Analysis of Security Overhead in Broadcast V2V Communications

Mujahid Muhammad[1(✉) [0000-0003-2177-5086]], Paul Kearney[1[0000-0002-6484-3344]],
Adel Aneiba[1[0000-0002-8021-445X]] and Andreas Kunz[2]

[1] Birmingham City University, Birmingham, UK
mujahid.muhammad@mail.bcu.ac.uk
{paul.kearney,adel.aneiba}@bcu.ac.uk
[2] Lenovo, Oberursel, Germany
akunz@lenovo.com

Abstract. This paper concerns security issues for broadcast vehicle to vehicle (V2V) messages carrying vehicle status information ((location, heading, speed, etc.). These are often consumed by safety-related applications that e.g. augment situational awareness, issue alerts, recommend courses of action, and even trigger autonomous action. Consequently, the messages need to be both trustworthy and timely. We explore the impact of authenticity and integrity protection mechanisms on message latency using a model based on queuing theory. In conditions of high traffic density such as found in busy city centres, even the latency requirement of 100ms for first generation V2V applications was found to be challenging. Our main objective was to compare the performance overhead of the standard, PKC-based, message authenticity and integrity protection mechanism with that of an alternative scheme, TESLA, which uses symmetric-key cryptography combine with hash chains. This type of scheme has been dismissed in the past due to supposed high latency, but we found that in high traffic density conditions it outperformed the PKC-based scheme. without invoking congestion management measures. Perhaps the most significant observation from a security perspective is that denial of service attacks appear very easy to carry out and hard to defend against. This merits attention from the research and practitioner communities and is a topic we intend to address in the future.

Keywords: V2V, Security, Performance, Queuing Theory.

1 Introduction

The term Intelligent Transportation System (ITS) covers a range of advanced road transport applications. These include safety related services employing direct radio communications between vehicles (V2V) and between vehicles and roadside infrastructure (V2I), which are included in the topic of Connected Vehicles (CV) in the US and Cooperative ITS (C-ITS) in Europe.

In the US, Society of Automotive Engineers (SAE) International has defined standard J2735 [1] covering the format, structure and contents of V2V and V2I messages.

Sixteen message types are listed, plus provision for regionally defined text messages. The main type of concern to this paper is the Basic Safety Message (BSM), which is broadcast by vehicles to provide status information (location, heading, speed, etc.) to other vehicles in the vicinity. The information is utilized by a variety of applications in receiving vehicles to (in conjunction with data from on-board sensors and other sources) to augment the driver's situational awareness, issue alerts, recommend courses of action, and potentially to trigger autonomous action. By default, BSMs are broadcast 10 times per second.

In Europe, ETSI has published two related standards:

- EN 302 637-2 [2] gives the specification of a Co-operative Awareness Basic Service including the syntax and semantics of the Cooperative Awareness Message (CAM).
- EN 302 637-3 [3] does likewise for a Decentralised Environment Notification (DEN) Basic Service and the associated DEN Message (DENM).

Like BSMs, CAMs provide vehicle status data and are broadcast periodically. The CAM transmission frequency can be varied between 1 and 10Hz depending on conditions. In contrast DENMs are alerts that are sent when particular events occur. Broadly speaking, BSMs and CAMs are comparable, and work is going on to align the two standards. The typical latency requirement for CAMs from first-generation use cases such as Forward Collision Warning and Emergency Vehicle Warning is better than 100ms. For next generation use cases such as Vehicle Platooning, the requirement reduces to 10ms and for Autonomous Driving, as low as 1ms.

Clearly, the utility of the applications built on the V2V messaging services depends critically on the timeliness and trustworthiness of the received messages. These two concerns are linked in that measures taken to protect and assure the integrity and authenticity of a message consume time and tie up resources, thereby increasing the time taken to deliver a message. The main purpose of this paper is to examine the trade-off between security and timeliness, and to compare the performance of different approaches to security. In particular, we address the question of whether utilization of a derivative of the Timed Efficient Stream Loss-tolerant Authentication (TESLA) protocol [4] is a viable alternative to the prevailing solutions based on Public-Key Cryptography (PKC). Note that this paper is only concerned with the per-message overhead of the schemes. We recognise that longer-timescale issues such as certificate distribution and revocation and/or renewal in the case of PKC and key-chain generation and distribution of initial commitments in the case of TESLA are also germane, and these will be addressed in subsequent papers.

