

Automated Security Analysis for Real-World IoT Devices

Lélio Brun

National Institute of Informatics
Tokyo, Japan
lelio_brun@nii.ac.jp

Yasushi Ono

Institute of Information Security
Yokohama, Japan
y.ono@iisec.ac.jp

Ichiro Hasuo

National Institute of Informatics
Tokyo, Japan
hasuo@nii.ac.jp

Taro Sekiyama

National Institute of Informatics
Tokyo, Japan
sekiyama@nii.ac.jp

ABSTRACT

Automatic security protocol analysis is a fruitful research topic that demonstrates the application of formal methods to security analysis. Several endeavors in the last decades successfully verified security properties of large-scale network protocols like TLS, sometimes unveiling unknown vulnerabilities.

In this work, we show how to apply these techniques to the domain of IoT, where security is a critical aspect. While most existing security analyses for IoT tackle individually either protocols, firmware or applications, our goal is to treat IoT systems as a whole. We focus our work on a case study, the Armadillo-IoT G4 device, highlighting the specific challenges we must tackle to analyze the security of a typical IoT device. We propose a model using the TAMARIN prover, that allows us to state certain key security properties about the device and to prove them automatically.

KEYWORDS

cryptographic protocols, Internet of Things, formal verification

ACM Reference Format:

Lélio Brun, Ichiro Hasuo, Yasushi Ono, and Taro Sekiyama. 2023. Automated Security Analysis for Real-World IoT Devices. In *Hardware and Architectural Support for Security and Privacy 2023 (HASP '23)*, October 29, 2023, Toronto, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3623652.3623667>

1 INTRODUCTION

Internet of Things (IoT) is a paradigm where many different objects can be queried and operated over communication networks. IoT devices are pervasive in fields like home automation, healthcare, transportation, energy management, manufacturing, etc. Because these applications can be either privacy or safety-critical, security is widely acknowledged as one of the major stakes of IoT growth [39]. This work is part of a project whose aim is the analysis and verification of security properties of IoT systems. In particular, we present the case study of the Armadillo-IoT G4¹

¹<https://armadillo.atmark-techno.com/armadillo-iot-g4>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

HASP '23, October 29, 2023, Toronto, Canada

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1623-2/23/10.

<https://doi.org/10.1145/3623652.3623667>

device, which embodies a typical architecture for secure IoT devices: it embeds a secure element and a main chip that features the ARM TrustZone technology [30] to provide a Trusted Execution Environment (TEE) [34]. TEEs are a standard mean in secure IoT systems [36] to provide both secure execution isolation for so-called Trusted Applications (TAs) and secure isolated storage, e.g., for secrets. The analysis of a typical platform like the Armadillo is challenging as it requires:

- Modeling the interactions between elements following distinct protocols and specifications;
- Modeling a notion of global shared state;
- Modeling the execution of TAs.

Most existing research on security analysis of IoT systems focuses independently either on communication protocols, firmware, or applications. However, independent analysis does not ensure that an entire IoT system is secure. For example, if a protocol does not meet a property assumed in the analysis of an application, the analysis results for the protocol and application cannot be combined on the whole system.

Our goal with this work is to approach security analysis of IoT systems as a whole and to model the interactions between these components. We follow the symbolic model-checking approach. Over the last decades, symbolic model-checking techniques have been designed to tackle the difficult challenges of security analysis. Tools like AVISPA [4] or DeepSec [10] achieve automated model checking by restraining the generally undecidable security problem to its decidable fragment. Other automated provers like ProVerif [6] or TAMARIN [11, 28, 35] give up on termination and decidability to allow analysis of a broader class of protocols without bounding the number of sessions. Researchers have used these two state-of-the-art tools to prove security properties such as secrecy or authentication for large-scale real-world protocols like TLS 1.3 [5, 13], ARINC823 [7], Signal [26], 5G AKA [12] or WPA2 [14]. Because these tools build on the general theoretical frame of state/transition systems, we show that their original scope restricted to network protocol analysis can be extended to model and analyze typical IoT systems.

Our contribution is a TAMARIN model that treats selected functional aspects of the Armadillo device. This model allows us to reason about and automatically prove certain critical security properties of the device. Our approach is general enough to be readily extended to other case studies. Namely, we prove the secrecy of its root of trust and the integrity of TAs execution.

We organize the paper as follows. We discuss related work in section 2. In section 3, we give an overview of the Armadillo device. Section 4 briefly presents the TAMARIN prover, and we detail our model in section 5. We discuss our results and conclude in section 6.

2 RELATED WORK

Formal methods approaches to security analysis in the domain of IoT are a recent research topic. Hofer-Schmitz and Stojanovi  [23] thoroughly review existing work in this vein, detailing the techniques and tools used as well as the specific treated aspects. In general, most of these works focus on analyzing the security of protocols that happen to be commonly used by IoT devices, like ZigBee, Z-Wave, or Sigfox². For example, Kim et al. [25] propose a TAMARIN model showcasing Denial-of-Service (DoS) attacks on Sigfox. Our approach is complementary: instead of focusing on one protocol, we are interested in the interactions of the different components using different protocols inside a typical IoT device.

There exist frameworks targeting security analysis for IoT. SOTERIA [8] extracts a state/transition system model from existing IoT applications' source code and performs model checking on the extracted model. IOTSAN [31] follows the same approach. While the scope of SOTERIA focuses on applications, IOTSAN also encompasses their interactions with other applications, sensors and actuators. Both frameworks are tightly bound to the Groovy source language and to the SmartThings ecosystem in general. We pursue the same goals as IOTSAN, while proposing a more agnostic approach. On the other hand their fine-grain analysis at the application level is not offered by our current solution: it is a topic for future work. IOTA [17] is another recent framework following a different approach. It builds logic Prolog models of IoT systems from their formal descriptions and generates attack graphs using MulVAL [33] for further analysis.

On the technical side, we take inspiration from the work of Kremer and K unnemann [27], who show how to extend TAMARIN to support global and shared state. Jacomme et al. [24] presented another extension even more relevant to us that provides support for execution in a TEE, with a *reporting* system that allows signing the output of trusted computation. These two generic extensions are part of SAPIC [9], an alternative front-end for TAMARIN. SAPIC proposes a syntax close to the applied π -calculus [2, 3] internally translated into TAMARIN's rewriting rules. SAPIC is still under development, and we only use the original TAMARIN formalism.

3 OVERVIEW OF THE ARMADILLO-IoT G4

The Armadillo-IoT G4 is part of a family of business card-sized embedded platforms mainly targeting high-performance Artificial Intelligence (AI) processing and machine learning. It is equipped with an Arm Cortex-A53 4-core SoC, the i.MX 8M Plus³, Gigabit Ethernet ports, USB 3.0, HDMI interfaces, and a Neural Processing Unit (NPU).

