Balancing Auditability and Privacy in Vehicular Networks

Jong Youl Choi
Dept. of Computer Science
Indiana Univ. at Bloomington
Bloomington, IN 47405, USA

Markus Jakobsson
School of Informatics
Indiana Univ. at Bloomington
Bloomington, IN 47408, USA

Susanne Wetzel
Dept. of Computer Science
Stevens Inst. of Tech.
Hoboken, NJ 07030, USA

ABSTRACT

We investigate how to obtain a balance between privacy and audit requirements in vehicular networks. Challenging the current trend of relying on asymmetric primitives within VANETs, our investigation is a feasibility study of the use of symmetric primitives, resulting in some efficiency improvements of potential value. More specifically, we develop a realistic trust model, and an architecture that supports our solution. In order to ascertain that most users will not find it meaningful to disconnect or disable transponders, we design our solution with several types of user incentives as part of the structure. Examples of resulting features include anonymous toll collection; improved emergency response; and personalized and route-dependent traffic information.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*

General Terms

Algorithms, Management, Design, Security, Verification

Keywords

Audit, incentive, light-weight, privacy, symmetric

1. INTRODUCTION

The wealth of information that could be obtained from vehicular networks is quite enormous, ranging from location and speed to emergency alerts and requests for roadside assistance. However, the very richness of information also threatens to cause deployment to come to a grinding halt if there are adverse consumer reactions to the technology. Thus, in order to deploy vehicular networks on a large scale, and provide personalized services beyond the most straightforward applications (such as emergency alerts), we believe it is necessary to ensure that the deployed system respects the privacy of users—both with respect to each other and the infrastructure.

The protection of privacy is an active research topic in computer security. Traditional public key based cryptography—exemplified by blind signatures (e.g., [6]), magic ink signatures [17], probabilistic encryption (e.g., El-Gamal encryption) and mix networks (e.g., [8, 12])—offers mathematically crisp answers to privacy needs, but suffers problems both in terms of usage and deployment costs. While a large body of research relies on public key techniques to offer privacy guarantees, significantly less work has been done in the area of symmetric key cryptography. The context of privacy within vehicular ad hoc networks (VANETs) is not an exception to this rule.

We believe that it is important to consider solutions based on symmetric key cryptography, and to evaluate the pros and cons of these. While symmetric key cryptography is less flexible in many ways than asymmetric (or public key) approaches, it is also well-known to allow for more compact representations, require less computational effort, and be less vulnerable to cryptanalytic advances. This is associated with benefits in an environment like the one we consider. For example, consider a situation where vehicles exchange information with another. Here, the exchange of information must be performed within a very short time frame, which limits both the possible message generation time and the available effective bandwidth. This suggests that there are benefits associated with using symmetric cryptographic techniques, as these typically result in smaller transcripts. At the same time, one could argue that the infrastructural advantages associated with a public key approach are limited, due to the enormous cost of maintaining an up-to-date certification infrastructure in which it sometimes may be impossible for a node to establish timely communication with a CA (Certificate Authority) or a CRL (Certificate Revocation List). We believe many local decisions have to be made without the benefits normally associated with a certificate infrastructure. Bluntly stated, we therefore believe that public key approaches may not be more suitable than symmetric key solutions in the context of many of the tasks of relevance.

Thus, this paper is a feasibility study on the topic of using symmetric cryptosystems to build vehicular networks with balanced privacy and auditability, but with lighter requirements on communication and computation than current proposals. While the need to base a solution on symmetric key primitives does not pose any unsurmountable challenge, it *does* restrict the types of functionality that can be obtained

Q2SWinet'05, October 13, 2005, Montreal, Quebec, Canada.
Copyright 2005 ACM 1-59593-241-0/05/0010 ...$5.00.

in a setting that relies on peer-to-peer communication for collection of data.

1.1 Incentives

Incentives are important independently of implementation issues. However, they are particularly important in contexts that do not allow for local verification of data, as is the case in symmetric models in which one does not establish pairwise keys between all peers. More specifically, we study incentive mechanisms suitable to prevent large-scale cheating that could otherwise be the result of the lack of public verifiability inherent in most[1] symmetric key designs. Without attractive incentives, large numbers of consumers simply disable their transponders. In particular, in order to be assured of consumer approval—and use—of any design, we must both be cautious not to make such a solution undesirable to rational users willing to disconnect equipment or limit its functionality.

One can argue that nobody benefits from having data collected that indicates that they were at fault in causing an accident to occur, and therefore, that everybody would opt out. We believe that this is not the case. In particular, if the undesired events are uncommon and unexpected, and consumer psychology denies the likely occurrence of such an event (such as being at fault) *to them*, then our solution will still be rational to use if there are sufficient benefits of some type that can be anticipated and experienced. However, even events of unexpected nature can serve as a motivator to deploy the system; for example, faster and better customized emergency response would serve as a reassurance even to consumers that do not believe they are likely to be in an accident. Another example of an incentive is anonymous collection of tolls; this avoids tracking of user-identifiable movement patterns, and is trivially achieved in our architecture. We describe how to achieve a handful of features to incentivize the use of our proposed methods; however, it should be noted that this is mostly for the purposes of illustration of the diversity of incentives that can be applied within our architecture. We anticipate that many more can be added by using the same basic mechanisms we base our solution on.

