Enforcing Replica Determinism in the Road Side Units of Fault-Tolerant Vehicular Networks

João Almeida[1,2](✉), Joaquim Ferreira[1,3], Arnaldo S.R. Oliveira[1,2],
Paulo Pedreiras[1,2], and José Fonseca[1,2]

[1] Instituto de Telecomunicações, Aveiro, Portugal
[2] DETI, Universidade de Aveiro, Aveiro, Portugal
[3] ESTGA, Universidade de Aveiro, Águeda, Portugal
{jmpa,jjcf,arnaldo.oliveira,pbrp,jaf}@ua.pt

Abstract. This paper presents a strategy to enforce replica determinism in the road-side units (RSUs) of wireless vehicular networks. An active replication scheme is used to enhance fault-tolerant behaviour in these RSU nodes, which are responsible for handling channel access and admission control policies in real-time vehicular communications protocols. The proposed solution guarantees consistency among all RSU replicas, by introducing a dedicated link shared exclusively by these units, allowing them to implement an atomic commit protocol of the packets received through the wireless medium. This strategy also has the advantage of reducing packet loss, since only one replica needs to successfully decode the packet in order for it to become available to the RSU group of replicas. It should be noticed however that this method increases the total delay of packet delivery to the upper layers of the communications protocol, so its impact on the real-time properties of the network needs to be further evaluated.

Keywords: Vehicular networks · Intelligent transportation systems · Real-time communications · Fault-tolerance mechanisms · Active replication scheme · Replica consistency · Atomic commit protocol

1 Introduction

Cooperative Intelligent Transportation Systems (C-ITS) aim to reduce road accidents, by extending vehicle's field of view and producing warnings alerts in case of dangerous traffic situations. In addition, these systems have the goal of decreasing CO_2 emissions, road congestions and energy consumption. Infotainment applications based on C-ITS can also be developed, in order to improve passenger's comfort and to provide a better riding and driving experience. Vehicular communications are the main enabling technology supporting these collaborative systems, since they allow vehicles to communicate among each others, to exchange information with the road-side infrastructure and eventually with pedestrians, cyclists and other objects located close to the roads. This is the

J. Ferreira and M. Alam (Eds.): Future 5V 2016, LNICST 185, pp. 3–12, 2017.
DOI: 10.1007/978-3-319-51207-5_1

essential concept behind Internet of Vehicles (IoV) [8], which constitutes a subgroup of the future Internet of Things (IoT). In the IoV, at least a significant percentage of vehicles will be equipped with these communications capabilities, and will be able for instance to disseminate an accident detection, avoiding further chain collisions. Another advantage is that vehicles will be connected with other networks, for example with the sensors network of the driver's home, allowing him/her to control the heat or air conditioning systems, the garage door, the lightning and surveillance systems, etc. before arriving home [7].

Vehicular networks are based on the IEEE WAVE protocol stack in USA and on the ETSI ITS-G5 in Europe. Both of them rely on the IEEE 802.11 standard [9] for the physical and medium access control (MAC) layers, the same standard employed in Wi-Fi technology, but with some different settings, namely the reduced channel bandwidth (10 MHz instead of 20 MHz) to mitigate the impacts of multipath and Doppler effects, and the absence of authentication and association procedures for faster link establishment due to the short connections in these very dynamic environments. The use of the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) method, defined in the standard for managing channel access, resulted in several studies [6,10] reporting multiple problems associated with this mechanism under congested traffic scenarios. Some of these issues are related with the high number of packet collisions and the large values for the end-to-end delay. In order to deal with this problem, ETSI for instance has proposed a Decentralized Congestion Control mechanism that still relies on the CSMA/CA scheme but enables a fine control of some variables that directly influence channel occupation, such as receiver sensitivity, transmit power level, etc. Other works in the literature [5,11] have proposed deterministic MAC protocols to support real-time wireless communications in these safety-critical environments. Due to the higher reliability they provide, solutions based on the support provided by the road-side infrastructure are typically more suitable for vehicular scenarios, where the network's topology changes very quickly. The reference nodes placed in fixed locations, named road-side units (RSUs), can handle admission control policies and the whole traffic scheduling from the on-board units (OBUs) placed inside the vehicles, as suggested in [11]. However, since in this scheme, the RSUs control all network communications, these nodes become critically important for the operation of the intelligent traffic system. As a result, a fault-tolerant vehicular network has been proposed [3], essentially targeting the role played by these master nodes and aiming to increase dependability of the overall network. An active replication scheme for the RSUs was deployed [2], enabling fast recovery of the RSU failed node, by using a backup replica to replace the operation of the primary unit. Nevertheless, in order to work properly, this mechanism assumes that both the primary and the backup replicas are synchronized with respect to the information both units hold about the network. For that to be true, all replicas need to receive and process the same packets send by all other vehicular nodes. Since the wireless medium is not totally reliable, even if the replicas are co-located, they may not be able to decode the same packets without errors. Consequently, this works focus on guaranteeing replica consistency by employing an atomic commit protocol to disseminate all the messages received by both RSU replicas.