There are two main families of network infrastructure that have been proposed to support V2V message transmission. The one that has been around longest and is arguably better-established uses WiFi technology based on the IEEE 802.11p standard running in the 5.9 GHz frequency band. Its use is specified by the Dedicated Short-Range Communication (DSRC) collection of standards in the US and by ITS-G5 within the European Cooperative ITS initiative in the EU. The second, which its proponents argue to be the better long-term bet, is known as Cellular-V2X (C-V2X) and is being defined by the 3GPP consortium as an extension to its mobile network architecture. In C-V2X, longer-range communications are sent via the cellular network, but shorter-range, low latency communications utilize the so-called PC5 interface (also known as the sidelink channel). PC5 messages are sent directly over the air and not via the cellular network. In this paper, we mainly consider C-V2X PC5 communications, but the basic issues and conclusions apply to both families.

The paper is structured as follows. After discussion related work, we outline the model based on queuing theory that we have used. We briefly describe the C-V2X PC5 mode with network managed resource allocation and derive an appropriate choice of model parameters to describe its behaviour. We then examine a traffic scenario corresponding to a busy city centre and the necessity for congestion management. The next sections compare the performance overheads of the standard, PKC-based authenticity and integrity protection scheme and one based on the TESLA protocol. The main finding is that there is no clear winner: PKC prevails at low message traffic densities, and TESLA when they are high. There then follows a discussion of the implications of the study and recommendations for future work. Of particular concern is the potential for disruption of V2V messaging by denial of service attacks that are simple to carry out and difficult to defend against.

2 Related Work

Queuing theory has been used widely to model telecommunication systems including vehicular networks. The authors of [5] proposed an analytical model describing the performance of periodic broadcasts in vehicular ad hoc networks, in terms of packet collision probability and average packet delay. A comprehensive M/M/∞ model of vehicular traffic dynamics over a roadway, with intermittently connected networks is presented in [6]. Also, the work of [7] describes analytical models to assess how queue length estimation at an intersection is influenced by the percentage of probe vehicles in the traffic stream. A discrete time D/M/1 model for analysing the performance periodic broadcast in VANETS is presented in [8]. The model shows numerical results of packet collision probability and average packet delay. In [9], the authors utilise an M/M/m queuing model to evaluate the probability that a vehicle finds all channels busy, and to derive the expected waiting times.

None of these works has modelled the security overhead for broadcast messages and its effect on the system performance. In the study presented, the overheads of PKC- and TESLA-based security mechanisms have been modelled and compared in a saturated vehicular traffic condition. A further difference from previous works is that we consider LTE-V2V as the network technology used, and specifically the variant exploiting in-coverage operation where the radio resources are assigned to transmitting vehicles by the infrastructure.

3 A Simple Queuing Theory Model

The delivery of a broadcast message has three main stages. First, the message is composed and formatted ready for transmission, then it is broadcast, and finally it is received and decoded/interpreted by all the receivers in range. All three steps involve shared use of finite resource. In the first and third steps, the resource is a processor assumed to be able to process one message at a time, and in the second step it is the wireless medium. There are numerous strategies for sharing the available bandwidth among transmitters, but ultimately its message-carrying capacity is finite.

If a message arrives at a resource and finds it is busy, then it must either be added to a queue to wait its turn, or else it will be lost. Conversely, if a resource finishes processing one message and finds its queue empty, then it will be idle until the next message arrives. Increasing the average message arrival rate makes it more likely that a given message will find the resource busy, and so the average queue length and time spent in the queue will grow. However, the greater the average queue length, the less likely that the resource will be idle, so that more messages will be processed per unit time. Provided that the average message arrival rate is less than the capacity of the resource, an equilibrium will be reached such that on average, input and output rates are equal. The higher the throughput, the greater will be the queue length and the longer the time taken.

The simplest model of a queuing system is a memoryless continuous time Markov chain denoted in so-called Kendall notation as an M/M/1 model. In an M/M/1 queuing system, the is no limit on the length of the queue, the queuing discipline is ‘first come first served’, the distributions of message arrival intervals and of time take to process a message are both exponential, and there is a single processing resource. In such a case, the average time a message spends in the system (including time spent being processed) and the average queue length are respectively:

$$T = 1/(\mu - \lambda) \text{ and } L = \lambda T = \lambda/(\mu - \lambda)$$

where λ is the average message arrival frequency and μ is the average rate at which messages can be processed by the resource. Notice that there is a singularity when $\lambda = \mu$ indicating that the steady state equilibrium model breaks down and the values of T for $\lambda \geq \mu$ have no physical meaning. Performance targets will not typically be expressed in terms of averages, but rather as expectations regarding exceptions to the norm. It is also useful, therefore, to consider the time within which a fraction x of messages is likely to be processed:

$$T_x = -\ln(1-x)/(\mu - \lambda) \tag{1}$$

Note that $T_{1-1/e} = T$.

