

Impact of Pseudonym Changes on Geographic Routing in VANETs

Elmar Schoch¹, Frank Kargl¹, Tim Leinmüller²,
Stefan Schlott¹, and Panos Papadimitratos³

¹ Ulm University, Media Informatics Institute
{elmar.schoch, frank.kargl, stefan.schlott}@uni-ulm.de

² DaimlerChrysler AG, Group Research
Tim.Leinmueller@DaimlerChrysler.com

³ École Polytechnique Fédérale de Lausanne
panos.papadimitratos@epfl.ch

Abstract. Inter-vehicle communication is regarded as one of the major applications of mobile ad hoc networks (MANETs). In these so called vehicular ad hoc networks (VANETs) security and privacy are crucial factors for successful deployment. In a scenario, where each vehicle would have a unique identifier, eavesdroppers could easily accumulate location profiles.

As a solution approach, several authors suggest using changeable pseudonyms as temporary vehicle identifiers. If a vehicle changes its pseudonym from time to time, long-term tracking can be avoided. However, as we show in this paper, changing identifiers has detrimental effects on routing efficiency and increases packet loss.

So, designers of VANET systems need to aim for a balance between privacy protection on the one and performance on the other hand. The results of this paper provide advise on how to achieve this balance.

1 Introduction

Vehicular ad hoc networks – often called VANETs – are one of the most promising application scenarios for mobile ad-hoc networks.

With the advent of car-to-car communication, both passenger safety and driving comfort can be improved significantly. A car detecting an icy road could inform follow-up vehicles and thereby prevent accidents. If an accident occurs anyway, inter-vehicle communication could support emergency relief units to reach the accident site faster by warning drivers blocking the road ahead or preemption of traffic lights. Regarding driving comfort, inter-vehicle communication could serve to exchange traffic flow information for improved navigation or intelligent adaptive cruise control.

Several research initiatives (e.g. projects like Fleetnet [1] or CarTALK [2]), both in Europe and the U.S., have already produced results in the investigation of vehicular ad hoc networks. For instance, geographic routing has been selected as routing scheme due to its compliance with application needs and its good performance under extremely dynamic network conditions [3].

Ongoing work is now taking the next steps. One step is the effort to define common standards among car manufacturers, resulting in initiatives like the Car2Car Communication Consortium (C2C-CC) [4] and the Vehicle Safety Communication Consortium (VSCC) [5]. Another important step is the research on security and privacy issues of VANETs, because consumers will definitely not accept attackable systems in their cars nor the ability to trace their itinerary. In Europe, the SEVECOM project [6] is specifically dedicated to that.

The importance of privacy is illustrated in Figure 1. Because both geographic routing as well as many VANET applications make extensive use of position information, locations of vehicles are constantly exposed on the wireless communication channel. For instance, several VANET routing and application protocols use beacon messages that are broadcasted periodically, containing the current position and perhaps also speed or other vehicle information. While the dissemination of these data usually does not cause any problem when considering only a single moment, information and place, the combination of several data over time and at different places can uncover privacy relevant information.

As an example, large petroleum companies may have an interest in detecting the routes which (potential) customers travel throughout the day. Using this knowledge, they could plan new petrol stations or adapt prices based on customer behavior. In order to gather these data, they would simply install C2C-ready communication devices at their petrol stations¹ and collect the beacons sent out by all cars that carry VANET equipment. Using information gathered from electronic payment at petrol stations, these companies might even link cars to individual persons and start targeted advertisement for specific customer groups.

Whereas this scenario may have only a modest impact on the privacy of each individual², other scenarios with a more severe background seem far more threatening. Government agencies could easily control where people go with their cars in a much more complete and reliable fashion as it is possible with video surveillance and automatic image recognition. Likewise, private investigators could track and trace cars easily through the cities by following the cars in 100 meters distance which is equivalent to placing a radio beacon on the car.

One has to consider, that these location profiles may be accumulated over years and that you might become a suspect of a crime, just because your car was detected near a crime scene three months ago. This may also allow behavior profiles, e.g. your boss may be interested in the fact that you visit the hospital twice a month.