2 Dependable Wireless Vehicular Networks

This section provides the background of this work, describing the operation of a real-time vehicular communications protocol for which a fault-tolerant infrastructure-based architecture was designed, aiming to improve the dependability attributes of the network. A detailed description of the proposed RSU replication scheme is provided, for which the replica consistency strategy presented in this paper was developed.

2.1 Real-Time Vehicular Communications Protocol

A deterministic MAC protocol for infrastructure-based vehicular networks was proposed in [11]. This protocol, named Vehicular Flexible-Time-Triggered (V-FTT), is based on a spatial TDMA scheme with a multi-master multi-slave architecture. A backhauling network interconnects the RSUs (master nodes), enabling them with a global vision of the road traffic situation. In V-FTT, time is divided in consecutive Elementary Cycles (ECs) with 50 ms duration, as depicted in Fig. 1. Each EC is further divided into three main windows. The first interval (Infrastructure Window) is used by RSUs to transmit warning messages concerning road hazards together with packets containing scheduling information to the OBUs (slave nodes). These scheduling messages include the time slot assignment for the transmission of OBUs' packets. After the Infrastructure Window, there is a congestion based interval named Free Period, during which vehicles can register with the road infrastructure and nodes non-compliant with V-FTT protocol can transmit. At the end of each EC (Synchronous OBU Window), the OBUs send safety messages according to the time slot scheduling previously disseminated. These packets include information about the vehicle's state (e.g. position, speed and orientation) and the perceived environment (e.g. dangerous traffic situations).

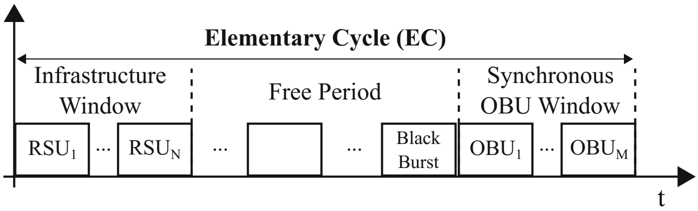

Fig. 1. V-FTT protocol.

2.2 Fault-Tolerant Network Architecture

Based on the need to provide strict real-time and dependability requirements for safety-critical vehicular applications, a fault-tolerant infrastructure-based network architecture was proposed in [3]. Given the key role played by the RSUs in the admission control and scheduling processes, a fail silence mechanism was designed [1], ensuring that these units can only fail by not sending any message to the network. After guaranteeing fail silent behaviour, an active replication scheme was developed [2], providing low-delay recovery procedure of the failed RSU nodes. On the OBU side, a medium guardian entity was devised in order to constrain packet transmission solely for the period inside the time slot previously allocated for that specific mobile unit. A diagram of the fault-tolerant architecture comprising these three major mechanisms is depicted in Fig. 2.