As shown in Fig. 1, we model each of the three stages mentioned above as such an M/M/1 queuing system, so that the average end-to-delay is given by:

$$T_{Total} = T_S + T_T + T_R + A; \tag{2}$$

$$T_S = 1/(\mu_S - \lambda);\ T_T = 1/(\mu_T - N\lambda);\ T_R = 1/(\mu_R - N\lambda);$$

Where the subscripts S, T and R stand for Sender, Transmission and Receiver respectively, N is the number of vehicles within reception range and A is a constant term added for generality. λ is the rate at which messages are generated by applications in each of the N vehicles. The multiplicative factor N is applied to the message traffic flowing through the transmission medium as it is shared by all vehicles, and to that for reception because each message is received by all vehicles.

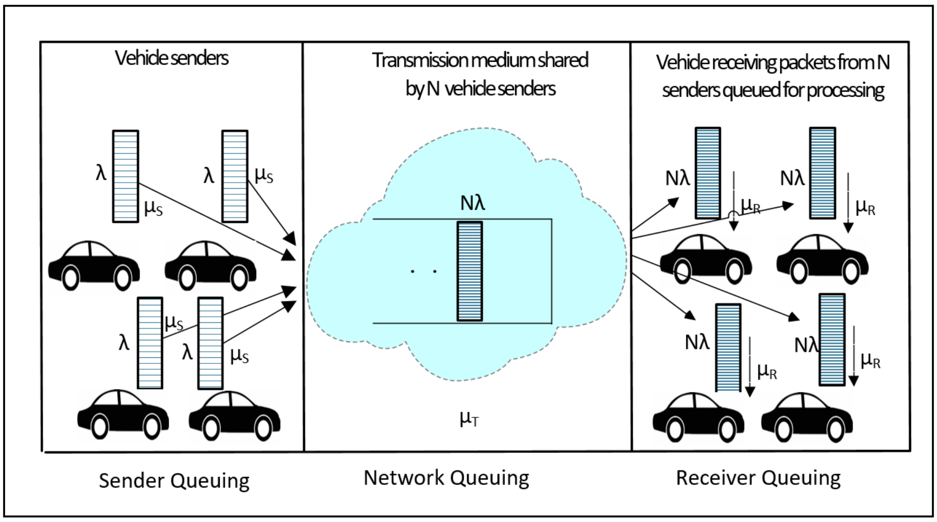
Vehicle senders
Transmission medium shared by N vehicle senders
Vehicle receiving packets from N senders queued for processing
λ
μ_S
Nλ
μ_R
μ_T
Sender Queuing
Network Queuing
Receiver Queuing

Fig. 1. End-to-End Queuing Delay Model

We assume that the message processing times have a component proportional to message length and a fixed component independent of length, so that:

$$\mu_i = 1/(\mathrm{l}r_i + c_i);\ i = S,T,R \tag{3}$$

where l is the message length in bytes, r*i* is the time to process one byte and c*i* is the additional per-message processing time. Message authenticity and integrity measures affect both message length and the per-message processing time.

4 Network-Specific Issues

3GPP started work on Vehicle to Everything (V2X) in its Release 14 [10], utilizing the Long Term Evolution (LTE) radio and Evolved Packet Core (EPC) for message transmission for the different scenarios Vehicle to Vehicle (V2V), Vehicle to Infrastructure (V2I), Vehicle to Network (V2N) and Vehicle to Pedestrian (V2P). V2V reused the previously defined PC5 interface for mission critical subscribers (Proximity Services [11]) and made it applicable for public usage in vehicles, but in a different radio band. Currently work in 3GPP is ongoing to specify V2X for the 5G system [12].

