DOI: http://dx.doi.org/10.14236/ewic/ICS2015.15

A System to Recognize Intruders in Controller Area Network (CAN)

Paul Carsten, Todd R. Andel, Mark Yampolskiy, Jeffrey T. McDonald, Samuel Russ
University of South Alabama
{pwc1221}@jagmail.southalabama.edu, {tandel, yampolskiy, jtmcdonald, sruss}@southalabama.edu

The existing automotive Controller Area Network (CAN) is vulnerable. The absence of sender verification in its communication means that an attacker acting as a node on its network will be treated like a legitimate node that has always been present, allowing attackers to transmit and receive messages freely. This paper proposes a system that will allow nodes using CAN to identify legitimate messages and reject those that come from illegitimate nodes. In addition, the system provides resistance against replay attacks that can clog up the safety-critical real time networks.

Keywords: controller area network (CAN), automotive security, security mechanism, network security

1. INTRODUCTION

Over the past 20 years, Controller Area Network (CAN) has become well-known as a robust, error-tolerant, and efficient protocol for in-vehicle communications. Despite its many safety features, CAN was not designed with security in mind, since automobiles were assumed to be isolated systems. However, modern in-vehicle networks are threatened by attackers with goals of stealing information, distorting information, or providing false information and potentially rendering the vehicles unsafe to occupants and their surroundings.

We propose a system by which CAN might detect external threats and distinguish between messages coming from within the vehicle's legitimate ECUs from those coming from unauthorized sources. We also seek to create a system that could resist replay attacks and thus resist the dangers of real-time network resource depletion.

2. BACKGROUND

2.1 CAN Behavior

Controller Area Network (CAN) is one of the most widely used of all modern automotive communication protocols. Originally designed in the mid-1980s by BOSCH, CAN operates as a broadcast network. Packets created by a node are distributed to all areas of the network and the receiving nodes determine for themselves whether to use or discard the packets. The rationale for this behavior is to keep the information consistent across the network (Texas Instruments-2008). The structure of a standard CAN packet is shown in Figure 1.

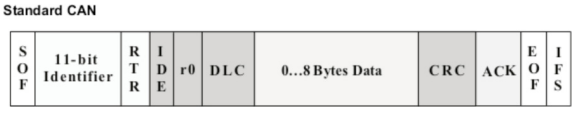

Figure 1 *Standard CAN Packet (Texas Instruments-2008)*

In addition to the standard length, there is an extended version which features a 29-bit identifier field which allows for a wider variety of priorities.

The individual nodes decide whether or not to process the packets by using the identifier (CAN ID) which identifies a message type.

2.2 CAN Vulnerabilities

Several potential vulnerabilities exist in CAN implementations. Receiving nodes decide to process a packet based on the CAN ID and all of the nodes that can operate on a packet will do so. Therefore individual nodes cannot tell if a packet is designed for them, nor do they know which node that packet came from. This results in another kind of vulnerability, namely CAN's inability to identify whether nodes are legitimate. CAN does not feature any sort of authentication of its nodes. Therefore an attacker could easily spoof messages through a compromised node. All of these concerns result in a system that is not only insecure but also ill-equipped to identify threats.

2.3 CAN Attack Surface

There are a surprising number of ways for an attacker to infiltrate a vehicle's network. The most straightforward approach is through the On-Board Diagnostics (OBD-II) port. It provides direct access to the CAN bus and can be used for gathering diagnostic packets (Checkoway-2011). However, if an attacker were to interact with the port, they would be treated as a node on the network. This would allow them to gather information and potentially transmit messages as well. Alternately the attacker could install some sort of remote device in the port and attack the vehicle remotely.

Recent trends in automobile manufacturers have moved towards the addition of wireless communications in vehicles. Systems like GM's OnStar, Toyota's SafetyConnect, and BMW's BMW Assist (Checkoway-2011) allow vehicles to communicate via broadband services. These systems provide new routes an attacker can use to interact with a vehicle.

3. GOALS

The purpose of our research is to produce some functional solution to the CAN's lack of security. At the heart of CAN's vulnerability is the protocol's inability to identify the source of messages.

In developing a solution, we wanted to come up with a method that could be compatible with the current model of CAN. We chose to design a mechanism that will allow CAN to recognize ECUs that are a part of the local vehicle's hardware (legitimate ECUs) from external attackers (malicious ECUs). In addition, we aim to create a system that will allow CAN to resist replay attacks by creating a structure that differentiates one frame from the next.

In order to achieve these goals, we set about creating a system that will allow CAN to recognize intruders and resist replay through the addition of a new field which will be added to the existing CAN frame. The mechanism will produce new values designed to be included in this field. One of these values is a hash value generated from the data included in the frame and a secret value. This hash value will act as an authentication value which can identify messages originating from the components of the vehicle present at its manufacture. The other value will serve as a time stamp which will allow the system to recognize replay attack activity.

