Portofoliu 10 lucrari reprezentative în domeniul de doctorat vizat:

Calculatoare si Tehnologia Informaţiei

Conf. Dr. Ing. Bogdan Ioan Groza

[1] Bogdan Groza, Bogdan Warinschi, "Client puzzles and DoS resilience, Revisited", ***Designs Codes and Cryptography***, Springer-Verlag, vol. 73, pp. 177-207, 2014.

[2] Bogdan Groza, Stefan Murvay, "Efficient Protocols for Secure Broadcast in Controller Area Networks", ***IEEE Transactions on Industrial Informatics***, vol.9, no.4, pp. 2034-2042, Nov. 2013.

[3] Bogdan Groza and Marius Minea. Bridging Dolev-Yao adversaries and control systems with time-sensitive channels. *In Conference on Critical Information Infrastructures Security (CRITIS)*. **Springer-Verlag**, **LNCS**, vol. 8328, pp. 167-178, 2013.

[4] Bogdan Groza, Stefan Murvay, Anthony van Herrewege, Ingrid Verbauwhede, LiBrA-CAN: a Lightweight Broadcast Authentication protocol for Controller Area Networks, *Proc. 11th International Conference on Cryptology and Network Security (CANS'12)*, **Springer-Verlag**, **LNCS**, vol. 7712, pp. 185-200, 2012. (Rank B).

[5] Bogdan Groza, Rene Mayrhofer, SAPHE - Simple Accelerometer based wireless Pairing with HEuristic trees, *Proc. 10th International Conference on Advances in Mobile Computing and Multimedia (MoMM'12)*, **ACM**, pp. 161-168, 2012. (Rank B).

[6] Bogdan Groza, Marius Minea, Formal modelling and automatic detection of resource exhaustion attacks. *Proc. 6th ACM Symposium on Information, Computer and Communications Security (ASIACCS'11)*, pp. 326-333, **ACM**, 2011. (Rank B).

[7] Bogdan Groza, Marius Minea, Customizing protocol specifications for detecting resource exhaustion and guessing attacks. *Proc. 9th International Symposium on Formal Methods for Components and Objects (FMCO'10)*, **Springer-Verlag**, **LNCS** vol. 6957, pp. 45-60, 2010.

[8] Groza, Bogdan, and Marius Minea. "A calculus to detect guessing attacks." Proc. 12th *Information Security Conference*. **Springer-Verlag**, **LNCS**, vol. 5735, pp. 59-67, 2009.

[9] Paula Vasile, Bogdan Groza, Stefan Murvay, Performance analysis of broadcast authentication protocols on CAN-FD and FlexRay, *10th Workshop on Embedded Systems Security (affiliated to ESWEEK'15)*, **ACM**, pp. 1-8, 2015.

[10] Stefan Murvay, Bogdan Groza, "Source Identification Using Signal Characteristics in Controller Area Networks", ***Signal Processing Letters***, **IEEE**, vol. 21, no. 4, pp. 395-399, 2014.

Des. Codes Cryptogr. (2014) 73:177–207
DOI 10.1007/s10623-013-9816-5

Cryptographic puzzles and DoS resilience, revisited

Bogdan Groza · Bogdan Warinschi

Received: 13 September 2012 / Revised: 20 January 2013 / Accepted: 22 March 2013 /
Published online: 5 April 2013

Abstract Cryptographic puzzles (or client puzzles) are moderately difficult problems that can be solved by investing non-trivial amounts of computation and/or storage. Devising models for cryptographic puzzles has only recently started to receive attention from the cryptographic community as a first step toward rigorous models and proofs of security of applications that employ them (e.g. Denial-of-Service (DoS) resistance). Unfortunately, the subtle interaction between the complex scenarios for which cryptographic puzzles are intended and typical difficulties associated with defining concrete security easily leads to flaws in definitions and proofs. Indeed, as a first contribution we exhibit shortcomings of the state-of-the-art definition of security of cryptographic puzzles and point out some flaws in existing security proofs. The main contribution of this paper are new security definitions for puzzle difficulty. We distinguish and formalize two distinct flavors of puzzle security which we call optimality and fairness and in addition, properly define the relation between solving one puzzle versus solving multiple ones. We demonstrate the applicability of our notions by analyzing the security of two popular puzzle constructions. We briefly investigate existing definitions for the related notion of security against DoS attacks. We demonstrate that the only rigorous security notion proposed to date is not sufficiently demanding (as it allows to prove secure protocols that are clearly not DoS resistant) and suggest an alternative definition. Our results are not only of theoretical interest: the better characterization of hardness for puzzles and DoS resilience allows establishing formal bounds on the effectiveness of client puzzles which confirm previous empirical observations. We also underline clear practical limitations for the

Communicated by C. Boyd.