LTE-V2V is based on the uplink Physical and Medium Access Control (MAC) layer radio network protocols of LTE. Thus, it utilizes Orthogonal Frequency Division Multiplexing at the Physical layer and Single Carrier Frequency Division Multiple Access at the MAC layer. A given LTE physical channel is divided into smaller fragments, both in time and frequency, which are referred to as frames. Every LTE frame is 10ms wide in the time domain and its length is equal to the system bandwidth in the frequency domain. LTE-V2V supports 10MHz and 20MHz channels, where each channel is divided into frames, Resource Blocks (RBs), and sub-channels. An RB is the smallest unit of frequency resources that can be allocated to an LTE user. It is 180 kHz wide in frequency (12 sub-carriers of 15 kHz) and one slot in time (i.e. 0.5ms). LTE-V2V defines a sub-channel as a group of RBs in the same sub-frame. Sub-channels are shared among vehicles for the transmission and reception of messages. The number of data bits carried by the group of RBs depends on the chosen Modulation and Coding Scheme (MCS). The number of RBs in each sub-channel for the transmission of messages depends on the available bandwidth and by configuration of the network. A typical LTE-V2V physical channel of 20MHz bandwidth can support a maximum data rate of 50Mbps (assuming 16QAM modulation scheme is used). This corresponds to a transmission of approximately 21,000 messages per second, given a safety message size of 300 bytes.

In PC5-based communication. RBs may either be allocated to a transmitting vehicle by a local element of radio access network infrastructure known as an eNodeB (eNB), or else selected autonomously by the vehicles using a distributed scheduling scheme [13]. The former is only possible when in network coverage. The following discussion applies to the case where RBs are allocated by an eNB.

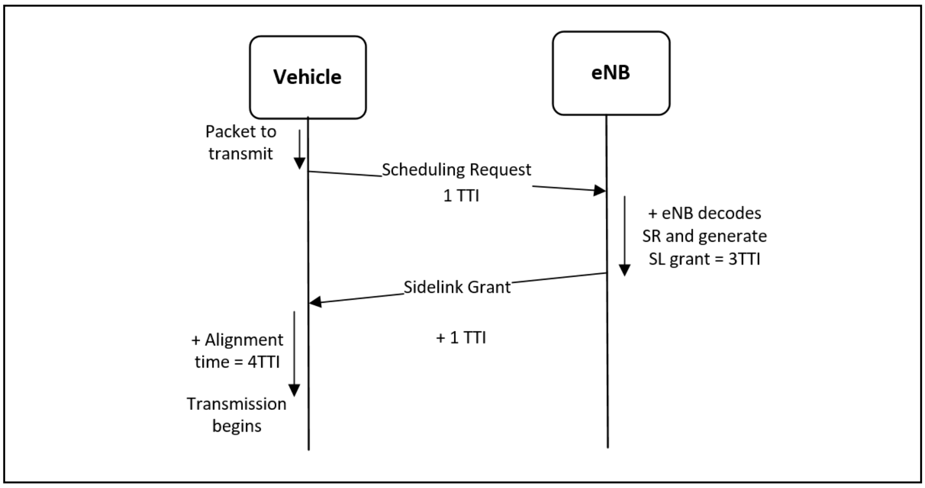

Fig. 2. Network-assisted sidelink resource allocation procedure

As shown in Fig. 2, in Mode 3, a vehicle V_S with data to transmit sends a scheduling request to the eNB and receives an allocation of slots in return. V*s* then begins transmission after a so-called alignment time, defined as the waiting time for decoding the received scheduling grants and processing the packets ready for transmission. Time advances in units referred to as the transmission time interval (TTI = 1ms). The total time elapsed between V_s sending a scheduling request and being able to begin transmission is 9TTI, i.e. 9ms. This contributes a fixed increment to the overall delay, i.e. it is part of A in equation (2) above. The value can, however, be reduced using a semi-persistent scheduling technique whereby the grant of a resource block remains valid for a period of time, so that a new request is not needed for every message.

5 Congestion Management

We used a simulation of a busy city centre scenario following a method detailed in [14] to obtain an initial ‘worst case’ assumption of 400 for the number of vehicles within the awareness range of a receiver. Combining this with the default message frequency of 10 messages per second per vehicle gives $N\lambda = 4000$ messages per second, meaning that both μ_T and μ_R must be much greater than this figure to avoid the rapid rise in latency and queue length as the singularities are approached. An estimate for transmission resource capacity of 21000 messages per second was given above, yielding $N\lambda/\mu_T \approx 0.2$. However, if the receiver capacity just meets the C-V2X requirement [15] of being able to process one message in 2ms, $N\lambda/\mu_R = 8$, which is well outside the region of validity of the queuing model.

