

A Touch of Evil: High-Assurance Cryptographic Hardware from Untrusted Components

Vasilios Mavroudis
University College London
v.mavroudis@cs.ucl.ac.uk

Andrea Cerulli
University College London
andrea.cerulli.13@ucl.ac.uk

Petr Svenda
Masaryk University
svenda@fi.muni.cz

Dan Cvrcek
EnigmaBridge
dan@enigmabridge.com

Dusan Klinec
EnigmaBridge
dusan@enigmabridge.com

George Danezis
University College London
g.danezis@ucl.ac.uk

ABSTRACT

The semiconductor industry is fully globalized and integrated circuits (ICs) are commonly defined, designed and fabricated in different premises across the world. This reduces production costs, but also exposes ICs to supply chain attacks, where insiders introduce malicious circuitry into the final products. Additionally, despite extensive post-fabrication testing, it is not uncommon for ICs with subtle fabrication errors to make it into production systems. While many systems may be able to tolerate a few byzantine components, this is not the case for cryptographic hardware, storing and computing on confidential data. For this reason, many error and backdoor detection techniques have been proposed over the years. So far all attempts have been either quickly circumvented, or come with unrealistically high manufacturing costs and complexity.

This paper proposes *Myst*, a practical high-assurance architecture, that uses commercial off-the-shelf (COTS) hardware, and provides strong security guarantees, even in the presence of multiple malicious or faulty components. The key idea is to combine protective-redundancy with modern threshold cryptographic techniques to build a system tolerant to hardware trojans and errors. To evaluate our design, we build a Hardware Security Module that provides the highest level of assurance possible with COTS components. Specifically, we employ more than a hundred COTS secure cryptoprocessors, verified to FIPS140-2 Level 4 tamper-resistance standards, and use them to realize high-confidentiality random number generation, key derivation, public key decryption and signing. Our experiments show a reasonable computational overhead (less than 1% for both Decryption and Signing) and an exponential increase in backdoor-tolerance as more ICs are added.

KEYWORDS

cryptographic hardware; hardware trojans; backdoor-tolerance; secure architecture

1 INTRODUCTION

Many critical systems with high security needs rely on secure cryptoprocessors to carry out sensitive security tasks (e.g., key generation and storage, legally binding digital signature, code signing) and provide a protection layer against cyber-attacks and security breaches. These systems are typically servers handling sensitive data, banking infrastructure, military equipment and space stations. In most cases, secure cryptoprocessors come embedded into Hardware Security Modules, Trusted Platform Modules and Cryptographic Accelerators, which are assumed to be both secure and reliable. This entails that errors in any of the Integrated Circuits (ICs) would be devastating for the security of the final system. For this reason, the design and fabrication of the underlying ICs must abide to high-quality specifications and standards. These ensure that there are no intentional or unintentional errors in the circuits, but more importantly ensure the integrity of the hardware supply chain. [52].

Unfortunately, vendors are not always able to oversee all parts of the supply chain [38, 60]. The constant reduction in transistor size makes IC fabrication an expensive process, and IC designers often outsource the fabrication task to overseas foundries to reduce their costs [35, 46, 99]. This limits vendors to run only post-fabrication tests to uncover potential defects. Those tests are very efficient against common defects, but subtle errors are hard to uncover. For instance, cryptoprocessors with defective RNG modules and hardware cipher implementations have made it into production in the past [31, 39].

Additionally, parts of the IC's supply chain are left vulnerable to attacks from malicious insiders [12, 63, 67, 84] and have a higher probability of introducing unintentional errors in the final product. In several documented real-world cases, contained errors, backdoors or trojan horses. For instance, recently an exploitable vulnerability was discovered on Intel processors that utilize Intel Active Management Technology (AMT) [49], while vulnerable ICs have been reported in military [59, 78] applications, networking equipment [40, 50], and various other application [2, 56, 76, 77]. Furthermore, the academic community has designed various types of hardware trojans (HT), and backdoors that demonstrate the extent of the problem and its mitigation difficulty [11, 18, 23, 54, 65, 91–93].

Due to the severity of these threats, there is a large body of work on the mitigation of malicious circuitry. Existing works have pursued two different directions: detection and prevention. Detection techniques aim to determine whether any HTs exist in a given circuit [3, 79, 95, 97], while prevention techniques either

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '17, October 30–November 3, 2017, Dallas, TX, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-4946-8/17/10...\$15.00
<https://doi.org/10.1145/3133956.3133961>

impede the introduction of HTs, or enhance the efficiency of HT detection [28, 69, 88, 89, 94]. Unfortunately, both detection and prevention techniques are brittle, as new threats are able to circumvent them quickly [94]. For instance, analog malicious hardware [98] was able to evade known defenses, including split manufacturing, which is considered one of the most promising and effective prevention approaches. Furthermore, most prevention techniques come with a high manufacturing cost for higher levels of security [19, 28, 94], which contradicts the motives of fabrication outsourcing. To make matters worse, vendors that use commercial off-the-shelf (COTS) components are constrained to use only post-fabrication detection techniques, which further limits their mitigation toolchest. All in all, backdoors being triggered by covert means and mitigation countermeasures are in an arms race that seems favorable to the attackers [94, 96].

In this paper, we propose *Myst* a new approach to the problem of building trustworthy cryptographic hardware. Instead of attempting to detect or prevent hardware trojans and errors, we proposed and provide evidence to support the hypothesis:

“We can build high-assurance cryptographic hardware from a set of untrusted components, as long as at least one of them is not compromised, even if benign and malicious components are indistinguishable.”

Our key insight is that by combining established privacy enhancing technologies (PETs), with mature fault-tolerant system architectures, we can distribute trust between multiple components originating from non-crossing supply chains, thus reducing the likelihood of compromises. To achieve this, we deploy distributed cryptographic schemes on top of an N-variant system architecture, and build a trusted platform that supports a wide-range of commonly used cryptographic operations (e.g., random number and key generation, decryption, signing). This design draws from protective-redundancy and component diversification [27] and is built on the assumption that multiple processing units and communication controllers may be compromised by the same adversary. However, unlike N-variant systems, instead of replicating the computations on all processing units, *Myst* uses multi-party cryptographic schemes to distribute the secrets so that the components hold only shares of the secrets (and not the secrets themselves), at all times. As long as one of the components remains honest, the secret cannot be reconstructed or leaked. Moreover, we can tolerate two or more non-colluding adversaries who have compromised 100% of the components.

Our proposed architecture is of particular interest for two distinct categories of hardware vendors:

- ❖ Design houses that outsource the fabrication of their ICs.
- ❖ COTS vendors that rely on commercial components to build their high-assurance hardware.

Understandably, design houses have much better control over the IC fabrication and the supply chain, and this allows them to take full advantage of our architecture. In particular, they can combine existing detection [8, 26, 68] and prevention techniques [25, 28, 69, 94] with our proposed design, to reduce the likelihood of compromises for individual components. On the other hand, COTS vendors have less control as they have limited visibility in the fabrication process

and the supply chain. However, they can still mitigate risk by using ICs from sources, known to run their own fabrication facilities.

To our knowledge, this is the first work that uses distributed cryptographic protocols to build and evaluate a hardware module architecture that is tolerant to multiple components carrying trojans or errors. The effectiveness of this class of protocols for the problem of hardware trojans has been also studied in previous theoretical works [6, 34].

To summarize, this paper makes the following contributions:

- ❖ **Concept:** We introduce *backdoor-tolerance*, where a system can still preserve its security properties in the presence of multiple compromised components.
- ❖ **Design:** We demonstrate how cryptographic schemes (§4) can be combined with N-variant system architectures (§3), to build high-assurance systems. Moreover, we introduce a novel distributed signing scheme based on the Schnorr blind signatures (§4.5).
- ❖ **Implementation:** We implement the proposed architecture by building a custom board featuring 120 highly-tamper resistant ICs, and realize secure variants of random number and key generation, public key decryption and signing (§5).
- ❖ **Optimizations:** We implement a number of optimizations to ensure the *Myst* architecture is competitive in terms of performance compared to single ICs. Some optimizations also concern embedded mathematical libraries which are of independent interest.
- ❖ **Evaluation:** We evaluate the performance of *Myst*, and use micro-benchmarks to assess the cost of all operations and bottlenecks (§6).

Related works and their relation to *Myst* are discussed in Section 7.

2 PRELIMINARIES

In this section, we introduce *backdoor-tolerance*, and outline our security assumptions for adversaries residing in the IC’s supply chain.

2.1 Definition

A *Backdoor-Tolerant* system is able to ensure the confidentiality of the secrets it stores, the integrity of its computations and its availability, against an adversary with defined capabilities. Such a system can prevent data leakages and protect the integrity of the generated keys and random numbers. Note that the definition makes no distinction between honest design or fabrication mistakes and hardware trojans, and accounts only for the impact these errors have on the security properties of the system.

2.2 Threat Model

We assume an adversary is able to incorporate errors (e.g., a hardware trojan) in some ICs but not all of them. This is because, ICs are manufactured in many different fabrication facilities and locations by different vendors and the adversary is assumed not to be able to breach all their supply chains. Malicious ICs aim to be as stealthy as possible, and conceal their presence as best as possible, while benign components may have hard to detect errors (e.g., intermittent faults) that cannot be trivially uncovered by post-fabrication tests.

Hence, malicious and defective components are indistinguishable from benign/non-faulty ones.

The adversary can gain access to the secrets stored in the malicious ICs and may also breach the integrity of any cryptographic function run by the IC (e.g., a broken random number generator). Moreover, the adversary has full control over the communication buses used to transfer data to and from the ICs. Hence, they are able to exfiltrate any information on the bus and the channel controller, and inject and drop messages to launch attacks. Additionally, they are able to connect and issue commands to the ICs, and if a system features more than one malicious ICs, the adversary is able to orchestrate attacks using all of them (i.e., colluding ICs). We also assume that the adversary may use any other side-channel that a hardware Trojan has been programmed to emit on – such as RF [5], acoustic [41] or optical [7] channels. Additionally, the adversary can trigger malicious ICs to cease functioning (i.e., availability attacks). We also make standard cryptographic assumptions about the adversary being computationally bound.

Finally, we assume a software developer or device operator builds and signs the applications to be run on the ICs. From our point of view they are trusted to provide high-integrity software without backdoors or other flaws. This integrity concern may be tackled through standard high-integrity software development techniques, such as security audits, public code repository trees, deterministic builds, etc. The operator is also trusted to initialize the device properly and choose ICs and define diverse quorums so that the probability of compromises is minimized.

3 ARCHITECTURE

In this section, we introduce the Myst architecture (build and evaluation in § 5 and 6), which draws from fault-tolerant designs and N-variant systems [27, 82]. The proposed design is based on the thesis that given a diverse set of k ICs sourced from k untrusted suppliers, a trusted system can be built, as long as at least one of them is honest. Alternatively, our architecture can tolerate up to 100% compromised or defective components, if not all of them collude with each other. As illustrated in Figure 1, Myst has three types of components: (1) a remote host, (2) an untrusted IC controller, and (3) the processing ICs.

Processing ICs. The ICs form the core of Myst, as they are collectively entrusted to perform high-integrity computations, and provide secure storage space. The ICs are programmable processors or microprocessors. They have enough persistent memory to store keys and they feature a secure random number generator. Protection against physical tampering or side-channels is also desirable in most cases and mandated by the industry standards and best practices for cryptographic hardware. For this reason, our prototype (Section 5), comprises of components that verify to the a very high level of tamper-resistance (i.e., FIPS140-2 Level 4), and feature a reliable random number generator (i.e., FIPS140-2 Level 3). Each implementation must feature two or more ICs, of which at least one must be free of backdoors. We define this coalition of ICs, as a *quorum*. The exact number of ICs in a quorum is determined by the user depending on the degree of assurance she wants to achieve (see also Subsection 3.2 and the Assessment Subsection 6).

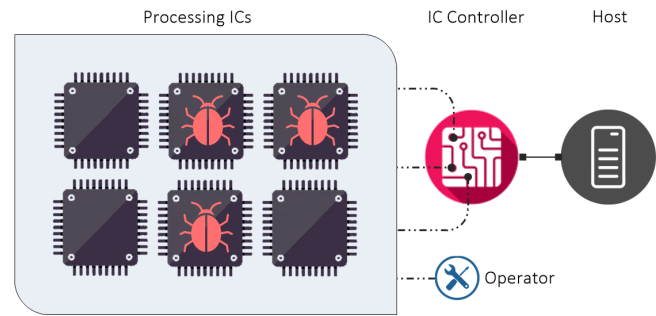


Figure 1: An overview of the Myst's distributed architecture, featuring all the integral components and communication buses. The gray area represents the cryptographic device, featuring several untrusted cryptoprocessors (ICs). As shown, the trusted operator interacts with individual ICs, while the host interacts with the device as a whole.