1.2 Privacy

It is important to realize that there are numerous interpretations of privacy; these are often highly application dependent. No previously developed privacy structure is suitable to our setting; however, we draw on many of these in order to achieve the appropriate privacy structure. Another contribution of ours is therefore a suitable model for privacy, matching reasonable needs and threats as well as the limitations posed by our infrastructure and trust model. In particular, we consider privacy against both *peers* and *authorities*—the latter correspond to either *front-end authorities* (typically base stations) that handle the bulk of the interaction with the nodes, and *back-end authorities* (the privacy representatives, or ombudsmen—these remain off-line for most of the time). The back-end authorities have an established trust relationship with nodes, while the front-end authorities do not. In more detail, we model privacy needs as follows:

Privacy against peers is achieved by use of short-lived pseudonyms, making correlation to reoccurring events and known identities difficult. It is worth noting here that peers will typically have a very *limited view* of the communicated information (as a function of both time and space), and thus, rather long pseudonym life spans (in the order of several hours) are reasonable from this perspective. On the other hand, although peers are considered limited in terms of their view of information, it is reasonable to model them as *malicious.*

Privacy against authorities is more complex. First of all, we believe that it is reasonable to model authorities as largely honest participants that occasionally may be corrupted by an adversary with full read access to their stored data. This corresponds to an insider attack or a malware attack with access to the records of either front-end or back-end authorities, but not both. This is a common assumption in the literature, and appears reasonable in the context of many known forms of corruption. Second—recognizing that the front-end authorities will have access to a large amount of data—we avoid storage of any identifying information with them. This information, instead, is stored with the back-end authorities, which only have access to traffic data in special cases in which tracing is necessary (and then still only in a “cleansed” form). In order to achieve a high degree of computational efficiency, we let front-end authorities be able to derive sequences of (to them linkable) pseudonyms from some piece of information obtained from the back-end authorities, which in turn helps them to handle hand-over to other front-end authorities in an autonomous fashion. This latter piece of information can be seen as a long-lived pseudonym; to avoid notational confusion, however, we will refer to it as a *handle.* The use of long-lived pseudonyms will allow the base stations to ensure normal communication and service of the operation, including handover and various audits and tracings. The use of long-lived pseudonyms allows the front-end authorities to link information to the extent that is necessary for efficient operation including data mining to support services that do not require the knowledge of the real identity of certain vehicles. This, for example, includes services such as navigation or infotainment. To the extent that these are pay-services (whether on a subscription basis or a per-use basis), the users can be anonymously billed, similar to how we will describe in the context of anonymous toll collection.

1.3 Communication

We consider two types of communication: inbetween nodes (where most nodes correspond to vehicles), and between a node and a base station. Both of these are one-hop wireless links in our protocol, with the exception of communication between some fixpoints and the base stations; this communication may be wired instead of wireless, and may therefore be a traditional multi-hop communication. (Either can be replaced by multi-hop communication, but that is an effort that is largely orthogonal to the emphasis of this paper.)

Communication between two nodes relates to information about the sender (such as “airbag deployment”, or “rapid braking”), while communication between a node and

[1] We avoid the use of primitives like TESLA [22], which rely on time-synchronization and delayed verification; such approaches are not suitable in environments with high mobility due to the computational demands associated with verifying the authenticity of messages from nodes that a given verifier has not recently interacted with.

a base station instead mostly[2] relates to the *context* of the sender. The contextual information consists of (potentially processed) messages from other nodes, and can be used for purposes of location and audit.

The incentives can be seen as aligned along the lines of these two types of communication as well. A node is encouraged to volunteer information about itself using one class of incentives (including emergency response) while contextual communication is associated with another incentive structure (exemplified by proving innocence in traffic accidents). Some types of services, such as location based services and response to car theft, relate to either type—depending on the design of the particular type of service.

1.4 Achievements

We propose an architecture addressing the needs of collecting data for purposes of increased security for drivers and passengers, and detail a corresponding solution that minimizes the bottleneck associated with data collection, audit, and occasional de-anonymization. The latter can only be performed by an active collaboration between local base stations and a protocol participant we may think of as a consumer representative, but which in practice may be implemented by a service provider to which the consumer has a trust relationship[3]. Our solution minimizes the required communication, and makes use of probabilistic collection techniques that we propose; these strive to ensure the rapid propagation of information likely to be useful for system audits and emergency response, and heuristically avoids the propagation of information of lesser value. We base our proposed solution on symmetric-crypto building blocks.

2. RELATED WORK

Vehicular Networks. Due to their enormous potential, vehicular networks have gained an increasing attention in both industry and academia. Research activities range from lower layer protocol design to applications and implementational issues. In the U.S., initiatives have received further support by means of the FCC dedicating 75MHz of spectrum of the 5.9 GHz band for Dedicated Short-Range Communication (implementational) [10].

While the need for security and privacy in vehicular networks was recognized early on (e.g., [19]), it was not until very recently that the numerous challenges have started to spur some increased research interest.

Zarki et al. [25] describe an infrastructure for driver assistance. In particular, they focus on immediate vehicle vicinity awareness and highway traffic conditions. In this context they discuss security requirements for the system which are met by introducing digital signatures and a public key infrastructure. Duri et al. [9] focus on assuring privacy and integrity of data in telematics applications. They present a comprehensive Data Protection Framework which integrates security components based on standard protocols like SSL or IPSec. In [5] Blum and Eskandarian propose the SecCar architecture which assumes a PKI infrastructure as well as a virtual network infrastructure to provide scalability and security for vehicular networks. The work by Golle, Greene and Staddon [11] focuses on detecting and correcting malicious data. In [14], Hubaux et al. discuss a number of security problems and attacks including DoS or the disabling or impersonation in case of electronic number plates. In [23, 24], Hubaux and Raya provide a comprehensive discussion of security and privacy issues in vehicular networks. They propose solutions based on the use of digital signatures and anonymous public keys.

The approach in this paper differs from previous work in that it is not based on public key cryptography. Instead, it uses a symmetric lightweight approach to minimize the computational overhead. This is combined with incentive mechanisms to address the problem of cheating, and to avoid having drivers disable the transponders of their cars.

Reputation. Reputation-based systems are generally used to foster cooperation or build trust in certain environments. Often, these environments lack a centralized infrastructure and thus reputation-based methods are used as alternative means. Different contexts range from e-marketing and trading (e.g., [1, 15]) to networking (e.g., [3, 4, 7, 21]).

