Computer Communications 122 (2018) 118–128

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Readjusting the privacy goals in Vehicular Ad-Hoc Networks: A safety-preserving solution using non-overlapping time-slotted pseudonym pools

David Eckhoff*,[a], Christoph Sommer[b]

[a] Robotics and Embedded Systems, Department of Computer Science, Technische Universität München, Germany
[b] Cooperative Mobile Systems, Heinz Nixdorf Institute and Dept. of Computer Science, Paderborn University, Germany

ABSTRACT

Current proposals for privacy measures in vehicular networking commonly suffer from either of three limitations: Many are so drastic that they compromise road traffic safety, a primary goal of vehicular networks. Others are more compliant, but (despite adding substantial overhead) are ineffective at protecting users' privacy against state-of-the-art attackers. The last group suffers from neither limitation, but undermine accountability and are thus in conflict with the requirements of future vehicular networks. As a consequence, workable privacy protection is often thought unattainable and privacy protection simply disregarded in the many field experiments, proposals, and standardization documents to date. In this work, we start fresh with a readjusted view on privacy goals and the capacities of a state-of-the-art attacker in mind, taking a structured approach to deriving a holistic solution for location privacy protection in Vehicular Ad-Hoc Networks (VANETs): We show that local privacy protection cannot be attained without compromising road traffic safety – nor does it have to be. Our approach is based on synchronized time-slotted pseudonym pools and the local announcing of pseudonym changes. By this, we overcome the privacy–safety problem while at the same time increasing privacy for all users. Our system is fully compatible with the requirements of vehicular networks and current standardization efforts.

1. Introduction

Vehicular networking, the wireless exchange of data between vehicles, is a key component of future intelligent transportation systems. Wirelessly sharing information among vehicles promises to improve road traffic safety and, as a pleasant side effect, can enable novel business concepts. Many vehicles in Japan can already rely on wireless short range communication technology and both the American IEEE and the European ETSI are finalizing standardization documents for the operation of future vehicular networks. Indeed, the US DOT has even announced its intent to make ad-hoc communication systems mandatory for new cars to reap the full benefits for road traffic safety that such a system can bring when deployed universally [1].

One of the key features targeted by the proposed systems is cooperative awareness, that is, vehicles communicating with each other to establish a virtual view of their surroundings for the sake of improving traffic safety. This is envisioned to be achieved with the help of periodic beacons, broadcast transmissions sent by a vehicle which include its current state. Beacons are commonly sent with a frequency of 1 Hz to 10 Hz. In IEEE WAVE and ETSI ITS-G5, each vehicle wirelessly informs all cars in its vicinity of state information that includes its current position, current heading, and current velocity.

As these transmissions are (and need to be) decodable by the general public, however, they can be received not just by other vehicles but by anyone with a receiver physically close enough to the sender (even at distances as far as 300 m to 800 m [2]). An adversary could exploit this and track a vehicle simply by linking consecutive transmissions. This can lead to a violation of location privacy of drivers, and through that also other types of privacy [3,4].

This privacy problem in vehicular networks has been understood from the very beginning [5]. The consensus in vehicular network privacy research is to use changing short-term identifiers, that is, *pseudonyms*, instead of static ones to complicate tracking for any eavesdropping adversary. An important challenge is to employ a suitable pseudonym change strategy, i.e., when (or where) a vehicle should change its pseudonym to maximize its location privacy. A broad range of these strategies [6], many of them aligned with the general privacy framework proposed in [7], has been proposed in the time since – and some of them even also considered in various field trials [8,9], albeit not always with the focus necessary to pave the way for a concrete strategy to become part of standardization documents.

A major obstacle in finding a suitable pseudonym changing strategy

* Corresponding author.
E-mail addresses: david.eckhoff@tum-create.edu.sg (D. Eckhoff), sommer@ccs-labs.org (C. Sommer).

https://doi.org/10.1016/j.comcom.2018.03.006
Received 30 April 2017; Received in revised form 29 January 2018; Accepted 11 March 2018
Available online 04 April 2018

is the fact that there seems to be no agreement in the parameters [10]: Strategies differ with regard to the adversary against which they are protecting, how they influence other applications such as traffic safety, and their compatibility with other system requirements, such as accountability or computational complexity. We believe that to make privacy protection a fundamental part of future vehicular networks, they have to take into account all these constraints and requirements. For example, safety applications rely on receiving and linking periodic messages; any privacy protection mechanism interfering with these applications is thus unlikely to be deployed.

In this work, we take a realistic look at the requirements of envisioned intelligent transportation systems and propose a holistic pseudonym-based solution that increases privacy without sacrificing safety:

- We make use of non-overlapping time-slotted pseudonyms to increase the overall privacy protection.
- At the same time, we advocate putting an end to chasing the goal of confusing eavesdropping adversaries, as this is completely opposite to the primary purpose of vehicular networks: allowing vehicles to track other nearby vehicles to avoid collisions.
- We support this claim with a detailed simulation study based on synthetic mobility and on real-world traces to show that confusing local adversaries is not possible without also affecting traffic safety.
- We present the underlying model of a state-of-the-art attacker using a multi-target tracking algorithm.

Our results give insights into the limitations of pseudonym changing strategies and consequently allow us to effectively tackle the privacy–safety trade-off. In addition, our solution is insusceptible to Sybil attacks and allows for efficient and privacy-preserving certificate revocation. It is also fully compatible with the upcoming North American IEEE and European ETSI families of standards.

