HIGHER LAYER AUTHENTICATION FOR BROADCAST IN CONTROLLER AREA NETWORKS

Bogdan Groza and Pal-Stefan Murvay

Faculty of Automatics and Computers, Politehnica University of Timisoara, Bd. V. Parvan, Timisoara, Romania

Keywords: Authentication, Broadcast, Controller area network.

Abstract: Controller Area Network (CAN) is a bus commonly used by controllers. The traditional view assumes that controllers operate in secure perimeters, but, as the degree of interconnectivity with the outside world increases, these networks may become open to intruders and CAN has no protection against Dolev-Yao adversaries. For this purpose one can implement security on higher layers. Here we design and implement a broadcast authentication protocol based on the well known paradigm of using one-way chains and time synchronization. In this way we can benefit from the use of symmetric primitives without the need of secret shared keys. As process control is a time critical operation, different to sensor networks where the life-time of the node is potentially the main limitation, here the authentication delay is the main optimization criteria. Several trade-offs are studied for this purpose in order to alleviate shortcomings on computational speed, memory, bandwidth and to assure a uniform bus-load. As for the experimental setup, we used S12 microcontrollers from Freescale to implement the proposed solution. To speed up cryptographic operations we also make use of the XGATE co-processor available on S12X.

1 INTRODUCTION AND RELATED WORK

Controller Area Network or simply CAN is a communication bus initially developed by BOSCH to be used by controller units in vehicular systems (BOSCH, 1991). The initial specifications are now superseded by the newer standard (ISO, 2003) while its area of application also extended outside vehicles to automation applications in general. Although high performance buses were developed in the last decade, such as FlexRay, because of its efficiency and reduced cost CAN is still the most commonly used communication bus in automotives today.

Traditionally, in control systems in general and in automotives in particular, reliability was a main concern but only with respect to natural phenomenons (electromagnetic disturbances, thermal noise, etc.) or accidents of various causes but not in front of Dolev-Yao adversaries. Thus CAN has very efficient mechanisms to deal with errors and to recover afterwards. In fact, the probability of an undetected error on CAN is extremely low, informally one undetected error occurs at about one thousand years for each vehicle that operates eight hours a day with an error each 0.7s. For the interested reader, an in-depth study of the performance of CAN error detection mechanism was done in (Charzinski, 1994).

However, in the last decade, industrial control systems and automotives become opened to malicious adversaries as well, and a significant part of the security community focused on this kind of issues. A recent comprehensive book for security and in particular cryptographic security in automotives is (Lemke et al., 2006) but a high amount of papers were published since then.

In this context, of malicious adversaries that can manipulate messages over the network, CAN does not have intrinsic support for any kind of security. Indeed, such kind of security is not needed if one sees CAN as operating in a secure perimeter. But, it is very likely that soon CAN-like networks will operate in environments that are opened for intruders. Recent research showed current automobiles to be unexpectedly vulnerable to external adversaries (Koscher et al., 2010) and it is likely that many other environments in which CAN operates are not completely isolated from the outside world. Security in front of such adversaries can be achieved by implementing this at the application level. In fact such improvements happened in the past, for example when deterministic delays were needed on the CAN bus with the

development of Time Triggered CAN (ISO, 2004). Still, to best of our knowledge there is no implementation available to assure authenticity in CAN networks. Thus, the main intention of this paper is to develop a higher layer implementation and to study several trade-offs to increase its efficiency. We analyze this both at a theoretical level by introducing the corresponding formalism and by designing an efficient protocol and at a practical level by following an efficient implementation. This is done on S12X microcontrollers from Freescale, a family of microcontrollers frequently found in the automotive industry, with the use of the XGATE co-processor available on S12X derivatives to speed up cryptographic functions.

As for the cryptographic mechanism that can be employed for this purpose, public-key cryptography is not the solution because of both the computational and communication overhead. As messages are short in CAN networks, usually fitting in the 64 bits of data carried by one CAN frame, using a public-key primitive such as the RSA will require thousands of bits and causes a significant overhead. Elliptic curves will significantly reduce the size of the messages, but still the computational overhead is too much to assure small authentication delays. While the computational overhead can be alleviated by dedicated circuits, such as ASICs and FPGAs, this will largely increase the cost of components, an issue that is largely avoided by manufacturers.

In contrast, symmetric primitives were efficiently employed for authentication in constrained environments such as sensor networks (Perrig et al., 2001b), (Liu and Ning, 2003), (Liu and Ning, 2004). Due to the broadcast nature of CAN, protocols similar to the well known TESLA protocol (Perrig et al., 2000), (Perrig et al., 2001a) can be used in this context as well. There is an extensive bibliography related to the TESLA protocol. Its history can be traced back to Lamport's scheme which uses one-way chains to authenticate users over an insecure network (Lamport, 1981). The work of Bergadano et al. (Bergadano et al., 2002) proposes several variants of one-way chain based protocols, with or without time synchronization. Previous work which inspired this family of protocols is the Guy Fawkes protocol from (Anderson et al., 1998). The context which is more related to our setting here is that of the application of such protocols in sensor networks. In particular, several trade-offs for sensor networks were studied in (Liu and Ning, 2003), (Liu and Ning, 2004) and several variants of the protocols are presented by Perrig as well in (Perrig et al., 2000), (Perrig et al., 2001a) .