The currently used CAN frames lack the capacity to accommodate the additional fields we are proposing. Therefore we have decided to use the new CAN FD (CAN with Flexible Data-Rate) that was recently proposed by BOSCH. This protocol features a variety of improvements over the CAN protocol, including a flexible data transmission rate and an extended data field. Whereas the traditional CAN format only allows for a maximum data length of 8 bytes, the CAN FD protocol can support messages up to 64 bytes (Hartwich-2012). In addition, the CAN FD protocol allows for a flexible data transmission rate (from the 1 mbps available in CAN to up to 15 mbps) which will enable the system to accommodate increased message lengths.

4. DESCRIPTION OF THE SYSTEM

The central concept of the system is the use of a small output cryptographic hash function. Cryptographic hash functions are considered one-way functions since it is computationally infeasible to reproduce the original message from the hashed output. In addition, cryptographic hash functions have collision resistance, meaning that it is computationally infeasible to produce two messages that generate the same hash value.

Hash functions have applications in both protection of message authenticity and sender verification. Since reproduction of the original value from the hash is very difficult, the contents of the material used to produce the hash are also protected from modification since they can be subsequently verified.

4.1 CAN Identification Number

The initial step of the system begins at the construction of the vehicle. During the process of manufacture of the vehicle, a particular value is associated with each vehicle. This value is unique to the particular vehicle and can be used as identification of that vehicle (similar to the Vehicle Identification Number (VIN) that is used in modern vehicles). Each of the ECUs in that particular vehicle will have this CAN Identification Number (CIN) hard-coded securely into their system. Access to this particular value will serve as an indication that the ECU was present at the construction of the vehicle. Maintenance procedures would have to allow for new authorized ECUs to be added using the original CIN from the manufacturer.

4.2 Structure of the Protected Frame

The mechanism will produce a new field called the CAN Message Authenticator (CMA) which will contain both the hash value and the timing information which will be used to generate the hash. The CMA will be used to verify the legitimacy of the message by a legitimate ECU upon receipt of a frame. This new frame will be positioned at the end of the current data frame.

4.3 Preparation of the Message Frame

During operation of the vehicle, an ECU will be required to perform the hashing algorithm anytime it wants to send a message. We have devised two different methods with which to obtain the hash, shown in Figure 2, which will be included in the new field along with the final frame to be sent.

Using a Universal Clock
The first version works on the assumption that the CAN system has access to a universal clock that can be monitored by all elements of the system. The clocking mechanism itself could be as simple as a monotonically increasing counter value, available to all ECUs. The initial message (IM) is concatenated (II) with the timing value (TV) and the CIN. This is then fed into the hash function (H), which produces a value which is dependent on all three elements. This value is then concatenated with the TV to produce the CMA. By including the current time at the point that the message was sent, the receiving ECU can replicate the hash (and thus verify its authenticity) on the assumption that it already has access to the CIN. The collision resistance present in cryptographic hash functions guarantees that only ECUs that have access to the CIN will be able to produce messages that hash properly.

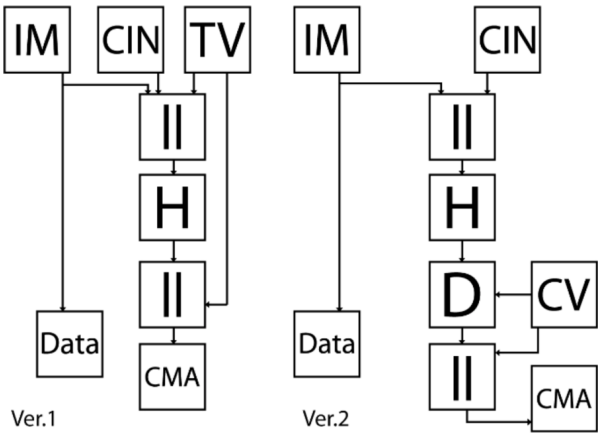

***Figure 2** Hashing Models*

Using Hash Division
The second version assumes that the universal clock is not present. Instead, each ECU will maintain a counter value (CV) which will be used for a hash block selection. The message and CIN are first concatenated (II) together and fed into the hash (H). The hash value that is produced is divided (D) into smaller sections. One of these subsections is selected based on the CV. The reduced hash and counter value are then concatenated resulting in the CMA. This method allows for the use of a smaller value for authentication and is more space efficient.

In both methods the frame is broadcast to every ECU on the network. In order to prevent any unnecessary hash checking, the ECUs will first check the CAN ID to determine if the message is relevant to the particular ECU. If an ECU is able to process the frame, it will first extract the authentication materials from the protected field. Authenticity is verified by using the message, the time value, and the CIN which was hard-coded at the vehicle's creation. If the hashes match, the message is considered valid and the message is processed. If not, the message is discarded.

5. BENEFITS OF THE SYSTEM

As was mentioned before, the current CAN system does not feature any mechanism to provide security for the distribution of its communications. After analysing the weaknesses of CAN, we concluded that the most basic weakness the system has is its inability to recognize intruders. This newly offered protection is the most important goal, but not the only benefit provided by the system's implementation.

