Ad-Hoc Networks," in *Proc. of WIT 2005*, March 2005, pp. 41-45.
- [15] Y.-C. Tseng, S.-Y. Ni, and E.-Y. Shih, "Adaptive Approaches to Relieving Broadcast Storms in a Wireless Multihop Mobile Ad Hoc Network," *IEEE Transact. on Computers*, vol. 52, no. 5, pp. 545-557, May 2003.
- [16] R. Baldessari, A. Festag, A. Matos, J. Santos, and R. Aguiar, "Flexible Connectivity Management in Vehicular Communication Networks," in *Proc. of WIT*, March 2006, pp. 211-216.