

# Challenges in Securing Vehicular Networks

Bryan Parno  
Carnegie Mellon University

Adrian Perrig  
Carnegie Mellon University

## ABSTRACT

In the near future, most new vehicles will be equipped with short-range radios capable of communicating with other vehicles or with highway infrastructure at distances of at least one kilometer. The radios will allow new applications that will revolutionize the driving experience, providing everything from instant, localized traffic updates to warning signals when the car ahead abruptly brakes. While resembling traditional sensor and ad hoc networks in some respects, vehicular networks pose a number of unique challenges. For example, the information conveyed over a vehicular network may affect life-or-death decisions, making fail-safe security a necessity. However, providing strong security in vehicular networks raises important privacy concerns that must also be considered. To address these challenges, we propose a set of security primitives that can be used as the building blocks of secure applications. The deployment of vehicular networks is rapidly approaching, and their success and safety will depend on viable security solutions acceptable to consumers, manufacturers and governments.

## 1. INTRODUCTION

The addition of short-range radios to both vehicles and highway infrastructure has the potential to significantly enhance the driving experience, providing increased safety and convenience. From a safety perspective, a car that informs other drivers of its sudden deceleration can reduce a ten-car pile-up to a fender-bender, or even prevent the accident entirely. Shared location information could eliminate traditional blindspots and assist drivers during lane changes or merges, potentially preventing thousands of accidents – in one year, lane changes and merges were responsible for over 630,000 crashes in the United States alone [6].

A convergence of forces from both the public and private sectors indicates that we are likely to see the birth of vehicular networks in the very near future. In 1999, the U.S. Federal Communications Commission (FCC) allocated a block of spectrum in the 5.850 to 5.925 GHz band for applications primarily intended to enhance the safety and efficiency of our highway system [8]. The FCC will not hesitate to reclaim unused spectrum, so vehicle manufacturers face very real pressures to make use of this allocation. Indeed, Audi, BMW, DaimlerChrysler, Fiat, Renault and Volkswagen have united to create the Car2Car Communication Consortium, an organization dedicated to developing industry standards for emerging wireless

---

This research was supported in part by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and by an NDSEG Fellowship from the Department of Defense. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of ARO, Carnegie Mellon University, Department of Defense, or the U.S. Government or any of its agencies.

technologies [3]. The consortium expects to have completed prototypes participating in field trials by March 2006. Since much of the impetus for the development of vehicular-network applications comes from vehicle manufacturers, we need to develop security techniques that operate in a distributed, ad hoc fashion. This approach will speed deployment, since it allows the manufacturers to incorporate these techniques without deploying additional infrastructure or deviating from their core business model.

To ensure the safety of vehicular networks, we must include security considerations from the very beginning. The very utility of these applications makes them likely targets for all forms of attackers and since many applications will affect life-or-death decisions, illicit tampering can have devastating consequences.

In this paper, we will focus on two potential vehicular network applications in order to ground the discussion in a concrete setting. First, we will consider a traffic congestion detection application designed to alert drivers to potential traffic jams, providing increased convenience and efficiency. In this application, vehicles detect when the number of neighboring vehicles exceeds a certain limit (and/or their average speed drops below a given threshold), and then relay this count to vehicles approaching the location of the congestion. The information can be relayed by vehicles traveling in either direction with the goal of propagating the information so that drivers heading towards the congestion receive it with sufficient time to choose alternate routes. In a city, the propagation may be a matter of blocks, but on major highways, it may be several kilometers.

Our second application is a deceleration warning system. Some of the worst traffic accidents involve tens or even hundreds of vehicles rear-ending each other after a single accident at the front of the line abruptly halts traffic. In this application, if a vehicle reduces its speed significantly,<sup>1</sup> it will broadcast a warning, along with its position and current velocity, to other vehicles. Recipients will relay the message to vehicles further behind, and any vehicles directly behind the vehicle in question will alert their drivers using visual and/or aural cues. Thus, vehicles separated from the warning originator by more than one or two vehicles will receive a warning before they see the chain reaction of brake lights.

**Contributions:** In this paper, we (1) analyze the security challenges specific to vehicular networks; (2) introduce a set of primitives for secure applications; (3) discuss vehicular properties that can support secure systems; and (4) present two security techniques, entanglement and reanonymizers, that leverage unique vehicular properties.