B. Groza (✉)
Faculty of Automatics and Computers, Politehnica University of Timisoara,
Bd. V. Parvan nr. 2, Timisoara, Romania
e-mail: bogdan.groza@aut.upt.ro

B. Warinschi
Computer Science Department, University of Bristol, Woodland Road,
Bristol BS8 1UB, UK
e-mail: bogdan@cs.bris.ac.uk

Efficient Protocols for Secure Broadcast in Controller Area Networks

Bogdan Groza, *Member, IEEE*, and Stefan Murvay, *Student Member, IEEE*

***Abstract*—Controller Area Network is a bus commonly used by controllers inside vehicles and in various industrial control applications. In the past controllers were assumed to operate in secure perimeters, but today these environments are well connected to the outside world and recent incidents showed them extremely vulnerable to cyber-attacks. To withstand such threats, one can implement security in the application layer of CAN. Here we design, refine and implement a broadcast authentication protocol based on the well known paradigm of using key-chains and time synchronization, a commonly used mechanism in wireless sensor networks, which allows us to take advantage from the use of symmetric primitives without the need of secret shared keys during broadcast. But, as process control is a time critical operation we make several refinements in order to improve on the authentication delay. For this we study several trade-offs to alleviate shortcomings on computational speed, memory and bandwidth up to the point of using reduced versions of hash functions that can assure ad hoc security. To prove the efficiency of the protocol we provide experimental results on two representative microcontrollers from the market: a Freescale S12X and an Infineon TriCore, both devices were specifically chosen as they are located somewhat on the extremes of computational power.**

***Index Terms*—Authentication, broadcast, CAN, S12X, TriCore.**

I. Motivation and Related Work

CONTROLLER Area Network or simply CAN is a communication bus frequently used in vehicular systems and also common in general purpose automations. Initially developed by BOSCH, the current specifications for CAN are found in the newer standard ISO-11898. Numerous further improvements appeared in the literature: dual channel architectures [9], flexible time-triggered communication [1], star topologies [4], [5], dynamic identifiers to improve timing requirements [8], domain specific adaptations in the aeronautical sector [26], etc.

In the last decade, industrial control and automation systems become open to malicious adversaries and a significant part of the security community focused on alleviating potential threats in such environments [13], [15]. Also, recent incidents of international level, such as the Stuxnet worm, have shown that embedded devices are not as isolated as once thought and can become vulnerable targets [14]. As for in-vehicle security, recent research showed current automobiles to be unexpectedly vulnerable to external adversaries [11], [18] and it is likely that many other environments in which CAN operates are not completely isolated from the outside world. Two comprehensive books for security in automotives and cryptography based solutions in particular are [20] and [17], both contain relevant chapters on intra-vehicle security. Also several research papers published in the last years address these issues [3]. Still, to best of our knowledge there is no implementation available to assure authenticity in CAN networks, while the secure perimeters in which CAN used to operate are becoming increasingly open to the outside world.

To withstand such threats, security can be implemented at the application level of CAN. For this we study several trade-offs to achieve a potentially optimal solution. Digital signatures provide an elegant method for signing broadcast data, but they are not the solution in our context because of both the computational and communication overhead. Elliptic curves can significantly reduce the size of the signatures, 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 increase the cost of components, an issue that is largely avoided by manufacturers. One alternative to digital signatures such as RSA, or ECDSA is the use of one-time signatures which were initially proposed by Merkle [24]. Although they are frequently mentioned in the literature, they are quite unused in practice mostly because of their one-time nature and less favourable re-initialization. In contrast, symmetric primitives were efficiently employed for authentication in constrained environments such as sensor networks [21], [22], [28]. Due to the broadcast nature of CAN, protocols similar to the well known TESLA protocol [27], [29] can be used in this context as well. Indeed, 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. While TESLA like protocols introduce delays that could be unsuitable for all real-time CAN based applications, there is a broad area of applications where they could be tolerated in exchange for security. In particular, delays in the order of milliseconds, or even below as proved to be achievable by our proof-of-concept implementation, are suitable for a broad area of real-time control tasks. Many examples of network control scenarios that can accommodate such delays can be found in the literature [23].