This manuscript constitutes an extended version of our previous work [11], now also including the proposed multi-target tracking algorithm used to model the attacker in our computer simulation studies and an in-depth description of the simulation setup.

The remainder of this manuscript is structured as follows: In Section 2 we describe the status quo of vehicular network privacy systems as envisioned in IEEE WAVE and ETSI ITS-G5. Here, and throughout the remainder of the manuscript, we will also refer to and discuss related work. Section 3 discusses the privacy threats of vehicular networks; Section 4 explains the constraints that privacy protection mechanisms must work within. In Section 5 we present our solution which we believe is a viable approach to coping with the location privacy challenges in VANETs without negatively impacting traffic safety. We back up central claims that motivated the construction of our solution using a novel model of a state-of-the-art attacker (Section 6) and investigating its leverage against state-of-the-art privacy solutions in a computer simulation study (Section 7).

2. Status Quo: vehicular PKI

Authenticity and integrity are essential security requirements in vehicular networks. Only authorized devices should be able to participate in the network and it must be guaranteed that forged messages can be detected as such. These security goals can be achieved by means of a Public Key Infrastructure (PKI) as described in standards of IEEE (1609.2-2016) and ETSI (102 941). In addition, this PKI is also the basis for privacy protection through the use of authenticated pseudonymous identifiers.

A (slightly simplified) explanation of the system is shown in Fig. 1. Vehicles are equipped with a base identity (or long-term identifier), consisting of a certificate and public-private key pair (Step 1). This identity is unique to a certain vehicle and must therefore never be used for car-to-car communication. It is only used to generate or request pseudonyms (in the form of pseudonymous certificates) from a Certificate Authority (CA) trusted by all vehicles (Step 2). If the identity is valid (as indicated by a signature of the CA) and the information in the pseudonym request is correct, the CA signs the pseudonyms and sends them back to the vehicle (Step 3). Each vehicle maintains a pool of pseudonyms and uses a selected pseudonym as its visible address, that is, to sign and send messages over the wireless channel (Step 4). Other vehicles will only consider received messages if signed with a valid pseudonym.

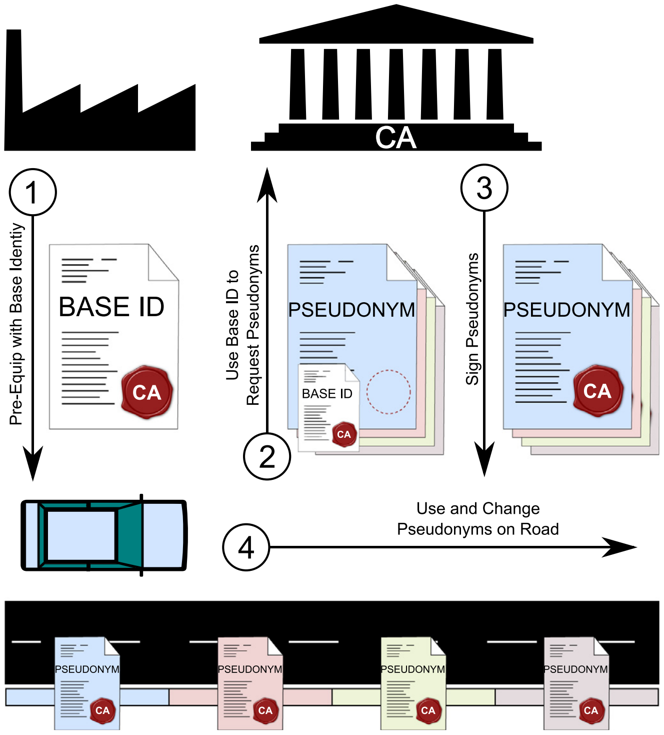

Fig. 1. A simplified vehicular Public Key Infrastructure (PKI).

It is, however, unclear how these pseudonym pools are organized and how vehicles should select which pseudonym to use for which transmission. For example, it was discussed that multiple (or even all) pseudonyms are valid at the same time and that the On-Board Unit (OBU) of the vehicle can choose freely or randomly which pseudonym to use. This introduces the problem of Sybil attacks [12,13], that is, one vehicle pretending to be many at the same time, thus subverting consensus-based approaches to credibility checks. Other vehicles would have no trivial method of identifying such an attack, as they cannot link different pseudonyms to the same vehicle. In earlier work, we have suggested the use of non-overlapping pseudonyms to avoid this problem [14].

Other proposals for privacy protection include the use of silent periods, that is, not transmitting beacons after a pseudonym change [15] or the use of group cryptography [16] to prevent eavesdropping. However, both of these proposals are not compatible with the upcoming standards as they interfere with traffic safety or conflict with the unencrypted transmissions of periodic beacons. Gerlach et al. have proposed to consider the context of a vehicle to determine when a pseudonym change can be effective [17], Freudiger et al. presented their concept of mix-zones, that is, geographic areas for pseudonym changes [18]. The results we present in Section 7 show that these proposals are not sufficient to protect the privacy of drivers.

As of today, the IEEE and ETSI family of standards do not recommend a specific pseudonym changing strategy, nor do they discuss existing solutions. The documents only mention the need to “use a pseudonym that cannot be linked to [⋯] the user’s true identity” (ETSI 102 893-v1.1.1) and suggest to change it frequently “[⋯] to avoid simple correlation between the pseudonym and the vehicle”