---

<sup>1</sup>The exact definition of “significantly” will depend on road conditions, the vehicle’s current speed, and other vehicle-specific parameters.

## 2. VEHICULAR NETWORK CHALLENGES

Vehicular network challenges include technical problems like key distribution as well as more abstract difficulties, such as the need to appeal simultaneously to three very different markets.

**Authentication versus Privacy.** In a vehicular network, we would like to bind each driver to a single identity to prevent Sybil [7] or other spoofing attacks. For instance, in the congestion avoidance scheme, we would like to prevent one vehicle from claiming to be hundreds in order to create the illusion of a congested road. Strong authentication also provides valuable forensic evidence and allows us to use external mechanisms, such as traditional law enforcement, to deter or prevent attacks on vehicular networks.

However, drivers value their privacy and are unlikely to adopt systems that require them to abandon their anonymity. For example, if we try to prevent spoofing in a manner that reveals each vehicle's permanent identity, then we may violate drivers' privacy expectations. Balancing privacy concerns with security needs will require codifying legal, societal and practical considerations. Most countries have widely divergent laws concerning their citizens' right to privacy. Since most vehicle manufacturers operate in multinational markets, they will require security solutions that satisfy the most stringent privacy laws, or that can be customized to meet their legal obligations in each market. Authentication schemes must also weigh societal expectations of privacy against practical considerations. Most drivers would resent a system that allows others to track their movements, but from a practical perspective, vehicles today are only partially anonymous. Each vehicle has a publicly displayed license plate that uniquely identifies it (and identifies the owner of the car, given access to the appropriate records). Thus, individual drivers have already surrendered a portion of their privacy while driving. Ideally, a secure vehicular network would build on these existing compromises instead of encroaching any further upon a driver's right to privacy.

**Availability.** For many applications, vehicular networks will require real-time, or near real-time, responses as well as hard real-time guarantees. While some applications may tolerate some margin in their response times, they will all typically require faster responses than those expected in traditional sensor networks, or even ad hoc networks. However, attempts to meet real-time demands typically make applications vulnerable to Denial of Service (DoS) attacks. In the deceleration application, a delay of even seconds can render the message meaningless. The problem is further exacerbated by the unreliable communication layer, since one potential way to cope with unreliable transmission is to store partial messages in the hopes that a second transmission will complete the message.

Current plans for vehicular networks rely on the emerging standard for dedicated short-range communications (DSRC) [2], based on an extension to the IEEE 802.11 technology [1]. Yin et al. provide a detailed, low-level evaluation of the performance of a simulated DSRC network and find that while the current DSRC standard provides an acceptable latency, the reliability is still lacking [22]. According to their simulations, on average, only 50-60% of a vehicle's neighbors will receive a broadcast message. Since vehicles moving in opposite directions will remain within communications range for only a few seconds, opportunities to retry a broadcast will be limited. On a positive note, DSRC features a high data rate.

**Low Tolerance for Errors.** Many applications use protocols that rely on probabilistic schemes to provide security. However, given the life-or-death nature of many proposed vehicular applications, even a small probability of error will be unacceptable. In fact, since the U.S. Bureau of Transportation Statistics estimates that there are over 200 million cars in the U.S. [21], even if only 5%

of vehicles use an application that functions correctly 99.99999% of the time, the application is still more likely to fail on at least one vehicle than function correctly on all vehicles. Thus, to provide the level of guarantees necessary for these scenarios, applications will have to rely on deterministic schemes or probabilistic schemes with security parameters large enough to make the probability of failure infinitesimally small.

Furthermore, for many applications, security must focus on prevention of attacks, rather than detection and recovery. In an ad hoc network, it may suffice to detect an attack and alert the user, leaving recovery and clean-up to the humans. However, in many safety-related vehicular network applications, detection will be insufficient, since by the time the driver can react, the warning may be too late. Instead, security must focus on preventing attacks in the first place, which will require extensive foresight into the types of attacks likely to occur (see Section 4).