There is an extensive bibliography related to the TESLA protocol, its history can be traced back to Lamport's authentication scheme [19]. The work of Bergadano *et al.* [7] proposes several variants of one-way chain based protocols, with or without

Manuscript received September 20, 2011; revised January 22, 2012; accepted November 29, 2012. Date of publication January 11, 2013; date of current version October 14, 2013. This work was supported in part by CNCSIS-UEFISCDI project number PNII-IDEI 940/2008 and by the strategic grant POSDRU 107/1.5/S/77265 and POSDRU/21/1.5/G/13798. Paper no. TII-11-543.

The authors are with the Faculty of Automatics and Computers, Politehnica University of Timisoara, Romania (e-mail: bogdan.groza@aut.upt.ro; stefan.murvay@gmail.com).

Digital Object Identifier 10.1109/TII.2013.2239301

Bridging Dolev-Yao Adversaries and Control Systems with Time-Sensitive Channels

Bogdan Groza and Marius Minea

Politehnica University of Timişoara and Institute e-Austria Timişoara*
bogdan.groza@aut.upt.ro, marius@cs.upt.ro

Abstract. Defining security objectives for industrial control scenarios is a challenging task due to the subtle interactions between system components and because security goals are often far from obvious. Moreover, there is a persistent gap between formal models for channels and adversaries (usually, transition systems) and models for control systems (differential or recurrent equations). To bind these two realms, we translate control systems into transition systems by means of an abstraction with variable time granularity and compose them with a channel model that is controlled by Dolev-Yao adversaries. This opens the road for automatic reasoning about the formal model of a control system using model checkers in a context where the communication channel is tampered with. We address a security objective that has so far largely eluded in models, namely freshness, which is highly relevant for control systems. Beyond the traditional resilience to replay attacks, we point out several flavours of freshness which are often overlooked, e.g., ordering and bounded lifespan. We formalize these notions and show that their absence can lead to attacks that subvert the control system. Finally, we build a proof-of-concept implementation that we use to determine attacks on a simple model which clearly shows that real-world scenarios are within reach.

Keywords: control system, formal modelling, freshness.

1 Introduction

CONTEXT AND REALISM. The generic image of a control system is that of a closed loop in which a controller regulates the behaviour of a process (usually called plant) as shown in Figure 1. The design of such systems is commonly based on intricacies known only to producers and insiders that use them. However, an important aspect reconfirmed time and again by incidents such as Stuxnet that is that the internal details of a system are hard to be kept secret to well motivated outsiders (likely, the worm exploited specific system details). Relying on security through obscurity is a fatal flaw. Since for cryptographic protocols, modeling has proved to be a crucial tool to assess security, it is quite obvious that modeling industrial control systems in their relation to the communication channels and adversaries is much more relevant today as they become exposed as parts of Internet-like structures, e.g., Internet of Things (IoT) [1]. Indeed, there has been constant attention in the previous years on attack surfaces and countermeasures for control

* This work is supported in part by FP7-ICT-2009-5 project 257876, SPaCIoS: Secure Provision and Consumption in the Internet of Services.

E. Luiijf and P. Hartel (Eds.): CRITIS 2013, LNCS 8328, pp. 167–178, 2013.

LiBrA-CAN: A Lightweight Broadcast Authentication Protocol for Controller Area Networks

Bogdan Groza[1], Stefan Murvay[1], Anthony van Herrewege[2], and Ingrid Verbauwhede[2]

[1] Faculty of Automatics and Computers, Politehnica University of Timisoara
{bogdan.groza,pal-stefan.murvay}@aut.upt.ro
[2] ESAT/COSIC - IBBT, KU Leuven, Belgium
{anthony.vanherrewege,ingrid.verbauwhede}@esat.kuleuven.be

Abstract. Security in vehicular networks established itself as a highly active research area in the last few years. However, there are only a few results so far on assuring security for communication buses inside vehicles. Here we advocate the use of a protocol based entirely on simple symmetric primitives that takes advantage of two interesting procedures which we call key splitting and MAC mixing. Rather than achieving authentication independently for each node, we split authentication keys between groups of multiple nodes. This leads to a more efficient progressive authentication that is effective especially in the case when compromised nodes form only a minority and we believe such an assumption to be realistic in automotive networks. To gain more security we also account an interesting construction in which message authentication codes are amalgamated using systems of linear equations. We study several protocol variants which are extremely flexible allowing different trade-offs on bus load, computational cost and security level. Experimental results are presented on state-of-the-art Infineon TriCore controllers which are contrasted with low end controllers with Freescale S12X cores, all these devices are wide spread in the automotive industry. Finally, we discuss a completely backward compatible solution based on CAN+, a recent improvement of CAN.