In the case of the industrial controllers, some of the constraints are similar. For example, computational power is also low and, although high speed microcontrollers are also available on the market, low speed microcontrollers are preferred to reduce costs. But while low computational power gives some similarities, other constraints are different. For example, energy consumption is a relevant issue in sensor networks, but usually for control units inside a car this is not a main concern since they do not strongly rely on small batteries. On the other side, a different constraint here, that is not so prevalent in sensor networks, is the size of the message which is limited to 64 bits on a CAN frame. Indeed, larger messages can be split in smaller messages but the overhead inflicted by the structure of the CAN frame is around 50%. This becomes prevalent in the case of one-way chain based protocols, where hash-functions are used to compute the chain elements and thus to send an element of the chain will require at least two exchanged messages (assuming the simplest hash function outputs 128 bits). To this, one will need to add the message authentication code as well, which again requires at least two exchanged messages, etc. Thus, at least four CAN frames are needed to transmit just the security elements of one frame with useful information. Still, the most critical part, in automotive communication and control systems in general, where this protocol is mostly used, are the authentication delays, i.e., how fast a packet can be deemed as authentic. For this purpose, the most relevant constraint to which we want to give a positive answer is the authentication delay. In particular we must assure that a node, if the bus is not taken by a higher priority message, is able to transmit the message and the message can be checked for authenticity as soon as possible. This condition is initially limited by the computational power, but as checking for authenticity can happen only as soon as the disclosure delay expires and the next element of the chain is committed, this also depends on the structure of the chain which is determined by the amount of memory, and also on the bandwidth. Using too large chains means too much time in the initialization stage and large amounts of memory, while too short chains means either high authentication delays or too frequent re-initializations, etc. Depicting an optimum in this context is not straight forward and we study this in detail in what follows. In particular, we used in our scheme several levels of one way chains. While three levels of one-way chains were reported to be close to optimal in sensor networks, due to memory constraints and to reduce initialization in some situations we used more levels. This is because of both the time horizon of the protocol and of the duration of the disclosure interval. In sensor networks the disclo-

sure interval was in the order of tens or hundreds of milliseconds, while here, to increase communication speed we want to reduce this as much as possible. Of course, we are finally limited by the bus speed and by the synchronization error, which in fault-tolerant CAN will not allow us to drop the disclosure delay under several milliseconds. Practical examples are given in the experimental results section.

The paper is organized as follows. Section 2 gives an overview of the protocol, starting from several details of the CAN protocol to the examination of the influence of chain lengths, structure and timings on performance. In section 3 we present experimental results concerning the implementation of the proposed protocol on S12X microcontrollers. This includes experimental results for the implementation of cryptographic primitives on the XGATE co-processor. Section 4 holds the conclusions of our paper.

2 THE PROTOCOL

From an upper view the design paradigm is the following. Memory, computational speed and bandwidth gives bounds on the length of the chains that can be used and the number of levels. This, along with the synchronization error, further bounds the authentication delay, as messages cannot be authenticated faster than the disclosure delays. Further, to improve on the delays, we allocate equidistant timings in order to avoid a non-uniform load on the CAN bus. Indeed, since initialization packets must have the highest priority, if timings are non-uniform, there will be periods when more chains need to be initialized and thus the bus will be heavily loaded by initialization packets. Before getting to the description of the protocol we enumerate briefly in what follows some relevant aspects of the CAN bus.

2.1 Short Description of the CAN Protocol

We are not interested here in typical aspects of the protocol such as error detection, synchronization, etc., so we will not mention them. CAN bus has a broadcast nature. Nodes are connected by a two wire bus topology, as shown in figure 1, and access to the bus is gained with priority based on a message identifier which forms the first part of a frame. This identifier has 29 bits in extended frames and 11 bits in standard frames. The structure of the CAN frame consists in the arbitration field (the identifier), 6 bit control field, 0-64 bits of data, 15 bit CRC and a 2 bit acknowledgement. Additionally 1 bit marks the start of frame and

7 bits mark its end. This structure is suggested in figure 2. Few words on arbitration are in order. The way of arbitrating is by judging the winner based on the state of a particular bit, namely recessive bits (value 1) are overwritten by dominant bits (value 0). So, if the case, all nodes can start to write a message at the same moment on the bus, but, whenever a node writes a recessive bit and reads a dominant one it means that it lost the arbitration and will stop, otherwise it can continue. After each 6 consecutive bits of identical values a stuffing bit of different value is added. The body of a message can have at most 8 bytes and is followed by a 15 bit CRC. In the worst case a frame can have 154 bits out of which only 64 bits are of actual useful information. Thus, the overhead is high from the basic design of the protocol, in the worst case exceeding 50%. But this is needed to achieve reliability as mentioned before. Two kinds of CAN nodes are commonly available on the market: fault tolerant low-speed nodes which operate at 128kBps and high-speed nodes that work up to 1MBps.


Figure 1: Generic CAN Bus topology.

Figure 2: Structure of a CAN data frame.

2.2 Overview of the Protocol

The generic structure of the key chains is depicted in figure 3. We use time synchronization and multiple levels of one-way chains in order to authenticate the broadcast from a particular node. Packets arriving on the communication bus are depicted as well, the dotted lines from an element of the chain to the packet denotes that the element was used as a key, and for the re-initialization packets in particular one element of the key chain was also used as a message. The picture denotes the case of more than one sender, e.g., $P_{i,j}$ is packet sent by principal i to initialize a chain from level j. In the next sections we will use the following notations:

- L - number of chain levels,

- $\lambda_i, i = 1..L$ - length of the chain on level i,

- $\delta_i, i = 1..L$ - disclosure delay on level i,

- B - bus speed, normalized to packets per second (where one packet can carry one key),

- M - available memory (measured in elements of the key chain that can be stored),

- S - computational speed (number of elements of the key chain that can be computed per second),

- T - time horizon of the protocol (one may view this as the disclosure delay of level 0),

- t - denotes actual time and is a positive integer, a subscript indicates particular details.

Based on these notations in the next section we discuss the optimal allocation of the broadcast parameters.