2.3 Active RSU Replication Scheme

The active replication scheme proposed in [2] for the RSUs of infrastructure-based vehicular networks guarantees that in case of failure in the primary unit, a backup replica will replace its operation within a small time interval ($\approx 25\ \mu$s). This way, there is no discontinuity of the traffic scheduling and the real-time guarantees provided by deterministic MAC protocols such as V-FTT, can still be delivered. The proposed mechanism works in the following way. The packet transmission of the backup RSU is always delayed by a small interval in relation to the primary transmission. In this sense, the backup RSU will sense the wireless medium and, if it is occupied, it knows that the primary replica is free of error. Otherwise, if the medium is free, the backup transmits the message, considering that the primary unit failed to send the message. This is only possible due to

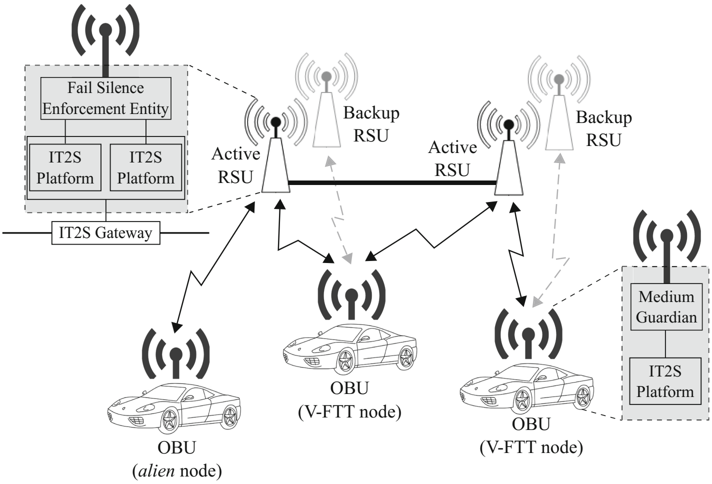

Fig. 2. Fault-tolerant vehicular network architecture [2].

the fail-silent behaviour of the RSU replicas, both in value and time domains. However, for the correct operation of this replication scheme, it is necessary that both replicas are synchronized in time, which is ensured by local GPS receivers, and with respect to the information they hold about the current road traffic situation, e.g. number and identification of vehicles in the coverage area. This last premise is not always valid, since due to the unpredictability of the wireless medium in such dynamic environments, the successfully received messages may not be the same in both replicas, even considering that they are co-located. Therefore, a new strategy to ensure replica consistency between the primary and backup units of each RSU node is proposed in this paper.

3 RSU Replica Determinism

The main goal behind ensuring replica determinism is to guarantee that in case of failure in one of the replicas, the remaining ones continue to deliver the same service without causing any inconsistency in the system. For that purpose, all the replicas need to be synchronized, sharing a consistent view of the environment in which they operate. In the case of infrastructure-based vehicular communications, the distinct RSU replicas must hold the same information about the current state of the road traffic network. For that purpose, they need to successfully decode the same packets and process these messages in a deterministic way. Otherwise, the vehicular network database residing in both replicas could differ and lead to inconsistent sequences of exchanged messages. Let's consider for instance, the situation in which a vehicle transmits a message to register with the network. Assume that this message is correctly received by the active RSU, but in the backup unit, there is an error in the decoding process. As a result, the active node will generate and send a new identification number to this OBU and will schedule a time slot for this new registered vehicle to transmit in the next Elementary Cycle. Then, imagine that the active RSU fails, being automatically replaced by the backup unit. In this scenario, the backup replica will not recognize the messages sent by this new OBU and will not include it in the future slot assignments, since it was not previously registered in the vehicular network database of the backup RSU. This situation is unacceptable and cannot be allowed in the operation of the real-time communications protocol.

In order to solve this problem, this work proposes the introduction of a dedicated wired link between the active and backup replicas, with the main goal of ensuring the correct dissemination of all messages successfully received by at least one of the replicas. This way, it is possible to implement a distributed atomic commitment of the packets received through the wireless medium without utilizing additional network resources (e.g. channel bandwidth). Figure 3 shows the proposed solution, in which correctly decoded packets are delivered to both replicas by employing an atomic commit protocol.