5.1 Invasion Resistance

The system allows CAN to recognize invaders to its system through the use of the one-way hash function and the CIN which is unique to each vehicle. In order for a message to be accepted, the hash included must be exactly the same as the hash produced upon receipt of the message. Due to the strong avalanche effect of the hash function, it would be computationally infeasible to produce a hash that would verify without knowledge of the CIN. In addition, the CIN could not be feasibly extracted from the hash produced. This would make it very difficult for a malicious ECU to pose as a legitimate ECU.

5.2 Replay Attack Resistance

In both versions of the hash function, the individual frames sent include a value which represents a sort of time stamp for that particular frame. This value is included in the hash as well, meaning that verification of a frame's authenticity is directly dependent on this value. In addition, this value could be retained in a temporary sense in order to verify the next packet received with that ID. If the same hash is received multiple times, that would indicate that the same value for time is being sent multiple times, thus indicating the likelihood of a replay attack.

6. LIMITATIONS / REQUIREMENTS

While this methodology would achieve the goals of adding security to the currently insecure CAN protocol, it does not come without limitations. Certain additions would need to be made to the

current system for it to be feasible. Also, its use would impose limitations on the system, possibly making it inadequate as a real-time system. The first additional requirement to make the system functional is the CIN. There is currently no process in the manufacture of automobiles to add such a value. Since so much of the security comes from the authenticity of this value, its uniqueness would need to be guaranteed. This unique value would be required to be maintained by the manufacturer as well. In addition to the creation of the CIN, the ECUs would need to be able to perform the hashing functions without any significant taxation of resources.

Perhaps the most substantial additional requirement would be the time stamp value required by both versions of the hash function. Version 1 of the hash function requires the ECUs to have access to some sort of universal time structure. There is currently no universal synchronized clock value available in CAN, so for version 1 to be adopted this hardware would have to be added into each vehicle. In addition, the current time would have to be available to all ECUs at all times. Depending on the method used to accomplish this task, the network buses might become overtaxed by the additional traffic.

If version 2 was adopted, the ECUs would need to have both a method to maintain the counter and also maintain an association of recent hashes received (to compare with incoming frames to identify replay attacks). This additional functionality would further tax the resources of the already limited ECU hardware and might render the methodology infeasible.

7. CONCLUSION

CAN's vulnerability to invasion and attack is unacceptable considering its use as a safety critical real-time system. The system we have presented in this paper is designed to improve the security of what is currently an insecure system. The methodology allows for CAN to distinguish frames received from within its own system from frames received from an intruder. It also includes a timing aspect which will allow it to resist replay attacks.

We believe that this system would be a valuable addition to the current CAN model. It would allow the system to recognize threats, whether they come from within or without, and take action to resist attack. We hope that such a system might be implemented in future vehicles so that this vulnerability can finally be addressed.

ACKNOWLEDGEMENTS

This material is based in part upon work supported by the National Science Foundation under Grant No. DUE-1241675.

REFERENCES

Checkoway, S. et al. (2011) Comprehensive experimental analyses of automotive attack surfaces. In: *USENIX Security Symposium*.

Hartwich, F. (2012) CAN with flexible data-rate. In: *13th International CAN Conference (iCC2012)*. Hambach, Germany.

Hoppe, T., Kiltz, S., and Dittmann J. (2008) Security threats to automotive CAN networks–practical examples and selected short-term countermeasures. In: *Computer Safety, Reliability, and Security*. Berlin Heidelberg, Germany: Springer. 235–248.

Koscher, K., et al. (2010) Experimental security analysis of a modern automobile. In: *IEEE Symposium Security and Privacy (SP)*.

Kleberger, P., Olovsson T., and Jonsson E. (2011) Security aspects of the in-vehicle network in the connected car. In: IEEE *Intelligent Vehicles Symposium (IV)*.

Larson, U. and Nilsson D. (2008) Securing vehicles against cyber attacks. In: *Proceedings of the 4th Intelligence Research: Developing Strategies Annual Workshop on Cyber Security and Information to Meet the Cyber Security and Information Intelligence Challenges Ahead*. ACM.

Miller, C. and Valasek C. (2013) *Adventures in automotive networks and control units*. Available from http://illmatics. com/car_hacking.pdf

Pagliery, J. (2014) *Your car is a giant computer-and it can be hacked*. CNN Money. http://money.cnn.com/2014/06/01/technology/security/car-hack/index.html?hpt=hp_t1 (Accessed 2 June 2014).

Phung, P. H. and Nilsson D. K. (2010) A model for safe and secure execution of downloaded vehicle applications. In: *IET Road Transport Information and Control Conference and the ITS United Kingdom Members' Conference*.

SLOA101A (2008) *Introduction to the controller area network (CAN)*. Texas Instruments, Dallas, TX.

Wolf M., Weimerskirch A., and Paar C. (2004) Security in automotive bus systems. In: *Workshop on Embedded IT-Security in Cars*. Bochum, Germany.