Previous work such as [7,8,9,10] suggests the use of randomly changing identifiers – so called pseudonyms – to prevent this kind of privacy intrusions. While it is still possible to collect data, associating it with identities over time gets much more complicated if node identifiers are used only for a short period of time.

However, when changing pseudonyms a number of new problems arise:

- *Traceability due to context*

A vehicle may be tracked despite of regular pseudonym changes because of certain circumstances. For instance, if the car changes its pseudonym while

¹ Depicted as circles in the figure.

² But it may have an influence on gas prices we have to pay.

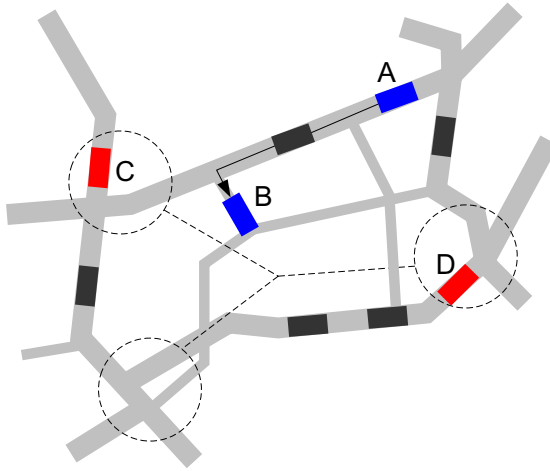


Fig. 1. Passive eavesdropper records beacons and/or application messages at important places and vehicle A tracks way of vehicle B

very few other vehicles are around, linking old and new pseudonym is rather simple by tracing its trajectory using beacons [7]. Similarly, if a car uses changing pseudonyms daily and is parked on the same reserved parking slot each day, the pseudonyms can also be related easily.

– *Traceability due to cross-layer influence*

Changing the pseudonym on one communication layer does not make sense if protocols on other, non-encrypted layers also use identifiers. In this case, node pseudonyms could be linked by the identifiers of other communication layers. So, changing pseudonyms must be coordinated between layers.

– *Security implications*

Anonymity has also drawbacks. Many security schemes that want to protect MANETs from selfish or malicious nodes propose mechanisms where these misbehaving nodes are first detected and then excluded from the network. With pseudonyms, misbehaving nodes can evade this exclusion by simply creating a new identity. Preventing this is a hard problem.

– *Problems with application protocols*

There are applications that need a long-term communication relationship between the involved parties. Examples include any type of file-transfer or interactive chat-sessions. Often, these protocols have an explicit session layer which controls authentication, association, stream control and similar issues. When identifiers change, it can become very complex and expensive to re-establish the session, as partners need to be re-authenticated³, some data may need to be replayed, etc.

– *Impact on communication protocols*

In most communication protocols, identifiers play a vital role. For example, beaconing is an important service for geographic routing as well as some

³ See [11] for a potential solution.

applications that deal with context-awareness in VANETs. While high frequency of changing pseudonyms improves privacy, it also complicates the design of communications protocols.

In this paper, we focus on the last aspect. Because geographic routing relies on stable identifiers of neighboring nodes, frequent pseudonym changes disturb proper routing functionality.

Changing the pseudonym once a day may be enough to prevent long-time tracking, but will not prevent a private investigator from following a car throughout the day. There are also applications where the car is needed to identify itself or to its communication partner, e.g. when you do electronic payment of tolls and the money is collected using bank transfers. Once this has happened, your movement profile for the whole day can be directly linked to your identity.

On the other hand changing the pseudonym only once every night while the car is parked at home in the garage has surely no significant influence on communication performance or on-going sessions.

On the other hand, changing the pseudonym every 10 milliseconds might increase privacy protection but will surely render most communication useless, as no node will be able to send you a packet as a reply to a previous packet you send earlier.

This paper analyzes the effects of privacy-enhancement mechanisms on the functionality of position-based routing protocols that forward packets hop-by-hop to the destination. In contrast to topology-based protocols, position-based routing is well suited to the specific characteristics of VANET scenarios [3], e.g. in terms of node mobility and application needs. From a privacy point of view, these protocols have the drawback that they link the position and identity of a node in every beacon message they send.