**Mobility.** Traditional sensor networks frequently assume a relatively static network, and even ad hoc networks typically assume limited mobility, often focusing on handheld PDAs and laptops carried by users. For vehicular networks, mobility is the norm, and it will be measured in miles, not meters, per hour. Also, the mobility patterns of vehicles on the same road will exhibit strong correlations. Each vehicle will have a constantly shifting set of neighbors, many of whom it has never interacted with before and is unlikely to interact with again. The transitory nature of interactions in a vehicular network will restrict the utility of reputation-based schemes. For example, rating other vehicles based on the reliability of their congestion reports is unlikely to prove useful, a specific driver is unlikely to receive multiple reports from the same vehicle. Furthermore, since two vehicles may only be within communication range for a matter of seconds, we cannot rely on protocols that require significant interaction between the sender and receiver.

**Key Distribution.** Key distribution is often a fundamental building block for security protocols. In vehicular networks, distribution poses several significant challenges. First, vehicles are manufactured by many different companies, so installing keys at the factory would require coordination and interoperability between manufacturers. If manufacturers are unable or unwilling to agree on standards for key distribution, then we could turn to government-based distribution. Unfortunately, in the U.S., most transportation regulation takes place at the state level, again complicating coordination. The federal government can impose standards, but doing so would require significant changes to the current infrastructure for vehicle registration, and thus is unlikely to occur in the near future. However, without a system for key distribution, applications like traffic congestion detection may be vulnerable to spoofing.

A potential approach for secure key distribution would be to empower the Department of Motor Vehicles (DMV) to take the role of a Certificate Authority (CA) and to certify each vehicle's public key. Unfortunately, this approach has many shortcomings. First, assuming the role of a CA is a challenging operation which is not in line with the DMV's current functionality. Extensive anecdotal evidence [16] suggests that even specialized CAs offer questionable security against dedicated attackers trying to obtain a certificate for another institution/entity. Second, vehicles from different states or different countries may not be able to authenticate each other unless vehicles trust all CAs, which reduces security. Finally, certificate-based key establishment has the danger of violating driver privacy, as the vehicle's identity is revealed during each key establishment.

**Incentives.** Successful deployment of vehicular networks will require incentives for vehicle manufacturers, consumers, and the government, and reconciling their often conflicting interests will prove challenging. For example, law-enforcement agencies would

quickly embrace a system in which speed-limit signs broadcast the mandated speed and vehicles automatically reported any violations. Obviously, consumers would reject such intrusive monitoring, giving vehicle manufacturers little incentive to include such a feature. Conversely, consumers might appreciate an application that provides an early warning of a police speed trap. Manufacturers might be willing to meet this demand, but law-enforcement is likely to object.

**Bootstrap.** Initially, only a small percentage of vehicles will be equipped with DSRC radios and little infrastructure will exist to support them. Thus, in developing applications for vehicular networks, we can only assume that a few other vehicles are able to receive our communications, and the applications must provide benefits even under these limited conditions (with increasing benefits as the number of DSRC-equipped vehicles increases).

### 3. ADVERSARIES

The nature and the resources of the adversary will largely determine the scope of the defenses needed to secure a vehicular network. A realistic assessment of the vehicular environment suggests the following classes of adversaries (in increasing order of threat severity):

**Greedy Drivers.** While we might hope that most drivers in the system could be trusted to follow the protocols specified by the application, some drivers will attempt to maximize their gains, regardless of the cost to the system. In our congestion avoidance system, a greedy driver might try to convince the neighboring vehicles that there is considerable congestion ahead, so that they will choose alternate routes and allow the greedy driver a clear path to his/her destination.

**Snoops.** This category of adversary encompasses everyone from a nosy next-door neighbor to a government agency attempting to profile drivers. A burglar might try to use a vehicular network to detect which garages (and hence which houses) are empty. Companies may want to track consumers' purchasing habits and use correlated data to alter prices and discounts. While data mining may be acceptable over aggregate data, it raises serious privacy concerns if one can extract identifying information for an individual.

**Pranksters.** Pranksters include bored teenagers probing for vulnerabilities and hackers seeking fame via their exploits. For example, a prankster targeting a collision-avoidance or platooning application might sit by the road and convince one vehicle to slow down while persuading the vehicle behind it to speed up. The need for real-time responses potentially leaves security mechanisms vulnerable to DoS attacks. A prankster could abuse this vulnerability to disable applications or prevent critical information from reaching another vehicle.