1 Motivation and Related Work

Vehicular network security established itself as an intense research topic in the last few years. Remarkable research papers from Koscher et al. [7] and later Checkoway et al. [4] showed vehicles to be easy targets for malicious adversaries.

While most of previous research was focused on vehicle to vehicle and vehicle to infrastructure communication there seem to be only a few results for assuring security on communication buses inside vehicles. There are several reasons behind this. First, the relevance of security inside vehicles was decisively shown only in the last two years [7], [4]. Second, the design principles used by manufacturers are somewhat out of reach for the academic community, being hard in this way to understand many assertions behind protocol design. Third, which is relevant for our research here, intra-vehicle communication is subject to constraints and specifications that are quite different from other well studied protocols. Most of the approaches advocate the use of secure gateways between different ECUs (Electronic Control Unit) or subnetworks [1], [13] and rely on basic building blocks from cryptography (encryptions, signatures, etc.). However, none

J. Pieprzyk, A.-R. Sadeghi, and M. Manulis (Eds.): CANS 2012, LNCS 7712, pp. 185–200, 2012.

SAPHE - Simple Accelerometer based wireless Pairing with HEuristic trees

Bogdan Groza
Faculty of Automatics and Computers
Politehnica University of Timisoara
Bd. V. Parvan nr. 2, Timisoara, Romania
bogdan.groza@aut.upt.ro

Rene Mayrhofer
Department for Mobile Computing
University of Applied Sciences Upper Austria
Softwarepark 11, Hagenberg, Austria
rene.mayrhofer@fh-hagenberg.at

ABSTRACT

Accelerometers provide a good source of entropy for bootstrapping a secure communication channel in autonomous and spontaneous interactions between mobile devices that share a common context but were not previously associated. We propose two simple and efficient key exchange protocols based on accelerometer data that use only simple hash functions combined with heuristic search trees. Using heuristics such as the Euclidean distance proves to be beneficial as it allows a more effective recovery of the shared key. While the first protocol seems to give just some performance improvements, the second, which we call *hashed heuristic tree*, is more secure than previous proposals since it increases the difference in protocol execution between benign and malicious parties. Nevertheless, the *hashed heuristic tree* is an entirely new approach which has the advantage of allowing different heuristics in the search, leaving plenty of room for future variants and optimizations.

Categories and Subject Descriptors

D.4.6 [**OPERATING SYSTEMS**]: Security and Protection—*Authentication*; K.6.5 [**MANAGEMENT OF COMPUTING AND INFORMATION SYSTEMS**]: Security and Protection—*Authentication*

General Terms

Security

Keywords

authentication, accelerometer, key-exchange

1. MOTIVATION

While today's mobile devices embed countless applications that boost up usability, they still heavily rely on manual input for performing day by day routines such as pairing, i.e., bootstrapping a communication channel. Examples for such manual input are the PIN code that needs to be entered for Bluetooth pairing, the WPS password for connecting to many WiFi networks, or a shared password for many secure chat, telephone, or file transfer applications. Depending on user data not only reduces usability, but inextricably gambles security since such input usually lacks the desired entropy for making a connection secure (as evidenced by the standard pairing PIN codes '0000' or '1234' for many Bluetooth headsets).

In contrast, using environmental information is an attractive mechanism for secure pairing between devices in autonomous and spontaneous interactions where previously shared information such as public key certificates does not exists. While manually entered information is time consuming and requires additional input interfaces, using motion, sound, light, etc., seems to be an attractive alternative that provides a fast and more scalable solution that can be also used by unskillful users. A complete survey on different techniques that can be used for this purpose can be found in [17].

In particular, motion characteristics acquired from accelerometers are considered to be a rich source of entropy, easy to produce or re-produce and difficult to guess by outsiders. Mayrhofer and Gellersen explore the use of shaking patterns to pair mobile devices in [22] and apply this technique to secure a bluetooth connection. Bluetooth pairing is known to be particularly vulnerable [13] to man-in-the-middle attacks caused by external adversaries and shaking patterns can be used to remove this problem. The same idea is also explored by Castelluccia and Mutaf in [5], Kirovski et al. in [16] and by Bichler et al. in [3]. All methods are distinct in the way key extraction from sensor data is performed. More specifically, in [22] data is extracted in the frequency domain by using a coherence function, initially used in [18] while the other two proposals extract data in the time domain. Our approach also uses the time domain but the way in which we extract data by using heuristic trees and hashed heuristic trees appears to be entirely distinct from any previous procedure. The closest method to our extraction procedure could be [3], but the concept of hashed heuristic tree that we introduce seems to be entirely new.