With our results, we aim at supporting the design of VANET systems that balance between privacy and operative requirements like performance or session stability. In our further analysis, we focus on the performance implications of pseudonym changes. In Section 2, we first describe a theoretical analysis of potential causes for packet loss and expected effects on routing. Later in Section 3, we support these findings by means of simulations. Before we finally summarize and conclude our results in Section 5, we give a short overview on related work in Section 4.

2 Effects of Pseudonym Changes on Geographic Routing

2.1 Routing Approach

We have based our analysis on the Cached Greedy Geocast (CGGC) routing protocol [12] which has been developed as part of the Fleetnet project. With CGGC, nodes periodically announce their identifier and current location using beacon messages. Nodes broadcast beacons every b seconds to all neighbors within reception range. Based on the information contained in beacons, nodes build up

neighbor tables. Table entries expire after t_o seconds and are removed from the table afterward.

If a node m generates a packet or receives one for forwarding, it searches its neighbor table for the node which is located closest to the destination:

$$\min(d(n, dest)) \forall n \in NT$$

where NT stands for neighbor table, $dest$ is the node identifier of the destination node, n is a node entry in the neighbor table, and $d(n, dest)$ is the Euclidean distance between n and $dest$.

If no such node is available (i.e. all $d(n, dest) >= d(m, dest)$ with m being the own node identifier), then the node simply stores the packet in a packet cache until a suitable neighbor becomes available due to node movement.

2.2 Analysis of Effects

For the analysis, we make the basic assumptions that beacons and data packets are sent at fixed intervals, but without any synchronization between each other. Likewise, nodes change pseudonyms at a fixed rate which is also not synchronized to the other intervals.

Data packets and beacons are sent as simple datagrams, there are no higher layer retransmission mechanisms (e.g. TCP) in place. This is a reasonable assumption, as most multi-hop applications which disseminate messages in VANETs work this way.

Let b be the beacon interval in seconds (e.g. 1s), p be the packet interval (e.g. 2s), and c be the pseudonym change interval (e.g. 10s).

For simplicity, we further assume $c > t_o$, i.e. there is at most one pseudonym change per beacon timeout interval. Changing pseudonyms more frequently is usually not reasonable. Further, we assume that $b \leq p$, because many VANET applications send simple information or warning messages at rather long intervals, so this assumption seems reasonable. However, we will also shortly discuss the $b > p$ case at the end.

Figure 2 shows an excerpt from a VANET scenario. Node m has created or received a packet for forwarding which is destined for node o . n is the the neighbor node with the smallest Euclidean distance to o and would be selected as next hop. n periodically broadcasts beacons with its identity and position which m uses to update its position table. But what happens when n changes its identity to n' ?

m still has n in its neighbor table and might send packets to n for forwarding. n' will however ignore these packets, because otherwise packets sent to n would be resent by n' so both identities could easily be correlated. Therefore such packets will get lost.

We are now interested in the percentage of packets that may get lost due to a pseudonym change in contrast to a similar scenario without pseudonym change. Figure 5 shows the potential cases and the respective probabilities.

Multiplying the probabilities for each path that leads from the root to a leaf where packets are lost and adding up the results gives us the overall probability




Fig. 2. Analysis scenario




Fig. 3. Timeline




Fig. 4.

of packet loss or in other words the expected percentage of packets that will be lost due to pseudonym changes. The result is that $\frac{t_o}{2c}$ of all packets will get lost.

We now explain the probability tree in detail from top to bottom.

Is a packet affected by pseudonym change? Looking at Figure 3, we see that only packets that follow a pseudonym change (PC) can be affected by this PC. One t_o interval after the last beacon with the old pseudonym (b) has been sent, this information will be removed from the neighbor table and only information from beacons containing the new pseudonym (b') will be available to node m . As the last beacon (b) was sent on average $b/2$ before the pseudonym change and as the t_o interval starts from this point on, the average timeframe for affected packets is $t_o - \frac{b}{2}$. Considering the whole time between two pseudonym changes, $\frac{t_o - \frac{b}{2}}{c}$ of all packets will be affected on average. The other $1 - \frac{t_o - \frac{b}{2}}{c}$ packets will be delivered regularly.