**Industrial Insiders.** Attacks by insiders are particularly insidious, and the extent to which vehicular networks are vulnerable will depend on other security design decisions. For example, if mechanics can update the software on a vehicle, they also have an opportunity to load malicious programs. If we allow vehicle manufacturers to distribute keys, then a single rogue employee at one manufacturer could create keys that would be accepted by all other vehicles.

**Malicious Attackers.** Malicious attackers deliberately attempt to cause harm via the applications available on the vehicular network. They may be individuals attempting to settle a score or terrorists attacking our infrastructure. In many cases, these attackers will have specific targets, and they will have access to more resources than the attackers described above. Terrorists might manipulate the deceleration warning system to create gridlock before detonating a bomb. Criminals might spoof the congestion avoid-

ance application to facilitate getaways. In general, while this class of attackers will hopefully be rarer than those outlined above, their combination of resources and directed malice makes them an important concern for any security system.

### 4. ATTACKS

While we obviously cannot anticipate every possible attack on vehicular networks, we can enumerate some of the more likely scenarios and ensure that applications are robust against this known set of potential attacks. While many of these attacks have appeared in other contexts, we list them here both for the sake of thoroughness and to examine the ramifications they may have in this new environment.

**Denial of Service (DoS).** If the attacker can overwhelm a vehicle's resources or jam the communication channel used by the vehicular network, then he can prevent critical information from arriving. Not only does this render the application useless, it could increase the danger to the driver if she has come to depend on the application's information. For instance, if a malicious adversary wants to create a massive pileup on the highway, he could provoke an accident and then use a DoS attack to prevent the appropriate deceleration warnings from reaching other drivers.

**Message Suppression Attacks.** In a more subtle attack, the adversary may use one or more vehicles to launch a suppression attack by selectively dropping packets from the network. A prankster might suppress congestion alerts before selecting an alternate route, thus consigning subsequent vehicles to wait in traffic.

**Fabrication Attacks.** An adversary can initiate a fabrication attack by broadcasting false information into the network. For example, a greedy driver might pose as an emergency vehicle to speed up his own trip. An attacker may also choose to fabricate her own information, including her identity, location, or other application-specific parameters. Defending against fabrication attacks in a vehicular network is particularly challenging, since the traditional remedy of using strong identities along with cryptographic authentication may conflict with the need to preserve the privacy of participants in the network.

**Alteration Attacks.** A particularly insidious attack in a vehicular network is to alter existing data. This includes deliberately delaying the transmission of information, replaying earlier transmissions or altering the individual entries within a transmission. For example, if the traffic congestion application requires a vehicle to collect "votes" from other vehicles at the site of the congestion, then an attacker might collect votes while traveling in normal traffic, but alter the locations and timestamps in the votes to make it appear that all of those vehicles were in the same place at the same time, deceitfully indicating a heavily congested highway. A malicious attacker might alter a message alerting vehicles to an obstacle ahead to persuade another vehicle that the lane is in fact clear. Clearly, applications on vehicular networks will need authentication of both the source of the data and the data itself.

### 5. PROPERTIES SUPPORTING SECURITY

While the previous discussion paints a grim picture of the security challenges facing vehicular networks, there are aspects of such networks that may aid the creation of secure applications.

**Regular Inspections.** In most U.S. states, all vehicles must pass inspection once a year. This yearly trip to the mechanic provides interesting possibilities for security maintenance in addition to the typical maintenance performed. For example, as part of the inspection, the mechanic might use SWATT [19] to verify the integrity of the software running on the vehicle's processor. This would also

be an opportune time to update/patch existing software, download new certificates, or receive the current list of revoked certificates. The inspection process affords the system an opportunity to return to a known, baseline state, and it serves as a firebreak against worm and virus infections. These inspections could also be included as part of the standard maintenance package anytime a driver brings the vehicle in for a tuneup or repair.

**Honest Majority.** Another advantage of vehicular networks is that the majority of drivers are likely to be honest. This will be reinforced by the fact that few people feel comfortable tinkering with their vehicles, so that most drivers will simply accept the default configurations. Assuming vehicles have adequate safeguards<sup>2</sup> against worms and viruses, the trustworthiness of drivers will translate into vehicles that correctly follow established protocols. Applications can take advantage of this property via polling and aggregation, with the expectation that correct responses will outweigh incorrect or malicious ones.