In a non-adversarial setting, accelerometer data can be directly used to distinguish whether or not two devices are carried together [14], [18], to recognize activities and gestures [15], [24], [19], [1] or an individual gait [2], [11], [10] by simply exchanging the information acquired from sensors. But it is not straight-forward to establish a common secret

MoMM2012, 3-5 December, 2012, Bali, Indonesia

Formal Modelling and Automatic Detection of Resource Exhaustion Attacks

Bogdan Groza
Politehnica University and Institute e-Austria
Timişoara, Romania
bogdan.groza@aut.upt.ro

Marius Minea
Politehnica University and Institute e-Austria
Timişoara, Romania
marius@cs.upt.ro

ABSTRACT

Many common protocols: TCP, IPSec, etc., are vulnerable to denial of service attacks, where adversaries maliciously consume significant resources of honest principals, leading to resource exhaustion. We propose a set of cost-based rules that formalize DoS attacks by resource exhaustion and can automate their detection. Our classification separates excessive but legal protocol use (e.g., flooding) from illegal protocol manipulation that causes participants to waste computation time without reaching the protocol goals. We also distinguish simple intruder intervention leading to wasteful execution from DoS attacks proper, which can be repeatedly initiated. Our rules can highlight attacks that are undetectable by the targeted honest agents, or by all protocol participants. We have successfully tested an implementation of the methodology in a validation platform on relevant protocol examples, in what to the best of our knowledge is the first formal automated analysis of DoS attacks.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Protocol verification*; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection; C.4 [**Performance of Systems**]: *Reliability, availability and serviceability*

General Terms

Security, Verification

Keywords

denial of service, formal modeling, automated verification

1. INTRODUCTION

Protocols that base their security on cryptographic primitives are indispensable nowadays. However, from a computational perspective, not all cryptographic primitives are cheap. While secret-key primitives can be executed in microseconds on modern computers, public-key primitives require a thousand time more computational steps and can cause resource exhaustion even on well equipped servers. The problem is far-reaching: on low computational power devices, such as sensors, mobile phones, embedded devices, etc., the unjustified execution of even simple cryptographic primitives can cause resource exhaustion.

In recent years, many protocols have been found vulnerable to this kind of attacks and modified variants or countermeasures have been proposed. In this paper we focus on the automatic detection of these attacks, where honest participants can be maliciously determined to perform expensive operations, such as public-key encryptions or signatures, while the adversary consumes significantly less resources.

From the point of view of protocol execution, we consider useful to separate resource exhaustion DoS attacks in two main categories:

- Resource exhaustion DoS attacks due to *excessive use*. These are attacks in which there is no abnormal use of the protocol, however the adversary as participant consumes significantly less resources than other principals thus being capable to cause a DoS. Typical examples are attacks on the server side, such as flooding, spam, etc., which do not violate the protocol specification but can exhaust resources of honest principals.

- Resource exhaustion DoS due to *malicious use*. In these attacks the adversary manages to bring the protocol to an abnormal state (principals are not aware of their correct identities, shared keys do not match, etc.) from which the protocol goals cannot be correctly met. Many protocols with such vulnerabilities exist, probably the best known is Lowe's attack [13] on the STS protocol [8], which we will use as one case study.

Of course, in general *denial of service* (proper) requires repetition, and one condition for this to take place is the ability of the intruder to control the initiation of a session. This condition is sufficient in the case of protocols vulnerable to excessive use since there is no abnormal protocol behaviour in this case, and thus an honest principal cannot detect being under attack (the only prevention is to limit the use of the protocol). A commonly used solution to protect the server side are proof-of-work protocols based on moderately hard one-way functions, known as cryptographic puzzles or client puzzles. In this context, several protocols have been augmented with such constructions, including e-mail [9], TCP [12], authentication protocols [3], etc.

ASIACCS '11, March 22–24, 2011, Hong Kong, China.