Anonymous Routing. The solution proposed in this paper is a special instance of anonymous routing. There is a wealth of work in the context of routing for different types of networks. Anonymous routing protocols strive for providing anonymity and unlinkability of nodes on a particular route. For example, while onion routing is based on public key methods [12, 18], the methods in [16] assumes an operator to share a secret key with each individual node in the system.

Auditing and Escrowing. In our solution we will use escrowing mechanisms to enforce that a participating nodes' identity will generally remain anonymous. The solution has similarities to escrow-like payment schemes and to the use of escrow in encrypted communication. Examples include work on magic ink signatures [17] and on key escrow for encryption purposes [20].

3. MODEL

3.1 Network

In this paper we consider vehicular networks which are hybrid in nature—i.e., networks with both an infrastructure-based and an ad hoc network component. Because of high mobility nature of vehicles, one might argue that the vehicular network should be implemented using a pure ad hoc architecture, allowing groups of nodes to exchange information without the need of (or potential privacy threat of) an authority or infrastructure. However, there are clear benefits associated with a hybrid solution, where a backbone service provider can provide various services. These will in turn act as incentives for collaboration among users, and for not disconnecting from the network. We do not consider "partial disconnection" of devices, in which users have some functionality disabled, but not all. Our motivation for this is

[2] Some of the information sent from a node to a base station also contains data relating to the sending node as opposed to its context.

[3] We will refer to this party simply as the *ombudsman*, and will for denotational simplicity assume that there is only one. Our solution can straightforwardly be augmented to allow for *multiple* and independent ombudsmen, at which time the competition between these will also serve to strengthen their trustworthiness.

the increased sophistication required for such an attack. We distinguish between two main categories of communication in our hybrid network:

Network Communication: Both node-to-base station (i.e., up-link) communication and base station-to-node (or down-link) communication rely on an infrastructure established and maintained by the service provider. communicated between the vehicle and the service provider at the time of occurrence. It furthermore assumes that a service provider and the vehicle support sufficient means to allow for communication with each other whenever deemed necessary. transmitted to the other party at a later time than it was collected. Reasons requiring the support of offline communication are manifold and range from choice of design and cost of the service to necessity.

Peer-to-Peer Communication: This communication mode supports data exchange between nodes, without relying on a pre-defined infrastructure. This type of communication can be used to improve on infrastructure-based communication, e.g., bridging gaps in the service infrastructure or providing alternative communication channels in case of an overloaded infrastructure.

The data exchanged in vehicular networks can range from information pertaining to the vehicle itself to timing information and information observed in close proximity or surroundings. The richness of data is meaningful to transmit in cases of emergency, as it may allow for faster dispatch of appropriate emergency services, or get precisely personalized service such as auditing.

We only require loose time synchronization (by the authorities) for the interpretation of collected data; these will infer the local times of nodes originating and forwarding data. For practical purposes, authorities may also push clock corrections to nodes it communicates with.

3.2 Adversarial Model

We consider an adversary that for any point in time can corrupt *either* front-end or back-end infrastructure nodes, but not both. We assume that the adversary has full read-access to the contents of memory of a corrupted node, but no write access. (This corresponds to a typical malware attack or insider attack on a system protected by intrusion detection capabilities based on the monitoring of changes.) Furthermore, we assume that messages between infrastructure nodes are authenticated using digital signatures, where keying material is proactively [13] modified in order to maintain security against permanent corruption. Furthermore, we assume that the adversary may corrupt vehicular nodes, in order to infer information about neighboring nodes. This is assumed to be a static corruption, i.e., not change over time. A corrupted vehicular node may either be *disconnected* (not able to communicate) or *snooping* (reporting all observed network traffic to the adversary.)

3.3 Authentication

In many cases, the computational limitations and associated battery limitations of mobile devices discourage the use of public key cryptography for authentication. For example, public key cryptography introduces potential DoS vulnerabilities if used in ad-hoc applications within cellular telephony, and all but makes operation impossible in typical sensor networks. For vehicular networks, clearly, this is not a concern: one may assume sufficient power and fast processors, largely on par with those on personal computers. However, symmetric key authentication remains more suitable, largely due to the necessary burstiness of communication: the window of time suitable for communication between two rapidly approaching vehicles is very short; this severely restricts the available bandwidth in the peer-to-peer communication mode. Moreover, the main advantage of publicly verifiable authentication would probably be that peers could remove invalid messages, instead of propagating these. This assumes the constant availability of a Certificate Revocation List, which may be an unreasonable assumption to make, thus calling into question the benefits of relying on a PKI in the first place.

For light-weight use of MAC authentication, we propose the use of short-lived pseudonyms instead of a static identity. Any static value (or pseudonym) assigned to the nodes can be considered as another form of identity, just as social security numbers (by themselves meaningless and contentless) can uniquely identify the people they are associated with. Similarly, the use of deterministic (or static) encryption techniques for hiding of identities is not possible, while the use of probabilistic encryption techniques is not suitable due to the resulting increase in computational cost in the context of audits and other network functions.

3.4 Audit and Incentives

The amount of information that is potentially available in vehicular networks is enormous and is expected to even further increase in the future. However, with the potential of a large-scale system which allows for auditing also come a number of problems—the main one being privacy which generally may be perceived as negative ramification of the system. In order to gain wide-spread user acceptance of vehicular networks it is therefore important to strike a delicate balance between auditing and privacy.

From this point of view, our proposed scheme can provide auditability as well as privacy by use of anonymized but traceable data collection. To this end, we distinguish between front-end and back-end authorities. Front-end authorities, which are typically base stations, handle the communication with the vehicles and collect data tagged with short-term pseudonyms. From the data collection, front-end authorities can know only the so-called *handles* of vehicles which are long-lived (but still changing) pseudonyms. Back-end authorities, on the other hand, do not process detailed communication transcripts, but only handle already aggregated data (such as billing information), along with identifying information.

To ensure compliance with protocols, application-level incentives can be provided to users of the system. We describe some such incentives herein to illustrate this aspect of the design, but note that this is a far from complete description.