The device is security oriented. It runs a dedicated Linux distribution designed to be minimal and uses containers to achieve

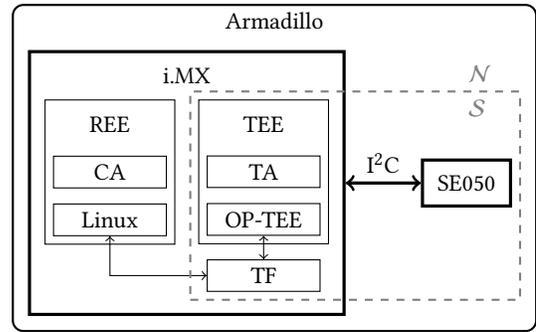


Figure 1: Architecture overview

separated sandboxed execution of applications. It benefits from the TrustZone technology⁴ offered by the i.MX to feature a TEE implemented by OP-TEE [1]. The device is equipped with a separate secure element, the EdgeLock SE050⁵, used to provide a root of trust and secure cryptographic operations. In this work, we focus on the interactions between the so-called *Normal world N*—also called Rich Execution Environment (REE) in TEE terminology, that is, the Linux-based OS and the contained applications—and the *Secure world S*—the SE050 element and OP-TEE.

Figure 1 gives a visual overview. On the left is the i.MX board: a non-secure contained Client Application (CA) running in the REE can request the execution of a secure TA in the TEE. The corresponding communication between Linux and OP-TEE is handled through the Trusted Framework (TF), using a sequence of messages and shared memory. On the right, establishing a secure channel over an I²C bus enables communication between the main board and the secure element SE050.

4 THE TAMARIN PROVER

The TAMARIN prover [11, 28, 35] is a tool that allows unbounded verification of network protocols in the symbolic model. Like MaudeNPA [16], it relies on *rewriting logic* to describe security protocols and their execution. A protocol is specified as a set of multiset rewriting rules defining a Labelled Transition System (LTS). The protocol can be analyzed by defining properties in a temporal first-order logic checked against a message deduction theory implementing the Dolev-Yao adversary model [15] and user-defined equational theories. The verification process builds on a constraint-based backward search algorithm and can run fully automated or be used in an interactive graphical mode. Termination is not guaranteed since the general security problem in the unbounded setting is undecidable [29]. We give in the following a brief overview of TAMARIN; the details are in the original presentation [35].

Term algebra. Messages are represented using a sorted term algebra. A term can be a public name of sort *pub* (e.g., a participant name), a fresh name (e.g., a nonce) of sort *fr*, a variable, or the application of a function symbol *f* with arity *n* from a given signature Σ

²<https://csa-iot.org/all-solutions/zigbee>, <https://www.z-wave.com>, <https://www.sigfox.com>

³<https://www.nxp.com/products/processors-and-microcontrollers/arm-processors/i-mx-applications-processors/i-mx-8-applications-processors/i-mx-8m-plus-arm-cortex-a53-machine-learning-vision-multimedia-and-industrial-iot:IMX8MPLUS>

⁴<https://www.arm.com/en/technologies/trustzone-for-cortex-a>

⁵<https://www.nxp.com/products/security-and-authentication/authentication/edglock-se050-plug-and-trust-secure-element-family-enhanced-iot-security-with-high-flexibility:SE050>

$$\frac{}{\text{Fr}(\sim x)} \quad \frac{\text{Out}(x)}{\text{K}(x)} \quad \frac{\text{K}(x)}{\text{In}(x)}[\text{K}(x)] \quad \frac{}{\text{K}(\$x)}$$

$$\frac{\text{Fr}(\sim x)}{\text{K}(\sim x)} \quad \frac{\text{K}(x_1) \cdots \text{K}(x_n)}{\text{K}(f(x_1, \dots, x_n))} \quad \forall (f, n) \in \Sigma$$

Figure 2: Built-in message deduction rules

(e.g., a cryptographic primitive) to n terms $f(t_1, \dots, t_n)$. To specify the sort of a variable, we use the following syntax: $\sim x$ is of sort *fr* and $\$x$ is of sort *pub*. If unspecified, the default sort is the message sort *msg*. Equations reflecting the properties of the functions define the equational theory (e.g., cancellation of encryption/decryption and arithmetic properties for Diffie-Hellman exponentiation).

Facts. In TAMARIN, the state of the LTS defined by a protocol is represented as a multiset of facts. A *linear* fact is a fact symbol (in uppercase by convention) applied to terms $F(t_1, \dots, t_n)$. A *persistent* fact is noted $!F(t_1, \dots, t_n)$. If a fact is linear, it can be consumed from the state when a transition fires, while persistent facts stay in the state forever. TAMARIN defines built-in facts: $\text{In}(m)$ denotes that a message m is received from the network, $\text{Out}(m)$ that m is sent over the network, $\text{Fr}(n)$ represents the generation of a fresh name n , and $\text{K}(m)$ models the fact that the adversary knows term m — K is always persistent and the $!$ modifier is omitted.

Rewriting rules and traces. A protocol is specified as a set of multiset rewriting rules. A rule is written as follows:

$$\frac{p}{c}[a] \quad \text{or} \quad \frac{p}{c} \quad \text{where} \quad \begin{cases} p \text{ is a premise multiset of facts} \\ c \text{ is a conclusion multiset of facts} \\ a \text{ is an action set of facts} \end{cases}$$

The LTS described by the rules is labeled by the action facts: the sequence of action facts resulting from an execution of the system is called a *trace*. A transition can occur non-deterministically as long as there exists a rule whose premises p are multiset-wise contained in the current state. As a result, the matching sub-multiset is rewritten into the conclusions c , and the set of action facts a is appended to the trace.

Figure 2 presents the built-in rules of TAMARIN associated with the facts In , Out , Fr and K , that implement the message deduction theory. The first rule models the generation of fresh nonces, which can always fire without any premise. The remaining describe the actual message deduction: the adversary can learn any term sent on the network, send any term they know, learn any public term, generate fresh nonces themselves, and apply any *public* function symbol to terms they know—note that TAMARIN offers the possibility of specifying *private* function symbols that the adversary cannot use, but we do not use this feature in this work.

Security properties. We can specify security properties about a protocol as formulas in a first-order temporal logic over traces. The additional *temp* sort characterizes *timepoints*, i.e., indexes of the trace, to reason about instants. The *temp* sort can be specified on variables using the $\#x$ syntax. The logic supports the universal

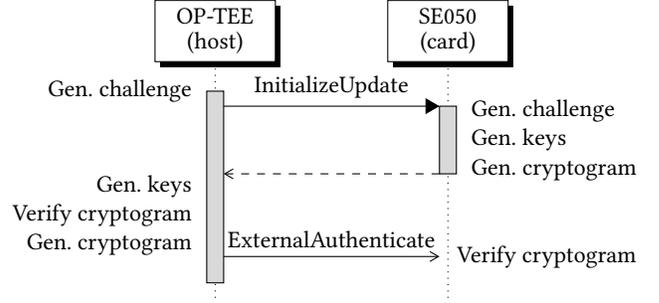


Figure 3: Binding sequence diagram

and existential quantifiers All and Ex , the implication \implies , conjunction $\&$, disjunction $|$, and negation not . One can assert that a fact occurred at a timepoint i : $F(m) \@ \#i$; that a timepoint i precedes another j : $\#i < \#j$; or that two variables are equal: $x = y$. A protocol as a set of rules P satisfies a formula ϕ , written $P \models \phi$, if all possible traces of P satisfy ϕ .