IC controller. The controller enables communication between the ICs, and makes Myst accessible to the outside world. It manages the bus used by the processing ICs to communicate with each other and the interface where user requests are delivered to. Specifically, it enables:

- ❖ *Unicast, IC-to-IC:* an IC A sends instructions an IC B, and receives the response.
- ❖ *Broadcast, IC-to-ICs:* an IC A broadcasts instructions to all other ICs, and receive their responses.
- ❖ *Unicast, Host-to-IC:* a remote client send commands to a specific IC, and receives the response.
- ❖ *Broadcast, Host-to-ICs:* a remote client broadcasts commands to all ICs, and receive their responses.

The controller is also an IC and is also untrusted. For instance it may drop, modify packets or forge new ones, in order to launch attacks against the protocols executed. It may also collude with one or more processing ICs.

Operator. The operator is responsible for setting up the system and configuring it. In particular, they are responsible for sourcing the ICs and making sure the quorums are as diverse as possible. Moreover, the operator determines the size of the quorums, and sets the security parameters of the cryptosystem (Section 4). They are assumed to make a best effort to prevent compromises and may also be the owner of the secrets stored in the device.

Remote Host. The remote host connects to Myst through the IC controller; it issues high level commands and retrieves the results. The remote host can be any kind of computer either in the local network or in a remote one. In order for the commands to be executed by the ICs, the host must provide proof of its authorization to issue commands, usually by signing them with a public key associated with the user's identity (e.g., certificate by a trusted CA). Each command issued must include: 1) the requested operation, 2) any input data, and 3) the host's signature (see also Section 3.1). We note that a corrupt host may use Myst to perform operations, but cannot extract any secrets, or forge signatures (see also § 3.1).

Communication Channels. At the physical level, the ICs, the controller, and the hosts are connected through buses and network interfaces. Hence, all messages containing commands, as well as their eventual responses are routed to and from the ICs through untrusted buses. We assume that the adversary is able to eavesdrop on these buses and tamper with their traffic (e.g., inject or modify packets). To address this, we use established cryptographic mechanisms to ensure integrity and confidentiality for transmitted data. More specifically, all unicast and broadcast packets are signed using the sender IC's certificate, so that their origin and integrity can be verified by recipients. Moreover, in cases where the confidentiality of the transmitted data needs to be protected, encrypted end-to-end channels are also established. Such encrypted channels should be as lightweight as possible to minimize performance overhead, yet able to guarantee both the confidentiality and the integrity of the transmitted data.

3.1 Access Control Layer

Access Control (AC) is critical for all systems operating on confidential data. In Myst, AC determines which hosts can submit service requests to the system, and which quorums they can interact with. Despite the distributed architecture of Myst, we can simply replicate established AC techniques on each IC. More specifically, each IC is provided with the public keys of the hosts that are allowed to have access to its shares and submit requests. Optionally, this list may distinguish between hosts that have full access to the system, and hosts that may only perform a subset of the operations. Once a request is received, the IC verifies the signature of the host against this public key list. The list can be provided either when setting up Myst, or when storing a secret or generating a key.

We note that it is up to the operator to decide the parameters of each quorum (i.e., size k , ICs participating), and provide the AC lists to the ICs. This is a critical procedure, as if one of the hosts turns out to be compromised, the quorum can be misused in order to either decrypt confidential ciphertexts or sign bogus statements. However, the secrets stored in the quorum will remain safe as there is no way for the adversary to extract them, unless they use physical-tampering (which our prototype (§5) is also resilient against). This is also true in the case where one of the authorized hosts gets compromised. For instance, a malware taking over a host can use Myst to sign documents, but it cannot under any circumstances extract any secrets.

3.2 Reliability Estimation

In the case of cryptographic hardware, in order for the operator to decide on the threshold k and the quorum composition, an estimation of the likelihood of hardware errors is needed. For this purpose we introduce *k-tolerance*, which given k foundries and an independent error probability, computes the probability that a quorum is secure.

$$\Pr[\text{secure}] = 1 - \Pr[\text{error}]^k \quad (1)$$

The quantification of the above parameters depends on the particular design and case, and as there is not commonly accepted way of evaluation, it largely depends on the information and testing

capabilities each vendor has. For instance, hardware designers that use split manufacturing [28, 69, 94] can estimate the probability of a backdoored component using the *k-security* metric [48]. On the other hand, COTS vendors cannot easily estimate the probability of a compromised component, as they are not always aware of the manufacturing technical details. Despite that, it is still possible for them to achieve very high levels of error and backdoor-tolerance by increasing the size of the quorums and sourcing ICs from distinct facilities (i.e., minimizing the collusion likelihood). It should be noted that as k grows the cost increases linearly, while the security exponentially towards one.

4 SECURE DISTRIBUTED COMPUTATIONS

In this section, we introduce a set of protocols that leverage the diversity of ICs in Myst to realize standard cryptographic operations manipulating sensitive keying material. More specifically, our cryptosystem comprises of a key generation operation (§4.1), the ElGamal encryption operation (§4.2), distributed ElGamal decryption (§4.3), and distributed signing based on Schnorr signatures (§4.5) [36, 37]. These operations are realized using interactive cryptographic protocols that rely on standard cryptographic assumptions. For operational purposes, we also introduce a key propagation protocol that enables secret sharing between non-overlapping quorums (§4.6).

Prior to the execution of any protocols, the ICs must be initialized with the domain parameters $T = (p, a, b, G, n, h)$ of the elliptic curve to be used, where p is a prime specifying the finite field \mathbb{F}_p , a and b are the curve coefficients, G is the base point of the curve, with order n , and h is the curve's cofactor. More details on the technical aspects of the initialization procedure are provided in the provisioning Section 6.4.

Distributed Key Pair Generation. The Distributed Key Pair Generation (DKPG) is a key building block for most other protocols. In a nutshell, DKPG enables a *quorum* Q of k ICs to collectively generate a random secret x , which is an element of a finite field and a public value $Y_{agg} = x \cdot G$ for a given public elliptic curve point G . At the end of the DKPG protocol, each player holds a share of the secret x , denoted as x_i and the public value Y_{agg} . All steps of the protocol are illustrated in Figure 2, and explained in turn below.

The execution of the protocol is triggered when an authorized host sends the corresponding instruction (Ⓢ). At the first step of the protocol, each member of Q runs Algorithm 4.1 and generates a triplet consisting of: 1) a share x_i , which is a randomly sampled element from \mathbb{Z}_n , 2) an elliptic curve point Y_i , and 3) a commitment to Y_i denoted h_i . (Ⓣ)

Algorithm 4.1: TripletGen: Generation of a pair and its commitment.

Input : The domain parameters λ
Output : A key triplet (x_i, Y_i, h_i)

- 1 $x_i \leftarrow \text{Rand}(\lambda)$
- 2 $Y_i \leftarrow x_i \cdot G$
- 3 $h_i \leftarrow \text{Hash}(Y_i)$
- 4 **return** (x_i, Y_i, h_i)

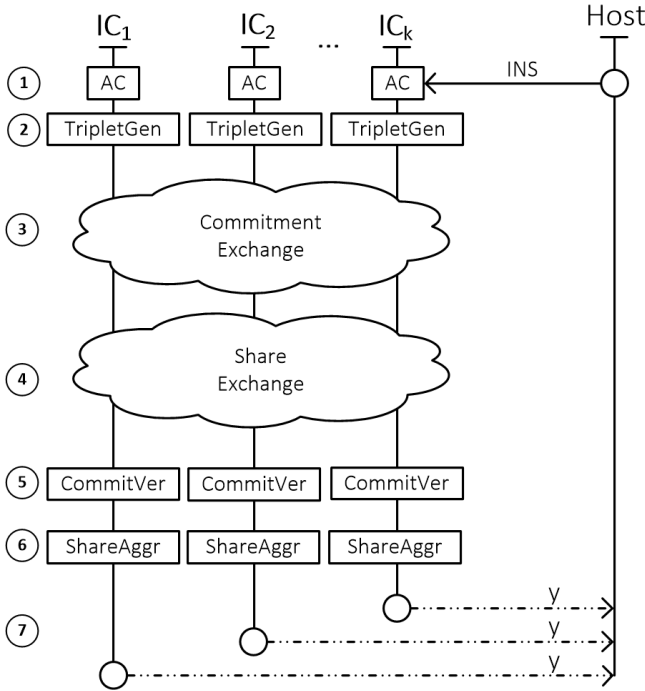


Figure 2: The interaction between the different participants during the execution of the Distributed Key Pair Generation (DKPG) protocol.

Upon the generation of the triplet, the members perform a pairwise exchange of their commitments (③), by the end of which, they all hold a set $H = \{h_1, h_2, \dots, h_t\}$. The commitment exchange terminates when $|H_q| = t \forall q \in Q$. Another round of exchanges then starts, this time for the shares of Y_{agg} (④) $Y = \{Y_1, Y_2, \dots, Y_t\}$. The commitment exchange round is of uttermost importance as it forces the participants to commit to a share of Y_{agg} , before receiving the shares of others. This prevents attacks where an adversary first collects the shares of others, and then crafts its share so as to bias the final pair, towards a secret key they know.

Algorithm 4.2: CommitVerify: Checks Y_i , against their respective commitments.

Input : (Y, H)
Output : Bool

```

1 for  $i \in \{1, |Y|\}$  do
2   if  $Hash(Y_i) \neq h_i$  then
3     return False
4   return True
```

Once each member of the quorum receives k shares (i.e., $|Y| = k$), it executes Algorithm 4.2 to verify the validity of Y 's elements against their commitments in H . (⑤) If one or more commitments fail the verification then the member infers that an error (either intentional or unintentional) occurred and the protocol is terminated. Otherwise, if all commitments are successfully verified, then the member executes Algorithm 4.3 (⑥) and returns the result to

the remote host (⑦). Note that it is important to return Y_{agg} , as well as the individual shares Y_i , as this protects against integrity attacks, where malicious ICs return a different share than the one they committed to during the protocol [42, 64]. Moreover, since Y_i are shares of the public key, they are also assumed to be public, and available to any untrusted party.

Algorithm 4.3: ShareAggr: Aggregates elements in a set of shares (e.g., ECPoints, field elements).

Input : Set of shares Q
Output : The aggregate of the shares q

```

1  $q \leftarrow 0$ 
2 for  $q_i \in Q$  do
3    $q \leftarrow q + q_i$ 
4 return  $q$ 
```

In the following sections, we rely on DKPG as a building block of more complex operations.

4.1 Distributed Public Key Generation

The distributed key generation operation enables multiple ICs to collectively generate a shared public and private key pair with shares distributed between all participating ICs. This is important in the presence of hardware trojans, as a single IC never gains access to the full private key at any point, while the integrity and secrecy of the pair is protected against maliciously crafted shares.

We opt for a scheme that offers the maximum level of confidentiality (t -of- t , $k = t$), and whose execution is identical to DKPG seen in Figure 2. However, there are numerous protocols that allow for different thresholds, such as Pedersen's VSS scheme [42, 64, 81]. The importance of the security threshold is discussed in more detail in Section 6.3.

Once a key pair has been generated, the remote host can encrypt a plaintext using the public key Y , request the decryption of a ciphertext, or ask for a plaintext to be signed. In the following sections we will outline the protocols that realize these operations.

4.2 Encryption

For encryption, we use the Elliptic Curve ElGamal scheme [24, 36] (Algorithm 4.4). This operation does not use the secret key, and can be performed directly on the host, or remotely by any party holding the public key, hence there is no need to perform it in a distributed manner.

Algorithm 4.4: Encrypts a plaintext using the agreed common public key.

Input : The domain parameters T , the plaintext m encoded as an element of the group \mathbb{G} , and the calculated public key Y_{agg}
Output : The Elgamal ciphertext tuple, (C_1, C_2)

```

1  $r \leftarrow Rand(T)$ 
2  $C_1 \leftarrow r \cdot G$ 
3  $C_2 \leftarrow m + r \cdot Y_{agg}$ 
4 return  $(C_1, C_2)$ 
```

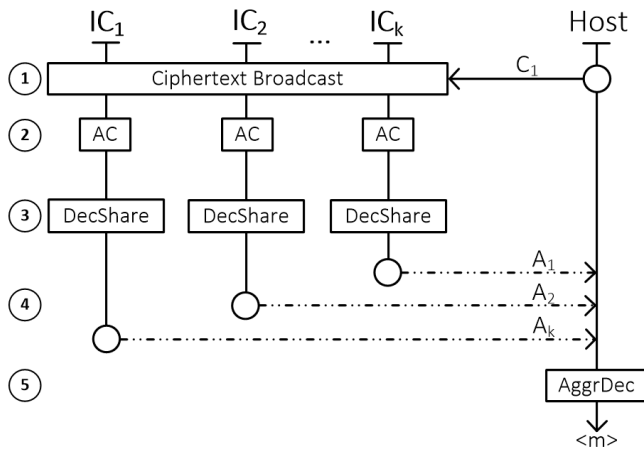


Figure 3: The interaction between the different ICs during the execution of the distributed decryption protocol.

4.3 Decryption

One of the fundamental cryptographic operations involving a private key is ciphertext decryption. In settings, where the key is held by a single authority, the decryption process is straightforward, but assumes that the hardware used to perform the decryption does not leak the secret decryption key. Myst addresses this problem by distributing the decryption process between k distinct processing ICs that hold shares of the key (Figure 3).