Possible incentives can be found in the following areas:

- **Improved Navigation Guidance:** Traditional navigation systems which are widely available to date provide navigation assistance solely based on the current position of the vehicle and the intended destination. In a system like the one we propose, it is possible to personalize guidance based on personal preferences, driving history, or current location and destination—without interfering with privacy issues.

- **Roadside and Emergency Assistance:** These are services that are already available in systems such as OnStar by GM [2]. The driver can contact the operator on the push of a button to report problems pertaining to the car or alternatively, events such as the deployment of an airbag will automatically trigger an emergency call to the service provider. The benefit of the service can be further enhanced by allowing, for example, a driver's personal information to be reported or kept on file.
- **Retrieval of Stolen Vehicles:** Theft of vehicles will become more difficult as transponders need to be disabled in order to avoid reporting of location information. Furthermore, in an enhanced reporting and audit system (some) vehicles will be equipped with cameras. In situations where there is a suspicion that a transponder has been disabled, they can thus provide information of their surroundings which can be easily analyzed (such as, e.g., color, size, and make of other vehicles).
- **Hassle-free Operation of Vehicles:** This includes services such as automatic toll payments or automatic reminders of necessary maintenance checks for the vehicle.

We wish to emphasize that it is not the goal of this paper to exhaustively describe all types of incentives, but merely to exemplify these in order to support the versatility of our collection structure in the face of necessary privacy protection mechanisms.

4. IMPLEMENTATION

We start with the presentation of our model by introducing the different participants. All participants share the following system parameters:

System Parameters. There are two types of time-related indices: one (T) for long time periods, and the other (t) for shorter time periods. Without loss of generality, we may assume that the length of a long time period is an integer multiple of the length of a short interval, and for simplicity, that t indicates the index of a given short time interval within a given time interval[4], while T indicates the index of the longer time interval. Thus, neither relate to the actual *length* of the intervals. The shorter of the intervals may be a minute long, and the longer interval twenty-four hours. However, either one could be of a length that does not exactly correspond to a standard unit of time. At the end of each short (resp. long) time interval, the corresponding pseudonym (resp. handle) is updated. Another common parameter is λ, the bit-length of a pseudonym, and κ represents the bit-length of a long-lived pseudonym, which we call a handle.

Consider the following protocol participants:

Node ($\mathcal{N}_k$). A node can be a mobile vehicle or a static infrastructure node, such as a traffic light or traffic sign. Most of our discussion is focused on mobile vehicles, which collect information transmitted by other nodes and base stations, and transmit information to the same in turn. Each node, say $\mathcal{N}_k$ for a certain index k, has the following parameters:

- ID_k**:** A unique identification number.
- SD_k**:** A node specific seed value for pseudonym generation.
- LT_k**:** The node's local time (which may differ from other participants' local time).
- PS_k^i, KS_k^i**:** A node's i^{th} short term pseudonym and corresponding session key respectively. A node can have a pair of multiple pseudonyms and session keys which are changed at each small time interval t.
- HD_k^j**:** The j^{th} handle to provide a way to correlate an identity to/from a pseudonym. Each node has multiple handles over time, since the handle changes with T.
- $\mathsf{DB}_k^{\mathcal{N}}$**:** A repository to store received packets (which in turn may include several messages). Whenever deemed necessary, a node can forward saved packets to other participants such as roadside base stations or other nodes. The database $\mathsf{DB}_k^{\mathcal{N}}$ consists of four columns: $\underline{\mathsf{LT}_{rcv}}$ (receiver's local time), PS (sender's pseudonym), MSG (received messages), and LT_{snd} (sender's local time). Note that the indexed column is underlined. In the context of our paper, indexed columns describes sorted columns which are used for later database searching. (Note that a more detailed discussion of the database structure is beyond the scope of this paper.)

Ombudsman ($\mathcal{OM}$). The ombudsman escrows associations between identities and pseudonyms. The $\mathcal{OM}$ may collaborate with a base station to reveal identities for given pseudonyms only after the fulfillment of specific conditions such as, for example, agreement, law enforcement, or order of court. The following is associated with an ombudsman:

- $f_{\mathcal{OM}}$**:** A publicly available one-way function which is used to generate long-term pseudonyms (or handles) from identities.
- $\mathsf{DB}^{\mathcal{OM}}$**:** A repository for storing and searching for node identities. The database $\mathsf{DB}^{\mathcal{OM}}$ has four columns: $\underline{\mathsf{HD}}$ (handle value), $\underline{\mathsf{T}}$ (large time index), ID (identification number), and SD (seed value). Again, the indexed columns are underlined as before. Since the handle changes according to the time index T, for example, the same pair of ID_k and SD_k of the node $\mathcal{N}_k$ may appear with different handle values HD_k^j and T. The database operations will be detailed later.

Base Station ($\mathcal{BS}_l$). Base stations are the stationary roadside infrastructure. A base station can send and receive network packets within its (limited) range of radio power. Each base station, say $\mathcal{BS}_l$ for a certain index l, has the following parameters:

[4] We assume that the internal clock of each node is synchronized with the corresponding clock of the base station within the one "long" time interval. This is, for all practical purposes, a reasonable assumption.

- $\mathsf{PK}_l, \mathsf{SK}_l$**:** A public key and a private key respectively.

- $f_{\mathcal{BS}}$**:** A publicly available one-way function which is used to generate short-term pseudonyms from long-term pseudonyms (handles).

- $\mathsf{DB}_l^{\mathcal{BS}}$**:** A database for storing "hello" messages which are received from nodes during initialization. The database $\mathsf{DB}_l^{\mathcal{BS}}$ consists of four columns: $\underline{\mathsf{PS}}$ (pseudonym), $\underline{\mathsf{KS}}$ (session key), $\underline{\mathsf{t}}$ (small time index), and HD (handle value). Again, indexed columns are underlined as before. As described before, the same handle can appear in multiple rows. The details on the database operations will be discussed later.