In principle we need two distinct protocols: an initialization protocol and a broadcast protocol. The role of the initialization protocol is to allow each unit to commit its initialization values and to achieve time synchronization with other participants. This part of the protocol can rely on more expensive operations required by public-key cryptography. In this stage we consider that each principal will authenticate itself by using a public key certificate that is signed by a trusted authority. Initial authentication based on public-key infrastructure is important to assure composability. This ensures that different components, from potentially different manufacturers, can be assembled in one system and is a common demand of the market today. For example, in a different context (that of communication alone), the latest state-of-the-art protocol, FlexRay, has communication segments that are preallocated such that different components from different providers can be bind into a system. Thus we require that each node must store the public key of a trusted authority. Although public key certificates are larger and will require more than one frame (which can carry at most 64 bits) in general it should not be a problem to transport them over CAN if this does not happen too often and just in the initialization stage.

The initialization protocol must also ensure time synchronization. This is done with respect to a central node, which will play the role of a communication master. As usual, synchronization between two nodes is loose and it requires a handshake and counting the round trip time until it is below a tolerance margin. This is usually achieved in two protocol steps as follows: $A \rightarrow B : N_A; B \rightarrow A : Sig_B(t_B, N_A)$. Here N_A denotes a nonce generated by principal A and t_B

denotes the current time at principal B when sending its response. However, in our scenario a digital signature costs tens, or hundreds of milliseconds, which will result in a high tolerance margin that will further require an even larger disclosure delay. Because of this, instead of a digital signature we will use a message authentication code which is several orders of magnitudes faster. In particular, in our experiments, the round-trip-time was less than 2 milliseconds. Afterwards, the round trip time ε_{AB} becomes the synchronization error. If the nonce was sent by A at time t_0 and now A's clock points to t_1 then A knows that the time on B side is in the interval $[t_B + t_1 - t_0, t_B + t_1 - t_0 + \varepsilon_{AB}]$. Further, the initialization values for the chains can be shared between each node and the central node who can broadcast them to all communication participants. We will not insist on details of the initialization protocol which is done in the initial phase with no constraints.

2.3 Optimal Allocation of Key Chain Lengths and Levels

For brevity we consider a homogeneous network with nodes that have equal computational abilities, thus the same computational delays and lengths for the chains can be handled by all of them. Otherwise, the protocol can be scaled according to the computational abilities of each client, but we want to keep the model as simple as possible.

To formalize the optimal allocation of chain lengths and levels, different to previous work, we use a tolerance relation to define the lengths of the chains and the disclosure delays at each level. The tolerance relation is formed by vectors which are defined as initialization values for each communication participant.

Definition 1. We say that a set of pairs $< \lambda_i, \delta_i >$ forms tolerance relation with respect to memory, computational speed, bandwidth and time, denoted as $< \Lambda, \Delta >_{M,S,B,T}$ if the following constraints hold:

$$\sum_{i=0,L} \lambda_i = \lambda_0 + \lambda_1 + ... + \lambda_L \leq M \qquad (1)$$

$$\sum_{i=1,L} \prod_{j=0,i} \lambda_j = \lambda_0 \cdot \lambda_1 + ... + \lambda_0 \cdot \lambda_1 \cdot ... \cdot \lambda_L \leq S \cdot T \quad (2)$$

$$\sum_{i=0,L} \prod_{j=0,i} \lambda_j = \lambda_0 + \lambda_0 \cdot \lambda_1 + ... + \lambda_0 \cdot \lambda_1 \cdot ... \cdot \lambda_L \leq B \cdot T$$
$$(3)$$

Relation (1) gives the memory bound, i.e., the sum of the lengths of the chains cannot exceed the total memory. Relation (2) and (3) are bounds on computational time and bandwidth, i.e., the total number of
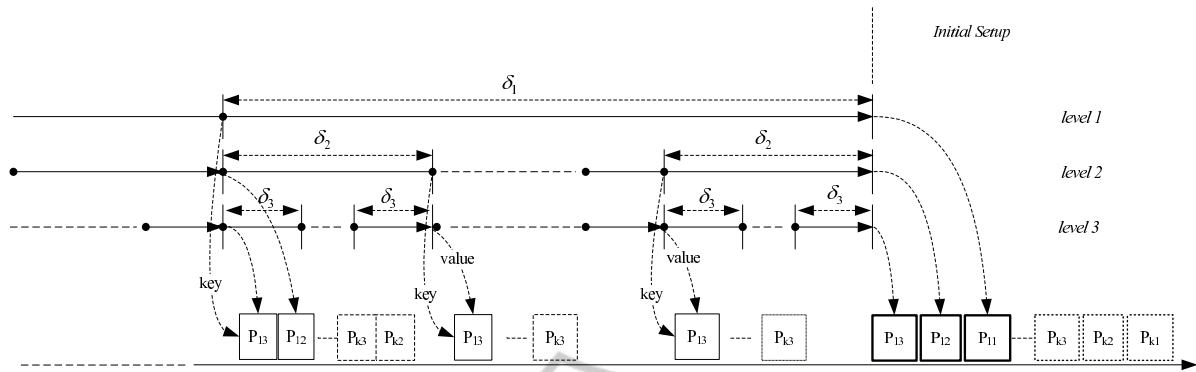
Figure 3: Basic structure of the one-way chains.

elements of the one-way chain cannot require more computational time than available over protocol lifetime T and cannot exceed the available bandwidth for transmission. Value λ_0 is missing from relation (2) as the first key chain is computed in the initialization stage, while λ_0 is present in relation (3) as the first key chain will be disclosed as well on the bus. Indeed, relations (2) and (3) can be further refined for the disclosure delays on all up to the last level since the re-initialization of each chain should be done in the disclosure delay of the previous level. Thus, we can write (2) as $\forall k > 1, \sum_{i=k,L} \prod_{j=k-1,i} \lambda_j = \leq S \cdot \delta_k$ and the same holds for bandwidth. We introduced this definition to keep our presentation formal, but it is obvious that defining good tolerance relations is a matter of good engineering.