Both the SAE and ETSI message standards provide for congestion management mechanisms For example, SAE J2945/1 [16] includes a congestion control algorithm that includes adaptive functions for calculation of Inter-Transmission Time (ITT) and transmission power for BSMs. To adjust the ITT, each vehicle keeps track of the number of other vehicles within 100m. If this exceeds a threshold value, then the ITT is increased linearly until a maximum value is reached. This corresponds to the message transmission frequency decreasing from 10Hz to 1.6Hz. Similarly, the transmission power decreases from 20 dBm to 10 dBm as the ‘channel busy ratio’ grows from 50% to 80%, which results in fewer vehicles being within reception range. The ETSI CAM standard is less prescriptive but provides for a reduction in CAM generation frequency from the nominal 10Hz to 1Hz. A reduction of the message generation frequency to 1.6Hz and of vehicles in range to 60 would result in $N\lambda/\mu_R \approx 0.2$.

6 The Overhead of Authenticity and Integrity Protection

Both the US and EU ITS schemes use PKC as described in IEEE standard 1609.2 to provide source authentication, data integrity and non-repudiation in vehicular communication. This solution involves signing a message using ECDSA (Elliptic Curve Digital Signature Algorithm), and attaching a public key certificate to each signed message to enable verification of the message at the receiving end. This process incurs a high computational cost per message for a) signature generation by the sender and b) verification by the receiver, although the cost is significantly reduced by use of a hardware security module (HSM) with support for ECDSA. The service rate or processing capacity of the vehicle on-board unit is reduced in consequence. There are two components to this overhead resulting respectively from:

1. The increased length of the message due to appending the signature and the certificate. This affects all three queuing systems, and

2. The per-message delay due to generation and verification of the signature by sender and receiver respectively.

Based on specification sheets of HSMs, the signature generation and verification times are estimated to be 0.125ms and 0.5ms. Taking 300 bytes as the basic message size, 64 bytes as the signature length and 194 bytes as the signature length, we can use (3) to estimate the per-byte processing speeds needed to meet the C-V2X requirements of $\mu_T \geq 1000$Hz and $\mu_R \geq 500$Hz as $r_T \approx 1.6\text{x}10^{-6}$s and $r_R \approx 2.7\text{x}10^{-6}$s. For simplicity we take $r_T = r_R = 1.5\text{x}10^{-6}$s.

The overhead for symmetric-key cryptography (SKC) is much lower than for PKC. This is because generation and verification of a symmetric-key Message Authentication Code (MAC, not to be confused with the Medium Access Control) requires much less computation than a PKC signature, the MAC itself is shorter than a signature, and no digital signature is included. However, MACs do not provide a solution to the message integrity and authenticity requirement without a secure and efficient means of sharing symmetric keys.

Hash chain techniques potentially offer a way to combine the best of both worlds. One commonly used protocol for broadcast authentication in wireless ad hoc networks is TESLA [4]. TESLA uses an SKC MAC algorithm to protect the integrity of messages but introduces the element of asymmetry by delaying the disclosure of the secret key used. A given key may only be used by a sender to generate MACs within a well-defined time window, after which it is made public and may be used by receivers to verify the integrity of messages sent within that window. A new key is then used for the next window. A sequence of keys used by a given sender is generated such that the Nth key used is the result of applying a hash function to the N+1th key. Thus, the hash function can be used to verify a sequence of keys used by a given sender, and hence the sequence of messages it sent, provided that the first key in the sequence can reliably be attributed to that sender. The main benefits of TESLA are low computation overhead, low communication overhead, and robustness to packet loss. However, TESLA also has some shortcomings: the basic version cannot provide non-repudiation which is crucial in V2V systems; the one-way key chain has a finite length, so new chains need to be created periodically; and there is a requirement for loose synchronization between sender and receivers. All of these can be addressed, e.g. access to trusted time-stamping permits non-repudiation, and there are various options for synchronisation, which is required by cellular protocols anyway. Such measures do add complications and overheads on longer timescales that must be weighed against the infrastructure required to support the competing PKC-based approach in a comprehensive comparison. Such topics will be covered in a future paper; here we focus on the per-message overhead. As a lightweight message authentication mechanism, TESLA has been employed in several research proposals to address the security problems in V2V systems [17], [18].