- $\mathsf{DB}_l^{\mathcal{N}}$**:** A repository for storing received network packets from nodes. The structure and operation are the same as for the node's repository $\mathsf{DB}_k^{\mathcal{N}}$.

4.1 Building Blocks

In our implementation, the following building blocks are frequently used:

1. **MAC generation/verification (**MAC_{ks}**):** The MAC is generated by means of encryption with a session key ks for given payload p, denoted $\mathsf{MAC}_{ks}(p)$. Since ks is symmetric, the same key ks is used for MAC verification.

2. **Public key encryption (**Enc_{pk}**):** An encryption is performed using the receiver's public key pk for a given message m, denoted by $\mathsf{Enc}_{pk}(m)$. Decryption for the given encryption ϵ is done using the receiver's secret key sk such as $\mathsf{Dec}_{sk}(\epsilon)$.

5. PROTOCOLS

In the following, we present our abstract implementation, which is event-driven. For each occurring event, the corresponding protocols will be launched.

5.1 Key Registration

This event occurs as part of the bootstrapping of a node. Registration includes generation of the identification number, say ID_k, of the node $\mathcal{N}_k$, and escrowing of the ID_k by the ombudsman $\mathcal{OM}$. The protocol is as follows:

1. The identification number ID_k and the seed value SD_k are randomly selected by $\mathcal{OM}$ and shared with the node $\mathcal{N}_k$.

2. The $\mathcal{OM}$ computes a set of handles, i.e.,

$$\left\{\mathsf{HD}_k^0, \mathsf{HD}_k^1, \cdots, \mathsf{HD}_k^\delta\right\},$$

where δ is some rather small integer value such that

$$\mathsf{HD}_k^j = f_{\mathcal{OM}}(\mathsf{ID}_k, \mathsf{SD}_k, \mathsf{T}_j) \qquad (1)$$

where $\mathsf{T}_j = \mathsf{T} + j$ for each $j = 0, 1, \cdots, \delta$. Each handle of the set is saved as an entry in a private database $\mathsf{DB}^{\mathcal{OM}}$, indexed by HD_k^j and T_j. Thus, multiple rows of $\mathsf{DB}^{\mathcal{OM}}$ have the same ID_k and SD_k. That is, for $\mathsf{T}_j = \mathsf{T}, \mathsf{T}+1, \cdots, \mathsf{T}+\delta$ the ombudsman $\mathcal{OM}$ can trace ID_k and SD_k based on the $(\delta + 1)$ pre-computed handles.

5.2 Initialization

This event occurs as a bootstrapping process whenever a node starts operation or enters a new area, i.e., when switching to a different base station.

1. A node $\mathcal{N}_k$ sends a hello message to the nearest base station $\mathcal{BS}_l$, which responds with its certified public key PK_l, whose validity is verified by the node.

2. The node $\mathcal{N}_k$ computes the j^{th} handle HD_k^j such that

$$\mathsf{HD}_k^j = f_{\mathcal{OM}}(\mathsf{ID}_k, \mathsf{SD}_k, \mathsf{T}_j) \qquad (2)$$

where ID_k and SD_k are known through key registration.

3. The node encrypts the handle HD_k^j and the time t with PK_l such that $\mathsf{Enc}_{\mathsf{PK}_l}(\mathsf{HD}_k^j, \mathsf{t})$ and transmits the ciphertext to $\mathcal{BS}_l$.

4. In case of identity critical services, the base station $\mathcal{BS}_l$ can check the validity of the given HD_k^j immediately by asking the ombudsman $\mathcal{OM}$ for help.

5. The base station $\mathcal{BS}_l$ decrypts the ciphertext. Using the one-way function $f_{\mathcal{BS}}$, the base station then computes a series of values

$$\mathsf{O}_i = f_{\mathcal{BS}}(\mathsf{t}'_i, \mathsf{HD}_k^j) \qquad (3)$$

for each $\mathsf{t}'_i = \mathsf{t} + i (0 \leq i \leq \gamma)$. Like δ before, γ is a small, system dependent integer representing the expected number of pseudonyms expected to be used by the node within the area of coverage of the base station. Then, let PS_k^i be the leftmost λ bits of O_i and KS_k^i be the remainder of O_i.

6. Finally, the base station stores γ rows with tuples of PS_k^i and KS_k^i for each i into its private database $\mathsf{DB}_l^{\mathcal{BS}}$. Multiple rows for the sets of $\{\mathsf{PS}_k^i\}$, $\{\mathsf{KS}_k^i\}$, and $\{\mathsf{t}'_i\}$ have the same handle value HD_k^j. That is, as in the key registration protocol, the base station can trace the handle HD_k^j with one of the multiple values PS_k^i and KS_k^i.

Note that λ may be small enough for collisions to be possible. Such collisions can later be resolved by means of verifying the authentication string of a given message with respect to the different candidate keys.

5.3 Handover

The pseudonyms are generated from the handle of the associated time period, along with the index of the shorter time interval. This is done on both sides—node and base station. When a node moves out of range of a base station, the base station performs a handover to the neighboring base station about to take over the communication with the node. The handover information contains state information about the node, including the current handle. From this, pseudonyms are generated by the new base station; the node can remain largely unaware of the transition.

5.4 Pseudonym Generation

For communication with the base station, each packet (except for "hello" messages) is tagged with the node's pseudonyms. Since the pseudonym will be changed at each time interval t, a node denoted by $\mathcal{N}_k$ should be able to generate a new i^{th} pseudonym as follows:

1. The node $\mathcal{N}_k$ computes the current time indices T and t, denoted by T_j and t_i respectively.

2. Then, the node computes O_i as:

$$\mathsf{O}_i = f_{\mathcal{BS}}(\mathsf{t}_i, f_{\mathcal{OM}}(\mathsf{ID}_k, \mathsf{SD}_k, \mathsf{T}_j)) \tag{4}$$

3. The i^{th} pseudonym PS_k^i is the leftmost λ bits of O_i and the corresponding i^{th} session key KS_k^i is the remaining bits of O_i. Now, the node $\mathcal{N}_k$ can use the new session key KS_k^i to generate a MAC as $\mathsf{MAC}_{\mathsf{KS}_k^i}$.