Is a packet the first packet after PC? For the first packet after each pseudonym change, there are two alternatives: either the packet p is sent before the next beacon b' or is sent after the next beacon b' . In the first case, there is no information available on the new pseudonym n' and the packet will be definitely lost. In the second case, the fate of the packet depends on the node movements, as we will see in the next step.

In Figure 4, we explain how to calculate the probability of first receiving a packet before a beacon. We still assume $b < p$. Depending on the parameters b and p , there is a varying amount of "complete" beacon intervals within each packet interval, on average $\lfloor \frac{p}{b} \rfloor$. When the pseudonym change happens inside one of these intervals, we will definitely receive another beacon before the next packet.

The "rest" in front and at the end of the packet interval has a varying length which depends on the offset of b at the beginning of the p interval. Depending on this offset, this "rest" ranges between zero and two beacon intervals. On average, it will be one beacon interval, as the packet sent always divides one

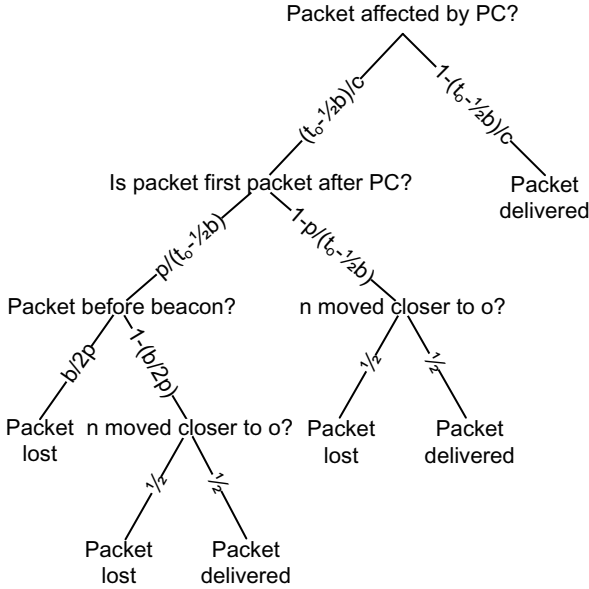


Fig. 5. Loss probability tree

interval which then contributes to the "rest" of the previous and next packet interval.

The relative amount of time of one beacon interval compared to one packet interval is $\frac{b}{p}$. On average, half of this time is located at the beginning of the packet interval and half of the time at the end. Only when the pseudonym change happens in the part the end of a packet interval, the packet is sent without a preceding beacon and the packet is lost (see Figure 4). The probability of loosing a packet because of this reason is therefore $\frac{b}{2p}$.

Has n moved closer to o while changing pseudonyms? For all packets that are affected by a pseudonym change and that are sent after a beacon with the new pseudonym (b') has been received, the following situation applies: the sending node has both information on n and n' in the neighbor table. As the sending node m cannot link n and n' , it will simply select the forwarding node based on its routing metric, i.e. the node which is closer to the destination o . Let $d(n, o)$ be the Euclidean distance between n and o . We assume that node n can move in the radius r within one beacon time as shown in Figure 2.

The last beacon before pseudonym change reported position p_n , the next beacon after pseudonym change reports position $p_{n'}$. As m cannot correlate n and n' , it assumes that there is a known node at position p_n and another node at position $p_{n'}$. If $d(n', o) < d(n, o)$, the new node will be preferred, otherwise m will try to send its packets to the previously known node.

If $d(n', o) < d(n, o)$, the packet will be received and forwarded, if $d(n', o) > d(n, o)$ the packet will not be received and gets lost. If $d(n', o) = d(n, o)$, m randomly selects one of the two nodes, which gives a 50% chance of success or loss.

If o is far away ($d(n, o) \gg 1$), the circle around n with radius r is divided in two halves, where all positions on the left are further away from o and all positions on the right are closer to o . Assuming random node movement of n , the chance of packet loss because of movement is therefore estimated to be $\frac{1}{2}$.

One might object that node n might even move to a position outside of transmission range of m and packets will then be lost with 100% probability. As this case can however occur with or without pseudonym change and packets sent to n are not received at the new position $p_{n'}$ anyway, the packet is lost, no matter if there is a pseudonym change or not. Since we are only interested in additional packet loss due to pseudonym change, this case will be neglected here.