Customizing Protocol Specifications for Detecting Resource Exhaustion and Guessing Attacks*

Bogdan Groza and Marius Minea

Politehnica University of Timişoara and Institute e-Austria Timişoara
bogdan.groza@aut.upt.ro, marius@cs.upt.ro

Abstract. Model checkers for security protocols often focus on basic properties, such as confidentiality or authentication, using a standard model of the Dolev-Yao intruder. In this paper, we explore how to model other attacks, notably guessing of secrets and denial of service by resource exhaustion, using the AVANTSSAR platform with its modelling language ASLan. We do this by adding custom intruder deduction rules and augmenting protocol transitions with constructs that keep track of these attacks. We compare several modelling variants and find that writing deductions in ASLan as Horn clauses rather than transitions using rewriting rules is crucial for verification performance. Providing automated tool support for these attacks is important since they are often neglected by protocol designers and open up other attack possibilities.

1 Introduction and Motivation

Formal verification tools provide an efficient means for automatic verification of security protocols, once models of these have been written, e.g., some variant of symbolic transition systems. Usually, the focus is on verification of standard security goals, such as authenticity and confidentiality. However, in many cases, satisfying these goals is not sufficient to consider a protocol safe and a more in-depth analysis to rule out other kinds of attacks is necessary.

This paper focuses on two such attacks which are not handled routinely by many protocol verifiers, namely guessing attacks and denial of service (DoS). Both of these attacks are a main concern in protocol design. Guessing attacks are relevant because users tend to choose weak passwords, and some values such as PIN codes have intrinsically low entropy. They can become the weakest link in more complex protocols, leading to other attacks as well. Resource exhaustion is relevant as a common source of DoS as well as from an economic point of view if one considers ruling out protocol designs that can be exploited to make honest participants spend unreasonable amounts of resources, time or memory.

* This work is supported in part by FP7-ICT-2007-1 project 216471, AVANTSSAR: Automated Validation of Trust and Security of Service-oriented Architectures and by strategic grant POSDRU/21/1.5/G/13798 of the Human Resources Development Programme 2007-2013, co-financed by the European Social Fund – Invest in People.

B.K. Aichernig, F.S. de Boer, and M.M. Bonsange (Eds.): FMCO 2010, LNCS 6957, pp. 45–60, 2011.

A Calculus to Detect Guessing Attacks*

Bogdan Groza[1] and Marius Minea[2]

[1] Politehnica University of Timişoara
bogdan.groza@aut.upt.ro
[2] Institute e-Austria Timişoara
marius@cs.upt.ro

Abstract. We present a calculus for detecting guessing attacks, based on oracles that instantiate cryptographic functions. Adversaries can *observe* oracles, or *control* them either on-line or off-line. These relations can be established by protocol analysis in the presence of a Dolev-Yao intruder, and the derived guessing rules can be used together with standard intruder deductions. Our rules also handle partial verifiers that fit more than one secret. We show how to derive a known weakness in the Anderson-Lomas protocol, and new vulnerabilities for a known faulty ATM system.

1 Introduction and Related Work

Analyzing vulnerability to guessing attacks is of high practical relevance. A value is deemed guessable if it has small entropy (is chosen from a small cardinality set), and the guess can be verified. An adversary can perform guessing by off-line computation, or on-line, exploiting the interaction with honest participants.

Conceptually, guessing involves two steps. Any protocol must have a *generation oracle* which computes some value (the *verifier*), given the secret as input. Next, a boolean *verification oracle* compares a verifier for the guess with one computed for the actual secret. We use the term oracle for an abstract object that produces a value, regardless of how the computation is done. In particular, the adversary might use other participants as on-line oracles for this purpose.

Separating the verifier generation from the verification itself, and modeling them as oracles is key to our analysis of guessing attacks in both off-line and on-line settings. It is often argued that on-line guessing can be blocked after a threshold of incorrect guesses. However, if the adversary's guesses are cached as valid protocol interactions, relying on blocking is not a justified defense.

Our analysis identifies various guessing situations with partial or complete view over inputs and outputs of oracles and with off-line, on-line or blockable on-line oracle access. We provide inference rules which can detect guessing attacks in these situations. Once such a vulnerability is detected, it is up to further review to decide if it can be removed by limiting protocol runs. We will also

* This work is supported in part by FP7-ICT-2007-1 project 216471, AVANTSSAR: Automated Validation of Trust and Security of Service-oriented Architectures.