The general relation between chain lengths and disclosure delays will now be the following:

$$\delta_i = \delta_{i-1}/\lambda_i, i > 0 \tag{4}$$

Intuitively, this also means that $\forall i, \lambda_i$ denotes the number of one-way functions computations that can be performed in time δ_{i-1} to generate a new chain. Thus λ_i can also be defined as a function of δ_{i-1}. One can use this to create sub-intervals by using the previous relation.

Remark 1. Given a tolerance relation $< \Lambda, \Delta >_{M,S,B,T}$ with respect to time, space and bandwidth the disclosure delay and the computational overhead have an inverse variation. Thus: the minimal disclosure delay is achieved if chains are of equal size, i.e., $\lambda_i = M/L$, while the minimal computational and communication overhead is achieved if upper level chains are smaller, i.e., $\lambda_0 < \lambda_1 < ... < \lambda_L$.

This is intuitively since the product of two values whose sum is fixed is maximal if the two values are equal and minimal if one of the values is 1. For example, assume $x + y = z$ then $\forall k \geq 1, z/2 \cdot z/2 > (z/2 - k) \cdot (z/2 + k) = z^2/4 - k^2$. Now, to achieve

the minimum delay, the product of the chain lengths $\lambda_0 \cdot \lambda_1 \cdot ... \cdot \lambda_L$ must be maximal. If we split this product exactly into the half left and half right terms, assuming an even number of terms which is wlog, then the maximum product is achieved if: the left and right are maximal and they are equal, and so on. To achieve minimal bandwidth and computational requirements we need $\lambda_0 \cdot \lambda_1 + ... + \lambda_0 \cdot \lambda_1 \cdot ... \cdot \lambda_L$ to be minimal. This can be written as $\lambda_0 \cdot (\lambda_1 + ... + \lambda_1 \cdot ... \cdot \lambda_L)$ and as right term cannot be equal to 1 the left term must be 1 in order to minimize the product and so on. Thus, the optimum choice of lengths with respect to delays, is to have all chains of equal size. For the purpose of generality, as well as for the fact that in practice small variations between chain lengths may occur when the amount of available memory is not directly divisible with the number of levels, we will keep the following exposition for the case when chains are of arbitrary sizes.

2.4 Optimal Allocation of Timings

By optimal allocation of the key disclosure time we want to achieve minimal delays for sending messages and authenticating them. Of course, the authentication delay is bounded by the disclosure delay, i.e., packets cannot be authenticated sooner than the disclosure delay from the last level. This bound was already fixed by the tolerance relation. However, the straight forward mechanism suggested in figure 3, in which chains are re-initialized successively, causes more overhead at the disclosure time of keys from upper layer chains (since at this time all chains from lower levels need to be re-initialized as well). To overcome this, we define a procedure which we call *equidistant timing* by which all packets are disclosed at periods of time separated by equal delays. More, different to the use in sensor networks where upper layer chains are used only to authenticate key commitments from the lower levels, we use chains from up-

per levels to authenticate information packets as well. Thus, we will normalize the disclosure time to:

$$\delta_{norm} = \frac{T}{\lambda_0 \cdot \lambda_1 \cdot \lambda_2 ... \lambda_L} \qquad (5)$$

In the forthcoming scheme, keys on all levels are released at δ_{norm} intervals. Of course, if we use relation (4), δ_{norm} is equal to δ_L but we prefer to use this notation to avoid confusion with a generic scheme in which this may not happen, i.e., not all keys are released at δ_{norm} but only the keys from the chain on level L. In what follows we need to establish three things:

- which is the disclosure time for a particular key k and its converse which follows,

- given a particular time t which key must be disclosed, or which is the last key that was disclosed upon t,

- given a particular packet, containing authentication information, what condition must be met on the receiver's side to deem this packet as on-time.

To determine all these, the easiest way to decide is by writing the time with respect to the vectors of the tolerance relation which is established by the next definition.

Definition 2. Given a discrete time value t we denote by $t_{<\Lambda,\Delta>} = < t_L t_{L-1}...t_0 >$ the discrete time decomposition of t with respect to the lengths of the chain and normalized time as a basis, i.e., the writing of the time value as:

$$\begin{aligned} t &= \delta_{norm} \cdot \sum_{i=L..1} (t_i \cdot \prod_{j=i..1} \lambda_j) + t_0 \qquad (6) \\ &= t_L \cdot \lambda_L \cdot ... \cdot \lambda_1 \cdot \delta_{norm} + ... + t_1 \cdot \lambda_1 \cdot \delta_{norm} + t_0 \end{aligned}$$

2.4.1 Sender's Perspective.

For the moment we consider the case of a single sender. Let t_{start} denote the time at which the broadcast was started and assume that it is started by the communication master which also gives time synchronization. Thus t_{start} is the exact time at which the broadcast started (no drifts for the sender).