The protocol runs as follows: Initially, the host broadcasts the decryption instruction along with the part of the ciphertext C_1 to the processing ICs (1). Upon reception, the ICs first verify that the request is signed by an authorized user (2), and then execute Algorithm 4.5 to generate their decryption shares A_i (3). Once the shares are generated they are sent back to the host, signed by the ICs and encrypted under the host’s public key (4). Once the host receives k decryption shares, executes Algorithm 4.6 to combine them and retrieve the plaintext (5).

Algorithm 4.5: DecShare: Returns the decryption share for a given ciphertext.

Input : A part of the Elgamal ciphertext (C_1) and the IC’s private key of x_i .

Output: The decryption share A_i , where i is the IC’s uid

- 1 $A_i \leftarrow -x_i \cdot C_1$
 - 2 **return** A_i
-

It should be noted that during the decryption process, the plaintext is not revealed to any other party except the host, and neither the secret key nor its shares ever leave the honest ICs. An extension to the above protocol can also prevent malicious ICs from returning arbitrary decryption shares, by incorporating a non-interactive zero knowledge proof [21] in the operation output.

Algorithm 4.6: AggrDec: Combines the decryption shares and returns the plaintext for a given ciphertext.

Input : The Elgamal ciphertext C_2 and the set of decryption shares A .

Output: The plaintext m

- 1 $D \leftarrow 0$
 - 2 **for** $A_i \in A$ **do**
 - 3 | $D \leftarrow D + A_i$
 - 4 $m \leftarrow (C_2 + D)$
 - 5 **return** m
-

4.4 Random Number Generation

Another important application of secure hardware is the generation of a random fixed-length bit-strings in the presence of active adversaries. The key property of such systems is that errors (e.g., a hardware backdoor), should not allow an adversary to bias the generated bitstring.

The implementation of such an operation is straightforward. The remote host submits a request for randomness to all actors participating in the quorum. Subsequently, each actor independently generates a random share b_i , encrypts it with the public key of the host, and signs the ciphertext with its private key. Once the host receives all the shares, he combines them to retrieve the b and then uses an one way function (e.g., SHA3-512 [17]) to convert it to a fixed length string.

4.5 Signing

Apart from key generation, encryption and decryption, Myst also supports distributed signing – an operation that potentially manipulates a long term signature key. Here we introduce a novel multi-signature scheme, based on Schnorr signature [73].

A multi-signature scheme allows a group of signers to distributively compute a signature on a common message. There has been a very long line of works on multi-signature schemes [14, 22, 55, 57, 58, 62] featuring several security properties (e.g. accountability, subgroup-signing) and achieving various efficiency trade-offs. A significant portion of work has been dedicated in reducing the trust in the key generation process. However, this often involves the use of expensive primitives or increases the interaction between parties [22, 55, 57]. In our construction, we rely on Myst’s DKPG for the key generation process.

Our multi-signature scheme is based on Schnorr signatures [73] and has similarities with Bellare and Neven [14] construction. One crucial difference between existing multi-signature schemes and ours, is that we utilize a host intermediating between signers (i.e., ICs). This intermediating host allows us to eliminate any interaction between the ICs and thus improve the efficiency of the construction. To further minimise the interaction between the host and ICs we adapt existing techniques used in proactive two-party signatures [61] into the multi-signature context.

We let $PRF_s(j)$ be a pseudorandom function with key s , that takes j as input and instantiates it as $Hash(s||j)$.

Initially, all k ICs cooperate to generate a public key y using the distributed key generation operation (Section 4.1), and store securely their own key share x_i . Moreover, each IC generates a

Algorithm 4.7: SigShare: Returns the signature share of the IC for a given plaintext and index j .

Input : The digest of the plaintext to be signed $H(m)$, the IC's private key of x_i and secret s , an index j , and the aggregated random EC point R_j .

Output: The signature share tuple $(\sigma_{ij}, \epsilon_j)$

- 1 $\epsilon_j \leftarrow \text{Hash}(R_j || \text{Hash}(m) || j)$
- 2 $r_{ij} \leftarrow \text{PRF}_s(j)$
- 3 $\sigma_{ij} \leftarrow r_{ij} - x_i \cdot \epsilon_j \pmod n$
- 4 **return** $(\sigma_{ij}, \epsilon_j)$

secret s for the PRF, and stores it securely. After these steps, the signing protocol can be executed. The protocol comprises of two phases: *caching* and *signing*.

In the *caching phase* (Figure 4), the host queries the ICs for random group elements R_{ij} , where i is the id of the IC and j an increasing request counter (❶). Once such a request is received, the IC verifies that the host is authorized to submit such a request and then applies the keyed pseudorandom function on the index j to compute $r_{i,j} = \text{PRF}_s(j)$ (❷). The IC then uses $r_{i,j}$ to generate a group element (EC Point) $R_{ij} = r_{i,j} \cdot G$ (❸), which is then returned to the host. Subsequently, the host uses Algorithm 4.3 to compute the aggregate (R_j) of the group elements (Algorithm 4.3) received from the ICs for a particular j , and stores it for future use (❹). It should be noted that the storage cost for R_j is negligible: for each round the host stores only 65 Bytes or 129 Bytes depending on the elliptic curve used (for R_j) and the corresponding round index j . This allows the host to run the caching phase multiple times in parallel, and generate a list of random elements that can be later used, thus speeding up the signing process.

The *signing phase* (Figure 5) starts with the host sending a Sign request to all ICs (❶). Such a request includes the hash of the plaintext $\text{Hash}(m)$, the index of the round j , and the random group element R_j corresponding to the round. Each IC then first verifies that the host has the authorization to submit queries (❷) and that the specific j has not been already used (❸). The latter check on j is to prevent attacks that aim to either leak the private key or to allow the adversary to craft new signatures from existing ones. If these checks are successful, the IC executes Algorithm 4.7 and generates its signature share (❹). The signature share $(\sigma_{i,j}, \epsilon_j)$ is then sent to the host (❺). Once the host has collected all the shares for the same j , can use Algorithm 4.3 on all the $\sigma_{i,j}$ to recover σ_j , obtaining the aggregate signature (σ_j, ϵ_j) (❻).

The recipient of $\langle (m, j), \sigma, \epsilon \rangle$ can verify the validity of the signature by checking if $\epsilon = \text{Hash}(R || \text{Hash}(m) || j)$, where $R = \sigma \cdot G + \epsilon \cdot Y$.

Security Analysis. The security of a multi-signature scheme requires that if at least one of the signers behaves honestly, then it is impossible to forge a multi-signature. In our context, the presence of a single honest IC guarantees the signature cannot be forged even in presence of a malicious host controlling all the remaining ICs. The key generation process is a crucial step for the security of multi-signature schemes due to the so called *rogue-key attacks* [14]. The latter enables an attacker to maliciously generate shares of the public key in such a way that is possible for her to forge multi-signatures. In Myst's DPKG process, malicious ICs cannot subvert

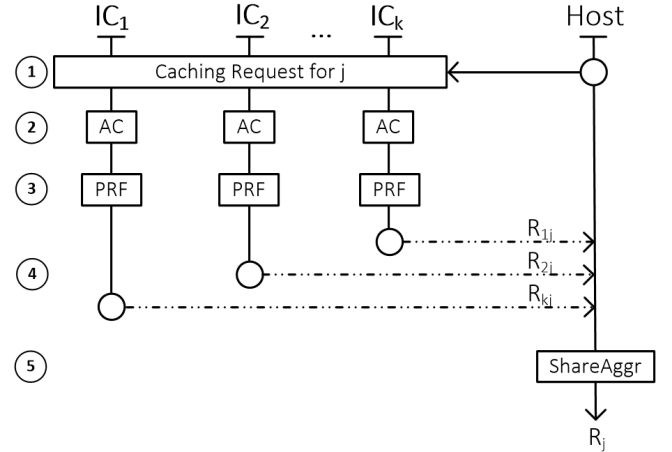


Figure 4: The interaction between the different players during the caching phase of the distributed signing protocol.

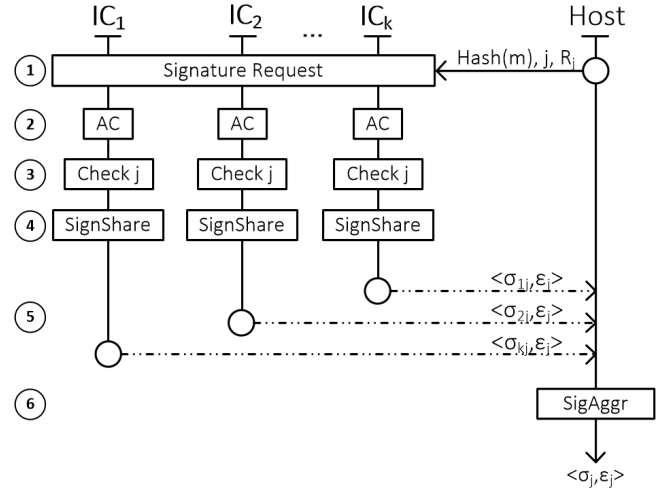


Figure 5: The figure illustrates the interaction between the different players when signing.

the key generation process as long as at least one IC is not colluding with the others, thus preventing rogue-key attacks. The security of our multi-signature scheme is proved in Theorem 4.1 assuming the hardness of on the one-more discrete logarithm problem [13]. We refer to Appendix A for the security definitions of multi-signatures and the proof of the following Theorem.

THEOREM 4.1. *If there exists a (q_S, q_H, ϵ) -forger \mathcal{F} for the multi-signature scheme described in Figures 4 and 5 interacting in $q_S = O(1)$ signature queries, making at most q_H Hash queries and succeeding with probability ϵ , then there exists an algorithm \mathcal{A} that solves the $(q_S + 1)$ -DL problem with probability at least*

$$\delta \geq \frac{\epsilon^2}{q_H^{q_S+1}} - \text{negl}(\lambda)$$

4.6 Key Propagation

In cases where more than one quorum is available it is useful to enable them all to handle requests for the same public key. This is of particular importance for both the system’s scalability and its availability, as we further discuss in Sections 6.2 and 6.4 respectively.

Once a quorum Q_1 has generated its public key y , (Section 4.1) the operator can specify another quorum Q_2 that is to be loaded with y . Each member q_i of Q_1 then splits its secret x_i in $|Q_2|$ shares and distributes them to the individual members of Q_2 . To do that q_i follows the secret sharing method shown in Algorithm 4.8. However, any t -of- t secret sharing schemes proposed in the literature [20, 64, 75] would do.

Algorithm 4.8: SecretShare: Returns a vector of shares from a secret.

Input : The domain parameters T , a secret s which is to be shared, and the number of shares k .

Output : A vector of shares \vec{v}_s

```

1 for ( $i = 0$  to  $k - 2$ ) do
2   |  $\vec{v}_s[i] \leftarrow \text{Rand}(T)$ 
3  $\vec{v}_s[k - 1] \leftarrow (s - \vec{v}_s[1] - \vec{v}_s[2] - \dots - \vec{v}_s[k - 2])$ 
4 return  $\vec{v}_s$ 
    
```

Once each member of Q_2 receives $|Q_1|$ shares, which they then combine to retrieve their share of the secret corresponding to y . Each member of Q_2 can retrieve its share by summing the incoming shares, modulo p (the prime provided in the domain parameters T). An additional benefit of such a scheme is that Q_1 and Q_2 may have different sizes.

It should be also noted that a naive approach of having each member of q_1 send their share of x to a member of q_2 is insecure, as malicious members from q_1 and q_2 can then collude to reconstruct the public key.

5 IMPLEMENTATION

In this section, we provide the implementation details of our Myst prototype. We first focus on the custom hardware we built, and outline its different components and capabilities. Thereafter, we discuss the development of the software for the untrusted ICs and the details of the remote host.




Figure 6: Overview of Myst’s components.

5.1 Hardware Design & Implementation

For our Myst prototype, we designed our own custom circuit board, which features 120 processing ICs (set to use 40 quorums of 3 smartcards from different manufacturers) interconnected with dedicated buses with 1.2Mbps bandwidth.




Figure 7: Myst’s smartcard board supports 120 ICs (60 on each side). Our configuration splits them in 40 quorums of 3 diverse ICs each.

The processing ICs are JavaCards (version 3.0.1), loaded with a custom applet implementing the protocols outlined in Section 4. JavaCards are a suitable platform as they provide good interoperability (i.e., applets are manufacturer-independent), which contributes to IC-diversification and prevents lock-in to particular vendors. Moreover, they also fulfill all the requirements listed in Section 2: (1) they are tamper-resistant (FIPS140-2 Level 4, CC EAL4) and can withstand attacks from adversaries with physical access to them [70], (2) they feature secure (FIPS140-2 compliant) on-card random number generators, (3) they offer cryptographic capabilities (e.g., Elliptic curve addition, multiplication) through a dedicated co-processor and (4) there are numerous independent fabrication facilities (Section 6). In addition to these, they have secure and persistent on-card storage space, ideal for storing key shares and protocol parameters.