Verification algorithm. To verify that $P \models \phi$, TAMARIN first performs a normalization phase on P and ϕ , then tries to produce a proof by contradiction. It performs a backward search in the LTS graph of P to try to find a counter-example, i.e., an attack, by using an incremental constraint-solving algorithm to solve the constraint system encoded from ϕ . If the algorithm terminates, it either found an attack or proved $P \models \phi$.

5 THE TAMARIN MODEL

The Armadillo-IoT G4 is a simple device, but reasoning formally about its security is an involved task since it comprises distinct hardware and software components that communicate with each other through different channels. In this section, we detail a restricted TAMARIN model which covers three key security aspects:

The binding process between the i.MX board and the SE050 secure element establishes a session-based private encrypted channel over the I²C bus. This is critical as it ensures that the secure element, which is used to perform cryptographic operations and to store the root of trust of the device, is accessed and operated securely.

The derivation of the Secure Storage Key (SSK) which is a secret key used by OP-TEE to encrypt storage operations. The secrecy of this key is crucial since it is the basis of the security of the storage in the *Secure world*.

The execution of TAs running in the TEE. To keep our model general, we focus here on the integrity of the executed code.

5.1 Binding the secure element

As a first security step, the Armadillo device uses *secure boot* to establish a cascading chain of trust, as described in the Armadillo documentation [38]. Secure boot ensures the authentication at boot-time of the firmware and operating system. It is a critical feature for IoT devices to prevent an attacker from booting its own malicious code. As part of this process, OP-TEE boots before Linux, and uses its own I²C driver to bind with the SE050. Afterward, as indicated

$$\begin{array}{l}
 1 \frac{!TEE(tee_l) \quad Fr(\sim cha^h)}{InitStTEE(\sim cha^h) \quad Out(\langle 'InitializeUpdate', \sim cha^h \rangle)} \\
 2 \frac{!SharedKeys(k_{enc}, k_{mac}) \quad Fr(\sim cha^c) \quad In(\langle 'InitializeUpdate', cha^h \rangle)}{InitStSE(cha^h, \sim cha^c, s_{enc}, s_{mac}) \quad Out(\langle 'InitializeUpdate_R', \sim cha^c, cry^c \rangle)} \quad \text{where } \begin{cases} s_{enc} = h(\langle k_{enc}, cha^h, \sim cha^c \rangle) \\ s_{mac} = h(\langle k_{mac}, cha^h, \sim cha^c \rangle) \\ cry^c = \text{cryptogram}(s_{mac}, 'card', \langle cha^h, \sim cha^c \rangle) \end{cases} \\
 3 \frac{!SharedKeys(k_{enc}, k_{mac}) \quad InitStTEE(cha^h) \quad In(\langle 'InitializeUpdate_R', cha^c, cry^c \rangle)}{!SessKeysTEE(s_{enc}, s_{mac}) \quad Out(\langle 'ExternalAuthenticate', cry^h \rangle)} [AuthTEE()] \quad \text{where } \begin{cases} s_{enc} = h(\langle k_{enc}, cha^h, cha^c \rangle) \\ s_{mac} = h(\langle k_{mac}, cha^h, cha^c \rangle) \\ cry^{h,c} = \text{cryptogram}(s_{mac}, 'card', \langle cha^h, cha^c \rangle) \\ cry^h = \text{cryptogram}(s_{mac}, 'host', \langle cha^h, cha^c \rangle) \\ cry^c \doteq cry^{h,c} \end{cases} \\
 4 \frac{InitStSE(cha^h, cha^c, s_{enc}, s_{mac}) \quad In(\langle 'ExternalAuthenticate', cry^h \rangle)}{!SessKeysSE(s_{enc}, s_{mac})} [AuthSE()] \quad \text{where } \begin{cases} cry^{c,h} = \text{cryptogram}(s_{mac}, 'host', \langle cha^h, cha^c \rangle) \\ cry^h \doteq cry^{c,h} \end{cases}
 \end{array}$$

Figure 4: Binding rewriting rules

in the OP-TEE documentation [18], subsequent communication between the main board and the secure element is managed securely through the REE's I²C driver, for performance reasons.

The binding process implements the Secure Channel Protocol 03 (SCP03) specified by GlobalPlatform⁶ [20]: it establishes an encrypted and MAC-authenticated channel. The communication between the two elements respects an Application Protocol Data Unit (APDU) interface specified by NXP [32]. Figure 3 represents the sequence diagram of the process. At the end of the sequence, both elements authenticated each other and have the same *session keys* used for subsequent encryption and MAC-ing.

The TAMARIN rules describing the binding process are shown in figure 4. As a convention, we will always typeset rules representing operations performed by OP-TEE with a light gray background. The first rule represents the generation of the host *challenge* as a fresh nonce $\sim cha^h$, and sending the *InitializeUpdate* message on the network. The syntax '*InitializeUpdate*' denotes a public name used for identifying the message. TAMARIN supports tuples through the $\langle t_1, \dots, t_n \rangle$ syntax. As a simplification, we do not model the I²C channel independently of the default network channel. The standard design pattern in TAMARIN to model protocols that execute in ordered steps is to use facts to represent the step sequence state: we save the host challenge for later using the *InitStTEE*($\sim cha^h$) fact. The fact *!TEE*(*tee_l*) is a state fact modeling initialization of OP-TEE, we can ignore the meaning of *tee_l* for now.

The second rule models the reception of the *InitializeUpdate* message by the secure element. We assume a persistent fact stating the existence of two Advanced Encryption Standard (AES) keys k_{enc} and k_{mac} that are static and shared by the host and the card prior to the binding. The secure element also generates a challenge and uses it with the received host challenge to derive the session keys s_{enc} and s_{mac} from k_{enc} and k_{mac} respectively. To model the Key Derivation Function (KDF) used for this derivation, we use

the TAMARIN built-in hashing function h . Then, the secure element generates a cryptogram from the s_{mac} key using the challenges and a static string 'card'. While this generation is specified to use the same KDF, we decided to model it without loss of generality using a user-specified function *cryptogram*. Finally, the secure element sends back its challenge and its cryptogram.

When the host receives the response message in the third rule, it proceeds to the exact same key and cryptogram derivation steps. To model the verification of the cryptogram, we use so-called *restrictions*: a restriction is a formula used to restrict the traces TAMARIN considers during the verification phase. This is the standard way in TAMARIN to filter out unwanted behaviors and to model branching or operation success. We denote the equality restriction using the \doteq operator: $cry^c \doteq cry^{h,c}$ means that we only consider traces where the two cryptograms are equal. Finally, the host generates its own cryptogram with the static string 'host' and sends it with the *ExternalAuthenticate* message to the card, which will perform a similar verification as modeled by the fourth rule. Upon verification success, the third and fourth rules both record authentication in the trace, and save the session keys s_{enc} and s_{mac} in permanent state facts, for subsequent encrypted and authenticated communication.