Definition 3. Let $\tau \in \Lambda^{i,j}$ denote that τ is a vector of $j - i + 1$ positive integers such that each element on position i is less or equal to λ_i. Given a tolerance relation $< \Lambda, \Delta >_{M,S,B,T}$, an indexed key K_τ is a key entailed by a vector $\tau \in \Lambda^{1..L}$. An indexed key chain $\mathcal{K}_{<\Lambda,\Delta>}$ is a collection of indexed keys, derived as follows: having a fixed fresh seed K_{master}, a key derivation process $\mathcal{KD}$ and a one-way function $\mathcal{F}$, then:

$$\forall i \in [1,L], \tau_i \in [1,\lambda_i], \tau_{\text{left}} \in \Lambda^{i+1,L}, \tau_{\text{right}} \in \Lambda^{1,i-1},$$
$$K_{\tau_{\text{left}}|0|\tau_{\text{right}}} = \mathcal{KD}(K_{master}, \tau_{\text{left}}, \tau_{\text{right}}),$$
$$K_{\tau_{\text{left}}|\tau_i|\tau_{\text{right}}} = \mathcal{F}(K_{\tau_{\text{left}}|\tau_i-1|\tau_{\text{right}}}) \quad (7)$$

Definition 4. Let $\mathcal{DT}(K_\tau)$ denote the disclosure time of the indexed key K_τ. Given a broadcast started at t_{start} the disclosure time of this key is:

$$\mathcal{DT}(K_\tau) = t_{start} + \delta_{norm} \cdot \sum_{i=L..1} (\tau_i \cdot \prod_{j=i..1} \lambda_j) \quad (8)$$

Definition 3 allows the generation of chains with the structure from figure 3 while definition 4 allows keys to be released at equal time intervals δ_{norm}.

Remark 2. The key released by the sender at time $t_{current}$ given a broadcast started at t_{start} is $K_{t_{<\Lambda,\Delta>}/t_0}$ where $t = t_{current} - t_{start}$. By $t_{<\Lambda,\Delta>}/t_0$ we denote all elements from t except the last term, i.e., t_0. This means that

- the key is from level $L - l$, where l is the first non-zero index from right to left, i.e., the maximum value such that $\forall i \leq l, t_i = 0$,

- the position of the key on the key chain from level l is t_l,

- the number of chains previously released at this level is j is $t_{l+1} \cdot \lambda_{l-1} + t_{l+2} \cdot \lambda_{l-1} \lambda_{l-2} + ... + t_L \cdot \prod_{i=l..L} \lambda_i$.

2.4.2 Reinitialization Packets and Efficiency

To avoid a non-uniform bus load, as discussed previously, re-initialization packets will be dropped *equidistantly* in the δ_{norm} intervals. Otherwise, packets carrying data must be delayed until all re-initialization packets are sent. This is because re-initialization packets must have priority on the bus, otherwise the protocol will succumb an will require passing through the initialization stage again which is even more expensive. Thus we can also use as an efficiency criteria the maximum delay until a new CAN frame can be send. For the basic scheme the maximum delay fluctuates with the number of initialization packets. This delay can be easily established for the basic scheme. Let $z_i(x)$ denote the number of consecutive zeros in vector x starting from position 1. At discrete time value t, given $t_{<\Lambda,\Delta>}$ the delay until the next packet can be sent is:

$$delay = \max[0, z((t - t_{start})_{<\Lambda,\Delta>})] \quad (9)$$

Indeed, the number of consecutive zeros at the end of the time value denotes how many chains need to be initialized at that particular time. This value becomes constant for the equidistant scheme and data packets

are delayed by at most one packet.

To complete the view on efficiency, we should also define the overhead induced by the authentication mechanism. The overhead has two distinct components, the *authentication overhead* which is the overhead inflicted by the authentication keys that are released on the bus, and the *re-initialization overhead* which is the overhead inflicted by the re-initialization material. Indeed, to this one may want to add the overhead induced by the message authentication codes, MACs, associated to each data packet that is send over the bus. We will not take this into account because this is application dependent, not protocol dependent, indeed in some applications the size of the data packets can be small, and thus adding a MAC to each data packet will greatly increase the overhead. In other applications it may be the reverse, and data packets can be large and the MAC will not significantly increase the overhead. Based on these we can also define the total overhead inflicted by the protocol.

Remark 3. Given a tolerance relation $< \Lambda, \Delta >_{M,S,B,T}$ the authentication and re-initialization overheads and the bus load induced by the protocol is:

$$OH_{auth} = (\lambda_0 + \lambda_0 \cdot \lambda_1 + \dots + \lambda_0 \cdot \lambda_1 \cdot \dots \cdot \lambda_L) \cdot B^{-1}$$
(10)

$$OH_{reinit} = (\lambda_0 \cdot \lambda_1 + \dots + \lambda_0 \cdot \lambda_1 \cdot \dots \cdot \lambda_L) \cdot B^{-1} \quad (11)$$

$$OH_{total} = OH_{auth} + OH_{reinit} \quad (12)$$

One may note that if the authentication delay is bigger the overhead also becomes lower since fewer elements of the chain are sent. Also, the re-initialization overhead increases with the number of levels. We give concrete examples for these in the experimental results.

Figure 4 shows the influence of chain length on bus overhead and disclosure delays. Plots are given for three and four levels of key chains. We note that the delays drop rapidly by increasing the number of levels, but in the same manner the overhead increases (at 100% the bus is locked and communication halted). Figure 5 shows the variation of memory requirements, which is the same as the initialization time, and of the number of re-initialization packets with the number of chain levels. All plots are taken for a time range of 24 hours, in figure 5 the delay is fixed to 5 ms.