In TESLA, the delayed key disclosure results in a delay before safety messages are verified. Each key is disclosed in the following safety message packet broadcast from the sending vehicle. This means that each receiver has to buffer the received messages for at least one time interval and wait for the corresponding key in the following safety message broadcast.

Comparing TESLA's performance with that of a PKC-based scheme, TESLA gains because it uses SKC (shorter messages and lighter-weight computation), but loses out because the receiving vehicle must wait a full time interval before receiving the key that enables it to verify and process a message. Clearly, the length of the time interval (Q) is crucial to whether TESLA is competitive with PKC. It must be chosen so that the vast majority of messages arrive at the receiver within the same interval in which they were sent. In order to choose an appropriate value for Q we drop the receiver term from (2) and apply the logarithmic factor from (1) with x = 0.99 to the S and T terms. This gives a conservative measure of the time taken for a message to reach the receiver that assumes 99th percentile delays at the sender and in the network and results on a choice of Q ≈ 0.012s. This value appears as an additional contribution to A in (2) when modelling TESLA.

Fig. 3 uses (2) to compare the average end-to-end delays as a function of λ for the following cases: no security (solid line), PKC security (dashed line), SKC security (dash-dot line) and TESLA (long dashes). Also shown (dotted line) is the curve corresponding to the 99th percentile sender and transmission terms for the SKC case that was used to derive a value for Q.

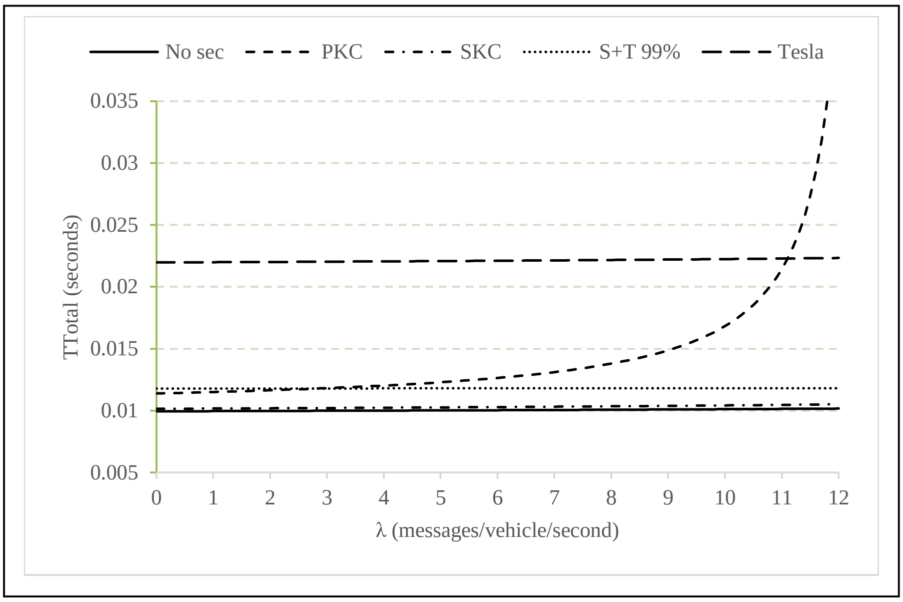

Fig. 3. Overhead from use of integrity and authenticity protection schemes

It is apparent that while the PKC overhead is small when message traffic is low, it increases rapidly when the λ/μ ratio approaches 1 for one of the queuing systems. The receiver queuing system appears to be the critical one, with a singularity at λ ≈ 12.5 for our estimated values. Notice that the TESLA curve is fairly flat whereas the PKC curve increases rapidly above λ ≈ 7s^{-1}. The two cross around λ = 11s^{-1}, which is close to the default rate of λ = 10s^{-1}. Below this rate, the PKC-based mechanism is preferable, but above it incurs a severe penalty. Given the many estimates made and the simplicity of the queuing theory model, it is not wise to infer much from the exact numbers. However, it appears that a TESLA-based scheme should not be ruled out without more detailed study.