5.5 Pseudonym Lookup

Since a pseudonym is coupled with a particular session key, a base station, say $\mathcal{BS}_l$, needs to find the session key KS_k^i corresponding to the pseudonym PS_k^i to allow for MAC verification or generation. Failure of the pseudonym lookup process results in dropping the packet.

Verification of MACs. When $\mathcal{BS}_l$ receives a network packet tagged by a pseudonym PS', it queries its database $\mathsf{DB}_l^{\mathcal{BS}}$ to find the corresponding session key KS' to verify the correctness of the MAC of the packet.

Generation of MACs. Packets sent from the base station to a node are MACed using the shared key KS', as described before. In many cases, KS' will already be stored in the database $\mathsf{DB}_l^{\mathcal{BS}}$. If the information is not available to the base station, the base station can compute it from the handle which corresponds to the time (as perceived by the recipient).

5.6 Peer-to-Peer Communication

This protocol is used for node-to-node communication. Let us assume that node $\mathcal{N}_s$ is sending a set of n messages $\{m_c\}$ $(1 \leq c \leq n)$ to node $\mathcal{N}_r$ as part of one network packet. The protocol is defined as follows:

1. If necessary (that is if the small time interval is expired), node $\mathcal{N}_s$ generates the new pseudonym PS_s^i and its corresponding session key KS_s^i(using the pseudonym generation protocol defined before). Otherwise, it will use the current pseudonym PS_s^i and session key KS_s^i.

2. The node $\mathcal{N}_s$ computes MAC_i by using the session key KS_s^i as

$$\mathsf{MAC}_i = \mathsf{MAC}_{\mathsf{KS}_s^i}(\mathsf{PS}_s^i \,||\, \mathsf{LT}_s^i \,||\, \{m_c\}) \tag{5}$$

where LT_s^i is the local time of the sender and the operation $||$ denotes concatenation.

3. $\mathcal{N}_s$ sends the network packet P_i with

$$\mathsf{P}_i = \mathsf{PS}_s^i \,||\, \mathsf{LT}_s^i \,||\, \{m_c\} \,||\, \mathsf{MAC}_i. \tag{6}$$

4. The receiving node $\mathcal{N}_r$ stores the received messages $\{m_c\}$ in its private database $\mathsf{DB}_r^{\mathcal{N}}$, i.e., a row is inserted for each received packet. Thus, the columns—PS, MSG, LT_{snd}—are filled with PS_s^i, $\{m_c\}$, LT_s^i respectively and the LT_{rcv} column with the receiver's local time LT_r. (Whenever deemed necessary, the node will send the saved data to a certain nearby base station using the up-link communication protocol described below.)

Note that the receiver $\mathcal{N}_r$ does not verify MAC_i of packet P_i. This is only verified by the base station.

5.7 Network Communication

This type of communication occurs between a node and a base station. In our proposal, we distinguish the up-link, in which a node sends a network packet to the base station, from the down-link, in which the base station sends a network packet to a node. Usually, the up-link is used for requesting or reporting something to the base station, while the down-link is used to provide information or some sort of service.

Up-link. A node sends a network packet with the short-term pseudonym to hide its identity. Then, the receiving base station queries its database to find the session key corresponding to the sender's pseudonym. At that moment the base station only knows pseudonyms—not the identity—of the sender. (We note that short-term pseudonyms can be linked to each other only by somebody who knows the corresponding handle; handles can not be linked by nodes other than the ombudsman, given that the one-way function used to derive these is keyed.)

Assuming that a node $\mathcal{N}_k$ reports a set of messages $\{m_i\}$ to a base station $\mathcal{BS}_l$, we can present the protocol as follows:

1. The node $\mathcal{N}_k$ follows the first step in the peer-to-peer communication protocol described previously.

2. Looking at PS_k^i included in the received packet P_i, the receiver $\mathcal{BS}_l$ runs the pseudonym lookup protocol as described before to find $\mathcal{N}_k$'s corresponding session key KS_k^i in its own database $\mathsf{DB}_l^{\mathcal{BS}}$.

3. The base station $\mathcal{BS}_l$ verifies the authenticity of packet P_k^i by verifying the MAC. If it is not, the packet is dropped; otherwise, the next step is executed.

4. Finally, the base station stores the received packet P_i in its database $\mathsf{DB}_l^{\mathcal{N}}$ which is the same as the database in the nodes; thus, operation is the same as in Step 4 in the peer-to-peer communication protocol described previously.

Down-link. When sending a packet to a node, the base station generates the message authentication code using the same key as is used as the receiver's session key. Since the session key is changed periodically in accordance with the pseudonym, the base station queries its database $\mathsf{DB}^{\mathcal{BS}}$ to find the receiver's pseudonym and session key for packet generation. Let us assume that base station $\mathcal{BS}_l$ sends messages $\{m_c\}$ to the node $\mathcal{N}_k$. Then, the protocol is as follows:

1. The base station $\mathcal{BS}_l$ queries its database $\mathsf{DB}_l^{\mathcal{BS}}$ to find the right pseudonym and corresponding session key, say PS_k^i and KS_k^i, of the node by means of the pseudonym lookup protocol for MAC generation.

2. Then, $\mathcal{BS}_l$ computes the message authentication code MAC_i by using the session key KS_i as

$$\mathsf{MAC}_i = \mathsf{MAC}_{\mathsf{KS}_k^i}(\mathsf{PS}_k^i \,||\, \mathsf{LT}_l \,||\, \{m_c\}) \tag{7}$$

where LT is the local time of the base station.

