CAN Transceivers with cyber security functions

NXP, the market-leading CAN transceiver manufacturer, has introduced ideas to secure the CAN lower layers that can be implemented in smart CAN transceivers.

The modern connected car with various internal and external communication interfaces, up to 150 electronic control units (ECUs) and 100 million lines of code, is a cyber-physical system rather than a simple mechanical system. One challenge of seamless connectivity to the Internet and end-user devices is the exposure of the vehicle to malicious exploitation of vulnerabilities, such as buffer overflow exploits, malware and Trojans. The connected car's potential for attack (its "attack surface") is increasing as the amount of connectivity, electronics and software continues to increase.

A common method to mitigate these risks is Defense-in-Depth (DiD). DiD is a concept in which multiple layers of security countermeasures are placed through a system to provide redundancy in the event a single security countermeasure fails or a vulnerability is successfully exploited. This is important as the attacker will need to circumvent multiple countermeasures to launch a successful attack.

The responsibility to define what level of security is required lies with the vehicle manufacturer. Current state of the art solutions are cryptographic-based with secure key exchange, authentication and possibly encryption. Cryptographic checks of message authenticity are adding message latency and requiring considerable computing power. Thus, the disruption of applying these kinds of solutions can be prohibitive or lead to only partial implementation for protecting solely a low percentage of the CAN messages in a network. NXP therefore proposes an additional layer in the DiD concept, either complementing state of the art security solutions, or as a standalone solution for less critical, low cost ECUs, providing a basic-level of protection and hack containment.

Proposed is a distributed intrusion detection methodology, based on CAN network specific parameters, like identifiers of the CAN messages and the contribution to the overall network busload of an ECU. This method helps contain network attacks like spoofing, remote frame tampering and denial of service (flooding).

The method described is implemented solely in a smart CAN transceiver, operating fully independent and isolated from the microcontroller (MCU) – providing an inherent level of security, without neither impacting the message latency nor increasing the processor load. It can be introduced into a network in a stepwise approach, without impacting other ECUs. Such smart transceivers can be provided as drop-in replacements with today's standard CAN transceivers avoiding further hardware and software changes on the ECU and do not affect the operation of other ECUs. This makes the proposed approach a fast, low-effort and highly cost-effective way to introduce a basic level of security or fortify state of the art security solutions with a last layer of defense.

Spoofing, tampering, and flooding

Spoofing a CAN-ID means that a compromised ECU attempts to use an ID that it is not intended to be send by this ECU. This can be useful to pretend to be another ECU. This technique has been used in practical attacks on modern cars, see Figure 1. Spoofing in the body and comfort domain can be become safety relevant, as sudden unexpected actions can distract the driver tremendously, e.g. set radio volume to maximum, or turn lights on/off.

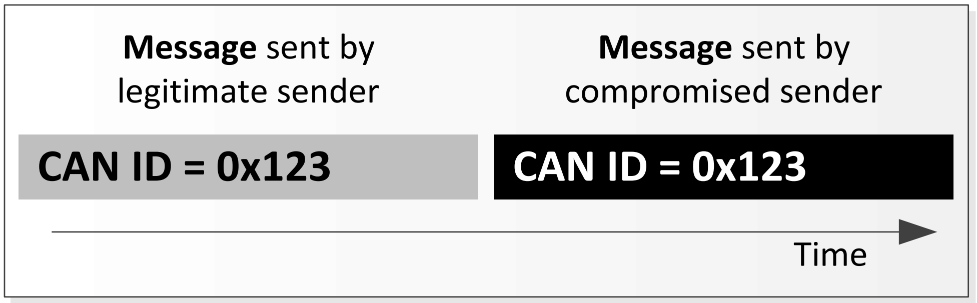
Message sent by legitimate sender
Message sent by compromised sender
CAN ID = 0x123
CAN ID = 0x123
Time

Figure 1: Spoofing attack.

For the tampering attack, the attacker aims to adjust a message, which another ECU is currently sending on the bus. The attacker must also adjust the cyclic redundancy check (CRC) to match the tampered data. Before a successful tamper attack can be accomplished, the legitimate sender must be forced into the Error-passive state, or else it will publish an active error on the bus when the attacker causes a bit flip. The attacker can put the legitimate sender in Error-passive state by intentionally publishing errors on the bus for several times. The tampering attack is useful since it gives the attacker the power to tamper with the messages that are being sent on the bus, which may be of critical operation for the car. This kind of attack has been presented at several conferences, see Figure 2. The effects in the network are like caused by spoofing.