Of course, movements of other nodes in the neighborhood of m might also change the potential forwarding node. This effect is also not considered here as it also happens independent of pseudonym change at the same rate.

If we now go back to the probability tree in Figure 5 we need to add the multiplied probabilities of all paths leading to a leaf node where packets get lost:

$$\begin{aligned} P_{loss} &= \frac{t_o - \frac{b}{2}}{c} \frac{p}{t_o - \frac{b}{2}} \frac{b}{2p} + \\ &\quad \frac{t_o - \frac{b}{2}}{c} \frac{p}{t_o - \frac{b}{2}} \left(1 - \frac{b}{2p}\right) \frac{1}{2} + \\ &\quad \frac{t_o - \frac{b}{2}}{c} \left(1 - \frac{p}{t_o - \frac{b}{2}}\right) \frac{1}{2} \\ &= \frac{t_o}{2c} \end{aligned}$$

What is interesting to see is that the loss probability is independent of the packet send rate and the beacon rate, but instead depends only on the relation between neighbor cache timeout and pseudonym change rate.

This result is valid only for our assumptions where $b \leq p \leq t_o \leq c$. Using a similar reasoning we can also show that for $b \geq p$, the loss probability is

$$P_{loss} = \frac{b^2 - c^2 + 4bp + 2t_ob}{4bc}$$

As you can see, in this case the situation gets more complex and all four parameters influence the loss probability.

In the next section, we will now present the results of simulations that analyze the effects of pseudonym changes on the packet delivery rate.

3 Simulation Results

The analysis in the previous section clearly points out which effects may occur when pseudonyms are changed. To estimate the order of magnitude of these effects on geographic routing, we conducted simulations with the network simulator ns-2, version 2.29 [13].

Table 1. Short overview on simulation parameters

| Parameter | Value |
|-------------------------------|-----------------|
| Number of nodes | 100 |
| Length of square node field | 1000 – 4000m |
| Max. node velocity | 10 – 50 m/s |
| Pause times | 0.0 s |
| Mobility model | Random Waypoint |
| Link-/MAC layer | IEEE 802.11b |
| Wireless transmission range | 250 m |
| Number of sent messages | 500 |
| Pseudonym change interval p | 5 – 60 s |
| Simulation time | 120 s |
| Simulation runs | 20 |
| Beacon interval b | 1 s |
| Neighbor cache timeout t_o | 6 s |

In these simulations, nodes are equipped with the previously described greedy-based geographic routing layer. Besides, every node changes its pseudonym with a defined frequency that is randomly jittered within ± 5 seconds. After having changed its pseudonym, just packets addressed to the node’s new identifier are accepted. Pseudonym change and beaconing intervals are completely independent, which means that there is no extra beacon after the pseudonym change. Both settings help keeping privacy - accepting packets for the old addresses or sending a beacon immediately after the pseudonym change would cut location privacy to some extent.

As data traffic, messages are generated and sent from randomly selected source nodes to random destinations using the described geographic routing protocol. This leads to larger distances between sender and destination when the network field size is increased. Moreover, packets are sent as geo-anycast, which assures that a destination is reachable regardless of changing identifiers. Detailed simulation settings like network topology, node mobility and composition are summarized in Table 1.

3.1 Basic Impact of Pseudonym Changes

A straight-forward measurement to quantify influences on routing is given by the number of packets that reach their destination. This delivery ratio directly reflects the performance of the routing protocol. Figure 6 shows the decrease in delivery ratio that is caused by pseudonym changes in relation to routing with stable identifiers, which is marked in the graph by the value ∞ as pseudonym change interval. Main insights of Figure 6 are on the one hand that delivery performance almost reaches the level without any identifier change when the interval is 60 seconds. On the other hand, a change interval of about five seconds seriously decreases delivery ratio. For instance, in the network with $2000m \times 2000m$ field size, geographic routing usually delivers over 60% of all messages,