The host is implemented using a computer that runs a python client application, which submits the user requests (e.g., Cipher-text Decryption) to Myst using a RESTful API exposed by the device. The incoming requests are handled by a Spring server, which parses them, converts them to a sequence of low-level instructions, and then forwards these to the IC controller, through an 1Gbps TCP/UDP interface. The ICs controller is a programmable Artix-7 FPGA that listens for incoming instructions and then routes them to the processing IC, through separate physical connections. We took special care that these buses offer a high bandwidth (over 400kbps), to prevent bottlenecks between controller and ICs even under extreme load. Once the ICs return the results, the controller communicates them back to the server, that subsequently forwards them to the host.

5.2 Software

We implement the protocols of Section 4 and provide the necessary methods for inter-component communication and system parameterization.

We develop and load the processing ICs with JavaCard applets implementing methods for 1) Card Management, 2) Key Generation, 3) Decryption, and 4) Signing. Although JavaCard APIs since version 2.2.2 specifies a BigInteger class, this class is either unsupported by real cards or provides only too small internal type length (typically at most 8 bytes). The only third-party BigInteger library available (i.e., BigNat¹) is unmaintained and lacked essential operations. Moreover, low-level elliptic curve operations are only

¹https://ovchip.cs.ru.nl/OV-chip_2.0

poorly supported by standard JavaCard API. The IC vendors often offer own proprietary APIs which provides the required operations (e.g., EC point addition) – but for the price of reduced portability. This made the implementation of our cryptographic schemes more complicated.

We extend BigNat to provide methods catering to our specific needs. Additionally, we develop own EC operations library based solely on public JavaCard API to support ICs where proprietary API is not available or cannot be used. Our EC library provides methods for EC point initialization, negation, addition, subtraction and scalar multiplication, and has been released as an independent project ².

Note, that although vendor’s proprietary API limits the portability, it usually provides better performance and better protection against various side-channel attacks in comparison to custom software-only implementations. For this reason, we structured the IC applet code for easy incorporation of proprietary API with minimal code changes. Our EC library is thus used only when no other option is available for target IC.

Our current implementation is based on the NIST P-256 [1, 86] curve that provides at least 128 bits of security. However, it can also be trivially adapted for any other curve.

Optimizations. We optimize our JavaCard applet for speed and to limit side-channel attacks. Although JavaCard applets are compiled with standard Java compiler, common Java implementation patterns (e.g., frequent array reallocation due to resizing) are prohibitively expensive on JavaCards. Therefore, we use the following optimization techniques based on good practices and performance measurements from real, non-simulated, smart cards [83]:

- ❖ We use hardware accelerated cryptographic methods from JavaCard API instead of custom implementations interpreted by JavaCard virtual machine where possible.
- ❖ We store the session data in the faster RAM-allocated arrays instead of persistent, but slower EEPROM/Flash
- ❖ We use copy-free methods which minimize the move of data in memory, and also utilize native methods provided by JCSys class for array manipulation like memory set, copy and erase.
- ❖ We made sure to pre-allocate all cryptographic engines and key objects during the applet installation. No further allocation during the rest of the applet lifetime is performed.
- ❖ Surplus cryptographic engines and key objects are used to minimize the number of key scheduling and initialization of engine with a particular key as these operations impose non-trivial overhead [83].
- ❖ We refrain from using single long-running commands which would cause other cards to wait for completion, thus increasing the latency of the final operation.

Finally, we optimized two fundamental operations commonly used in cryptographic protocols: 1) integer multiplication, and 2) the modulo operation optimized for 32 byte-long EC coordinates. This was necessary, as the straightforward implementation of these two algorithms in JavaCard bytecode is both slow and potentially vulnerable to side-channel timing attacks. Instead, we implemented

both operations so as to use the native RSA engine and thus have constant-time run-time.

The integer multiplication of a and b can be rewritten as $a \cdot b = ((a + b)^2 - a^2 - b^2)/2$. The squaring operation (e.g., a^2) can be quickly computed using a pre-prepared raw RSA engine with a public exponent equal to 2 and a modulus n , that is larger than the sum of the lengths of both operands. On the other hand, the integer modulo of two operands a (64 bytes long) and b (32 bytes long) is not so straightforward. We exploit the fact that b is always the order of the fixed point G in the P-256 Elliptic Curve [1, 86], and transform $a \bmod b = a - (((a \cdot x) \gg z) \cdot x)$ where x and z values are pre-computed offline [44]. As a result, a modulo operation can be transformed to two RSA operations, one shift (with z being multiple of 8) and one subtraction. Note that we cannot directly use RSA with a public exponent equal to 1 as operands are of different length and also shorter than smallest RSA length available on the card.

5.3 System States

Initially, the system is in a non-operational state, where the processing ICs do not respond to user requests. To make it operational, a secure initialization process has to be carried out. During the initialization the processing ICs and the other components described in 3 are loaded with verified application packages, and the domain parameters for the distributed protocols are set. Moreover, the ICs are provided with their certificates that they will later use to sign their packets and establish secure communication channels.

Once the initialization process has been completed, the system switches to an operational state and is available to serve user requests. Depending on the configuration, the system may be brought back to an uninitialized state, in order to load new software or change the protocol parameters. We further discuss the importance of the initialization process in Section 6.4.

6 EVALUATION

In this section, we evaluate Myst by examining both its performance, and its qualitative properties.

Experimental Setup. All evaluations were performed using the setup illustrated in Figure 6. The host is a single machine with a CentOS 7 OS (3.10.0 Linux kernel), featuring a 3.00GHz Intel(R) Core i5-4590S CPU, 16GB of RAM, and uses a python client application to submit service requests to Myst, through a 1Gbps Ethernet interface (as described in Section 5). Upon reception, the server uses the Java Spring framework [51] to parse them, and then forward the instructions to the Artix-7 FPGA, which subsequently routes them to the individual smart cards. In all experiments, we collect response-time measurements from both the host and the Spring server. On average the round-trip from the host to the server takes 5ms. For accuracy, we use the host measurements in the rest of the section.

6.1 Performance Impact

This subsection evaluates the performance impact of Myst, and compares its throughput and latency with that of a single-IC system. Moreover, it examines the impact of our optimizations on the overall performance of the system.

²<https://OpenCryptoJC.org>

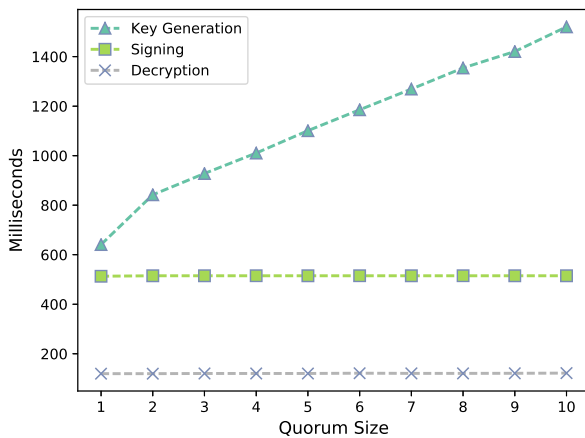


Figure 8: The average response time for each distributed operation of the cryptosystem, in relation to the quorum (i.e., a coalition of multiple ICs) size.

Methodology. To better understand the overhead that the use of a distributed architecture entails we run experiments that measure the latency as the size of the protocol quorum grows. We first submit 1,000 requests for each cryptosystem operation (Section 4) in one JavaCard and measure the response time. We then extend the experiment to larger quorums with sizes ranging from 2 to 10, and measure the latency in completing the same 1,000 operations. Simultaneously, to gain a more in-depth understanding of the impact that each low-level instruction type has, we micro-benchmark the response time for all intra-system communications.

Results. Figure 8 plots the average response time for performing Key Generation, Decryption and Signing using IC quorums of different sizes. To begin with, Decryption is the fastest (119ms), since it implements a single round protocol. Moreover, when we compare the runtime between the single-IC run, and the runs with multiple ICs, we observe that the latency is stable and the overhead remains always smaller than 0.8%. Hence, we conclude that the Decryption time does not increase with the size of the quorum, due to the ICs performing the decryption computations simultaneously. This highlights that Decryption is only CPU bound, and the network capacity of our prototype does not pose a bottleneck; and demonstrates that Myst can essentially provide increased assurance, with negligible impact on the decryption runtime. It should be noted that, high-throughput decryption is of extreme importance in applications such as secure key derivation in mix-network infrastructures [33] that heavily rely on decryption. Similarly, the runtime for signing remains the constant (~ 517ms) regardless of the quorum size. This is because our multi-sig signing protocol does not require one-to-one interaction between the ICs. The runtime difference between decryption and signing is mainly because of JavaCard API limitations. Specifically, we were forced to perform some of the mathematical operations for signing in the general purpose processor of the card, and not in the cryptographic coprocessor. The caching phase of the signing protocol (takes on average 169ms/operation) was performed offline for thousands of




Figure 9: Breakdown of the runtime for low-level instructions that comprise the key generation operation, in relation to the quorum size. The horizontal reference line represent the cost of using a single IC.

indices j and is not considered in the measurements, as it is ran only occasionally.

As with signing and decryption, the runtime for random number generation is also constant, as the protocol is single round and can be executed in parallel.

On the other hand, Key Generation (DKPG) requires two phases (commitment and revelation) and this adds significant latency. In particular, as seen in Figure 8, each additional IC results in an runtime increase of approximately 90ms. Figure 9 examines the timings of low-level operations involved in the key generation protocols. When quorums are small the cost of key generation is dominated by the “INS_KEYGEN_INIT” operation that generates a secret share, and derives a public key (624ms). However, as the IC quorums grow the operations related to exchanging public keys (“INS_KEYGEN_STORE_PUBKEY”) and commitments (“INS_KEYGEN_STORE_HASH”) become significant, and for a quorum of 10 ICs, nearly doubles the cost of key generation. However, for quorums of 3 the impact on runtime is merely 303ms, compared to a single IC. Other low-level operations used in DKPG have negligible cost, regardless of the quorum size.

6.2 Scalability & Extensibility

This section examines how the throughput of our prototype grows as the more processing power (i.e., quorums) is added. The absence of bottlenecks in the system is important to be able to provide high-availability in production environments.

Methodology. To determine how efficiently our design scales in practice, we run a series of experiments that measure Myst’s throughput, for varying numbers of processing quorums. As described in Section 5, our board supports up to 120 processing ICs which can be divided into multiple quorums and serve requests simultaneously. To benchmark the scalability of the system, on each iteration of the experiment we submit 10,000 requests for each




Figure 10: The average system throughput in relation to the number of quorums ($k = 3$) that serve requests simultaneously. The higher is better.

high-level operation supported by our cryptosystem, and measure its throughput. However, this time we fix the quorum size k to 3, and on each iteration we increase the number of quorums serving the submitted requests by one, until we involve 40 quorums, which is the maximum number of 3-IC quorums that our prototype can support. For simplicity, we assign each processing IC to only one quorum. However, it is also technically possible for an IC to participate in more than one quorums, and store more than one secret shares.

Results. Figure 10 illustrates the throughput of the Myst system (in operations per second) as more of the ICs are used for processing transactions. The maximum throughput of Myst was 315ops/sec for Decryption and 77ops/sec for Signing, when using all 40 quorums. We also observe that as the number of quorums increases, the performance increases linearly, at a rate of ~ 9 requests per second per additional quorum for the Decryption operation, and ~ 1.9 requests per second for Signing. This illustrates that the system is CPU bound, and the communication fabric between ICs and the host is not a bottleneck. Consequently, a system with our proposed architecture can be easily tailored for different use cases to provide the throughput needed in each of them. It should be also noted that using a lower threshold $k < t$ does not affect the performance of the system. However, this may result in some ICs being idle longer. For this purpose, it would be beneficial if ICs participated in more than one quorum, thus minimizing their idle time.

6.3 Tolerance levels

Myst is tolerant against both fabrication-time and design-time attacks (assuming independent foundries and design houses). Alternatively, if there is only one design house, then Myst protects against fabrication-time attacks only. In this section, we examine the relationship between the system parameters and the tolerance levels provided depending on the attack type (Table 1). The reported tolerance levels are those provided by the schemes outlined in Section 4,

Parameters	Leakage	Denial-of-Service	IC Failures
<i>Single IC</i>	0	0	0
$k = t$	$t - 1$	0	$n - 1$
$k < t$	$k - 1$	$t - k$	$(t - k) * n$

Table 1: Number of malicious/faulty ICs that different setups can tolerate in each abnormal scenario. The system is assumed to feature n identical quorums of size t , with a secret-sharing threshold k .

and were also empirically verified using our prototype introduced in Section 5.