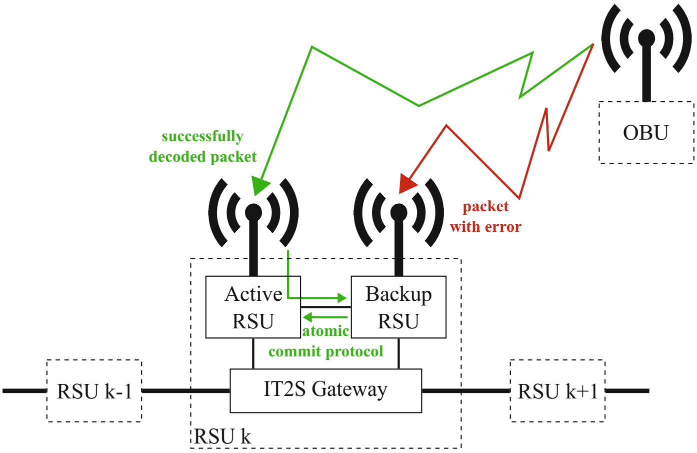

Fig. 3. Replica consistency strategy in infrastructure-based vehicular networks.

3.1 Distributed Atomic Commitment

In distributed systems, an atomic commitment is traditionally achieved by utilizing the two-phase or the three-phase commit protocols or some variant of these. The two-phase commit (2PC) protocol [4, Chap. 7] consists of two communication phases that are coordinated by one of the participants. Initially, during the prepare or voting phase, the coordinator sends a VOTE-REQ (i.e., vote request) message to all other participants. Upon reception of a VOTE-REQ message, the participant replies with a YES or NO message, depending if it is ready or not to commit the transaction. If the answer is negative, the participant immediately aborts the transaction. In the second phase, also known as commit phase, if the coordinator has received at least one negative reply, it will abort the transaction and send an ABORT message. On the other hand, if the coordinator has only received affirmative responses from all the participants, it will commit the transaction and will send a COMMIT message. In both cases (COMMIT or ABORT), the participants will respond with an acknowledgement message after having acted accordingly to the received command.

Despite its efficiency, the two-phase commit protocol has a main drawback. In case of failure of the coordinator during the transition period between the two phases, all the participants will remain blocked waiting for the decision. In order to solve this problem, the three-phase commit (3PC) protocol [12] was proposed, providing a non-blocking solution. In the 3PC protocol, an additional intermediate phase is introduced between the prepare and commit steps of the two-phase commit protocol. This extra phase is usually designated as pre-commit and has the main objective of delaying the final decision until all currently active participants are aware of which resolution the coordinator is about to take. In this scheme, the coordinator sends a PRE-COMMIT message to all participants if the prepare phase is successful. Then, after another round of acknowledgements, the coordinator finally moves to the commit phase. In case of a missing acknowledgement message, the coordinator will abort the transaction. On the other hand, if a participant does not receive the final COMMIT message after a successful pre-commit phase, it will automatically commit the transaction. This way, the blocking situation that can occur in the two-phase commit protocol will be avoided by introducing this intermediate phase.

3.2 Atomic Commit Protocol

The distributed atomic commitment protocol proposed in this paper is inspired in the previously described solutions. However, given the specific characteristics of the scenario presented in Fig. 3, some aspects of the proposed protocol are different from the traditional ones. For instance, in the developed RSU replication scheme there are only two replicas, which means that in each transaction (corresponding to a packet received through the wireless interface), besides one coordinator, there is only one participant in the atomic commit protocol. Additionally, the scenario depicted in Fig. 3 poses an additional challenge, giving the fact that there are two interfaces by which the RSU replicas can communicate, i.e. the wireless and the cabled one. Consequently, several situations may occur, since at a given instant both interfaces can be working perfectly, both can have failed or one can be running while the other is down. Moreover and in order to comply with the real-time requirements of vehicular environments, the implemented replication strategy follows a low latency recovery procedure, which leads to the presence of an active and a backup replica with distinct functional behaviours. For example, since the active RSU is the only replica to transmit messages over the wireless medium, failures in the wireless communications interface of the backup unit can not be detected by the active one.