3. $\mathcal{BS}_l$ sends the packet P_i constructed as

$$\mathsf{P}_i = \mathsf{PS}_k^i \,||\, \mathsf{LT}_l \,||\, \{m_c\} \,||\, \mathsf{MAC}_i. \quad (8)$$

4. The receiver $\mathcal{N}_k$ may store the received messages $\{m_c\}$ in its private database $\mathsf{DB}_k^{\mathcal{N}}$ for future purposes. The operation is the same as Step 4 in the peer-to-peer communication protocol described previously.

5.8 Auditing

Sometimes, it may be necessary to determine the *identity* of a node at a given place and time, as opposed to simply its pseudonym. Alternatively, it may be necessary to determine the location at a certain time of a node of a given identity. Yet other alternatives exist, such as determining whether two vehicles (indexed by their identities) were within the same geographical area within a given interval of time. Each of these tracing operations require collaboration between base stations and the ombudsman. We describe some of these auditing mechanisms herein:

Pseudonym Auditing. Let us assume a base station $\mathcal{BS}_l$ wants to send information to vehicles which were observed in some specific area or conditions; for example, cars passed some toll gates an hour ago or cars parked in the roadside for a long time. Only need for the base station is the set of pseudonyms $\{\mathsf{PS}'_n\}$. The operation is as follows:

1. The base station $\mathcal{BS}_j$ queries its database $\mathsf{DB}_j^{\mathcal{N}}$ with various conditions such as the message column MSG includes the specific location, time, or car's conditions. The database query returns the corresponding rows. These contain a set of related pseudonyms $\{\mathsf{PS}_n\}$.

2. For each output PS_n, the base station $\mathcal{BS}_j$ will use the down-link protocol.

By design, the identity auditing to link pseudonyms to identities is only possible by involvement of the ombudsman which the drivers trusted for escrowing their identity.

Basic Identity Auditing. Let us assume a base station $\mathcal{BS}_l$ intends to trace a pseudonym PS'_x. The basic operation for auditing is as follows:

1. The base station $\mathcal{BS}_l$ queries its database $\mathsf{DB}_l^{\mathcal{BS}}$ to find the handle HD'_x corresponding to the given pseudonym PS'_x. The query returns one or a number of rows for PS matching PS'_x. Since each row contains handle column HD, the base station sends handle HD'_x to the ombudsman $\mathcal{OM}$.

2. For the received handle HD'_x, the ombudsman queries its database $\mathsf{DB}^{\mathcal{OM}}$; i.e., for the given handle HD'_x, the query returns one or a number of rows with HD matching the given HD'_x. Then, the ombudsman gets the row containing the identity ID'_x of the handle HD'_x.

Application of Identity Auditing. Let us assume a base station needs to find some node at a certain time or at a certain place. Then, the base station $\mathcal{BS}_l$ follows these steps:

1. The base station $\mathcal{BS}_j$ queries its database $\mathsf{DB}_j^{\mathcal{N}}$ according to a number of conditions. That is, the query returns the rows for which LT matches the given time or the column MSG includes some specific information. Then, the output rows contain a set of related pseudonyms $\{PS_y\}$.

2. For each PS_y of outputs, the base station $\mathcal{BS}_j$ can apply the identity auditing described above.

An Example—Anonymous Toll Payments. Now we introduce a simple example of building *Anonymous Toll Payments* system to demonstrate the benefit of vehicular networks and how our protocols can be combined. Assume there is a toll gate with a base station $\mathcal{BS}_t$ in close proximity.

1. Whenever passing the toll gate, the vehicles send a packet including their short-term pseudonyms to the base station $\mathcal{BS}_t$.

2. The base station $\mathcal{BS}_t$ stores only packets with valid pseudonyms verified by the ombudsman $\mathcal{OM}$ into its database $\mathsf{DB}_t^{\mathcal{N}}$. The sender of invalid packets, for example, including fake handles, will be detected and punished later on—much like it is currently done in cases where road tolls are avoided.

3. At the end of the billing cycle (e.g., every week or month), the base station $\mathcal{BS}_t$ makes a list of saved pseudonyms and their matching handles, $\{\mathsf{PS}_z, \mathsf{HD}_z\}$, as described in the section on basic identity auditing (see Section 5.8). The base station sends the set of $\{\mathsf{HD}_z\}$ to the ombudsman.

4. The ombudsman $\mathcal{OM}$ finds $\{\mathsf{ID}_z\}$ from the given handles $\{\mathsf{HD}_z\}$ and produces the billing information.

Note that the base station $\mathcal{BS}_t$ does not know any of the identities. That is, toll fees are accounted for anonymously.

6. SECURITY ANALYSIS

In the following we will briefly describe how the new protocol provides for privacy and authentication at a lower computational cost than other exiting solutions. A more complete analysis can be found in the full version of the paper.

Privacy. The privacy in *peer-to-peer* communication is achieved by using short-lived pseudonyms (which are continually changing at every short time interval t) instead of real identities of the nodes. The pseudonyms can not be predicted from the previous values due to the use of the one-way function $f_{\mathcal{BS}}$. Assuming collision resistance of the one-way function $f_{\mathcal{BS}}$ will render a pseudonym forgery attack impossible.

The same applies to the *up-link* and *down-link* communications. The nodes' identities are obscured by the periodically changing pseudonyms. Meanwhile, the handles which are provided to the base stations to trace identities are also periodically changing pseudonyms at every time interval T and untraceable due to the one-wayness of the function $f_{\mathcal{OM}}$. Thus, the handles, like short-lived pseudonyms, obscure the nodes' identities.

6.1 Authentication

To reduce computational overhead for the nodes, our proposal does not make use of authentication in *peer-to-peer* communication.[5] MAC verification is performed by the recipient (base station resp. node) in both the *up-link* and the *down-link* communication protocols. For that, the shared session keys are computed independently by the nodes as well as the base station, using the pre-shared handles. Similarly as in pseudonym generation, the security of the session key also depends on the collision freeness of the one-way function $f_{\mathcal{BS}}$.