In practice, the threshold k and the size of the quorums t express the trade-off between confidentiality and robustness for each particular quorum. The relationship between these two parameters determines how many ICs can cease to function, before the quorum fails: When k equals the number of processing ICs t , then secrets are safe in the presence of $t - 1$ compromised and colluding ICs. On the other hand, a more “relaxed” threshold ($k < t$) enables the quorum to remain fully functional unless more than $t - k$ ICs fail (maliciously or through wear). Alternatively, ($k = t$)-systems can withstand multiple ICs failing (due to wear) using the technique introduced in Section 4.6. This technique enables several quorums to handle requests for the same public key, and thus even if multiple ICs (and consequently the quorums they belong to) fail, the remaining quorums can still handle incoming requests for those public keys. It should be noted that this technique provides robustness only in the presence of faulty & defective ICs, but does not mitigate denial of service attacks. This is because, all n quorums are identical (same manufacturer diversity) and thus a malicious IC from a breached supply chain will participate in all quorums. In contrast, defective ICs will fail with a probability less than 1 (post-fabrication tests detect always failing ICs) and thus several quorums will remain functional.

From all the above, the security guarantees offered by Myst are determined by the threshold k , the IC diversity and the number of quorums. In our prototype, we chose to maximize confidentiality, and resist the presence of $t - 1$ actively malicious ICs. Malicious ICs launching denial-of-service attacks are easier to detect and replace, compared to those subtly leaking secrets or weakening keys. In cases where increased robustness and availability are paramount, the security level can be adjusted in favor of redundancy using the appropriate threshold schemes [64].

6.4 Other Considerations

In this section, we consider several qualitative properties commonly used in the industry.

Physical Security & Diversity. SmartCards forms the core of our prototype and have multiple benefits as they were designed to operate in a hostile environment that is assumed to be fully controlled by the adversary [70]. For this reason, they come with very-high tamper-resistance (FIPS140-2, Level 4) and secure storage capabilities, that are constantly evolving. Another benefit is that there are several manufacturers owning foundries including

NXP-semiconductors, ST Microelectronics, Samsung, Infineon and Athena. Moreover, there are numerous independent facilities used by different vendors to fabricate their ICs [30, 74, 80, 85], which also contributes to the necessary diversity for the quorums.

Code & Parameter Provisioning. Crucial component of the security of Myst. If the code on all the ICs, or the cryptographic parameters contain errors, then the security of the system is not guaranteed. We propose two strategies to ensure secure provisioning. First, we may assume that provisioning occurs at the factory. This leverages our assumption that some of the factories are honest, and ensures that an adversary would have to compromise all fabrication facilities to extract any secrets. The second strategy is to assume the existence of a trusted off-line provisioning facility that is not under the control of an adversary. This facility needs only to guarantee high-integrity, as no secrets are involved in the provisioning step (secrets are generated within the ICs as part of protocols executed later).

Formal Security Validations. Any off-premise use of ICs raises a question of trust. To address this problem, independent security validations (e.g., FIPS140-2, or Common Criteria) have been introduced. These validations are performed by a third party, which verifies the claims of hardware vendors. However, these validations are a serious bottleneck when introducing new IC design. Myst does not need to go through this process as it executes provably secure cryptographic schemes on already validated ICs. As a result, even if one of the ICs has passed an independent validation, the whole instance of the Myst carries this assurance.

Real-world Practicality. Myst provides a platform for generating legally binding digital signatures (eIDAS [29]) under the sole control of the user. Moreover, due to the tamper-resistance properties of SmartCards, ICs can be also stored remotely thus making Myst a highly practical system that is able to support some non-typical real-world use-cases. For instance, Myst can perform code-signing of mobile applications for apps stores (e.g., AppStore or Google Play), by sharing the signing keys between developers' laptops, managers' computers, and securely stored ICs providing protection against internal enterprise/company attackers without compromising developers' control over the signing process. Finally, another practical feature of Myst is that each party maintains its own independent audit logs, thus ensuring non-repudiation.

7 RELATED WORK

This section examines existing literature on hardware trojans and countermeasures, outlines relevant fault-tolerant designs and architectures, and discusses how Myst compares to prior works.

Hardware Trojans & Countermeasures. To better understand the extend of the Hardware Trojans threat, numerous attacks and exploitation techniques have been proposed in the literature in the last decade. For instance, the authors in [53] design two hardware trojans and implement a proof-of-concept attack against the PRINCE cipher [23]. The novelty of their attacks is that they use dopant polarity changes (first introduced in [11]), to create a hard-to-detect fault-injection attack backdoor. Moreover, [65] also introduces a hardware trojan attacking RSA applications. In this attack, the adversary is able to use power supply fluctuations to trigger the trojan,

which then leaks bits of the key through a signature. Another very hard to detect class of trojans (inserted by fabrication-time attackers) was introduced by K Yang et al. in [98]. Such trojans leverage analog circuits and require only a single logic gate to launch their attack (e.g., switch the CPU's mode). Apart from these, detection evasion and stealthy triggering techniques have been proposed in [18, 54, 91–93].

As discussed in Section 1, malicious components carrying errors have been also observed in commercial and military hardware [2, 40, 50, 56, 59, 76–78], while a subset of those incidents also involved misbehaving cryptographic hardware [43, 45, 71, 72]. In all these cases, the errors were eventually attributed to honest design or fabrication mistakes, but the systems were left vulnerable to attacks regardless. For instance, one popular and well-studied example of attacks against weak cryptographic hardware is [16]. In this work, Bernstein et al. study the random number generators used in smartcards and found various malfunctioning pieces, that allowed them to break 184 public keys using in “Citizen Digital Certificates” by Taiwanese citizens.

To address the aforementioned threats, different approaches have been proposed. The most common ones attempt to either detect malicious circuitry, or prevent insertions. In particular, detection techniques aim to determine whether any HTs exist in a given circuit and feature a wide range of methods such as side-channel analysis [3, 79, 95, 97], logic testing [26], and trust verification [8, 68, 90, 102]. On the other hand, prevention techniques aim to either impede the introduction of HTs, or make HT easier to detect, such approaches are Split manufacturing [28, 69, 94] which tries to minimize the circuit/design exposure to the adversary, logic obfuscation [25] and runtime monitoring [47, 88]

Moreover, there is also a smaller body of work which attempts to tackle the even harder problem of inferring additional information about the HT such as its triggers, payload, and exact location [95, 97]. Other works considered verifiable computation architectures (such as [87]), which provide guarantees for the correctness of the computation on untrusted platforms. However, they come with a computation overhead and do not guarantee secure handling of secrets or protection from side-channel attacks. On top of the above, previous works [6, 34] have also theoretically discussed using multi-party computation protocols to distribute trust between untrusted manufacturers during the fabrication process.

Myst follows an alternative approach that leverages a diverse set of untrusted ICs to minimize the risk of compromises by distributing trust between them. For this reason, all the above countermeasures remain applicable and would even increase the security of the final system. In other words, our proposed approach is not an alternative to existing techniques, as it provides a way to combine ICs fabricated by different manufacturers, in various locations and featuring a wide-range of protection techniques in one system.

Fault-Tolerant Systems. Component-redundancy and component-diversification are both key concepts used in N-variant systems that aim to achieve high fault-tolerance [27]. An example of such a system is the Triple-Triple Redundant 777 Primary Flight Computer [100, 101], that replicates the computations in three processors and then performs a majority voting to determine the final result. The applications of N-variance in security scenarios has been studied in only few works aiming

to protect systems against software attacks. In particular, [32] introduces a method to generate randomized system copies, that will have disjoint exploitation sets, thus achieving memory safety. On the other hand, [4] proposes a N-variant method for IC diversification, aiming again to achieve disjoint exploitation sets. However, this method is not effective against fabrication-time attacks and protects only against (potentially exploitable) unintentional errors. Finally, heterogeneous architectures with COTS components have been also proposed in [9, 10]. However, the practicality of these works is very limited as the computations are simply replicated between the different components, thus not protecting against confidentiality attacks.

8 CONCLUSION

High-assurance cryptographic hardware, such as hardware security modules, is used in production under the assumption that its individual components are secure and free from errors. However, even though post-fabrication testing is capable of detecting defective ICs with high-accuracy, there are certain error classes that are hard to uncover (e.g., RNG defects). Unfortunately, these errors are also detrimental to the security of high-assurance hardware such as hardware security modules. Moreover, hardware trojans and malicious circuitry have been extensively studied, and there is an abundance of designs, mitigation techniques and countermeasure evasion methods. This line of work assumes that not all errors can be detected and that due to the arms race, between trojan horses and mitigation techniques, countermeasures will never be 100% effective against all possible threats.

To resolve this, we introduce Myst, which brings the adversary into the unfavorable position of having to compromise 100% of the hardware components to gain any advantage. By employing threshold schemes and a redundancy-based architecture, Myst is able to distribute both secrets and cryptographic computations among multiple, diverse integrated circuits. Consequently, an adversary aiming to breach the confidentiality or the integrity of the system, must be able to compromise all the ICs. This is not a trivial task when the ICs are manufactured from different vendors, in different premises. To evaluate Myst, we build a custom board featuring 120 Smart Cards controlled by an Artix-7 FPGA. The maximum throughput for distributed decryption is 3150ops/sec, while for signing is 77ops/sec. Both come with an overhead of less than < 1% compared to a typical single-point of failure system. All in all, our results show that Myst is highly scalable, and provides strong guarantees for the security of the system thus making it suitable for production.

9 ACKNOWLEDGEMENTS

This work was supported by the European Commission through the H2020-DS-2014-653497 PANORAMIX project and the European Research Council via the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307937, and the Czech Science Foundation under project GA16-08565S.

REFERENCES

- [1] Mehmet Adalier. 2015. Efficient and Secure Elliptic Curve Cryptography Implementation of Curve P-256. (2015).
- [2] Sally Adee. 2008. The hunt for the kill switch. *IEEE Spectrum* 45, 5 (2008), 34–39.
- [3] Dakshi Agrawal, Selcuk Baktir, Deniz Karakoyunlu, Pankaj Rohatgi, and Berk Sunar. 2007. Trojan detection using IC fingerprinting. In *Security and Privacy, 2007. SP'07. IEEE Symposium on*. IEEE, 296–310.
- [4] Younsa Alkabani and Farinaz Koushanfar. 2008. N-variant IC design: methodology and applications. In *Proceedings of the 45th Design Automation Conference, DAC 2008, Anaheim, CA, USA, June 8–13, 2008*. 546–551. <https://doi.org/10.1145/1391469.1391606>
- [5] Jacob Appelbaum, Judith Horchert, and Christian Stöcker. 2013. Shopping for spy gear: Catalog advertises NSA toolbox. *Der Spiegel* 29 (2013).
- [6] Giuseppe Ateniese, Aggelos Kiayias, Bernardo Magri, Yiannis Tselekounis, and Daniele Venturi. 2016. Secure Outsourcing of Circuit Manufacturing. *Cryptology ePrint Archive*, Report 2016/527. (2016). <http://eprint.iacr.org/2016/527>.
- [7] Michael Backes, Markus Dürmuth, and Dominique Unruh. 2008. Compromising reflections-or-how to read LCD monitors around the corner. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*. IEEE, 158–169.
- [8] Chongxi Bao, Yang Xie, and Ankur Srivastava. 2015. A security-aware design scheme for better hardware Trojan detection sensitivity. In *IEEE International Symposium on Hardware Oriented Security and Trust, HOST 2015, Washington, DC, USA, 5–7 May, 2015*. 52–55. <https://doi.org/10.1109/HST.2015.7140236>
- [9] Mark Beaumont, Bradley Hopkins, and Tristan Newby. 2012. Safer path: Security architecture using fragmented execution and replication for protection against trojaned hardware. In *Proceedings of the Conference on Design, Automation and Test in Europe*. EDA Consortium, 1000–1005.
- [10] Mark Beaumont, Bradley Hopkins, and Tristan Newby. 2013. Hardware trojan resistant computation using heterogeneous COTS processors. In *Proceedings of the Thirty-Sixth Australasian Computer Science Conference-Volume 135*. Australian Computer Society, Inc., 97–106.
- [11] Georg T Becker, Francesco Regazzoni, Christof Paar, and Wayne P Burleson. 2013. Stealthy dopant-level hardware trojans. In *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 197–214.
- [12] Georg T. Becker, Francesco Regazzoni, Christof Paar, and Wayne P. Burleson. 2014. Stealthy dopant-level hardware Trojans: extended version. *J. Cryptographic Engineering* 4, 1 (2014), 19–31. <https://doi.org/10.1007/s13389-013-0068-0>
- [13] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. 2003. The One-More-RSA-Inversion Problems and the Security of Chaum's Blind Signature Scheme. *J. Cryptology* 16, 3 (2003), 185–215.
- [14] Mihir Bellare and Gregory Neven. 2006. Multi-signatures in the plain public-Key model and a general forking lemma. 390–399.
- [15] Mihir Bellare and Phillip Rogaway. 1993. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3–5, 1993*. 62–73.
- [16] Daniel J Bernstein, Yun-An Chang, Chen-Mou Cheng, Li-Ping Chou, Nadia Heninger, Tanja Lange, and Nicko Van Someren. 2013. Factoring RSA keys from certified smart cards: Coppersmith in the wild. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 341–360.
- [17] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. 2009. Keccak sponge function family main document. *Submission to NIST (Round 2)* 3 (2009), 30.
- [18] Swarup Bhunia, Michael S. Hsiao, Mainak Banga, and Seetharam Narasimhan. 2014. Hardware Trojan Attacks: Threat Analysis and Countermeasures. *Proc. IEEE* 102, 8 (2014), 1229–1247. <https://doi.org/10.1109/JPROC.2014.2334493>
- [19] Swarup Bhunia, Michael S Hsiao, Mainak Banga, and Seetharam Narasimhan. 2014. Hardware Trojan attacks: threat analysis and countermeasures. *Proc. IEEE* 102, 8 (2014), 1229–1247.
- [20] George Robert Blakley. 1979. Safeguarding cryptographic keys. *Proc. of the National Computer Conference 1979* 48 (1979), 313–317.
- [21] Manuel Blum, Paul Feldman, and Silvio Micali. 1988. Non-interactive zero-knowledge and its applications. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*. ACM, 103–112.
- [22] Alexandra Boldyreva. 2003. Threshold Signatures, Multisignatures and Blind Signatures Based on the Gap-Diffie-Hellman-Group Signature Scheme. In *Public Key Cryptography - PKC 2003*. 31–46.
- [23] Julia Borghoff, Anne Canteaut, Tim Güneysu, Elif Bilge Kavun, Miroslav Knezevic, Lars R Knudsen, Gregor Leander, Ventsislav Nikov, Christof Paar, Christian Rechberger, et al. 2012. PRINCE—a low-latency block cipher for pervasive computing applications. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 208–225.
- [24] Felix Brandt. 2005. Efficient Cryptographic Protocol Design Based on Distributed El Gamal Encryption. In *Information Security and Cryptology - ICISC 2005, 8th International Conference, Seoul, Korea, December 1–2, 2005, Revised Selected Papers (Lecture Notes in Computer Science)*, Dongho Won and Seungjoo Kim (Eds.), Vol. 3935. Springer, 32–47. https://doi.org/10.1007/11734727_5
- [25] Rajat Subhra Chakraborty and Swarup Bhunia. 2009. Security against hardware Trojan through a novel application of design obfuscation. In *Proceedings of the 2009 International Conference on Computer-Aided Design*. ACM, 113–116.