5.2 Deriving the Secure Storage Key

OP-TEE is an open-source implementation of a TEE that respects GlobalPlatform specification. Specifically, the TEE Internal Core API [21] defines a notion of *Trusted Storage* (also called *Secure Storage*). This storage should guarantee the confidentiality and integrity of the data, authentication, and separation between TAs.

The default implementation chosen by OP-TEE relies on encrypted file operations performed in the TEE on the Normal world file system of the REE. Data encryption is handled by a cascade of derived keys. This section focuses on SSK, a per-device key generated at OP-TEE boot time, according to the following scheme:

$$SSK = \text{HMAC}_{\text{SHA256}}(HUK, \text{Chip_ID})$$

⁶<https://globalplatform.org>

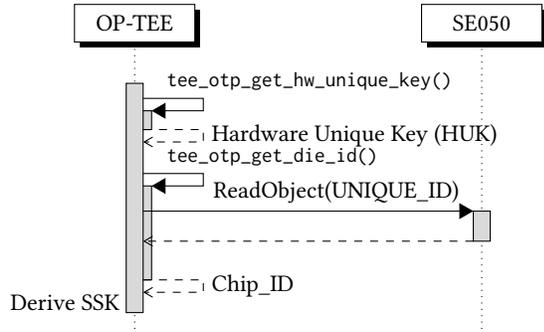


Figure 5: SSK derivation sequence diagram

The HUK and Chip_ID are platform-dependent, but the documentation recommends that HUK must be secret and not readable directly from software. In OP-TEE source code, the two following functions to access HUK and Chip_ID respectively are stubbed and need to be implemented by each platform:

```

TEE_Result tee_otp_get_hw_unique_key(...);
int tee_otp_get_die_id(...);

```

According to the source code, the Chip_ID is fetched from the SE050 secure element after the binding process and the establishment of the SCP03 secure channel. Figure 5 shows this ideal sequence diagram. The HUK is obtained from the i.MX board's own secure components. The Chip_ID is requested by OP-TEE using the established session keys s_{enc} and s_{mac} to encrypt and authenticate the APDU message *ReadObject* sent to the secure element, as required by the specification [32].

The corresponding TAMARIN rules are shown in figure 6. The first rule initializes OP-TEE and generates the secret HUK $\sim huk$. Again, we can ignore the facts mentioning $\sim tee_l$ for now. Similarly, the second rule initializes the SE050 secure element and generates the secret Chip_ID $\sim chipID$.

The third rule models sending the APDU message *ReadObject* to request the Chip_ID, named *UNIQUE_ID* in the secure element. The message is encrypted and MAC-ed using the session keys established previously by the SCP03 binding process described in section 5.1, as represented by the premise fact $!SessKeysTEE(s_{enc}, s_{mac})$. We use a user-supplied function symbol *mac* to represent the MAC operation, while the (symmetric) encryption uses the built-in TAMARIN function *senc*.

In the fourth rule, we use *pattern matching* to model the successful decryption and un-MAC-ing of the APDU message by the secure element. Pattern matching allows writing more compact rules that model the success of operations without relying on restrictions and function destructors, e.g., the *sdec* destructor symbol that models explicit decryption. In turn, the secure element answers with an encrypted and MAC-ed APDU message carrying the Chip_ID value.

Finally, on receiving this last message, OP-TEE can perform the SSK derivation $ssk = h(\langle huk, chipID \rangle)$. Again, pattern matching models successful decryption and un-MAC-ing.

5.3 Trusted Application execution

After the booting process, the REE and the TEE run side-by-side. The typical execution flow is that a Client Application (CA) running in the REE requests the execution of a Trusted Application (TA) by the TEE, to perform secure operations. In OP-TEE, this execution flow complies with TEE Client API Specification [19] which defines fundamental concepts:

Context: represents the connection between the CA and the TEE. Any number⁷ of concurrent contexts can be initialized.

Session: represents the connection between the CA and the TA, identified by its Universally Unique Identifier (UUID). Within a context, any number⁷ of concurrent sessions can be opened.

Command: represents the operation to be executed by the TA. Within a session, any number⁷ of concurrent commands can be called.

Shared Memory: allocated by the CA. It is used to transfer data between the CA and the TA.

Moreover, OP-TEE uses the shared memory to allow dynamic loading of the code of TAs. Figure 7 presents the sequence diagram corresponding to the loading and execution of a TA. We simplified the diagram: when an arrow crosses several components, it abstracts a sequence of messages between those components into one. The three labeled blocks *Contexts*, *Sessions*, and *Commands* represent the unbounded number of corresponding sub-sequences as loops. When a CA opens a session, OP-TEE will ask the REE to fetch the corresponding Executable and Linkable Format (ELF) binary code stored in its storage into the shared memory. Then upon verification of the signature of the code, OP-TEE loads the code from the shared memory into its own memory, for further execution of commands. Finally, sessions and contexts are closed.

The TAMARIN rules for our model are shown in figure 8. First, we simplify the execution of a CA, stating that the REE can initialize a context with the TEE at anytime, as long as the TEE has been previously initialized (modeled by the $!TEE(tee_l)$ fact). The second rule represents the initialization of a context by OP-TEE, identified by a fresh identifier $\sim ctx$. The action fact $InitContext(\sim ctx)$, along with the $LoopContext(ctx)$, $InitSession(ctx, uuid, sess)$ and $LoopSession(ctx, uuid, sess)$ action facts in further rules, will be used for proof purposes, as explained in the next section.

After the initialization of a context, the CA can request the opening of a session with a specific TA identified by its identifier *uuid*. We use the network channel to bind this identifier, to model the fact that the adversary might know it. In the fourth rule, OP-TEE receives the *OpenSession* message, opens a session identified by a fresh identifier $\sim sess$, and asks the REE to load the code of the corresponding TA into the shared memory.

When the REE receives the *LoadTA* message from OP-TEE, it loads the signed binary code in ELF format from its storage to the shared memory. The permanent fact $!Storage('TA', uuid, elf_s)$ models the binding of the identifier *uuid* to the code *elf_s* in a map called 'TA' inside the storage. Similarly, $SharedMem(sess, 'TA', uuid, elf_s)$ linearly models the binding of *uuid* to *elf_s* in a map called 'TA' representing a shared memory area allocated for the session *sess*. The

⁷Implementation-defined [19].