**Additional Input.** Unlike traditional sensor networks, each node in a vehicular network has a presumed intelligent operator available. Unlike ad hoc networks, the interaction with the operator must be minimal, since the application should not distract the driver. However, even minimal information from the driver may enable additional security techniques. For example, in the collision avoidance application, the vehicle can use the driver's reaction to a warning as input to a reinforcement learning algorithm. If the driver ignores the warning and no crash results, the vehicle may decide that the warning was erroneous and adjust its parameters accordingly. If the driver does brake, it supports the vehicle's decision to issue the warning. Information from the driver may also be supplemented with information from other sensors in the vehicle (e.g., proximity sensors).

**Central Registration.** Unlike ad hoc networks, all of the nodes in a vehicular network must register with a central authority, such as the state government (in the U.S.). These institutions have already developed extensive infrastructure for tracking and administering these registrations, and in the long term, we could potentially leverage this existing bureaucracy to strengthen the security of vehicular applications. However, given the vast nature of the existing infrastructure, any changes will take years of effort and require considerable funding. Thus, we cannot rely on this infrastructure to help bootstrap the system, but we can be prepared to take advantage of it when it comes online.

**Controlled Access.** Many portions of the transportation system already have access control mechanisms. For example, toll roads and many bridges have controlled entry and exit points, so it may be practical to equip these locations with the infrastructure necessary to distribute ephemeral identities and/or keys to allow vehicles to communicate while they remain on the controlled stretch of road. Since all vehicles must pass through the limited set of entrances, we can assume that any DSRC-equipped vehicle encountered in that segment of the highway will have received an appropriate key that will interoperate with everyone else's key. We discuss other aspects of key distribution in Section 6.

**Existing Enforcement Mechanisms.** Finally, we have the presence of existing law enforcement to aid the security of vehicular networks. In many cases, an adversary will need to be within physical proximity of the victim to launch an attack. If the victim can

<sup>2</sup>These safeguards will likely extend existing work on worm and virus detection and can leverage the primitives discussed in Section 6. Trusted Platform Modules (TPMs) [20] would also help secure vehicular software. At present, a TPM may constitute an excessive expense, but as production increases, they may become economically feasible.



**Figure 1: Secure Relative Localization** Vehicle B can use broadcasts from vehicles C, D and E to determine A's location.

unambiguously identify the attacker to local law enforcement, we can use existing legal remedies to provide disincentives for attacks. Also, since liability laws typically hold vehicle manufacturers to a far higher standard than software vendors, the manufacturers have a strong incentive to deploy secure systems that will resist attacks.

## 6. SECURITY PRIMITIVES

To secure vehicular networks, we propose the development of a set of security primitives, which can serve as the building blocks for secure applications. The ideal primitives should be specific to the vehicular setting, while remaining sufficiently general and composable to allow for reuse in a wide variety of protocols. Below, we describe some of the primitives that will be necessary for secure vehicular applications.

**Authenticated Localization of Message Origin.** Location will play a significant role in many vehicular applications, making it critical to determine that a message did indeed originate at a given location. For example, this primitive would prevent an attacker sitting on the side of the road from claiming to be a vehicle traveling on the highway. It would also prevent an adversary from using another communication medium (e.g., cellular) to replay a message heard in one location as though it had originated in a different location, thus preventing the types of attacks enabled by wormholes in sensor networks [9]. Such a primitive would assist in securing the congestion detection application, since the principal characteristic of congestion is the presence of multiple vehicles in a similar location.

One possible approach to authenticating message origin would be to deploy beacons capable of broadcasting their location, along with a timestamp and a signature. Public key signatures would likely require trusting a central authority but would simplify deployment. Vehicles could include the beacon's packet within their message to prove they were at the beacon's location at the time specified. However, this approach would require extensive deployment of additional infrastructure, and as we discussed in Section 1, ad hoc protocols that can operate between vehicles are more likely to be deployed. As an alternative, we could attempt to extend the protocols that have been proposed for secure localization in sensor networks [5, 12, 13, 14, 15] to this new setting. Unfortunately, these protocols focus on allowing a sensor to securely determine its own position (rather than the positions of its neighbors) or rely on the presence of multiple base stations. Instead, we propose to leverage the properties of the vehicular environment to provide a new method of secure relative localization.