P. Samarati et al. (Eds.): ISC 2009, LNCS 5735, pp. 59–67, 2009.

Performance analysis of broadcast authentication protocols on CAN-FD and FlexRay

Paula Vasile
paula.vasile10@gmail.com

Bogdan Groza
bogdan.groza@aut.upt.ro

Stefan Murvay
stefan.murvay@gmail.com

Politehnica University of Timisoara, Faculty of Automatics and Computers
Department for Automatics and Applied Informatics
Bd. V. Parvan nr. 2, Timisoara, Romania

ABSTRACT

In the light of the numerous reported attacks, designing cryptographic protocols for in-vehicle embedded networks was a constant preoccupation in the past few years. While several research proposals appeared, a concrete performance analysis of such protocols over a realistic network configuration is still absent from the literature. In this work we address the performance for various authentication protocols that were recently proposed for the two most prominent vehicular buses: FlexRay and CAN-FD. While a real-world vehicular network is still out of reach for our work, we achieve a first step in this direction by using a CANoe based simulation for these protocols over state-of-the-art automotive buses. This allows us to draw a more realistic perspective on the efficiency of existing proposals for bus authentication. Our results suggest that sharing symmetric keys between groups of nodes is the most realistic proposal as it creates a balance between bandwidth efficiency and security level.

Categories and Subject Descriptors

[**Security and privacy**]: Security in hardware—*Embedded systems security*

General Terms

Security

Keywords

broadcast authentication, embedded networks, CAN-FD, FlexRay

1. INTRODUCTION

Contemporary vehicles are the result of an evolutionary process that augmented mere mechanical devices with complex electronics and intricate software counterparts. Recent vehicles have dozens of ECUs (Electronic Control Units) that implement a plethora of functions dedicated for safety critical tasks, e.g., braking and stability control, or mere informational purposes to make a more pleasurable driving experience. Of course, behind these miniature computers called ECUs, a complex network infrastructure needs to be deployed that connects these ECUs via cables and unavoidably exposes information to the outsiders via various ports, e.g., OBD (On-Board Diagnostics), or even wireless channels, e.g., WiFi, Bluetooth, 3G, etc. There are little doubts that in terms of attacks and defenses vehicular networks will evolve in a similar manner to computer networks. A solid proof for this are the countless attacks that were published in the recent years [9], [2].

The academic research community was quick to react with various proposals for assuring security on vehicular buses (these are all surveyed in a forthcoming section). However, a comprehensive performance analysis of these solutions on a real-world vehicular network is still missing. The reason behind this is quite simple, in-vehicle infrastructures are still subject to many proprietary solutions and architectural details that are not fully accessible to academic researchers. In this work we take a first step in this direction by using the industry standard CANoe tool in order to simulate a realistic state-of-the-art network based on FlexRay and CAN-FD.

1.1 State-of-the-art in-vehicle buses: FlexRay and CAN-FD

The reasons for choosing these buses are clear: CAN is the most common bus inside cars and FlexRay was designed as its successor. Due to inherent expenses, FlexRay had not yet entered in all vehicles and a direct successor of CAN also emerged, i.e., CAN-FD (Controller Area Network with Flexible Data-rate). We give next a brief description of these two communication layers.

FlexRay is an automotive communication protocol developed by the FlexRay Consortium which gathers important players in the automotive industry. It was designed as a faster and more reliable alternative to other automotive communication protocols existing at that time, e.g., CAN or LIN. The protocol supports data rates up to 10 Mbit/s and a data payload length of 254 bytes. The data is transmitted between ECUs in the form of frames which have the structure presented in Figure 1. The bit-length of the fields are displayed below, the only exception being the data field for which the length is represented in bytes. FlexRay frames can be either time-triggered (static frames) or event-triggered (dynamic frames). FlexRay accomplishes the use of both static and dynamic frames by employing a predefined communication cycle (Figure 1) which defines specific segments

WESS'15, October 04-09, 2015, Amsterdam, Netherlands

DOI: http://dx.doi.org/10.1145/2818362.2818369 .