7 Discussion

It is clear from the study that V2V message latency is a serious issue for ITS, particularly for scenarios with high traffic density. For the parameter values chosen here, the main concern is with the ability of receivers to process messages sufficiently quickly. Sender performance is of much less concern as a vehicle only processes its own messages when sending, but it receives messages from all vehicles within range. Improving the receiver's per-byte processing performance helps, but when PKC is used, the time required to verify the digital signature is still large enough that the receiver delay remains the critical term. However, even if signature verification performance is considerably better than that assumed, this merely shifts the problem to one of the network being the bottleneck. With our model, the transmission singularity occurs at $N\lambda \approx 11000s^{-1}$ in the PKC case. In the absence of congestion management $\lambda=10s^{-1}$ and our estimate for N in a busy city centre is 400, yielding $N\lambda=4000s^{-1,}$ which is uncomfortably close. This is a more serious problem as network performs is determined by standards rather than by equipment.

This confirms the motivation for looking at alternative schemes for message integrity and authenticity protection and assurance. TESLA benefits from using symmetric key cryptography, but suffers the penalty of the receiver needing to wait for receipt of new keys to allow the preceding time interval's messages to be verified. This penalty, while significant, is a fixed overhead that does not depend on λ or N, so that TESLA wins out eventually over PKC at high traffic densities and/or message frequencies. Based on our choice of parameters, the cross-over point occurs with the λ range of interest for V2V communications. Much of TESLA's fixed overhead is due to the 9ms required by the LTE network-assisted resource allocation procedure. This affects all options examined, of course, but TESLA suffers a double dose. There is a mechanism that avoids the need for a sender to request a new resource allocation for each message that should reduce the average delay. It has not been taken into account in the results presented here, and if the reduction is significant it would render TESLA more competitive with PKC. The 9ms will cause problems in any case for next generation use cases such as Vehicle Platooning and Autonomous Driving that possess latency requirements of 10ms and 1ms respectively.

Our model assumes infinite queues, which is obviously not realistic. In reality, messages would be dropped when queues reach their limit. This would alleviate the load on the receivers' resources and would reduce the latency experienced by the delivered messages at the expense of infinite latency for the dropped messages. Given the safety-relevance of many V2V applications it seems unwise to leave it to chance to decide which messages get through and which do not.

It is clear that standards organisations SAE and ETSI recognise that there is a problem as both provide mechanisms for congestion management, albeit their main concern is with the transmission resources rather than the receiving vehicles. Whether reduction of the message transmission frequency and/or transmitter power based on vehicle density is acceptable depends on the application requirements. For example, what really matters is not the message frequency or latency, but whether the receiving application has sufficiently accurate and up-to-date information. Fortunately, those scenarios in which vehicle densities are highest, such as busy city centres, are also likely to feature the lowest vehicle speeds and hence require less frequent updating of information. Intermediate density scenarios such as busy highways may offer the most challenging combination of message traffic volumes and application requirements.

The SAE and ETSI congestion management mechanisms rely on sending vehicles being co-operative — the senders must detect high vehicle density and respond by reducing transmission frequency and/or power. They can, therefore, be thwarted by senders that are non-compliant because they implement the standard badly or not at all, because they malfunction, or because they are intentionally selfish. In this last category, consider that a small minority of vehicles that continue to broadcast at high frequency and power, while the majority follow the congestion management algorithms. The selfish minority benefit from the congestion control while the lawful majority suffer the costs.

Beyond selfish behaviour there is considerable potential for malicious disruption of V2V applications. Consider a malicious agent that simply broadcasts messages at high frequency and power in order to fill up the message queues of receiving vehicles causing delays and lost messages. Conventional authenticity protection cannot help, because the cost of verifying the authenticity of messages is contributing to the vulnerability being exploited by the attacker. This type of attack is likely to become common because of the low cost and skill requirement and is extremely difficult to defend against. Until a solution is found, the deployment of safety related applications that rely on timely information from other vehicles rather than simply benefiting from it may be judged to incur too high a risk. As an example, consider vehicle platooning. If the vehicles in a convoy rely on message exchange to enable shorter inter-vehicle distances and higher convoy speeds, then a sudden denial of service could well result in a collision.

8 Conclusions

In the study presented here we have used a simple three-stage queuing theory model to explore performance issues in broadcast V2V communication. Despite its simplicity it has proved valuable in developing a qualitative understanding of performance. The model is capable of representing a variety V2V broadcast technologies. The parameter values used here were based on the LTE-V2V mode with network-allocated resource blocks. We need to extend the study to the LTE-V2V autonomous allocation mode and also to US and European wifi-based solutions, but we expect broadly-similar results.

A first observation is that in conditions of high traffic density such as found in busy city centres, even the latency requirement of 100ms for first generation V2V applications appears challenging without invoking congestion management measures. The critical element is the queuing system representing a vehicle acting in message-reception role. This has to process messages from all vehicles within range whereas the sending queuing system only has to process outbound messages from its parent vehicle.