In this scheme, a vehicle's relative location is defined by its *entanglement* with other vehicles. Each vehicle will regularly broad-

cast its identity (a public key) along with its signature of a current timestamp. When a vehicle receives such a broadcast, it signs the other vehicle's ID and rebroadcasts it. In other words, when vehicle  $A$  receives public key  $K_B$  from vehicle  $B$ , it adds a signature  $\{K_B\}_{K_A^{-1}}$  with its private key  $K_A^{-1}$  to its regular broadcast. When vehicles pass each other traveling in opposite directions, this will allow both streams of traffic to perform relative localization (see Figure 1). If vehicle  $B$  hears vehicle  $C$  rebroadcast  $A$ 's identity before it rebroadcasts  $B$ 's identity, then  $B$  can conclude that  $A$  is ahead of him/her. Vehicle  $B$  can aggregate multiple indicators (i.e., from vehicles  $D$  and  $E$ ) to provide further assurance of  $A$ 's position. Furthermore, vehicle  $B$  can evaluate the entanglement data for those vehicles as well to determine how much weight to give their reports. We describe this protocol not as a final solution, but rather to illustrate how we can use the properties of vehicular networks to our advantage when designing security services.

**Anonymization Service.** An anonymization service would allow us to resolve some of the tension between authentication and privacy. It would rely on the observation that for almost all of the applications we envision, a vehicle does not need to authenticate the exact identity of the other vehicle sending the information, but only the connection between the information sent and a vehicle present on the road. Drivers could use their permanent identity to authenticate to an anonymization service. The service would then provide the driver with a temporary identification that cannot be traced back to the driver (although this could be modified to allow authorized entities to trace the connection between the temporary id and the original driver, either through sealed records or some form of escrow service). This primitive could help prevent spoofed identities, while still preserving drivers' privacy expectations.

As a sample implementation, consider a toll highway. Since all drivers must pass through the toll booth to enter the highway, they could also authenticate themselves to an anonymization service hosted in the booth. The toll booth would provide the driver with a temporary ID (i.e., a public key pair  $\{K, K^{-1}\}$ ) that could be used for the duration of the trip. The system could optionally make use of *reanonymizers* positioned at regular intervals in stoplights or mile markers to further enhance driver privacy. A reanonymizer would provide a new identity certificate to any vehicle that can prove that it already possesses a temporary identity. Each certificate will be issued with a short lifetime, so that the certificate expires shortly after a vehicle passes the next reanonymizer. This prevents an adversary from accumulating anonymous identities for use in a Sybil attack.

When a vehicle approaches a reanonymizer, the reanonymizer would broadcast a random nonce,  $N$ . The vehicle would sign the random nonce using its secret key and broadcast the signature along with its public key. After verifying the signature, the reanonymizer would broadcast a new certificate encrypted with the vehicle's old public key. The certificate would contain the vehicle's new identity,  $\{K'_V, K'^{-1}_V\}$ , a timestamp ( $T$ ), and the reanonymizer's signature. The exchange between the vehicle ( $V$ ) and the reanonymizer ( $R$ ) is summarized below:

$$\begin{aligned} R \rightarrow V &: N \\ V \rightarrow R &: K_V, \{N\}_{K_V^{-1}} \\ R \rightarrow V &: \{K'_V, K'^{-1}_V, T, \{K'_V, T\}_{K_R^{-1}}\}_{K_V} \end{aligned}$$

Thus, every time a vehicle passes a reanonymizer, it can acquire a new identity independent of its previous identity. The anonymization service would help secure the congestion detection application, since it would provide each vehicle with a single strong identity without compromising the vehicle's anonymity.

As an alternative to a dynamic anonymization service, we could

consider preloading each vehicle with a year's worth of anonymous keys, as Raya and Hubaux propose [18]. The supply could be refreshed by the certifying authority. Unfortunately, this approach provides an attacker with a ready host of "legitimate" identities for use in a Sybil attack [7]. Tamper-resistant devices would make this attack more difficult but not impossible and would add to the expense of the vehicular network.