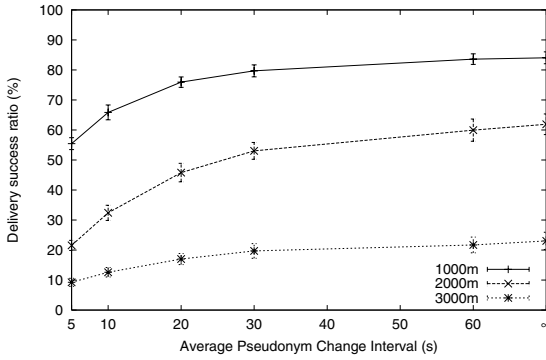


Fig. 6. Packet delivery ratio with influence of pseudonym change at different frequencies and network sizes




Fig. 7. Reduction of packet delivery ratio with pseudonym change compared to normal forwarding (100%) in dependence of network size and pseudonym change frequency




Fig. 8. Packet delivery ratio with different node velocities and both with and without pseudonym changes

whereas only little more than 20% reach their destination when nodes change their identifiers every five seconds.

3.2 Influence of Node Density

Another observation from Figure 6 is the fact that networks with higher node density (like $1000m \times 1000m$) can cope better with pseudonym changes. In this configuration, the additional packet loss is less than 30%, but over 40% in case of the $2000m \times 2000m$ sized field. Figure 7 gives a more detailed insight of the influence of node density. It shows the relation between successfully delivered packets when using changing pseudonyms in contrast to permanent identifiers. As we can see, even with a pseudonym change frequency of 30s, packet delivery ratio may decrease almost 20%. The fact that there is a peak decrease at about $2000m$ to $2500m$ field side lengths is due to decrease of delivery ratio without changing identifiers when node density gets low ($\sim 20\%$ at $3000m \times 3000m$, see also Figure 6). Therefore, only packets with short trip reach their destination in low-density networks anyway and thus are also likely to face no pseudonym change, too.

Regarding privacy, changing the pseudonym only every 30 seconds may perhaps be a too long time already. Though such a change interval surely avoids being tracked in a global scope, an attacker following another vehicle may be able to figure out into which direction the haunted vehicle has turned off at an intersection. On the other hand, re-identifying the tracked car after it has actually changed its pseudonym could be a difficult task for the attacker if the pseudonym change is done carefully.

3.3 Influence of Node Velocity

Node velocity is a crucial parameter in VANETs. Geographic routing has shown to cope well with this requirement [14]. As depicted in Figure 8, when using pseudonyms, delivery success ratio does not decrease much with higher node velocity. Interestingly, this contrasts to the decrease of successfully delivered packets when routing can rely on stable node identifiers. Particularly in the scenario with lower node density ($2000m \times 2000m$), delivery ratio decreases notably from 60% to 30%, whereas the difference with 10s pseudonym change interval is only about 10% between $10m/s$ and $50m/s$ maximum node velocity. Hence, the effect of changing pseudonyms decreases with higher node velocities.

3.4 Comparison with Theoretical Analysis

The loss probability of $\frac{t_o}{2c}$ that we found in the analysis in section 2 is independent of beacon and packet intervals. As this is a result that one might not expect, we explicitly verified it in our simulations. Figure 9 shows the packet delivery ratio with different values for p , c , b and t_o . Though single result values differ about 5%, at large, the graph confirms that b and p are not relevant for loss probability. Besides, also the order of magnitude of losses coheres with the analytical result.




Fig. 9. Packet delivery ratio at different pseudonym change intervals, with varied beacon intervals and message numbers

Taking a look at Figure 7, we see that the loss due to pseudonym change corresponds to on average 50% with $c = 5s$, 30% with $c = 10s$, and 5 – 10% with $c = 30s$. Table 2 compares this to the expected results using our formula $P_{loss} = \frac{t_a}{2c}$.

As a summary, we can conclude that we have a good correspondence of simulation findings and analytical results which both support our claim that pseudonym changes can lead to a significant reduction in routing performance under certain circumstances. It is interesting to see that the packet and beacon rates do not influence this loss.