2.4.3 Receiver's Perspective

We consider the case of a sender $\mathcal{S}$ with synchronization error ε_S and a receiver $\mathcal{R}$ with synchronization
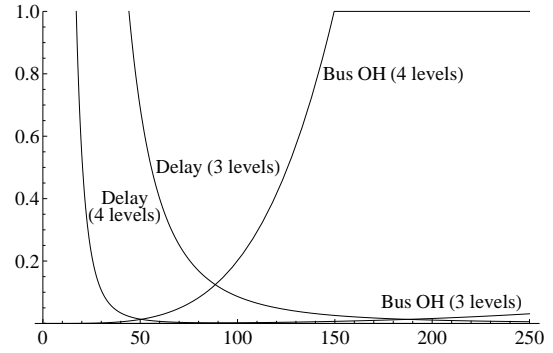

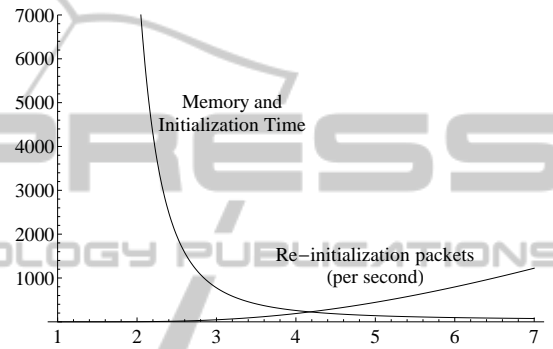Figure 4: Delay and overhead variation with chain length.

Figure 5: Memory, initialization time and re-initialization packets variation with number of levels.

error $\varepsilon_{\mathcal{R}}$. Now we define the security condition that must be met by all packets that contain authentication information, i.e., MAC codes, produced with an indexed key K_τ.

Definition 5. Given synchronization errors ε_S and $\varepsilon_{\mathcal{R}}$ for sender and receiver, an authentication packet indexed by τ must discarded unless:

$$\mathcal{T}_{rec}(P_\tau) \leq \mathcal{D}\mathcal{T}(\tau) - \varepsilon_S - \varepsilon_{\mathcal{R}} \quad (13)$$

Here $\mathcal{T}_{rec}(P_\tau)$ denotes the time at which packet P_τ was received. Here we assume that the time on the receiver side is common to that on the sender's side, since this is after the synchronization. Thus $\mathcal{T}_{rec}(P_\tau)$ is the minimum time value on the sender's side, while the synchronization errors on the right side ensure a safe margin for the maximum time value.

2.4.4 Protocol Description

We can now summarize previous notions in one definition for the entire protocol.

Definition 6. Given tolerance relation $< \Lambda, \Delta >_{M,S,B,T}$, an indexed key chain $\mathcal{K}_{<\Lambda,\Delta>}$ and the two roles sender and receiver denoted by $\mathcal{S}$, $\mathcal{R}$ each with synchronization errors ε_S, $\varepsilon_{\mathcal{R}}$ with respect to a common clock, protocol Broadcast$[\mathcal{S}, \mathcal{R}, < \Lambda, \Delta >_{M,S,B,T}$

$, \mathcal{K}_{<\Lambda,\Delta>}]$ is formed by the following two rules for the two roles:

1. $\mathcal{S}$ sends K_τ at $\mathcal{DT}(K_\tau)$ and the message M and its corresponding $MAC(K_\tau, M)$ in any empty time-slot with the condition that $MAC(K_\tau, M)$ is released no latter than $\mathcal{DT}(K_\tau) + \xi$,

2. $\mathcal{R}$ discards all $MAC(K_\tau, M)$ received later than $\mathcal{DT}(K_\tau) + \varepsilon_\mathcal{S} + \varepsilon_\mathcal{R}$ and deems authentic all other messages for which $MAC(K_\tau, M)$ can be verified with an authentic key. A key $K_{\tau_{\text{left}}|\tau_i|\tau_{\text{right}}}$ is authentic only if $K_{\tau_{\text{left}}|\tau_i|\tau_{\text{right}}} = \mathcal{F}(K_{\tau_{\text{left}}|\tau_i-1|\tau_{\text{right}}})$ and $K_{\tau_{\text{left}}|\tau_i-1|\tau_{\text{right}}}$ is a previously received/computed authentic key.

Here ξ denotes a tolerance margin until the sender can send a MAC with a particular key. Indeed, sending MACs too close to the disclosure time may be useless because the receiver may have to discard them if the security condition cannot be met. Thus ξ must be fixed as initial parameter for the protocol and it must hold that $\varepsilon_\mathcal{S} + \varepsilon_\mathcal{R} << \xi$. In time interval $[\mathcal{DT}(K_\tau), \mathcal{DT}(K_\tau) + \xi]$ the sender can safely disclose any kind of data packet, but not MACs.

$\text{Broadcast}[\mathcal{S}, \mathcal{R}, <\Lambda, \Delta>_{M,S,B,T}, \mathcal{K}_{<\Lambda,\Delta>}]$ is a secure broadcast authentication protocol. The security of this family of protocols is well established, the informal argument is that the adversary cannot construct $MAC(K_i, M)$, until K_i is released. As function $\mathcal{F}$ is one-way he cannot derive K_i from $\mathcal{F}(K_i)$ and by the time K_i is released any MAC with K_i that is further received will be discarded. Formal proofs for such protocols can be found in (Perrig et al., 2001a), (Bergadano et al., 2002).

2.5 The Case of Multiple Senders

The case of k senders can now be easily derived from the previous formalism. To preserve the *equidistant* time schedule we use the nominal disclosure time δ_{norm} and divide it by the number of senders k. Let us define the next sender delay as:

$$\delta_{next} = \frac{\delta_{norm}}{k} \qquad (14)$$

Definition 7. Let $\mathcal{DT}(K_\tau)$ denote the disclosure time of an indexed key by the by the k^{th} sender. Given a broadcast started at t_{start} we have:

$$\mathcal{DT}(K_\tau, k) = \mathcal{DT}(K_\tau) + k \cdot \delta_{next} \qquad (15)$$

The security conditions which has to be verified by receivers must also add the $k \cdot \delta_{next}$ term to the disclosure time of the k-th sender.