Instead, we could consider a dynamic key distribution system, in which vehicles could create a new anonymous key pair every day. Using a scheme similar to that suggested for creating anonymous keys for use with a TPM [20], each vehicle would have a vehicle identity public key pair  $\{K_V, K_V^{-1}\}$  along with a certificate  $C$  issued by the vehicle's manufacturer. To create an anonymous identity, the vehicle generates a new public key pair  $\{K, K^{-1}\}$  and sends a request for a new certificate for the public key  $K$  to a Certificate Authority (CA). The vehicle would sign the request with its identity key  $K_V$  and include the certificate  $C$  with the request. Assuming the CA trusts the vehicle's manufacturer, it can verify the signatures and issue a limited-lifetime certificate for  $K$  that is unlinkable (except by the CA) to the vehicle's actual identity. The drawbacks to this approach include the need for an online CA and the need for the vehicle to regularly communicate with the CA. In addition, the CA should not issue overlapping anonymous identities to the same vehicle (to prevent a Sybil attack), so creating a decentralized system may be challenging. Some drivers may also want additional anonymity beyond that granted by a daily key change, which places additional demands on the CA.

Another alternative for providing anonymity in vehicular networks is the use of group signatures. A group signature scheme allows one member of the group to sign a message such that other members of the group have the ability to verify that the message originated from a group member but not to identify the actual sender. Variants of group signature schemes allow linkability between two messages signed by the same group member, while still preserving sender anonymity [4]. Others allow unlinkability across larger time scales. An optional group manager may be endowed with the ability to link signatures with signers. This property may be desirable from the government's perspective, but it also raises a number of privacy issues.

**Secure Aggregation Techniques.** Applications running on vehicular networks can also benefit from secure aggregation primitives. These would allow, for example, one vehicle to count the number of vehicles it passes and report the sum to subsequent vehicles. After authenticating the count, these vehicles could use the data to estimate the amount of traffic ahead. This application lends itself to secure polling techniques, such as those developed by Kuhn [11], but other applications may require more general aggregation techniques, such as those by Przydatek et al. for sensor networks [17].

**Additional Primitives** Vehicular networks will also require primitives for performing key establishment and message authentication. In the context of sensor and ad hoc networks, key establishment continues to be an active research area. Unfortunately, these approaches do not readily translate to the vehicular setting, because the security requirements in vehicular networks are much more stringent than in either ad hoc or sensor networks, and the trust assumptions are different. Message authentication is another important primitive for vehicular networks. Secure message authentication prevents an external attacker from injecting malicious messages into the network, and can prevent a relaying node from altering the message. In conjunction with message freshness and appropriate mechanisms, authentication prevents an attacker from replaying an old message.

## 7. RELATED WORK

Few researchers have examined the problem of security in vehicular networks. Zarki et al. [23] present the DAHNI (Driver Ad Hoc Networking Infrastructure) system for providing driver assistance. They show how they can use a vehicular network to track nearby vehicles and report potential hazards to the driver. In contrast to their work, we argue that privacy and key establishment are two vital issues that require additional work before vehicular networks can be securely deployed. Hubaux et al. describe some of the attacks vehicular networks may face and propose a mechanism for providing secure positioning; they also suggest the congestion detection application discussed in this work [10]. In another work, Raya and Hubaux consider the issues involved with key management for vehicular networks, as well as the use of anonymous public keys. They also analyze the feasibility of using a PKI to support the security requirements of vehicular networks [18].

## 8. CONCLUSION

To make vehicular networks viable and acceptable to consumers, we need to establish secure protocols that satisfy the stringent requirements of this application space. Designing secure protocols is complicated by the seemingly conflicting requirements of consumers, automobile manufacturers, and government, particularly when trying to provide strong vehicle identification while protecting driver privacy. Fortunately, the properties of vehicular networks provide new approaches for these challenges, allowing us to develop new primitives based on, for example, the entanglement of vehicle trajectories and the use of simple reanonymizers. We anticipate that the challenges outlined in this paper and the new opportunities for solutions in vehicular networks will encourage other researchers to start studying this important and exciting research area.

## Acknowledgements

The authors gratefully acknowledge the help of D. Baker, J.-P. Hubaux, P. Koopman, T. Parno, R. Rajkumar, A. Seshadri, D. Seymour, and D. Song.