Our main objective was to compare the performance overhead of the standard, PKC-based, message authenticity and integrity protection mechanism with that of an alternative scheme, TESLA, which uses symmetric-key cryptography combine with hash chains. This type of scheme has been dismissed in the past due to supposed high latency, but we found that in high traffic density conditions it outperformed the PKC-based scheme. Simulation-based studies combined with benchmarking of representative equipment are needed to confirm this result and explore where the performance crossover occurs. Subject to this confirmation, the result indicates that TESLA merits deeper consideration. Subsequent papers will explore options for integration of TESLA with cellular communications standards and infrastructure. Various shortcomings of TESLA were noted in Section 6, and we intend as far as possible to leverage network capabilities to address them.

Perhaps the most significant observation from a security perspective arising from this study is that denial of service attacks appear very easy to carry out and hard to defend against. This merits attention from the research and practitioner communities and is a topic we intend to address in the future.

References

1. SAE International, Dedicated Short Range Communications (DSRC) Message Set Dictionary, V2X Core Technical Committee, https://doi.org/10.4271/J2735_200911
2. ETSI EN 302 637-2 V1.3.2 (2014-11), Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service
3. ETSI EN 302 637-3 V1.2.2 (2014-11), Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specifications of Decentralized Environmental Notification Basic Service
4. The TESLA Broadcast Authentication Protocol, Adrian Perrig, Ran Canetti, J. D. Tygar, Dawn Song, CryptoBytes, 5:2, Summer/Fall 2002, pp. 2-13
5. Yang Q, Zheng J, Shen L. Modeling and performance analysis of periodic broadcast in vehicular ad hoc networks. In2011 IEEE Global Telecommunications Conference-GLOBECOM 2011 2011 Dec 5 (pp. 1-5). IEEE.
6. Khabbaz MJ, Fawaz WF, Assi CM. A simple free-flow traffic model for vehicular intermittently connected networks. IEEE Transactions on Intelligent Transportation Systems. 2012 Sep;13(3):1312-26.
7. Comert G, Cetin M. Analytical evaluation of the error in queue length estimation at traffic signals from probe vehicle data. IEEE Transactions on Intelligent Transportation Systems. 2011 Jun;12(2):563-73.

8. Khabbaz MJ, Fawaz WF, Assi CM. Modeling and delay analysis of intermittently connected roadside communication networks. IEEE Transactions on Vehicular Technology. 2012 Jul;61(6):2698-706.
9. Fowler S, Häll CH, Yuan D, Baravdish G, Mellouk A. Analysis of vehicular wireless channel communication via queueing theory model. In2014 IEEE International Conference on Communications (ICC) 2014 Jun 10 (pp. 1736-1741). IEEE.
10. 3GPP TS 23.285 “Architecture enhancements for V2X services”, Release 16, V16.0.0., March 2019
11. 3GPP TS 23.303, “Proximity-based services (ProSe); Stage 2”, Release 15, V15.1.0., June 2018
12. 3GPP TS 23.287 “Architecture enhancements for 5G System (5GS) to support Vehicle-to-Everything (V2X) services”, Release 16, V0.3.0. April, 2019
13. 3GPP TS 36.331, “Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification”, Release 15, V15.5.1., April, 2019
14. Bazzi A, Masini BM, Zanella A. How many vehicles in the LTE-V2V awareness range with half or full duplex radios?. In2017 15th International Conference on ITS Telecommunications (ITST) 2017 May 29 (pp. 1-6). IEEE.
15. 3GPP TS 22.186 “Enhancement of 3GPP support for V2X scenarios; Stage 1”, Release 16, V16.1.0. December, 2018
16. SAE International (2016). On-Board System Requirements for V2V Safety Communications, version 201603.
17. C. Lyu, D. Gu, Y. Zeng, and P. Mohapatra, “Pba: Prediction-based authentication for vehicle-to-vehicle communications,” IEEE Transactions on Dependable and Secure Computing, vol. 13, no. 1, pp. 71–83, 2016.
18. C. Lyu, D. Gu, X. Zhang, S. Sun, and Y. Tang, “Efficient, fast and scalable authentication for vanets,” in Wireless Communications and Networking Conference (WCNC), 2013 IEEE. IEEE, 2013, pp. 1768–1773.