3.5 Improvement with MAC/LL Callback

The basic reason why packets get lost is because they starve out in link layer interface queues due to neighbors that are no more available. Therefore, applying a direct callback mechanism from link layer to routing layer is likely a method to overcome the problem of outdated neighbor table entries due to pseudonym change. In an enhancement of our simulation implementation, we tested this mechanism in conjunction with pseudonym changes. Thus, the link layer informs the geographic routing about unreachable neighbors immediately after a transmission failed and passes back the packet. The routing then takes up the packet again, determines the next hop the packet was originally sent to and removes it from its neighbor table. After the update, the packet is re-enqueued for routing.

Figure 10 depicts the aggregated number of MAC callbacks that occurred during simulations. In consistence with the previous results, most link failures

Table 2. Comparison of analytical results with simulation findings

| c | Simulation | Analytical |
|----|------------|---------------------------------------|
| 5s | ≈ 50% | $P_{loss} = \frac{5s}{2*5s} = 50\%$ |
| 10 | ≈ 30% | $P_{loss} = \frac{5s}{2*10s} = 25\%$ |
| 30 | ≈ 5 – 10% | $P_{loss} = \frac{5s}{2*30s} = 8,3\%$ |

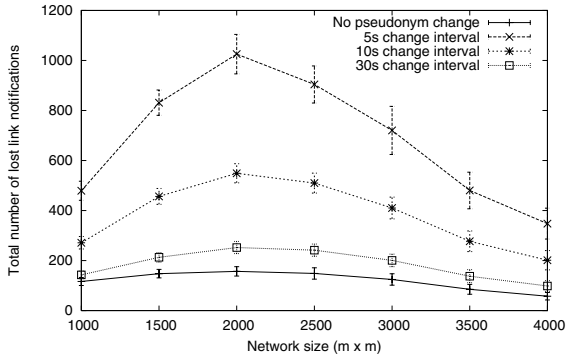


Fig. 10. Aggregated number of lost link notifications

occur in mid-sized networks. More importantly, the simulation results also show that the MAC callback mechanism is able to reduce the performance decrease almost to zero even with pseudonym change frequency of 5s.

Unfortunately in reality, wireless links usually are rather unstable. Thus, if a node removes a neighbor immediately after a single transmission failed, links to neighboring nodes may be removed though they are only temporarily unavailable. To meet this problem, direct MAC callback has to be used carefully, e.g. only after a set of retries.

4 Related Work

With progressing research on VANETs, the quest for privacy has emerged as a crucial factor. Several authors, for instance Hubaux et al. [15] or Dötzer in [8], addressed that topic. They argue, that cars are personal devices that are usually kept for a rather long time and even innocent looking data may become privacy-relevant when evaluated over a longer period of time. Therefore, they propose to use changing pseudonyms as temporary identifiers to preserve privacy. In [15], Hubaux et al. also review entropy as a metric to quantify the effectiveness of pseudonym changes. In [16], Sampigethaya et al. take up the idea of selecting certain nodes as mix nodes. All nodes that belong to one cluster communicate only through their mix node and thus manage to stay private.

On the other hand, [8] also states that vehicles are expected to work reliably, implying that applications of inter-vehicle communication have to face this requirement as well. Unfortunately, from the point of view of security, detection and exclusion of malicious nodes usually relies on the ability to identify nodes. Thus, there is a clear tradeoff between security and privacy. For example, as Golle et. al stated in [7], higher pseudonym change frequency leads to smaller margins for detection and correction of malicious behavior.

The proposed solution to the problem in [8] is to deploy a trusted third party, that issues a limited number of pseudonyms per vehicle and records the

corresponding, real identity. In case of problems, the issuer can withdraw the pseudonyms and disclose the real identity if necessary.

Further work was done in the field of location privacy in pervasive computing. In [17], Beresford and Stajano propose so-called mix-zones to overcome the linkability problem when nodes change pseudonyms arbitrarily. Schlott et al. investigate attacks on random pseudonym change schemes using some side channel information in [18].

5 Summary and Conclusion

In this paper, we focused on the effects of pseudonym changes on the performance of geographic routing that is intended to be used in VANETs. The analysis shows, where pseudonym changes affect routing procedures and result in packet losses. Both analytical results and simulation confirm serious performance decreases in case of less dense network connectivity and high pseudonym change intervals ($< 30s$).