## 9. REFERENCES

- [1] IEEE 802.11 WG, part 11: Wireless LAN Medium Access Control (MAC) and physical layer (PHY) specifications, August 1999.
- [2] Standard specification for telecommunications and information exchange between roadside and vehicle systems - 5 GHz band dedicated short range communications (DSRC) Medium Access Control (MAC) and physical layer (PHY) specifications, September 2003.
- [3] Car2Car Communication Consortium. <http://www.car-2-car.org/>, 2004.
- [4] Jan Camenisch and Anna Lysyanskaya. Efficient non-transferable anonymous multi-show credential system with optional anonymity revocation. In *Proc. of Advances in Cryptology - Eurocrypt*, 2001.
- [5] Srđan Čapkun and J.P. Hubaux. Secure positioning of wireless devices with application to sensor networks. In *Proc. of INFOCOM*, March 2005.
- [6] John Chovan, Louis Tijerina, Graham Alexander, and Donald Hendricks. Examination of lane change crashes and potential IVHS countermeasures. [http://www.itsdocs.fhwa.dot.gov/JPODOCS/REPTS\\_TE/61B01!.PDF](http://www.itsdocs.fhwa.dot.gov/JPODOCS/REPTS_TE/61B01!.PDF), March 1994.
- [7] John R. Douceur. The Sybil attack. In *First International Workshop on Peer-to-Peer Systems (IPTPS)*, March 2002.
- [8] FCC. FCC allocates spectrum in 5.9 GHz range for intelligent transportation systems uses. [http://www.fcc.gov/Bureaus/Engineering\\_Technology/News\\_Releases/1999/nret9006.html](http://www.fcc.gov/Bureaus/Engineering_Technology/News_Releases/1999/nret9006.html), October 1999.
- [9] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Packet leases: A defense against wormhole attacks in wireless networks. In *Proc. of IEEE INFOCOM*, April 2003.
- [10] Jean-Pierre Hubaux, Srđan Čapkun, and Jun Luo. The security and privacy of smart vehicles. *IEEE Security & Privacy Magazine*, 2(3):49–55, May–June 2004.
- [11] Markus Kuhn. Probabilistic counting of large digital signature collections. In *Proc. of USENIX Security Symposium*, 2000.
- [12] L. Lazos and R. Poovendran. SeRLoc: Secure range-independent localization for wireless sensor networks. In *Proc. of ACM Wireless Security Workshop (WiSe)*, October 2004.
- [13] L. Lazos, R. Poovendran, and Srđan Čapkun. ROPE: Robust position estimation in wireless sensor networks. In *Proc. of IPSN*, April 2005.
- [14] Zang Li, Wade Trappe, Yanyong Zhang, and Badri Nath. Robust statistical methods for securing wireless localization in sensor networks. In *Proc. of IPSN*, April 2005.
- [15] Donggang Liu, Peng Ning, and Wenliang Kevin Du. Attack-resistant location estimation in sensor networks. In *Proc. of IPSN*, April 2005.
- [16] PKI Forum. Verisign fraudulent certificates. <http://www.pkiforum.com/resources/verisigncerts.html>, Accessed on January 2005.
- [17] Bartosz Przydatek, Dawn Song, and Adrian Perrig. SIA: Secure information aggregation in sensor networks. In *Proc. of ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.
- [18] Maxim Raya and Jean-Pierre Hubaux. The security of vehicular ad hoc networks. In *Proc. of ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, November 2005.
- [19] Arvind Seshadri, Adrian Perrig, Leendert van Doorn, and Pradeep Khosla. SWATT: Software-based attestation for embedded devices. In *Proc. of IEEE Symposium on Security and Privacy*, May 2004.
- [20] Trusted Computing Group. Trusted platform module main specification, Part 1: design principles, Part 2: TPM structures, Part 3: Commands. <http://www.trustedcomputinggroup.org>, October 2003. Version 1.2, Revision 62.
- [21] U.S. Department of Transportation, Bureau of Transportation Statistics. Transportation Statistics Annual Report, 2003.
- [22] Jijun Yin, Tamer ElBatt, Gavin Yeung, Bo Ryu, Stephen Habermas, Hariharan Krishnan, and Timothy Talty. Performance evaluation of safety applications over DSRC vehicular ad hoc networks. In *Proc. of ACM workshop on Vehicular Ad Hoc Networks (VANET)*, 2004.
- [23] Magda El Zarki, Sharad Mehrotra, Gene Tsudik, and Nalini Venkatasubramanian. Security issues in a future vehicular network. In *European Wireless*, February 2002.