$$\begin{array}{c}
 \begin{array}{c}
 1 \frac{\text{Fr}(\sim huk) \quad \text{Fr}(\sim tee_l)}{\text{!HUK}(\sim huk) \quad \text{!TEE}(\sim tee_l)} [\text{InitTEE}(\sim huk, \sim tee_l)] \qquad \qquad \qquad 2 \frac{\text{Fr}(\sim chipID)}{\text{!ChipID}(\sim chipID)} [\text{InitSE}(\sim chipID)] \\
 \\
 3 \frac{\text{!SessKeysTEE}(s_{enc}, s_{mac})}{\text{ChipIDStTEE}() \quad \text{Out}(\text{mac}(\text{senc}(\langle \text{'ReadObject'}, \text{'UNIQUE_ID' \rangle), s_{enc}), s_{mac})} \\
 \\
 4 \frac{\text{!ChipID}(chipID) \quad \text{!SessKeysSE}(s_{enc}, s_{mac}) \quad \text{In}(\text{mac}(\text{senc}(\langle \text{'ReadObject'}, \text{'UNIQUE_ID' \rangle), s_{enc}), s_{mac})}{\text{Out}(\text{mac}(\text{senc}(\langle \text{'ReadObject_R'}, \text{'UNIQUE_ID'}, chipID \rangle), s_{enc}), s_{mac})} \\
 \\
 5 \frac{\text{!HUK}(huk) \quad \text{!SessKeysTEE}(s_{enc}, s_{mac}) \quad \text{ChipIDStTEE}() \quad \text{In}(\text{mac}(\text{senc}(\langle \text{'ReadObject_R'}, \text{'UNIQUE_ID'}, chipID \rangle), s_{enc}), s_{mac})}{\text{!SSK}(ssk)} [\text{DeriveSSK}(ssk)] \quad \text{where } ssk = h(\langle huk, chipID \rangle)
 \end{array}
 \end{array}$$

Figure 6: SSK derivation rewriting rules

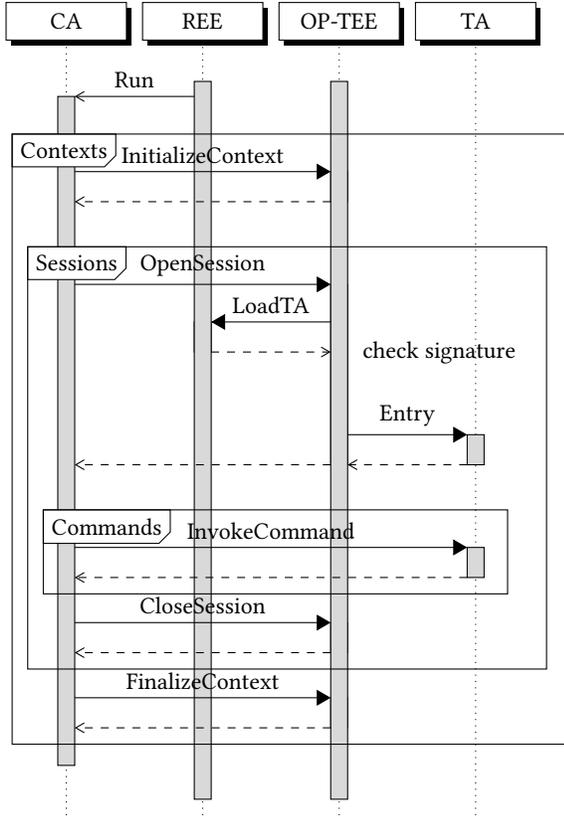


Figure 7: TA execution sequence diagram

sixth rule models the loading of the TA code from the shared memory to the OP-TEE memory, modeled by the Running fact. The verification of the signature of the binary code uses the built-in functions `revealSign`, `getMessage`, `revealVerify`, `pk` and `true`:

$$\begin{cases}
 \text{revealVerify}(\text{revealSign}(m, k_s), m, \text{pk}(k_s)) = \text{true}() \\
 \text{getMessage}(\text{revealSign}(m, k_s)) = m
 \end{cases}$$

The variable m represents the data to sign with the private key k_s . The public part of the key pair is given by $\text{pk}(k_s)$. In our rule, we again use a restriction to assert the success of the signature verification.

The seventh rule models the request from the CA to execute a given TA command, identified by its identifier cmd , with a given parameter $payload$. In full generality, depending on their size, the CA can pass parameters to the TA through the shared memory, but we do not model this feature in this work. We use the fact `ExecCmd(ctx, uuid, ~sess, cmd, payload)` as a state fact to keep track of the state of the TA waiting for the command to complete.

To model the execution of a command in the eighth rule, we use an abstract user-defined function `exec` that takes the ELF code of the TA, the session identifier, the command identifier and the input parameter, and returns the output. Following ideas from Jacomme et al. [24], we rely on *locations* to model the integrity of OP-TEE. The idea is to model the origin of computation using TAMARIN terms. This is akin to a signature-based authentication system. OP-TEE does not directly feature reporting, but our model adapt these ideas to model its integrity. The abstract function `report` produces a report proof that the command output has been produced at location tee_l . In turn, in the ninth rule, the CA checks the report using a restriction and the function `checkReport` which satisfies the following:

$$\text{checkReport}(\text{report}(m, l), m, l) = \text{true}()$$

The four last rules are straightforward: they model closing sessions and finalizing contexts.

We complete our model with rules in figure 9. The first one provisions the initial shared keys used during the binding process. The second one establishes the secret TA signing private key, and the last one installs the binary code, represented by the public name `'bin_code'`, in the `'TA'` map of the storage with key `'42'`.

5.4 Security properties and verification

5.4.1 Liveness properties. The first properties that we need to verify are *liveness*, or *executability* properties. Such properties ensure that the model is actually productive and that further security properties do not hold vacuously. To ensure liveness, it suffices to show

$$\begin{array}{c}
1 \frac{!TEE(tee_l)}{\text{Out}(\text{'InitializeContext'})} \quad 2 \frac{\text{In}(\text{'InitializeContext'}) \quad \text{Fr}(\sim ctx)}{\text{Context}(\sim ctx)} \text{[InitContext}(\sim ctx)\text{]} \\
3 \frac{\text{Context}(\sim ctx) \quad \text{In}(uuid)}{\text{Context}(\sim ctx) \quad \text{Out}(\text{'OpenSession'}, \sim ctx, uuid)} \text{[LoopContext}(\sim ctx)\text{]} \\
4 \frac{\text{In}(\text{'OpenSession'}, ctx, uuid) \quad \text{Fr}(\sim sess)}{\text{Session}(ctx, uuid, \sim sess) \quad \text{Out}(\text{'LoadTA'}, uuid)} \text{[InitSession}(ctx, uuid, \sim sess)\text{]} \\
5 \frac{\text{Session}(ctx, uuid, sess) \quad \text{In}(\text{'LoadTA'}, uuid) \quad !\text{Storage}(\text{'TA'}, uuid, elf_s)}{\text{Session}(ctx, uuid, sess) \quad \text{SharedMem}(sess, \text{'TA'}, uuid, elf_s)} \left[\begin{array}{l} \text{LoadedTA}(sess, uuid) \\ \text{LoopSession}(ctx, uuid, sess) \end{array} \right] \\
6 \frac{\text{Session}(ctx, uuid, sess) \quad !\text{TASig}(k_{ta_sig})}{\text{SharedMem}(sess, \text{'TA'}, uuid, elf_s) \quad \text{Running}(ctx, sess, elf)} \text{[LoopSession}(ctx, uuid, sess)\text{]} \quad \text{where} \left\{ \begin{array}{l} elf = \text{getMessage}(elf_s) \\ \text{revealVerify}(elf_s, elf, \text{pk}(k_{ta_sig})) \doteq \text{true}() \end{array} \right. \\
7 \frac{\text{Session}(ctx, uuid, \sim sess) \quad \text{In}(\langle cmd, x \rangle)}{\text{Session}(ctx, uuid, \sim sess) \quad \text{ExecCmd}(ctx, uuid, \sim sess, cmd, x)} \text{[LoopSession}(ctx, uuid, \sim sess)\text{]} \\
\quad \text{Out}(\text{'InvokeCommand'}, \sim sess, cmd, x) \\
8 \frac{\text{Running}(ctx, sess, elf) \quad !TEE(tee_l) \quad \text{In}(\text{'InvokeCommand'}, sess, cmd, x)}{\text{Running}(ctx, sess, elf) \quad \text{Out}(\text{'InvokeCommand_R'}, sess, cmd, y, rpt)} \quad \text{where} \left\{ \begin{array}{l} y = \text{exec}(elf, sess, cmd, x) \\ rpt = \text{report}(y, tee_l) \end{array} \right. \\
9 \frac{!TEE(tee_l) \quad \text{ExecCmd}(ctx, uuid, sess, cmd, x)}{\text{In}(\text{'InvokeCommand_R'}, sess, cmd, y, rpt)} \text{[ExecTA}(uuid, cmd, x, y)\text{]} \quad \text{where} \text{checkReport}(rpt, y, tee_l) \doteq \text{true}() \\
10 \frac{\text{Session}(ctx, uuid, \sim sess)}{\text{Session}(ctx, uuid, \sim sess) \quad \text{Out}(\text{'CloseSession'}, \sim sess)} \left[\begin{array}{l} \text{ClosingSession}(\sim sess) \\ \text{LoopSession}(ctx, uuid, \sim sess) \end{array} \right] \\
11 \frac{\text{Context}(\sim ctx)}{\text{Context}(\sim ctx) \quad \text{Out}(\text{'FinalizeContext'}, \sim ctx)} \left[\begin{array}{l} \text{FinalizingContext}(\sim ctx) \\ \text{LoopContext}(\sim ctx) \end{array} \right] \\
12 \frac{\text{Session}(ctx, uuid, sess) \quad \text{In}(\text{'CloseSession'}, sess)}{\quad} \quad 13 \frac{\text{Context}(ctx) \quad \text{In}(\text{'FinalizeContext'}, ctx)}{\quad}
\end{array}$$