3 APPLICATION SETTING AND EXPERIMENTAL RESULTS

As stated, for the implementation of the previously described protocol, we used a Freescale 16-bit automotive grade microcontroller (MC9S12XDT512) from the S12X family on a SofTec Microsystems ZK-S12-B development board. One special feature of this family is the presence of an incorporated co-processor called XGATE which can be used to increase the microcontroller's data throughput (Freescale, 2004). We made use of this module to increase the efficiency of our implementation by assigning it the task of computing the underlying cryptographic functions.

The S12X microcontrollers used in our experiments have 512Kbytes of FLASH memory and 20Kbytes of RAM. Both FLASH and RAM memories are banked as a consequence of the 16 bits wide address bus which is not sufficient to access all memory locations. Thus, a total of 8Kbytes of RAM space can be used for continuous allocation while the rest of the RAM can be accessed in 4Kbyte windows. The amount of RAM memory that can be used for storing key chains is relevant as it determines the maximum number of chain levels and their lengths.

The maximum bus frequency that can be set using the PLL module is, according to the data-sheet, 40MHz. We configured the PLL for frequencies beyond this specified value and were able to go up to speeds of 80MHz. After assessing that the behaviour of the microcontroller at this overclocked frequency is normal we used it in our tests and compared the results with the ones obtained for 40MHz.

The on-chip CAN module implements version 2.0A/B of the CAN protocol and supports a programmable bit-rate up to 1 MBps. A limitation for the maximum achievable CAN speed comes from the on board low speed fault tolerant transceiver which can only run at speeds up to 125Kbaud.

3.1 XGATE Module

The XGATE module has a built in RISC core with instructions for data transfers, bit manipulations and basic arithmetic operations. The RISC core can access the internal memories and peripherals of the microcontroller without blocking them from the main S12X CPU. The S12X CPU always has priority when the two CPUs access the same resource at the same time. To assure data consistency, the access priority can be controlled by using the hardware semaphores available on the microcontroller.

Interrupts can be routed to the XGATE module in order to decrease the interrupt load of the main S12X

CPU. Additionally, up to 8 software triggered channels can be used by the S12X CPU to request software execution on XGATE.

In order to obtain the maximum XGATE CPU speed, the code can be executed from RAM. Because RAM is a volatile memory, XGATE code is being stored into the FLASH memory and copied into RAM after each reset. Having a better RAM access speed and a speed-optimized instruction set, a typical function can run up to 4.5 times faster on XGATE than on the S12X CPU (Mitchell, 2004). Because it was designed for quick execution of small code requested by interrupts, the flash memory available for storing XGATE code is limited. For controllers in the S12XD family this limit is 30 Kbytes.

3.2 Implementation Details

To decrease the communication overhead that could be introduced by computing cryptographic primitives we assigned this task to the XGATE module. In order to evaluate the computational performance of the microcontroller we measured the execution speed of three hash functions: MD5, SHA1 and SHA-256. Measurements were done on S12X and XGATE for different input lengths using both the maximum specified frequency and the overclocked frequency. The measurements presented in tables 1, 2 and 3 show that on average the hash functions were performed approximately 2.12 times faster on XGATE than on S12X. The overclocking also increases the speed with a factor of 2.

Table 1: Performance of S12X and XGATE in computing MD5.

| Length (bytes) | Execution time | | | |
| | @ 40MHz | | @ 80MHz | |
	S12X	XGATE	S12X	XGATE
0	$732\mu s$	$312.5\mu s$	$371\mu s$	$156.2\mu s$
1	$736\mu s$	$314.5\mu s$	$373\mu s$	$157.4\mu s$
3	$737\mu s$	$314.5\mu s$	$373\mu s$	$157.2\mu s$
14	$738\mu s$	$316.0\mu s$	$374\mu s$	$158\mu s$
26	$739\mu s$	$317.5\mu s$	$374.5\mu s$	$158.9\mu s$
62	$1414\mu s$	$605\mu s$	$717\mu s$	$303\mu s$
80	$1374\mu s$	$592\mu s$	$697\mu s$	$296.5\mu s$

We chose MD5 for building the one-way chains which are needed by the protocol. Due to the small disclosure delay we consider that using MD5 is safe for our scenario. Each chain element will thus be a 128 bit value which is used to perform an HMAC over the message that has be sent at a certain point in time.

All cryptographic computations are done on the XGATE module following a software request. For passing data between the two processing units, a com-

Table 2: Performance of S12X and XGATE in computing SHA1.

| Length (bytes) | Execution time | | | |
| | @ 40MHz | | @ 80MHz | |
	S12X	XGATE	S12X	XGATE
0	$2.285ms$	$1.000ms$	$1.144ms$	$500\mu s$
1	$2.290ms$	$1.002ms$	$1.146ms$	$501\mu s$
3	$2.290ms$	$1.002ms$	$1.146ms$	$502\mu s$
14	$2.290ms$	$1.004ms$	$1.148ms$	$502\mu s$
26	$2.295ms$	$1.004ms$	$1.148ms$	$503\mu s$
62	$4.510ms$	$1.976ms$	$2.255ms$	$988\mu s$
80	$4.470ms$	$1.962ms$	$2.235ms$	$982\mu s$

Table 3: Performance of S12X and XGATE in computing SHA-256.

| Length (bytes) | Execution time | | | |
| | @ 40MHz | | @ 80MHz | |
	S12X	XGATE	S12X	XGATE
0	$5.51ms$	$3.155ms$	$2.755ms$	$1.578ms$
1	$5.52ms$	$3.155ms$	$2.760ms$	$1.580ms$
3	$5.52ms$	$3.155ms$	$2.760ms$	$1.580ms$
14	$5.50ms$	$3.150ms$	$2.755ms$	$1.578ms$
26	$5.49ms$	$3.145ms$	$2.750ms$	$1.578ms$
62	$10.86ms$	$6.24ms$	$5.44ms$	$3.125ms$
80	$10.80ms$	$6.22ms$	$5.40ms$	$3.115ms$

mon memory area is used. Each time a hash needs to be computed, the S12X writes the input data in the common memory area and makes a software request to the XGATE module. While the hash is computed on the XGATE side, the main CPU is free to execute other tasks such as, receiving messages or sending messages that have been already built. The XGATE module can signal the end of the function execution by issuing an interrupt to the S12X CPU.