Given these circumstances, a careful protocol design must be followed in order to ensure a reliable and consistent operation of the road-side infrastructure under

Algorithm 1. Active RSU's algorithm

```
if packet received through the wireless interface then
    # 1 → commit packet information
    # 2 → send packet to the backup RSU
    # 3 → wait for an acknowledgment message from the backup replica
    if timeout expired before acknowledgment message received then
        # 4 → stop sending any type of messages to the backup RSU
    end if
else if packet received through the dedicated wired link then
    # 1 → commit packet information
    # 2 → send acknowledgement message to the backup RSU
else if no packet sent to the backup RSU in the last EC period then
    # 1 → send keepAlive message
else if no packet received from the backup RSU in the last EC period then
    # 1 → inform upper layers about the faulty link/replica
end if
```

Algorithm 2. Backup RSU's algorithm

```
if packet received through the wireless interface then
    # 1 → send packet to the active RSU
    # 2 → wait for an acknowledgment message from the active replica
    if acknowledgment message received then
        # 3 → commit packet information
    else if timeout expired before acknowledgment message received then
        # 3 → inform upper layers about the faulty link/replica
        if active RSU is still transmitting in the wireless interface then
            # 4 → stop backup replica operation
        else
            # 4 → commit packet information
        end if
    end if
else if packet received through the dedicated wired link then
    # 1 → commit packet information
    # 2 → send acknowledgement message to the active RSU
else if no packet sent to the active RSU in the last EC period then
    # 1 → send keepAlive message
else if no packet received from the active RSU in the last EC period then
    # 1 → inform upper layers about the faulty link/replica
    if active RSU is still transmitting in the wireless interface then
        # 2 → stop backup replica operation
    end if
end if
```

all different fault scenarios. To simplify the design process, it is assumed that each RSU replica presents fail-silent behaviour not only in the wireless interface [1], but also in the dedicated wired link. This condition is not guaranteed by the previously developed RSU architecture, however it can be easily achieved by implementing a fail silence enforcement entity, similar to the one devised in [1], for the output traffic of the replica's cabled connection. This means that only valid packets will be transmitted within specified time intervals. In addition, it is also assumed that all packets correctly delivered by the atomic commit protocol are posteriorly handled in an upper layer of the protocol stack, before the information contained in the received messages is included in the RSU replica's database. In this step, it is verified for instance if there is a duplication of the packets received both from the wireless interface and the dedicated link between the replicas. Under perfect circumstances, all the packets would be received through both interfaces, so it is necessary to discard the duplicated messages.

The operation of the proposed atomic commit protocol, both in the active and in the backup replicas, is depicted by the Algorithms 1 and 2, respectively. In the absence of faults, all the packets received through the wireless interfaces of both replicas will be first committed in the active RSU, in order to avoid any inconsistency in case of failure of the active replica. Therefore, it should be noticed that if the active RSU stops working, the backup replica will only replace its operation, if it has delivered the same packets to the upper layers of the protocol stack. This is the reason why in Algorithm 1, all the packets received from both the wireless interface and the dedicated link in the active RSU, are immediately committed. Then, if the packet was received through the wireless interface, it will be sent to the backup RSU using the dedicated wired link. The active RSU will await for an acknowledgement message, that if not arrives within a maximum time limit (and after some retransmissions), will cause the active replica to completely stop sending messages to the backup unit. As a result, this is will force the backup RSU to stop its operation (Algorithm 2). In the other the hand, if the packet was received through the wired link, the active replica will just reply with an acknowledgement. Moreover and in order to verify the connectivity status between the replicas when there are no messages transmitted in the wireless medium, the active replica will transmit a *keepAlive* message at least once per Elementary Cycle (EC). On the other hand, if no packet is received from the backup unit during an entire EC, the active RSU will inform the upper layers that there is a problem with the wired link or backup replica, information that will be forwarded to the backhauling network (via IT2S Gateway).