Figure 8: TA execution rewriting rules

$$\frac{\frac{\text{Fr}(\sim k_{enc}) \quad \text{Fr}(\sim k_{mac})}{! \text{SharedKeys}(\sim k_{enc}, \sim k_{mac})} \quad \frac{\text{Fr}(\sim k_{ta_sig})}{! \text{TASig}(\sim k_{ta_sig})}}{! \text{Storage}(\text{'TA'}, \text{'42'}, \text{revealSign}(\text{'bin_code'}, k_{ta_sig}))} ! \text{TASig}(k_{ta_sig})$$

Figure 9: Initialization rewriting rules

that there exists at least one trace where the considered event appears.

LEMMA 5.1 (SCP03 LIVENESS). *There exists an execution where the binding process is successful, that is, there exists a trace satisfying:*

$$\text{Ex } \#i \#j. (\text{AuthTEE}() @ \#i) \& (\text{AuthSE}() @ \#j)$$

LEMMA 5.2 (SSK LIVENESS). *There exists an execution where the SSK is successfully derived, that is, there exists a trace satisfying:*

$$\text{Ex } ssk \#i. \text{DeriveSSK}(ssk) @ \#i$$

LEMMA 5.3 (TA EXECUTION LIVENESS). *There exists an execution where a command of the TA identified by '42' is successfully executed, that is, there exists a trace satisfying:*

$$\text{Ex } cmd \ x \ y \ \#i. \text{ExecTA}(\text{'42'}, cmd, x, y) @ \#i$$

TAMARIN automatically discharges lemmas 5.1 and 5.2, lemma 5.3 requires more work. The model describing the execution of the command of a TA is prone to entail the non-termination of TAMARIN's algorithm. We explain three techniques we applied to mitigate this issue, described in more details in TAMARIN's documentation [37].

Sort annotations. As a pre-computation step, TAMARIN tries to construct the derivation chains that are the source of premise facts of all the protocol rules. When it fails to find such sources, the *partial deconstruction chains* remaining might lead to non-termination. In

our model, some rules are responsible for partial deconstruction chains, but we apply a simple trick to fix them. The idea is to guide TAMARIN's pre-computation process by providing *atomicity* hints as sort annotations. For example, in the third rule of figure 8, we explicitly annotate the sort of $\sim ctx$ as fresh, to prevent TAMARIN from unifying the variable ctx with compound terms in the pre-computation process. We apply this technique to specific rules in a sound way in which we have the insurance that such annotated variables are never instantiated by non-atomic terms.

Restrictions. We already described restrictions as a way to model successful operations. We can also use them to make our model more precise and avoid unwanted looping behaviors leading to non-termination. In our model, we use a common pattern specifying as a restriction that a rule with an action fact F can only fire once:

$$\begin{aligned} & \text{All } x_1 \dots x_k \#i \#j. \\ & (F(x_1, \dots, x_k) @ \#i) \& (F(x_1, \dots, x_k) @ \#j) \implies \#i = \#j \end{aligned}$$

We apply this pattern to figure 8's fifth, tenth, and eleventh rules.

Induction. Some rules exhibit looping behavior, e.g., figure 8's third and fifth rules. Permanent facts are the standard way of dealing with this problem, but we need the Context and Session facts to be linear. In such cases, TAMARIN offers basic induction principles. The idea is to make TAMARIN first prove auxiliary lemmas using those induction principles, used later on for proving security lemmas whose direct verification could otherwise not terminate. In our case, the auxiliary lemmas are straightforward: we want to prove that a specific rule originally produces each *looping* fact. That is, we want to prove that the occurrences of such facts are well-founded.

LEMMA 5.4 (CONTEXTS ARE WELL-FOUNDED). *For all executions, any recurring occurrence of a context is preceded originally by its creation, that is, all traces satisfy:*

$$\begin{aligned} & \text{All } ctx \#i. \\ & \text{LoopContext}(ctx) @ \#i \implies \\ & \text{Ex } \#j. (\text{InitContext}(ctx) @ \#j) \& (\#i < \#j) \end{aligned}$$

LEMMA 5.5 (SESSIONS ARE WELL-FOUNDED). *For all executions, any recurring occurrence of a session is preceded originally by its creation, that is, all traces satisfy:*

$$\begin{aligned} & \text{All } ctx \text{ uuid } sess \#i. \\ & \text{LoopSession}(ctx, \text{uuid}, sess) @ \#i \implies \\ & \text{Ex } \#j. (\text{InitSession}(ctx, \text{uuid}, sess) @ \#j) \& (\#i < \#j) \end{aligned}$$

5.4.2 *Security properties.* TAMARIN can automatically prove basic secrecy properties about our model.