Source Identification Using Signal Characteristics in Controller Area Networks

Pal-Stefan Murvay and Bogdan Groza

***Abstract*—The CAN (Controller Area Network) bus, i.e., the de facto standard for connecting ECUs inside cars, is increasingly becoming exposed to some of the most sophisticated security threats. Due to its broadcast nature and ID oriented communication, each node is sightless in regards to the source of the received messages and assuring source identification is an uneasy challenge. While recent research has focused on devising security in CAN networks by the use of cryptography at the protocol layer, such solutions are not always an alternative due to increased communication and computational overheads, not to mention backward compatibility issues. In this work we set steps for a distinct approach, namely, we try to take authentication up to unique physical characteristics of the frames that are placed by each node on the bus. For this we analyze the frames by taking measurements of the voltage, filtering the signal and examining mean square errors and convolutions in order to uniquely identify each potential sender. Our experimental results show that distinguishing between certain nodes is clearly possible and by clever choices of transceivers and frame IDs each message can be precisely linked to its sender.**

***Index Terms*—CAN bus, physical fingerprinting, source identification.**

I. Motivation and Related Work

THE Controller Area Network (CAN) is a broadcast serial bus initially designed for in-vehicle communication. The ever growing design complexity of automotive embedded systems makes it difficult for the manufacturers to anticipate all possible threat scenarios. As a result, vulnerabilities in automotive systems are highlighted by an increasing number of recent papers [1], [4], [6]. All these prove that in-vehicle communication, in the absence of source authentication, is an easy target even in front of some of the most basic attacks, e.g., replays, packet injections, etc.

The CAN physical layer is typically implemented as a two wire differential bus as presented in Fig. 1. Each frame begins with an identifier (ID) which determines the priority of the frame and carries up to 64 bits of data followed by a standard 15-bit CRC. The 64 bit data field of the CAN frame gives the first hints on why assuring cryptographic security is uneasy. Clearly, it is not feasible to fit an authentication tag plus the message in the 8 byte block that is carried by each frame. Adding separate authentication frames increases the bus load and CAN is limited to 1 Mbps, an upper-bound that is already reached in many practical scenarios. Message Authentication Codes (MACs), the cryptographic tool for assuring message authentication, are usually in the order of 128 bits and while truncating them to a particular length is an option one still needs to fit them along with the message within the 64 bits (clearly, this is not possible). Van Herrewege *et al.* propose in [8] the use of CAN+ in order to hide the authentication bits within the bits of a regular CAN frame. But CAN+ capable controllers are not yet available on the market and it is not clear if they will be produced in the near future (currently, as a more expensive alternative, the industry is migrating to other layers such as CAN-FD or FlexRay). In [7] Szilagy and Koopman allow each node to vote on the authenticity of the message and each vote consists in several MACs that are truncated in order to fit them inside each frame. The suggested value from [7] is 8 bits for each MAC, this is clearly too low to assure real-world security.

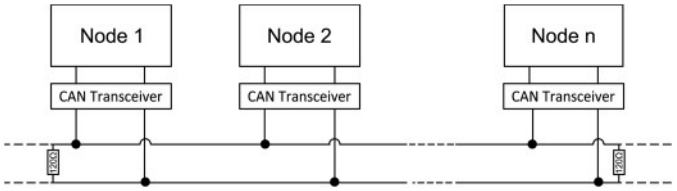

Fig. 1. Typical CAN bus topology.

Obviously, the alternatives for assuring source identification on CAN are limited. To alleviate this, here we take an entirely distinct approach by trying to identify the nodes based on the signal patterns. The CAN specification allows great freedom in the implementation of the physical layer. As a consequence, signals produced by transceivers from different manufacturers are not identical. Moreover, as each electronic component gathers unique physical characteristics, even signals generated by transceivers from the same manufacturer show up unique peculiarities that can help to distinguish between senders.

Related work. Several lines of work were already focused on physical layer security. Hall *et al.* [3] used radio frequency fingerprinting for intrusion detection in wireless networks. Beamforming and artificial noise were used in wireless networks in a physical layer approach to provide secure communication in the presence of eavesdroppers [5]. Investigations were done in the case of wired buses as well. The work of Gerdes *et al.* [2] is focused on identifying Ethernet cards by studying the synchronization signal (found at the beginning of each Ethernet frame) with the help of a matched filter. Reported experimental results show that Ethernet cards of different

Manuscript received October 10, 2013; revised December 23, 2013; accepted January 19, 2014. Date of publication January 31, 2014; date of current version February 06, 2014. This work was supported in part by the strategic grant POSDRU 107/1.5/S/77265, inside POSDRU Romania. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Yan Lindsay Sun.

The authors are with the Faculty of Automatics and Computers, Politehnica University of Timisoara, Timisoara, Romania (e-mail: stefan.murvay@gmail.com; bogdan.groza@aut.upt.ro).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/LSP.2014.2304139