In the backup RSU (Algorithm 2), if a packet is received through the wireless interface, it will be first sent to the active replica, and it will only be committed after receiving the acknowledgement message. If this acknowledgement does not reach the backup unit, the information about the faulty link or replica will be forwarded to the upper layers. In this case, if the active RSU is still transmitting in the wireless interface, the backup replica will stop its operation, since the active unit is free of faults. However, if the wireless interface of the active RSU is silent, that means the backup is the only replica operating and therefore it can commit the packet information. When receiving a packet through the dedicated wired link, the backup RSU will simply deliver it to the upper layers and send an acknowledgement message to the active replica. The rest of the algorithm is similar to the one in the active RSU, except for the situation when no packets are received from the dedicated link during an entire EC. In this case, the backup replica will verify if the active unit is still transmitting in the wireless interface and it will stop working under those circumstances, in order to avoid any inconsistency in the event of a failure in the active RSU.

4 Conclusions and Future Work

In this paper, an atomic commit protocol is proposed in order to enforce replica determinism between the active and the backup units of an RSU node in infrastructure-based vehicular networks. The goal of this protocol is to disseminate the messages exchanged in the wireless medium through a dedicated wired link connecting both replicas. This way, all the packets received by at least one of the replicas will be delivered to the upper layers in both units. Given the distinct roles played by the active and backup replicas, the algorithms running in both units for the implementation of the atomic commit protocol will be different, as explained in the paper. As future work, the protocol needs to be implemented in a practical vehicular communications system and the timing overhead introduced by the operation of the protocol needs to be evaluated, as well as its impact in the performance of the real-time vehicular network.

Acknowledgements. This work is funded by National Funds through FCT - Fundação para a Ciência e a Tecnologia under the PhD scholarship ref. SFRH/BD/52591/2014 and the project PEst-OE/EEI/LA0008/2013.

References

1. Almeida, J., Ferreira, J., Oliveira, A.: Fail silence mechanism for dependable vehicular communications. Int. J. High Perform. Comput. Networking (2017) (in press)
2. Almeida, J., Ferreira, J., Oliveira, A.S.R.: An RSU replication scheme for dependable wireless vehicular networks. In: 12th European Dependable Computing Conference (EDCC 2016), Gothenburg, Sweden, 5-9 September 2016, pp. 229–240 (2016). doi:10.1109/EDCC.2016.11
3. Almeida, J., Ferreira, J., Oliveira, A.S.R.: Fault tolerant architecture for infrastructure based vehicular networks. In: Alam, M., Ferreira, J., Fonseca, J. (eds.) Intelligent Transportation Systems. SSDC, vol. 52, Chap. 8, pp. 169–194. Springer, Cham (2016)
4. Bernstein, P.A., Hadzilacos, V., Goodman, N.: Concurrency Control, Recovery in Database Systems. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (1987). ISBN: 0-201-10715-5
5. Böhm, A., Jonsson, M.: Real-time communication support for cooperative, infrastructure-based traffic safety applications. Int. J. Veh. Technol. **2011**, 1–17 (2011)
6. Bilstrup, K., et al.: On the ability of the 802.11p MAC method, STDMA to support real-time vehicle-to-vehicle communication. EURASIP J. Wireless Commun. Networking **2009**(1), 1–13 (2009)
7. BMW Group. BMW ConnectedDrive at the IFA 2015 consumer electronics show in Berlin. https://www.press.bmwgroup.com/global/article/detail/T0232769EN/bmwconnecteddrive-at-the-ifa-2015-consumer-electronics-show-inberlin. Accessed 28 May 2016
8. Gerla, M., et al.: Internet of vehicles: from intelligent grid to autonomous cars and vehicular clouds. In: IEEE World Forum on Internet of Things (WF-IoT), pp. 241–246 (2014)
9. IEEE Standard for Information Technology - Telecommunications, information exchange between systems Local, metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. In: IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007), pp. 1–2793 (2012)
10. Kloiber, B., et al.: Performance of CAM based safety applications using ITS-G5A MAC in high dense scenarios. In: Intelligent Vehicles Symposium (IV), pp. 654–660. IEEE (2011)
11. Meireles, T., Fonseca, J., Ferreira, J.: The case for wireless vehicular communications supported by roadside infrastructure. In: Perallos, A., et al. (eds.) Intelligent Transportation Systems Technologies and Applications, pp. 57–82 (2015)
12. Skeen, D., Stonebraker, M.: A formal model of crash recovery in a distributed system. IEEE Trans. Softw. Eng. **SE–9**(3), 219–228 (1983)