LEMMA 5.6 (HUK SECURITY). *The adversary never learns the HUK, that is, all traces satisfy:*

$$\begin{aligned} & \text{All } huk \text{ tee}_1 \#i. \\ & \text{InitTEE}(huk, \text{tee}_1) @ \#i \implies \text{not } (\text{Ex } \#j. K(huk) @ \#j) \end{aligned}$$

LEMMA 5.7 (CHIP_ID SECURITY). *The adversary never learns the Chip_ID, that is, all traces satisfy:*

$$\begin{aligned} & \text{All } chipID \#i. \\ & \text{InitSE}(chipID) @ \#i \implies \text{not } (\text{Ex } \#j. K(chipID) @ \#j) \end{aligned}$$

We can also ensure that the adversary cannot produce reports impersonating the TEE.

LEMMA 5.8 (TEE REPORTING AUTHENTICITY). *The adversary never learns the TEE location, that is, all traces satisfy:*

$$\begin{aligned} & \text{All } huk \text{ tee}_1 \#i. \\ & \text{InitTEE}(huk, \text{tee}_1) @ \#i \implies \text{not } (\text{Ex } \#j. K(\text{tee}_1) @ \#j) \end{aligned}$$

6 DISCUSSION AND CONCLUSION

We presented a TAMARIN model supporting analysis of critical features of a security-oriented IoT device. TAMARIN automatically verifies our model of about 400 lines in about 15 s, on a laptop equipped with a 10-core 12th Gen. Intel Core i7-1250U processor and 15.25 GiB RAM. TAMARIN proves all our security properties correct, meaning that it could not find any attack invalidating these properties on the model. Generally speaking, anticipating performance cost is hard during the design phase. For example, we use an additional auxiliary lemma proved by induction specifically for the Running recurring fact in the eighth rule of figure 8. This lemma is unnecessary for ensuring termination, but without it, the verification time increases by one order of magnitude (150 s).

We designed our model to be simple enough to emphasize on feasibility and generality and we only proved fundamental but basic security properties. The main limitations reside in the TA execution representation. We want to model higher-order execution in full generality, but to our knowledge, none of the automatic tools for security analysis offers this kind of feature. In this paper, we adapt ideas of Jacomme et al. [24], but modeling, e.g., encrypted file operations by a TA involves further challenges that we want to tackle as future work. The alternative SAPIC front-end of TAMARIN implements these ideas for modeling TEEs, along with concepts about global and shared state [27], and we believe that we could adapt our model to this language once its development reaches a more stable state. Another related limitation is the discussed potential vulnerability in the communication with the secure element SE050 through the REE's I²C driver. To verify the safety of this communication, we would need to model the interaction of a TA with the SE050 and the details of the routing of the communication between the TA, OP-TEE, the REE's driver, and the SE050. In future work, we could also allow the adversary to tamper the storage and the shared memory directly.

Our model is general enough to be relevant for a large class of secure IoT devices following the same standard TEE-based architecture using a secure element. In the future, we will extend this first proof-of-concept model to support more features and protocols of the device. In such cases, the ability of using *composition* techniques would be highly beneficial to the automatic approach taken by TAMARIN and other similar tools. This interesting research topic has seen recent advances [22] that we also want to investigate.

ACKNOWLEDGMENTS

This work was supported by JST, CREST Grant Number JPMJCR21M3, and by ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603).

REFERENCES

- [1] 2023. OP-TEE. Trusted Firmware.
- [2] Martín Abadi, Bruno Blanchet, and Cédric Fournet. 2017. The Applied Pi Calculus: Mobile Values, New Names, and Secure Communication. *J. ACM* 65, 1 (Oct. 2017), 1:1–1:41. <https://doi.org/10.1145/3127586>