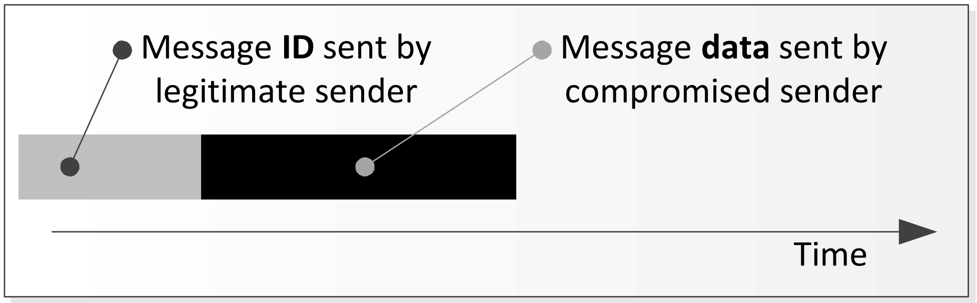
Message **ID** sent by legitimate sender
Message **data** sent by compromised sender
Time

Figure 2: Tamper attack.

Flooding the bus by continuously pumping the bus full of messages is a way to deny service, see Figure 3. This makes the bus unusable for all other ECU, which forces the entire vehicle into an emergency operating mode.

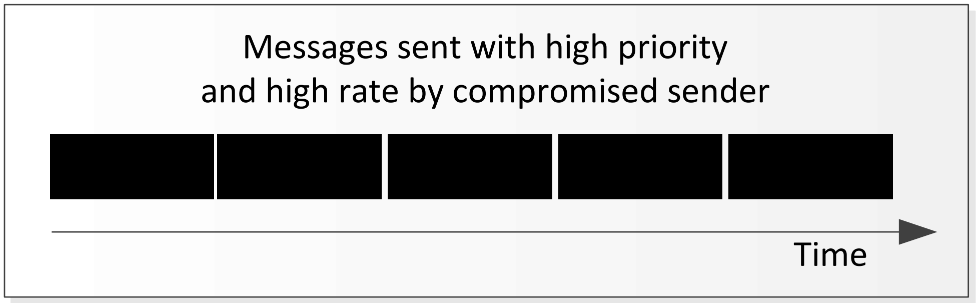
Messages sent with high priority and high rate by compromised sender
Time

Figure 3: Flooding attack.

Countermeasures

The methodology proposed by NXP can be implemented in smart CAN transceivers. All the countermeasures are based on parameters that the transceiver can perceive and are executed independently from the host, which might be compromised.

The first countermeasure, filtering messages based on CAN-IDs in the transmit path, is a way for the transceiver to protect the bus from a compromised ECU. If the ECU tries to send a message with an ID that is originally not assigned to it, the smart CAN transceiver can refuse to transmit this message on the bus by invalidating the message and deny subsequent transmissions. CAN ID-based filtering can be done using a white list of IDs that is user-configurable. For example, the IDs for Unified Diagnostic Services (UDS) as specified in ISO 14229 for off-board testers may be excluded from the whitelist. This would prevent a compromised ECU from starting a diagnostic session with another ECU in the vehicle to, for example, manipulate calibration values.

The second countermeasure against spoofing is the monitoring and invalidating messages on the bus based on the CAN-ID. This method enables every ECU to protect its own IDs in case a rouge ECU is not prevented from sending this ID; e.g. in case of an aftermarket device that is not under control of the car OEM and thus does not have a smart CAN transceiver with a configured transmission whitelist. When any ECU sends a message on the bus, the smart CAN transceiver of the legitimate ECU can actively invalidate that message by writing an active error frame to the bus. It can do this based of the same white list as the filtering in the transmit path. The compromised sender will repeat the spoofed message 16 times before Suspend-transmission behavior kicks in, limiting the bus load contribution, and finally another 16 repetitions will occur before the attacking ECU enters Bus-off state.

Preventing spoofing makes transferring a stolen cryptographic key to a rogue ECU useless, as the ECU cannot send the CAN IDs of the messages that it could authenticate with the stolen key!

Invalidating messages on the CAN network can also be used to prevent tampering. The smart CAN transceiver can check whether there was a valid message on the network, for which the local node has won arbitration, but stopped transmission (due to receiving a dominant bit while sending recessive). This is a clear sign that a compromised ECU has stepped into the transmission.

Limiting the number of transmitted messages per ECU of time can prevent flooding the network, when implemented at the sender side. In certain applications, a burst of messages on the CAN network is desirable, but this should only last for a certain amount of time. To prevent flooding, a leaky bucket mechanism can be used. In order, not to hamper diagnostic services, e.g. for uploading data, the contribution of messages with low priority IDs is neglected when filling the bucket. Flooding protection increases the availability of the network, also in case of babbling idiots.

Smart CAN transceivers with cyber security features are available

The proposed methodology is deployed on smart CAN transceivers, isolated from the host MCU and intended to be configured one-time at the Tier-1 production site and then locked, preventing future reconfiguration. An additional advantage of implementing in a CAN transceiver is it exploits the pervasiveness of the CAN transceivers in the in-vehicle network, enabling a fast and cost-effective security upgrade of existing ECUs without touching the MCU and/or software.

NXP has developed a demonstrator to prove the concept. It is based on demo silicon in an SO8 package with standard transceiver pin-out.

Bernd Elend, Georg Olma, Tony Adamson, Thierry Walrant – NXP Semiconductors