- [26] Rajat Subhra Chakraborty, Francis G. Wolff, Somnath Paul, Christos A. Papachristou, and Swarup Bhunia. 2009. MERO: A Statistical Approach for Hardware Trojan Detection. In *Cryptographic Hardware and Embedded Systems - CHES 2009, 11th International Workshop, Lausanne, Switzerland, September 6-9, 2009, Proceedings*. 396–410. https://doi.org/10.1007/978-3-642-04138-9_28
- [27] Liming Chen and Algirdas Avizienis. 1978. N-version programming: A fault-tolerance approach to reliability of software operation. In *Digest of Papers FTCS-8: Eighth Annual International Conference on Fault Tolerant Computing*. 3–9.
- [28] Zhang Chen, Pingqiang Zhou, T. Y. Ho, and Y. Jin. 2016. How secure is split manufacturing in preventing hardware trojan?. In *2016 IEEE Asian Hardware-Oriented Security and Trust (AsianHOST)*. 1–6. <https://doi.org/10.1109/AsianHOST.2016.7835561>
- [29] European Commission. 2016. Questions & Answers on Trust Services under eIDAS. <https://ec.europa.eu/digital-single-market/en/news/questions-answers-trust-services-under-eidas>. (February 2016).
- [30] Semiconductor Manufacturing International Corporation. 2017. Embedded Non-Volatile Memory for Smart Card & MCU. http://www.smics.com/eng/foundry/technology/tec_envm.php. (May 2017).
- [31] Nicolas T Courtois. 2009. The dark side of security by obscurity and cloning Mifare Classic rail and building passes, anywhere, anytime. (2009).
- [32] Benjamin Cox and David Evans. 2006. N-Variant Systems: A Secretless Framework for Security through Diversity. In *Proceedings of the 15th USENIX Security Symposium, Vancouver, BC, Canada, July 31 - August 4, 2006*. <https://www.usenix.org/conference/15th-usenix-security-symposium/n-variant-systems-secretless-framework-security-through>
- [33] George Danezis, Claudia Diaz, and Paul Syverson. 2009. Systems for anonymous communication. *Handbook of Financial Cryptography and Security, Cryptography and Network Security Series* (2009), 341–389.
- [34] Stefan Dziembowski, Sebastian Faust, and François-Xavier Standaert. 2016. Private circuits III: Hardware Trojan-Resilience via testing amplification. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 142–153.
- [35] Don Edenfeld, Andrew B. Kahng, Mike Rodgers, and Yervant Zorian. 2004. 2003 Technology Roadmap for Semiconductors. *IEEE Computer* 37, 1 (2004), 47–56. <https://doi.org/10.1109/MC.2004.1260725>
- [36] Taher ElGamal. 1985. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE transactions on information theory* 31, 4 (1985), 469–472.
- [37] Amos Fiat and Adi Shamir. 1986. How to prove yourself: Practical solutions to identification and signature problems. In *Conference on the Theory and Application of Cryptographic Techniques*. Springer, 186–194.
- [38] Task Force. 2005. High Performance Microchip Supply. (2005).
- [39] Bastian Fredriksson. 2016. A case study in smartcard security Analysing Mifare Classic Rev. (2016).
- [40] Sean Gallagher. 2014. Photos of an NSA “Jupgrade” factory show Cisco router getting implant. *Ars Technica* 14 (2014).
- [41] Daniel Genkin, Adi Shamir, and Eran Tromer. 2014. RSA key extraction via low-bandwidth acoustic cryptanalysis. In *International Cryptology Conference*. Springer, 444–461.
- [42] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. 2007. Secure Distributed Key Generation for Discrete-Log Based Cryptosystems. *J. Cryptology* 20, 1 (2007), 51–83. <https://doi.org/10.1007/s00145-006-0347-3>
- [43] Dan Goodin. 2013. ‘We cannot trust’ Intel and Via’s chip-based crypto’, FreeBSD developers say. <http://arstechnica.com/security/2013/12/we-cannot-trust-intel-and-vias-chip-based-crypto-freebsd-developers-say/>. (December 2013).
- [44] Torbjörn Granlund and Peter L Montgomery. 1994. Division by invariant integers using multiplication. In *ACM SIGPLAN Notices*, Vol. 29. ACM, 61–72.
- [45] FreeBSD Security Working Group. 2013. FreeBSD Developer Summit: Security Working Group. <https://wiki.freebsd.org/201309DevSummit/Security>. (December 2013).
- [46] Stefan Heck, Sri Kaza, and Dickon Pinner. 2011. Creating value in the semiconductor industry. *McKinsey & Company* (2011).
- [47] Matthew Hicks, Murph Finnicum, Samuel T. King, Milo M. K. Martin, and Jonathan M. Smith. 2010. Overcoming an Untrusted Computing Base: Detecting and Removing Malicious Hardware Automatically. In *31st IEEE Symposium on Security and Privacy, S&P 2010, 16-19 May 2010, Berkeley/Oakland, California, USA*. 159–172. <https://doi.org/10.1109/SP.2010.18>
- [48] Frank Imeson, Ariq Emtenan, Siddharth Garg, and Mahesh V Tripunitara. Securing Computer Hardware Using 3D Integrated Circuit (IC) Technology and Split Manufacturing for Obfuscation.
- [49] Intel Inc. 2017. About the Intel manageability firmware critical vulnerability. <http://www.intel.com/content/www/us/en/architecture-and-technology/intel-amt-vulnerability-announcement.html>. (May 2017).
- [50] Yier Jin and Yiorgos Makris. 2010. Hardware Trojans in wireless cryptographic ICs. *IEEE Design & Test of Computers* 27, 1 (2010).
- [51] Rod Johnson et al. 2005. Introduction to the spring framework. *TheServerSide.com* 21 (2005), 22.
- [52] Samuel T. King, Joseph Tucek, Anthony Cozzie, Chris Grier, Weihang Jiang, and Yuanyuan Zhou. 2008. Designing and Implementing Malicious Hardware. In *First USENIX Workshop on Large-Scale Exploits and Emergent Threats, LEET '08, San Francisco, CA, USA, April 15, 2008, Proceedings*. http://www.usenix.org/events/leet08/tech/full_papers/king/king.pdf
- [53] Raghavan Kumar, Philipp Jovanovic, Wayne P Burleson, and Ilia Polian. 2014. Parametric Trojans for Fault-Injection Attacks on Cryptographic Hardware. *IACR Cryptology ePrint Archive* 2014 (2014), 783.
- [54] Sebastian Kutzner, Axel York Poschmann, and Marc Stöttinger. 2013. Hardware trojan design and detection: a practical evaluation. In *Proceedings of the Workshop on Embedded Systems Security, WESS 2013, Montreal, Quebec, Canada, September 29 - October 4, 2013*. 1:1–1:9. <https://doi.org/10.1145/2527317.2527318>
- [55] Steve Lu, Rafail Ostrovsky, Amit Sahai, Hovav Shacham, and Brent Waters. 2013. Sequential Aggregate Signatures, Multisignatures, and Verifiably Encrypted Signatures Without Random Oracles. *Journal of cryptology* 26, 2 (2013), 340–373.
- [56] John Markoff. 2009. Old trick threatens the newest weapons. *The New York Times* 27 (2009).
- [57] Silvio Micali, Kazuo Ohta, and Leonid Reyzin. 2001. Accountable-subgroup multisignatures: extended abstract. In *CCS 2001, Proceedings of the 8th ACM Conference on Computer and Communications Security*. 245–254.
- [58] Markus Michels and Patrick Horster. 1996. On the Risk of Disruption in Several Multiparty Signature Schemes. In *Advances in Cryptology - ASIACRYPT '96, International Conference on the Theory and Applications of Cryptology and Information Security, Kyongju, Korea, November 3-7, 1996, Proceedings*. 334–345.
- [59] Subhashish Mitra, H-S Philip Wong, and Simon Wong. 2015. The Trojan-proof chip. *IEEE Spectrum* 52, 2 (2015), 46–51.
- [60] Inez Miyamoto, Thomas H Holzer, and Shahryar Sarkani. 2017. Why a counterfeit risk avoidance strategy fails. *Computers & Security* (2017).
- [61] Antonio Nicolosi, Maxwell N. Krohn, Yevgeniy Dodis, and David Mazières. 2003. Proactive Two-Party Signatures for User Authentication. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2003, San Diego, California, USA*.
- [62] Kazuo Ohta and Tatsuaki Okamoto. 1991. A Digital Multisignature Scheme Based on the Fiat-Shamir Scheme. In *Advances in Cryptology - ASIACRYPT '91*. 139–148.
- [63] United States. Defense Science Board. Task Force on High Performance Microchip Supply. 2005. *Defense science board task force on high performance microchip supply*. Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics.
- [64] Torben P. Pedersen. 1991. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*. 129–140. https://doi.org/10.1007/3-540-46766-1_9
- [65] Andrea Pellegrini, Valeria Bertacco, and Todd Austin. 2010. Fault-based attack of RSA authentication. In *Proceedings of the conference on Design, automation and test in Europe*. European Design and Automation Association, 855–860.
- [66] David Pointcheval and Jacques Stern. 2000. Security Arguments for Digital Signatures and Blind Signatures. *Journal of cryptology* 13, 3 (2000), 361–396.
- [67] Miodrag Potkonjak, Ani Nahapetian, Michael Nelson, and Tammara Massey. 2009. Hardware Trojan horse detection using gate-level characterization. In *Proceedings of the 46th Design Automation Conference, DAC 2009, San Francisco, CA, USA, July 26-31, 2009*. 688–693. <https://doi.org/10.1145/1629911.1630091>
- [68] Jeyavijayan JV Rajendran and Siddharth Garg. 2017. Logic Encryption. In *Hardware Protection through Obfuscation*. Springer, 71–88.
- [69] Jeyavijayan JV Rajendran, Ozgur Sinanoglu, and Ramesh Karri. 2013. Is split manufacturing secure?. In *Proceedings of the Conference on Design, Automation and Test in Europe*. EDA Consortium, 1259–1264.
- [70] Wolfgang Rankl and Wolfgang Effing. 2004. *Smart card handbook*. John Wiley & Sons.
- [71] RT. 2013. ‘We cannot trust them anymore’: Engineers abandon encryption chips after Snowden leaks. <https://www.rt.com/usa/snowden-leak-rng-randomness-019/>. (December 2013).
- [72] Bruce Schneier. 2013. Surreptitiously Tampering with Computer Chips. <https://www.schneier.com/blog/archives/2013/09/surreptitiously.html>. (November 2013).
- [73] Claus-Peter Schnorr. 1991. Efficient signature generation by smart cards. *Journal of cryptology* 4, 3 (1991), 161–174.
- [74] Hua Hong Semiconductor. 2017. Hua Hong Semiconductor Limited. <http://www.huahonggrace.com/html/about.php>. (May 2017).
- [75] Adi Shamir. 1979. How to share a secret. *Commun. ACM* 22, 11 (1979), 612–613.
- [76] Thomas Shrimpton and R Seth Terashima. 2015. A provable-security analysis of Intel’s secure key RNG. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 77–100.
- [77] S Skorobogatov. 2012. Hardware assurance and its importance to national security. Available On-line: <http://www.cl.cam.ac.uk/sp32/secnews.html> (2012).