7. CONCLUSION AND FUTURE WORK

We study the feasibility of using symmetric key constructions to realize secure VANETs, concluding that it measures up well to asymmetric designs. We argue that insufficient time has been spent evaluating the suitability of symmetric designs, which have several notable advantages over asymmetric designs—most notably a lesser reliance on availability of bandwidth. In particular, we suggest that asymmetric cryptographic models may not provide functional benefits to match their deployment costs, especially so in networks where the constant availability of certification authorities and certification revocation lists can not be taken for granted. To this end, we have proposed a new model which provides a high degree of efficiency coupled with auditability and privacy. This is achieved by combining symmetric authentication with the use of short-lived pseudonyms.

The development of security primitives for VANETs is an area that has not received a lot of attention to date. We believe there is a great need for a careful modeling of likely threats, and the development of matching security mechanisms. The latter may to some extent be of a provably correct nature (as is common in the field of cryptography), or be based on heuristics to detect fraudulent or malicious behavior; use of heuristics are common approaches within banking and electronic warfare. At the same time, we believe it may be worthwhile to consider the potential threat associated with an increased reliance on wireless communication for the smooth flow of traffic; for example, it may be important to study the potential impacts of DoS attacks on any VANET system; this aspect emphasizes the importance of using light-weight cryptographic constructions.

8. REFERENCES

[1] http://www.ebay.com

[2] http://www.onstar.com

[3] N. Ben Salem, J.-P. Hubaux, and M. Jakobsson. Reputation-based WiFi Deployment Protocols and Security Analysis. In *Proceedings of the 2nd ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots*, 2004.

[4] S. Buchegger and J.-Y. Boudec. Performance Analysis of the CONFIDANT Protocol: Cooperation of Nodes - Fairness in Distributed Ad Hoc Networks. In *Proceedings of MobiHOC*, 2002.

[5] J. Blum and A. Eskandarian. The Threat of Intelligent Collisions. In *IT Professional*, 6(1):24-29, Jan.-Feb. 2004.

[6] D. Chaum. *Untraceable electronic mail, return addresses, and digital pseudonyms.* Communications of the ACM, volume 24, pp. 84–88, Feb. 1981.

[7] Z. Desptovic and K. Aberer. Trust and Reputation in P2P Networks. In *1st Interdisciplinary Symposium on Online Reputation Mechanisms*, 2003.

[8] G. Danezis, R. Dingledine, and N. Mathewson. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, pp. 2–15, May 2003.

[9] S. Duri, M. Gruteser, X. Liu, P. Moskowitz, R. Perez, M. Singh, and J. Tang. Framework for Security and Privacy in Automotive Telematics. In *Proceedings of the 2nd International Workshop on Mobile Commerce*, Atlanta, Georgia, USA, pp. 25–32, 2002.

[10] Federal Communications Commission. FCC 99-305. FCC Report and Order, October 1999.

[11] P. Golle, D. Greene, and J. Staddon. Detecting and Correcting Malicious Data in VANETs. In *Proceedings of the First ACM Workshop on Vehicular Ad Hoc Networks*, pp. 29–37, 2004.

[12] D. Goldschlag, M. Reed, and P. Syverson. Onion routing. *Commun. ACM*, 42(2):39–41, ACM press, New York, February 1999.

[13] A. Herzberg, M. Jakobsson, S. Jarecki, H. Krawczyk and M. Yung, Proactive Public Key and Signature Systems. In *ACM Conference on Computer and Communications Security*, pp. 100–110, 1997.

[14] J.-P. Hubaux, S. Capkun, and J. Luo. The Security and Privacy of Smart Vehicles. In *IEEE SECURITY & PRIVACY*, Vol. 2, No. 3, pp. 49–55, 2004.

[15] D. Houser and J. Wooders. Reputation in Auctions: Theory and Evidence from eBay. *Working Paper 00-01, University of Arizona*, 2001.

[16] M. Jakobsson, S. Capkun, and J. P. Hubaux. Secure and Privacy-Preserving Communication in Hybrid Ad hoc Networks. *Technical Report IC/2004/10, EPFL-DI-ICA*, January 2004.

[17] M. Jakobsson and M. Yung. Distributed Magic Ink DSS Signatures. In *Proceeding of Eurocrypt '97*, 1997.

[18] J. Kong and X. Hong. Andor: Anonymous on Demand Routing with Untraceable Routes for Mobile Ad-Hoc Networks. In *Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'03)*, pp. 291–302, 2003.

[19] P. Karger and F. Yair. Security and Privacy Threats to ITS. In *The Second World Congress on Intelligent Transport Systems*, pp. 2452–2458, November 1995.

[20] S. Micali. Fair Cryptosystems. In *Proceedings of Crypto*, 1992.

[21] P. Michiardi and R. Molva. Core: A Collaborative Reputation Mechanism to Enforce Node cooperation in Mobile Ad Hoc Networks. In *Proceedings of the 6th IFIP Communications and Multimedia Security Conference*, 2002.

[22] A. Perrig, R. Canetti, J. D. Tygar, and D. Song. The TESLA Broadcast Authentication Protocol. *RSA CryptoBytes*, 5:(2):2–13, 2002.

[23] M. Raya and J.-P. Hubaux. The Security of Vehicular Networks. In *EPFL Technical Report IC/2005/009*, 2005.

[24] M. Raya and J.-P. Hubaux. Security Aspects of Inter-Vehicle Communications. In *Swiss Transport Research Conference(STRC) 2005*, March 2005.

[25] M. Zarki, S. Mehrotra, G. Tsudik, and N. Venkatasubramanian. Security Issues in a Future Vehicular Network. In *European Wireless*, 2002.

[5] Instead of relying on authentication methods to ensure honest behavior, we let the various services provided by base stations incentivize nodes to participate honestly. This is so since, under our assumption, drivers/owners may disable transponders, but are not able to modify the behavior of the communication devices in a more granular manner; we further assume that the services provided will not make complete disconnection a favorable option. These services, or incentives, were briefly described before, and will be elaborated on in the full version of the paper.