We suggest introducing a callback mechanism which informs the routing about failed transmissions. The routing can then cope better with pseudonym changes. On the other hand, such a callback mechanism needs to be implemented carefully because links are rather unstable in highly dynamic ad hoc networks like VANETs.

In conclusion, our work shows that operational and privacy requirements need to be balanced in VANETs. This can be achieved by choosing appropriate pseudonym change intervals and implementing a "soft" callback from link layer, for instance if a transmission failed several times.

Currently, both simulation and analysis focus on the case where beacons are sent more often than data packets ($b < p$) because we estimate this to be very important for many eSafety applications. We are now about to also investigate the opposite case with $b > p$.

These results will help us to develop a privacy protection mechanism for VANETs, which is one of the objectives of the SEVECOM project.

Acknowledgements

Parts of this work have been carried out in contribution to the SEVECOM project [6] that is supported by the European Commission e-Safety initiative under contract no. IST-027795. We also would like to thank Matthias Gerlach for his comments.

References

1. Franz, W., Wagner, C., Maihöfer, C., Hartenstein, H.: Fleetnet: Platform for inter-vehicle communications. In: Proc. 1st Intl. Workshop on Intelligent Transportation, Hamburg, Germany (2004)
2. CarTalk 2000 Project. <http://www.cartalk2000.net> (2004)

3. Mauve, M., Widmer, J., Hartenstein, H.: A survey on position-based routing in mobile ad-hoc networks. *IEEE Network* **1** (2001) 30–39
4. Car2Car Communication Consortium. (<http://www.car-to-car.org/>)
5. US Vehicle Safety Communication Consortium. (<http://www-nrd.nhtsa.dot.gov/pdf/nrd-12/CAMP3/pages/VSCC.htm>)
6. SEVECOM - Secure Vehicle Communications Project. (<http://www.sevecom.org/>)
7. Golle, P., Staddon, D.G.J.: Detecting and correcting malicious data in vanets. In: *Proceedings of the First ACM Workshop on Vehicular Ad Hoc Networks (VANET)*, Philadelphia, USA (2004)
8. Doetzer, F.: Privacy issues in vehicular ad hoc networks. In: *Workshop on Privacy Enhancing Technologies*, Cavtat, Croatia (2005)
9. Aijaz, A., Bochow, B., Dötzer, F., Festag, A., Gerlach, M., Kroh, R., Leinmüller, T.: Attacks on inter vehicle communication systems - an analysis. In: *Int'l Workshop on Intelligent Transportation (WIT)*. (2006)
10. Raya, M., Hubaux, J.P.: The security of vehicular ad hoc networks. In: *Proc. of Third ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN 2005)*, Alexandria, USA (2005)
11. Schlott, S., Kargl, F., Weber, M.: Re-identifying anonymous nodes. In: *International Workshop on Location- and Context-Awareness (LoCA 2006)*, Dublin, Ireland (2006)
12. Maihöfer, C., Eberhardt, R., Schoch, E.: CGGC: Cached Greedy Geocast. In: *Proc. 2nd Intl. Conference Wired/Wireless Internet Communications (WWIC 2004)*. Volume 2957 of *Lecture Notes in Computer Science*, Frankfurt (Oder), Germany, Springer Verlag (2004)
13. ns 2, N.S. <http://www.isi.edu/nsnam/ns/> (2004)
14. Füssler, H., Mauve, M., Hartenstein, H., Käsemann, M., Vollmer, D.: A comparison of routing strategies for vehicular ad hoc networks. *Technical Report TR-3-2002*, Department of Computer Science, University of Mannheim (2002)
15. Hubaux, J.P., Čapkun, S., Luo, J.: The security and privacy of smart vehicles. *IEEE Security and Privacy* **4** (2004) 49–55
16. Sampigethaya, K., Huang, L., Li, M., Poovendran, R., Matsuura, K., Sezaki, K.: Caravan: Providing location privacy for vanet. In: *Proceedings of Embedded Security in Cars (ESCAR)*. (2005)
17. Beresford, A.R., Stajano, F.: Location privacy in pervasive computing. *IEEE Pervasive Computing* **2** (2003) 46–55
18. Schlott, S., Kargl, F., Weber, M.: Random ids for preserving location privacy. (2005)