- [3] Martin Abadi and Cédric Fournet. 2001. Mobile Values, New Names, and Secure Communication. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '01)*. Association for Computing Machinery, New York, NY, USA, 104–115. <https://doi.org/10.1145/360204.360213>
- [4] A. Armando, D. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuellar, P. Hankes Drielsma, P. C. Héam, O. Kouchnarenko, J. Mantovani, S. Mödersheim, D. von Oheimb, M. Rusinowitch, J. Santiago, M. Turuani, L. Viganò, and L. Vigneron. 2005. The AVISPA Tool for the Automated Validation of Internet Security Protocols and Applications. In *Computer Aided Verification (CAV'05)*, Kousha Etessami and Sriram K. Rajamani (Eds.). Springer, Berlin, Heidelberg, 281–285. https://doi.org/10.1007/11513988_27
- [5] Karthikeyan Bhargavan, Bruno Blanchet, and Nadim Kobeissi. 2017. Verified Models and Reference Implementations for the TLS 1.3 Standard Candidate. In *2017 IEEE Symposium on Security and Privacy (SP'17)*. 483–502. <https://doi.org/10.1109/SP.2017.26>
- [6] Bruno Blanchet. 2001. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *Proceedings of the 14th IEEE Workshop on Computer Security Foundations (CSFW '01)*. IEEE Computer Society, USA, 82–96. <https://doi.org/10.1109/CSFW.2001.930138>
- [7] Bruno Blanchet. 2017. Symbolic and Computational Mechanized Verification of the ARINC823 Avionic Protocols. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF'17)*. 68–82. <https://doi.org/10.1109/CSF.2017.7>
- [8] Z. Berkay Celik, Patrick McDaniel, and Gang Tan. 2018. SOTERIA: Automated IoT Safety and Security Analysis. In *Proceedings of the 2018 USENIX Conference on Usenix Annual Technical Conference (USENIX ATC '18)*. USENIX Association, USA, 147–158.
- [9] Vincent Cheval, Charlie Jacomme, Steve Kremer, and Robert Künnemann. 2022. SAPIC+: Protocol Verifiers of the World, Unite!. In *31st USENIX Security Symposium (USENIX Security 22) (USENIX'22)*. USENIX Association, Boston, MA, 3935–3952.
- [10] Vincent Cheval, Steve Kremer, and Itsaka Rakotonirina. 2018. DEEPSEC: Deciding Equivalence Properties in Security Protocols Theory and Practice. In *2018 IEEE Symposium on Security and Privacy (SP'18)*. 529–546. <https://doi.org/10.1109/SP.2018.00033>
- [11] Cas Cremers, David Basin, Jannik Dreier, Simon Meier, Ralf Sasse, and Benedikt Schmidt. 2021. Tamarin Prover.
- [12] Cas Cremers and Martin Dehnel-Wild. 2019. Component-Based Formal Analysis of 5G-AKA: Channel Assumptions and Session Confusion. *Network and Distributed System Security Symposium* (2019). <https://doi.org/10.14722/ndss.2019.23394>
- [13] Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. 2017. A Comprehensive Symbolic Analysis of TLS 1.3. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17)*. Association for Computing Machinery, New York, NY, USA, 1773–1788. <https://doi.org/10.1145/3133956.3134063>
- [14] Cas Cremers, Benjamin Kiesel, and Niklas Medinger. 2020. A Formal Analysis of IEEE 802.11's WPA2: Countering the Kracks Caused by Cracking the Counters. In *Proceedings of the 29th USENIX Conference on Security Symposium (SEC'20)*. USENIX Association, USA, 1–17. <https://doi.org/10.5555/3489212.3489213>
- [15] Danny Dolev and Andrew C. Yao. 1983. On the Security of Public Key Protocols. *IEEE Transactions on Information Theory* 29, 2 (March 1983), 198–208. <https://doi.org/10.1109/TIT.1983.1056650>
- [16] Santiago Escobar, Catherine Meadows, and José Meseguer. 2009. Maude-NPA: Cryptographic Protocol Analysis Modulo Equational Properties. In *Foundations of Security Analysis and Design V: FOSAD 2007/2008/2009 Tutorial Lectures*, Alessandro Aldini, Gilles Barthe, and Roberto Gorrieri (Eds.). Springer, Berlin, Heidelberg, 1–50.
- [17] Zheng Fang, Hao Fu, Tianbo Gu, Pengfei Hu, Jinyue Song, Trent Jaeger, and Prasant Mohapatra. 2022. IOTA: A Framework for Analyzing System-Level Security of IoTs. In *2022 IEEE/ACM Seventh International Conference on Internet-of-Things Design and Implementation (IoTDI) (IoTDI'22)*. 143–155. <https://doi.org/10.1109/IoTDI54339.2022.00017>
- [18] Trusted Firmware. 2023. *OP-TEE documentation*. <https://optee.readthedocs.io/en/3.22.0/>
- [19] GlobalPlatform. 2010. *TEE Client API Specification*. <https://globalplatform.org/specs-library/tee-client-api-specification/>
- [20] GlobalPlatform. 2020. *Secure Channel Protocol '03'*. <https://globalplatform.org/specs-library/secure-channel-protocol-03-amendment-d-v1-2/>
- [21] GlobalPlatform. 2021. *TEE Internal Core API Specification*. <https://globalplatform.org/specs-library/tee-internal-core-api-specification/>
- [22] Andreas V. Hess, Sebastian A. Mödersheim, and Achim D. Brucker. 2023. Stateful Protocol Composition in Isabelle/HOL. *ACM Transactions on Privacy and Security* (Jan. 2023). <https://doi.org/10.1145/3577020>
- [23] Katharina Hofer-Schmitz and Branka Stojanović. 2020. Towards Formal Verification of IoT Protocols: A Review. *Computer Networks* 174 (June 2020), 107233. <https://doi.org/10.1016/j.comnet.2020.107233>
- [24] Charlie Jacomme, Steve Kremer, and Guillaume Scerri. 2017. Symbolic Models for Isolated Execution Environments. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P) (EuroS&P'17)*. 530–545. <https://doi.org/10.1109/EuroSP.2017.16>
- [25] Jun Young Kim, Ralph Holz, Wen Hu, and Sanjay Jha. 2017. Automated Analysis of Secure Internet of Things Protocols. In *Proceedings of the 33rd Annual Computer Security Applications Conference (ACSAC '17)*. Association for Computing Machinery, New York, NY, USA, 238–249. <https://doi.org/10.1145/3134600.3134624>
- [26] Nadim Kobeissi, Karthikeyan Bhargavan, and Bruno Blanchet. 2017. Automated Verification for Secure Messaging Protocols and Their Implementations: A Symbolic and Computational Approach. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P'17)*. 435–450. <https://doi.org/10.1109/EuroSP.2017.38>
- [27] Steve Kremer and Robert Künnemann. 2014. Automated Analysis of Security Protocols with Global State. In *35th IEEE Symposium on Security and Privacy (S&P'14) (S&P'14)*, IEEE Computer Society (Ed.). San Jose, United States, 163–178. <https://doi.org/10.1109/SP.2014.18>
- [28] Simon Meier, Benedikt Schmidt, Cas Cremers, and David Basin. 2013. The TAMARIN Prover for the Symbolic Analysis of Security Protocols. In *Computer Aided Verification (CAV'13)*, Natasha Sharygina and Helmut Veith (Eds.). Springer, Berlin, Heidelberg, 696–701. https://doi.org/10.1007/978-3-642-39799-8_48
- [29] J Mitchell, A Scedrov, N Durgin, and P Lincoln. 1999. Undecidability of Bounded Security Protocols. In *Workshop on Formal Methods and Security Protocols*.
- [30] Bernard Ngabonziza, Daniel Martin, Anna Bailey, Haehyun Cho, and Sarah Martin. 2016. TrustZone Explained: Architectural Features and Use Cases. In *2016 IEEE 2nd International Conference on Collaboration and Internet Computing (CIC) (CIC'16)*. 445–451. <https://doi.org/10.1109/CIC.2016.065>
- [31] Dang Tu Nguyen, Chengyu Song, Zhiyun Qian, Srikanth V. Krishnamurthy, Edward J. M. Colbert, and Patrick McDaniel. 2018. IotSan: Fortifying the Safety of IoT Systems. In *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '18)*. Association for Computing Machinery, New York, NY, USA, 191–203. <https://doi.org/10.1145/3281411.3281440>
- [32] NXP. 2021. *SE050 APDU Specification*. <https://www.nxp.com/docs/en/application-note/AN12413.pdf>
- [33] Xinming Ou, Sudhakar Govindavajhala, and Andrew W. Appel. 2005. MulVAL: A Logic-Based Network Security Analyzer. In *Proceedings of the 14th Conference on USENIX Security Symposium - Volume 14 (SSYM'05)*. USENIX Association, USA, 8.
- [34] Mohamed Sabt, Mohammed Achemlal, and Abdelmadjid Bouabdallah. 2015. Trusted Execution Environment: What It Is, and What It Is Not. In *2015 IEEE TrustCom/BigDataSE/ISPA (TrustCom'15, Vol. 1)*. 57–64. <https://doi.org/10.1109/TrustCom.2015.357>
- [35] Benedikt Schmidt, Simon Meier, Cas Cremers, and David Basin. 2012. Automated Analysis of Diffie-Hellman Protocols and Advanced Security Properties. In *2012 IEEE 25th Computer Security Foundations Symposium (CSF'12)*. 78–94. <https://doi.org/10.1109/CSF.2012.25>
- [36] Carlton Shepherd, Ghada Arfaoui, Iakovos Gurulian, Robert P. Lee, Konstantinos Markantonakis, Raja Naeem Akram, Damien Sauveron, and Emmanuel Conchon. 2016. Secure and Trusted Execution: Past, Present, and Future - A Critical Review in the Context of the Internet of Things and Cyber-Physical Systems. In *2016 IEEE TrustCom/BigDataSE/ISPA (TrustCom'16)*. 168–177. <https://doi.org/10.1109/TrustCom.2016.0060>
- [37] The Tamarin Team. 2023. *Tamarin-Prover Manual*. <https://tamarin-prover.github.io/manual/master/tex/tamarin-manual.pdf>
- [38] Atmark Techno. 2023. *Armadillo-IoT ゲートウェイ G4 セキュリティガイド*. https://manual.atmark-techno.com/armadillo-iot-g4/armadillo-base-os-security-guide_ja-1.3.1/
- [39] Teng Xu, James B. Wendt, and Miodrag Potkonjak. 2014. Security of IoT Systems: Design Challenges and Opportunities. In *2014 IEEE/ACM International Conference on Computer-Aided Design (ICCAD) (ICCAD'14)*. 417–423. <https://doi.org/10.1109/ICCAD.2014.7001385>