- [78] Sergei Skorobogatov and Christopher Woods. 2012. Breakthrough Silicon Scanning Discovers Backdoor in Military Chip. In *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9–12, 2012. Proceedings*. 23–40. https://doi.org/10.1007/978-3-642-33027-8_2
- [79] Oliver Soll, Thomas Korak, Michael Muehlberghuber, and Michael Hutter. 2014. EM-based detection of hardware trojans on FPGAs. In *Hardware-Oriented Security and Trust (HOST), 2014 IEEE International Symposium on*. IEEE, 84–87.
- [80] StarChip. 2017. Smart CARD ICs. <http://www.starchip-ic.com/en/smart-card-chips/>. (May 2017).
- [81] Douglas R. Stinson and Reto Strohli. 2001. Provably Secure Distributed Schnorr Signatures and a (t, n) Threshold Scheme for Implicit Certificates. In *Information Security and Privacy, 6th Australasian Conference, ACISP 2001, Sydney, Australia, July 11–13, 2001, Proceedings (Lecture Notes in Computer Science)*, Vijay Varadharajan and Yi Mu (Eds.), Vol. 2119. Springer, 417–434. https://doi.org/10.1007/3-540-47719-5_33
- [82] Lorenzo Strigini. 2005. Fault tolerance against design faults. (2005).
- [83] Petr Svenda. 2014. Nuances of the JavaCard API on the cryptographic smart cards—JCAIAGTest project. (2014).
- [84] Mohammad Tehranipoor and Cliff Wang. 2011. *Introduction to hardware security and trust*. Springer Science & Business Media.
- [85] Taiwan Semiconductor Manufacturing Company Limited TSMC. 2017. Value Chain Aggregator - KM211. http://www.tsmc.com/english/dedicatedFoundry/services/value_chain_aggregator_km211.htm. (May 2017).
- [86] Sean Turner, Russ Housley, Tim Polk, Daniel RL Brown, and Kelvin Yiu. 2009. Elliptic curve cryptography subject public key information. (2009).
- [87] Riad S. Wahby, Max Howald, Siddharth Garg, abhi shelat, and Michael Wal-fish. 2016. Verifiable ASICs. In *IEEE Security and Privacy (Oakland) 2016, eprint/2016/1243*.
- [88] Adam Waksman and Simha Sethumadhavan. 2010. Tamper evident microprocessors. In *Security and Privacy (SP), 2010 IEEE Symposium on*. IEEE, 173–188.
- [89] Adam Waksman and Simha Sethumadhavan. 2011. Silencing hardware backdoors. In *Security and Privacy (SP), 2011 IEEE Symposium on*. IEEE, 49–63.
- [90] Adam Waksman, Matthew Suozzo, and Simha Sethumadhavan. 2013. FANCI: identification of stealthy malicious logic using boolean functional analysis. In *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4–8, 2013*. 697–708. <https://doi.org/10.1145/2508859.2516654>
- [91] Xinmu Wang. 2014. *Hardware trojan attacks: Threat analysis and low-cost countermeasures through golden-free detection and secure design*. Ph.D. Dissertation. Case Western Reserve University.
- [92] Xinmu Wang, Tatini Mal-Sarkar, Aswin Raghav Krishna, Seetharam Narasimhan, and Swarup Bhunia. 2012. Software exploitable hardware Trojans in embedded processor. In *2012 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, DFT 2012, Austin, TX, USA, October 3–5, 2012*. 55–58. <https://doi.org/10.1109/DFT.2012.6378199>
- [93] Xinmu Wang, Seetharam Narasimhan, Aswin Krishna, Tatini Mal-Sarkar, and Swarup Bhunia. 2011. Sequential hardware trojan: Side-channel aware design and placement. In *Computer Design (ICCD), 2011 IEEE 29th International Conference on*. IEEE, 297–300.
- [94] Yujie Wang, Pu Chen, Jiang Hu, and Jeyavijayan Rajendran. 2016. The cat and mouse in split manufacturing. In *Proceedings of the 53rd Annual Design Automation Conference, DAC 2016, Austin, TX, USA, June 5–9, 2016*. 165:1–165:6. <https://doi.org/10.1145/2897937.2898104>
- [95] S. Wei and M. Potkonjak. 2012. Scalable Hardware Trojan Diagnosis. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 20, 6 (June 2012), 1049–1057. <https://doi.org/10.1109/TVLSI.2011.2147341>
- [96] Sheng Wei and Miodrag Potkonjak. 2013. The undetectable and unprovable hardware trojan horse. In *Proceedings of the 50th Annual Design Automation Conference*. ACM, 144.
- [97] Sheng Wei and Miodrag Potkonjak. 2014. Self-Consistency and Consistency-Based Detection and Diagnosis of Malicious Circuitry. *IEEE Trans. VLSI Syst.* 22, 9 (2014), 1845–1853. <https://doi.org/10.1109/TVLSI.2013.2280233>
- [98] Kaiyuan Yang, Matthew Hicks, Qing Dong, Todd Austin, and Dennis Sylvester. 2016. A2: Analog malicious hardware. (2016).
- [99] Age Yeh. 2012. Trends in the global IC design service market. *DIGITIMES research* (2012).
- [100] Ying C Yeh. 1996. Triple-triple redundant 777 primary flight computer. In *Aerospace Applications Conference, 1996. Proceedings., 1996 IEEE*, Vol. 1. IEEE, 293–307.
- [101] Ying C Yeh. 1998. Design considerations in Boeing 777 fly-by-wire computers. In *High-Assurance Systems Engineering Symposium, 1998. Proceedings. Third IEEE International*. IEEE, 64–72.
- [102] Jie Zhang, Feng Yuan, Lingxiao Wei, Yannan Liu, and Qiang Xu. 2015. VeriTrust: Verification for Hardware Trust. *IEEE Trans. on CAD of Integrated Circuits and Systems* 34, 7 (2015), 1148–1161. <https://doi.org/10.1109/TCAD.2015.2422836>

A MULTI-SIGNATURE SCHEME PROOF

In this Appendix we discuss the security of our construction introduced in Section 4.5. Before moving to the proof of the Theorem 4.1 we recall the definition of the one-more discrete logarithm problem [13] and security definitions for multi-signatures schemes.

Definition A.1 (N-DL). A group generator $\mathcal{G}(1^\lambda)$ is a probabilistic polynomial time algorithm that on input a security parameter λ returns a pair (p, G) where p is a λ -bit prime and G is a random generator of a cyclic group \mathbb{G} of order p .

An algorithm \mathcal{A} to solve the N -DL problem is a probabilistic polynomial time algorithm which receives as input an instance $(p, G) \leftarrow \mathcal{G}(1^\lambda)$ and can access two oracle O_{ch}, O_{dlog} . Upon request, the former returns a random group element in \mathbb{G} . The latter gets as input a group element C and returns its discrete logarithm with respect to the generator G , i.e. such that $C = x \cdot G$. Let C_1, C_2, \dots, C_N be the challenges returned by O_{ch} . We say that adversary \mathcal{A} wins if he returns $c_1, \dots, c_N \in \mathbb{Z}_p$ satisfying $C_i = c_i \cdot G$ by using a number of queries to O_{dlog} strictly smaller than N .

Next, we recall definitions of multi-signatures of [14], adapted to our settings.

Definition A.2 (Multi-Signature). A multi-signature scheme is a tuple $(KeyGen, Sign, Verify)$ of algorithms.

- ❖ *KeyGen*: This is an interactive protocol between n parties IC_i to jointly compute a common shared verification key Y and n individual signing keys x_i associated with Y .
- ❖ *Sign*: This is an interactive protocol between the n ICs that given input a common message m and their individual secret keys x_i allows to compute a shared signature Σ on m .
- ❖ *Verify*: This is a deterministic algorithm that given in input the common verification key Y , a signature Σ and a message m returns 1 if the signature is valid and 0 otherwise.

In the construction presented in Section 4.5 the key generation protocol is handled by the DKPG; *Sign* is described by Algorithm 4.7 and Figures 4 and 5; *Verify* consists of the verification algorithm of standard Schnorr signatures on the aggregated signature using the shared verification key Y .

Definition A.3 (Security Game). We consider and adversary \mathcal{F} attempting to forge a multi-signature. The attack is modelled as a game in three phases.

- ❖ *Setup*: We assume the key generation protocol among n parties to be successfully executed and returning a shared public key Y and a set of n secret keys x_i .
- ❖ *Attack*: We assume the forger \mathcal{F} to corrupt $n - 1$ ICs and learn their own individual secret keys x_i . Without loss of generality, we assume user IC_1 to be the only honest user in the system. The forger \mathcal{F} interacts as the host over q_S interactive signing sessions with IC_1 and arbitrarily deciding on the messages to be signed. Let Q to be the set of messages m used in the interactive signing sessions.
- ❖ *Forgery*: At the end of its execution, the forger \mathcal{F} returns a signature Σ on a message m which was not used on a signing session with IC_1 , namely $m \notin Q$. The forger \mathcal{F} wins the game if $Verify(Y, m, \Sigma) = 1$.

We model the security of the scheme in the random oracle model [15, 37]. This means that we assume the hash function to behave as an ideal random function. This is modelled by giving to the adversary \mathcal{F} access to an oracle O_{Hash} returning random values in a range $\{0, 1\}^\lambda$.

The advantage $adv_{ms}(\mathcal{F})$ of algorithm \mathcal{F} in forging a multi-signature is defined to be the probability that \mathcal{F} wins the above game, where the probability is taken over the random coins of \mathcal{F} , IC_1 , DKGP and the random oracle. We say that $\mathcal{F}(q_S, q_H, \epsilon)$ -breaks the multi-signature scheme if it participates in q_S signing sessions with IC_1 , makes at most q_H queries to O_{Hash} and the advantage $adv_{ms}(\mathcal{F}) \geq \epsilon$.

We now restate and prove Theorem 4.1.

THEOREM 4.1. *If there exists a (q_S, q_H, ϵ) -forger \mathcal{F} for the multi-signature scheme described in Figures 4 and 5 interacting in $q_S = O(1)$ signature queries, making at most q_H Hash queries and succeeding with probability ϵ , then there exists an algorithm \mathcal{A} that solves the $(q_S + 1)$ -DL problem with probability at least*

$$\delta \geq \frac{\epsilon^2}{q_H^{q_S+1}} - \text{negl}(\lambda)$$

PROOF. We start by describing the idea behind the reduction. Assume for a moment that a forger \mathcal{F} is able to produce two valid signatures (σ, ϵ) and (σ', ϵ') on the same message, i.e.

$$\epsilon = \text{Hash}(\sigma \cdot G + \epsilon \cdot Y || H(m) || j) \quad \epsilon' = \text{Hash}(\sigma' \cdot G + \epsilon' \cdot Y || m || j) \quad (2)$$

and such that

$$R = \sigma \cdot G + \epsilon \quad R = \sigma' \cdot G + \epsilon' \cdot Y \quad (3)$$

Dividing the two equations we obtain

$$(\sigma - \sigma') \cdot G = (\epsilon' - \epsilon) \cdot Y$$

then we get the discrete logarithm of Y with respect to G by computing $(\sigma - \sigma') * (\epsilon' - \epsilon)^{-1} \bmod p$.

Given a forger \mathcal{F} to output forged signatures, we construct an adversary \mathcal{A} for solving the $(q_S + 1)$ -DL problem. In the process, \mathcal{A} has to simulate signatures as produced by IC_1 and answer to the random oracle queries made by \mathcal{F} during the attack, as it happen in the security game for multi-signatures. In case \mathcal{F} succeeds in forging a first signature, then adversary \mathcal{A} rewinds \mathcal{F} and replay him reusing the same coin tosses used in the first execution. However in the second run of \mathcal{F} , the adversary replaces the answer of the random oracle query corresponding to the forgery produced in the first execution with a fresh random string in the range of the hash function. By applying a forking lemma type of argument [66] one can then argue that with good probability the replay of \mathcal{F} will return a new forgery which has the same R as in the first forged signature but different (σ', ϵ') . Once adversary \mathcal{A} obtains two signatures of this kind, he will be able to compute the discrete logarithm of Y as shown above³.

The description of \mathcal{A} is given in Algorithm A.1. Throughout the execution, adversary \mathcal{A} keeps track of the random oracle queries made by \mathcal{F} using a list $L[\cdot]$. Without loss of generality we assume

³The adversary described in the proof actually has to compute the discrete logarithm of Y_1 , the public key of IC_1 . This is trivial to do once the adversary \mathcal{A} has obtained the discrete logarithm of Y , the shared verification key.

the adversary \mathcal{F} checks all the signature queries made to her oracle, as well as the signature he attempts to forge.

We can summarise the interaction of adversary \mathcal{A} with \mathcal{F} in the following phases.