After protocol implementation, the total RAM memory left for storing key chains can hold 1216 elements (16 bytes each). Having this upper limit for M, L and λ have to be determined for best performances based on the bus speed and the desired disclosure delay. If we consider packets of 16 bytes in size, the measured bus speed for sending packets is 578 packets/second (one packet each $1.730ms$).

For example, if we decide to use three levels to assure authentication over a period of one day with a speed of 578 packets/second, we would have 233 elements on each level and the key disclosure time δ would be 6.8ms. The bus overhead for this situation is 25,2% and the time needed to initialize the key chains is approximately 700ms. Increasing the number of levels to 5 leads to chains of 26 elements so less memory is necessary and the time needed for initialization is reduced to 131ms. The cost of these improvements is a bus overhead of 26.9%. The bus load grows exponentially with the increase in the number

of levels used while the disclosure delay depends on the transmission speed and the total communication duration. The duration of the communication also affects the number of elements on each level which is upper bounded to M/L ($1216/L$ in our setting).

To further improve performance, using small scale variants of hash functions, as proposed in (Macchetti and Rivard, 2005) can be an interesting alternative. However, as collisions on these functions were already reported in (Steurer, 2006), a more in depth analysis is needed before using them in practice and we leave this as potential future work.

4 CONCLUSIONS

A protocol for assuring broadcast authenticity on the CAN bus was provided. We studied different trade-offs in order to depict the optimal choice of parameters. In particular we concluded that the main limitation is the bus speed (limited to 128kBps in fault-tolerant CAN), followed by memory and last by computational power. This is also due to the performance of the XGATE co-processor which is about two times faster than the regular S12 processor. In some cases, to reduce memory consumption and to shorten the initialization time, chains with more than three levels were also efficient. The theoretical estimations are entailed by experimental results on the S12X processor, a commonly used automotive grade microcontroller. As future work we plan to extend this paper with experimental results on other microcontrollers and to use a real-world communication scenario for an automotive network. By this research we hope that we give a first analysis on the feasibility of such a solution.

ACKNOWLEDGEMENTS

REFERENCES

Anderson, R., Bergadano, F., Crispo, B., Lee, J.-H., Manifavas, C., and Needham, R. (1998). A new family of authentication protocols. *SIGOPS Oper. Syst. Rev.*, 32:9–20.

Bergadano, F., Cavagnino, D., and Crispo, B. (2002). Individual authentication in multiparty communications. *Computers & Security*, 21(8):719 – 735.

BOSCH (1991). *CAN Specification Version 2.0.* Robert BOSCH GmbH.

Charzinski, J. (1994). Performance of the error detection mechanisms in can. In *Proceedings of the 1st International CAN Conference*, pages 20–29.

Freescale (2004). *MC9S12XDP512 Data Sheet, Rev. 2.21, October 2009.* Freescale.

ISO (2003). *ISO 11898-1. Road vehicles - Controller area network (CAN) - Part 1: Controller area network data link layer and medium access control.* International Organization for Standardization.

ISO (2004). *ISO 11898-4. Road vehicles - Controller area network (CAN) - Part 4: Time triggered communication.* International Organization for Standardization.

Koscher, K., Czeskis, A., Roesner, F., Patel, S., Kohno, T., Checkoway, S., McCoy, D., Kantor, B., Anderson, D., Shacham, H., and Savage, S. (2010). Experimental security analysis of a modern automobile. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 447 –462.

Lamport, L. (1981). Password authentication with insecure communication. *Commun. ACM*, 24:770–772.

Lemke, K., Paar, C., and Wolf, M. (2006). *Embedded Security in Cars Securing Current and Future Automotive IT Applications.* Springer Verlag.

Liu, D. and Ning, P. (2003). Efficient distribution of key chain commitments for broadcast authentication in distributed sensor networks. In *Proc. of the 10th Annual Network and Distributed System Security Symposium*, pages 263–276.

Liu, D. and Ning, P. (2004). Multilevel μtesla: Broadcast authentication for distributed sensor networks. *ACM Trans. Embed. Comput. Syst.*, 3:800–836.

Macchetti, M. and Rivard, P. (2005). Small-scale variants of the secure hash standard. In *ECRYPT workshop on RFID and lightweight cryptography*.

Mitchell, R. (2004). *Tutorial: Introducing the XGATE Module to Consumer and Industrial Application Developers, March 2006.* Freescale.

Perrig, A., Canetti, R., Song, D., and Tygar, J. D. (2001a). Efficient and secure source authentication for multicast. In *Network and Distributed System Security Symposium, NDSS '01*, pages 35–46.

Perrig, A., Canetti, R., Song, D., and Tygar, J. D. (2001b). Spins: Security protocols for sensor networks. In *Seventh Annual ACM International Conference on Mobile Computing and Networks (MobiCom 2001)*, pages 189–199.

Perrig, A., Canetti, R., Tygar, J., and Song, D. X. (2000). Efficient authentication and signing of multicast streams over lossy channels. In *IEEE Symposium on Security and Privacy*, pages 56–73.

Steurer, M. E. (2006). Multicollision attacks on iterated hash functions. Technical report.