- ❖ **Setup Phase:** \mathcal{A} initialises an empty list L . She picks a random string ρ for the randomness used by adversary \mathcal{F} as well as q_H random values π_i from the range of the hash function, hereby set to be $\{0, 1\}^\lambda$. The adversary then simulates the multi-signature key generation process: she generates pairs of public and secret keys for the corrupted cards IC_2, \dots, IC_n , as honest cards would do, and generates the public key of IC_1 by querying his own challenge oracle $Y_1 \leftarrow O_{ch}$. Finally, she generates the shared public key Y as in the key generation protocol (Algorithm 4.3). Then she starts adversary \mathcal{F} on inputs the shared verification key Y , the list of verification key-shares and the $n - 1$ signing key corresponding to the corrupted ICs.
- ❖ **Caching Phase:** the adversary \mathcal{F} , acting as the hosts, initiates the signature process by sending *Init* to IC_1 . In this phase the adversary emulates the caching phase of IC_1 by making q_S queries to O_{ch} and returning the list of group elements to \mathcal{F} .
- ❖ **Random Oracle Queries:** whenever \mathcal{F} makes a new query (*Hash*, M) to its random oracle, adversary \mathcal{A} answers it by picking the next unused random string in $\{\pi_1, \pi_2, \dots, \pi_{q_H}\}$. She keeps track of \mathcal{F} requests to answer consistently.
- ❖ **Signing Queries:** in this phase the forger submits signature queries as the host in the system. The adversary \mathcal{A} hashes the message, consistently with the list of random oracle queries and uses her oracle O_{dlog} to compute the requested signature.
- ❖ **Replay:** \mathcal{F} eventually attempts to output a forged multi-signature (σ, ϵ) . If the signatures verifies, \mathcal{A} rewinds adversary \mathcal{F} to the beginning and replay it on input the same random coins ρ and same inputs and reusing the same response in the caching phase. \mathcal{A} answers the random oracle and signing queries as in the first execution apart from the random oracle query which generated ϵ . \mathcal{A} answers this query by picking a random string $\pi'_i \in \{0, 1\}^\lambda$. The forger attempts to produce a new signature (σ', ϵ') . If conditions (2) and (3) are met, then \mathcal{A} computes the discrete logarithm of Y , and thus the discrete logarithm of Y_1 . Given the latter and the responses of O_{dlog} she can also compute the discrete logarithm of the $R_{1,j}$, the outputs of O_{ch} .

In order to succeed in her game, adversary \mathcal{A} has to simulate the interaction between \mathcal{F} and IC_1 as described in the security game of multi-signatures. The above adversary perfectly simulates the multi-signature key generation phase, the caching phase of IC_1 and the random oracle queries of \mathcal{F} . Extra care needs to be taken for signature queries simulating the interaction of IC_1 and \mathcal{F} .

Simulation of Signature Queries. Note that differently from standard Schnorr signatures, we cannot simulate a signature by programming the random oracle. This is because the caching phase of IC_1 fixes the randomness used in the signatures before \mathcal{A} knows the messages corresponding to them, preventing him to program the random oracle accordingly. Thus, in order to simulate IC_1 's

Algorithm A.1: Adversary $\mathcal{A}_{\mathcal{F}}^{O_{ch}, O_{dlog}}$ against $(q_S + 1)$ -DL

Input : (p, G) .
Output : $q_S + 1$ discrete logarithms using q_S calls to O_{dlog}

- 1 **Setup Phase**
- 2 $ctr = 0, j = 1$
- 3 Set $L[\cdot] = \emptyset, J[\cdot] = \emptyset$
- 4 Set random string ρ for adversary \mathcal{F} .
- 5 Pick $\pi_1, \pi_2, \dots, \pi_{q_H} \leftarrow \{0, 1\}^\lambda$
- 6 $Y_1 \leftarrow O_{ch}$
- 7 **for** $i = 2$ **to** n **do**
- 8 $x_i \leftarrow \mathbb{Z}_p$
- 9 $Y_i = x_i \cdot G \leftarrow O_{ch}$
- 10 $Y = \sum_{i=1}^n Y_i$
- 11 **Start** $\mathcal{F}(Y, (Y_1, \dots, Y_n), (x_2, \dots, x_n); \rho)$
- 12 **Caching Phase:** $Init \leftarrow \mathcal{F}$
- 13 $R_{1,1}, R_{1,2}, \dots, R_{1,q_S} \leftarrow O_{ch}$
- 14 **Run** $\mathcal{F}(R_{1,1}, R_{1,2}, \dots, R_{1,q_S})$
- 15 **Random Oracle Queries:** $(Hash, M) \leftarrow \mathcal{F}$
- 16 **if** $L[M] = \perp$ **then**
- 17 $ctr+ = 1$
- 18 $L[M] = ctr$
- 19 $k = L[M]$
- 20 **Run** $\mathcal{F}(\pi_k)$
- 21 **Signing Queries:** $(Sign, (R, H(m), j')) \leftarrow \mathcal{F}$
- 22 **if** $j' \neq j$ **then**
- 23 **return** \perp to \mathcal{F}
- 24 **else**
- 25 $j+ = 1$
- 26 **if** $L[R||H(m)||j'] = \perp$ **then**
- 27 $ctr+ = 1$
- 28 $L[R||H(m)||j'] = ctr$
- 29 $k = L[[R||H(m)||j']]$
- 30 $J[j'] = k$
- 31 $\epsilon_{j'} = \pi_k$
- 32 $\sigma_{j'} \leftarrow O_{dlog}(R_{1,j'} - h \cdot Y)$
- 33 **Run** $\mathcal{F}((\sigma_{j'}, \epsilon_{j'}))$
- 34 $(Forgery, (m, (\sigma, \pi_i))) \leftarrow \mathcal{F}$
- 35 **if** $Verify(Y, m, (\sigma, \pi_i)) = 0$ **then**
- 36 **return** \perp
- 37 **Replay** \mathcal{F}
- 38 $(Forgery, (m, (\sigma', \pi'_i))) \leftarrow \mathcal{F}$
- 39 **if** $Verify(Y, m, (\sigma', \pi'_i)) = 0$ **then**
- 40 **return** \perp
- 41 $x_1 = (\sigma - \sigma') * (\pi'_i - \pi_i)^{-1} - \sum_{j=2}^n x_j \pmod p$
- 42 **for** $j = 1$ **to** q_S **do**
- 43 $k = J[j]$
- 44 $r_{1,j} = s_j + x_1 * \pi_k \pmod p$
- 45 **return** $(x_1, r_{1,1}, \dots, r_{1,q_S})$

signatures \mathcal{A} takes advantage of her O_{dlog} oracle. As \mathcal{A} has to make fewer queries to O_{dlog} than to O_{ch} , she is allowed one call to the discrete logarithm oracle for each signature in the first round. The problem arises when \mathcal{A} replays the forger \mathcal{F} . In this execution we want \mathcal{F} to behave exactly as in the first round up to the point he queries (a second time) the random oracle on the forged message. This means that in the caching phase \mathcal{A} is forced to reuse the same random group elements as in the first execution. Since \mathcal{A} cannot make new queries to the challenge oracle, she cannot make any more queries to the O_{dlog} during the second execution of \mathcal{F} . Therefore, in the second execution, \mathcal{A} has to simulate signatures (on potentially different messages) without the help of her oracle. Assume for a moment that the index of hash queries used for the creation of signatures is the same in both executions. In this case, the response of the random oracle on the corresponding hash queries will be the same, even if the message has changed during the two execution. Since the the output of the hash is the same in both runs, \mathcal{A} is able to answer the signature queries by simply recycling the signature produced in the first execution.

Computation of Discrete Logarithm. Under the condition that \mathcal{A} simulates \mathcal{F} 's environment, the above adversary succeeds to extract the discrete logarithm of Y_1 as long as the index of the hash query associated to the forged signatures is the same in both runs. In this case conditions (2) and (3) are met with overwhelming probability and \mathcal{A} can proceed to compute discrete logarithm as illustrated above.

From the above observations we can summarise two conditions which make adversary \mathcal{A} to succeeds in her game:

- (1) The index i of the random oracle query corresponding to the forged signature is the same in both executions of \mathcal{F} .
- (2) The set J of indexes of random oracle queries corresponding to signature queries is the same in both execution of \mathcal{F} .

By assuming that \mathcal{F} forges signatures with non-negligible probability ϵ , the following lemma shows the above conditions holds with non-negligible probability. This concludes the proof as it gives a bound δ , the probability of success of \mathcal{A} . \square

The following lemma is an adaptation of the generalized forking lemma of [14] and Lemma 5 of [61].

LEMMA A.4. *Let \mathcal{F} be an adversary producing a multi-signature forgery with probability ϵ . Assuming that \mathcal{F} makes at most q_H hash queries and interacts in $q_S = O(1)$ signatures, then by replaying \mathcal{F} , \mathcal{A} succeeds with probability $\delta \geq \frac{\epsilon^2}{q_S+1} - \text{negl}(\lambda)$.*

PROOF. Let $i \in [q_H]$ be the index of the random oracle query associated to the forgery returned by \mathcal{F} and let $J \subseteq [q_H], |J| \leq q_S$ be the set of indexes of oracle queries associated with the signatures queries made by \mathcal{F} . Call i and J the *target* indexes. Note that J and $\{i\}$ are disjoint as a valid forgery cannot be on a message queried on the signature oracle.

For any fixed i, J let $\epsilon_{i,J}$ be the probability that \mathcal{F} produces a forgery on target indexes i, J . The probability of \mathcal{F} forging a signature is given by $\epsilon = \sum_{i,J} \epsilon_{i,J}$.

Consider now all the random inputs given to \mathcal{F} . These includes the random coins ρ , his inputs, the response of the caching phase, the responses of the random oracle queries and the signature

queries. We split these into: the i -th response to the random oracle query, π_i , and everything else, which we call R . Consider now a matrix $M_{i,J}$ with one column for every possible answer to π_i and one row for any possible string R . The entries of the matrix are 1, in case R and π_i will make \mathcal{F} produce a forgery for target indexes i, J , and 0 otherwise. For any i, J , let $\epsilon_{i,J,R}$ be the probability of the adversary \mathcal{F} to produce a forgery given randomness R . This corresponds to the density of the row R in the matrix $M_{i,J}$. Similarly, $\epsilon_{i,J}$ corresponds to the density of the matrix $M_{i,J}$, namely

$$\epsilon_{i,J} = \frac{1}{|R|} \sum_R \epsilon_{i,J,R}$$

In a similar fashion we define $\delta_{i,J}$ to be the probability of producing two forgeries on the same target indexes i, J while replaying the adversary on same randomness and replacing π_i with a random $\pi'_i \leftarrow \{0, 1\}^\lambda$. Similarly, we also define $\delta_{i,J,R}$ for which we have

$$\delta_{i,J} = \frac{1}{R} \sum_R \delta_{i,J,R}$$

We now relate probability δ of adversary \mathcal{A} to \mathcal{F} . The probability $\delta_{i,J,R}$ corresponds to the event of sampling randomness R for the adversary \mathcal{F} , hitting a first 1 in the row R of $M_{i,J}$ and then probe another random column in the same row and hitting another 1. Since the two runs are independent, the probability of succeeding in the replay attack is

$$\delta_{i,J,R} = \epsilon_{i,J,R} * \left(\epsilon_{i,J,R} - \frac{1}{2^\lambda} \right)$$

Replacing the above in the expression of $\delta_{i,J}$ and applying the Cauchy-Schwartz inequality give the following

$$\begin{aligned} \delta_{i,J} &= \frac{1}{|R|} \sum_R \delta_{i,J,R} \\ &= \frac{1}{|R|} \left(\sum_R \epsilon_{i,J,R}^2 - \frac{\epsilon_{i,J,R}}{2^\lambda} \right) \\ &\geq \left(\frac{1}{|R|} \sum_R \epsilon_{i,J,R} \right)^2 - \frac{\epsilon_{i,J}}{2^\lambda} \\ &= \epsilon_{i,J}^2 - \frac{\epsilon_{i,J}}{2^\lambda} \end{aligned}$$

Given the above we get the following bound on δ

$$\begin{aligned} \delta &= \sum_{i,J} \delta_{i,J} \geq \sum_{i,J} \left(\epsilon_{i,J}^2 - \frac{\epsilon_{i,J}}{2^\lambda} \right) \\ &= \sum_{i,J} \epsilon_{i,J}^2 - \frac{\epsilon}{2^\lambda} \\ &\geq \frac{1}{q_H^{q_S+1}} \left(\sum_{i,J} \epsilon_{i,J} \right)^2 - \frac{\epsilon}{2^\lambda} \\ &= \frac{\epsilon^2}{q_H^{q_S+1}} - \frac{\epsilon}{2^\lambda} \end{aligned}$$

where the last inequality is obtained by applying again the Cauchy-Schwartz inequality. \square

B PROTOTYPE EXTENSIONS

Our Myst prototype can also be extended to use multiple ICs boards, if higher throughput is needed. Figure 11 depicts an instantiation with 240 ICs: two boards holding 120 ICs each (60 ICs on each side) and an Artix-7 FPGA to facilitate the inter-IC communication within each board.




Figure 11: Myst's prototype with 240 JavaCards fitted into an 1U rack case.