



# Trojan-Resilience Without Cryptography

Suvradip Chakraborty<sup>2</sup>, Stefan Dziembowski<sup>1</sup>, Małgorzata Gałązka<sup>1(✉)</sup>,  
Tomasz Lizurej<sup>1</sup>, Krzysztof Pietrzak<sup>2</sup>, and Michelle Yeo<sup>2</sup>

<sup>1</sup> University of Warsaw, Warsaw, Poland

<sup>2</sup> IST Austria, Klosterneuburg, Austria

**Abstract.** Digital hardware Trojans are integrated circuits whose implementation differ from the specification in an arbitrary and malicious way. For example, the circuit can differ from its specified input/output behavior after some fixed number of queries (known as “time bombs”) or on some particular input (known as “cheat codes”).

To detect such Trojans, countermeasures using multiparty computation (MPC) or verifiable computation (VC) have been proposed. On a high level, to realize a circuit with specification  $\mathcal{F}$  one has more sophisticated circuits  $\mathcal{F}^\circ$  manufactured (where  $\mathcal{F}^\circ$  specifies a MPC or VC of  $\mathcal{F}$ ), and then embeds these  $\mathcal{F}^\circ$ 's into a *master circuit* which must be trusted but is relatively simple compared to  $\mathcal{F}$ . Those solutions impose a significant overhead as  $\mathcal{F}^\circ$  is much more complex than  $\mathcal{F}$ , also the master circuits are not exactly trivial.

In this work, we show that in restricted settings, where  $\mathcal{F}$  has no evolving state and is queried on independent inputs, we can achieve a relaxed security notion using very simple constructions. In particular, we do not change the specification of the circuit at all (i.e.,  $\mathcal{F} = \mathcal{F}^\circ$ ). Moreover the master circuit basically just queries a subset of its manufactured circuits and checks if they're all the same.

The security we achieve guarantees that, if the manufactured circuits are initially tested on up to  $T$  inputs, the master circuit will catch Trojans that try to deviate on significantly more than a  $1/T$  fraction of the inputs. This bound is optimal for the type of construction considered, and we provably achieve it using a construction where 12 instantiations of  $\mathcal{F}$  need to be embedded into the master. We also discuss an extremely simple construction with just 2 instantiations for which we conjecture that it already achieves the optimal bound.

---

T. Lizurej—Stefan Dziembowski, Małgorzata Gałązka, and Tomasz Lizurej were supported by the 2016/1/4 project carried out within the Team program of the Foundation for Polish Science co-financed by the European Union under the European Regional Development Fund.

K. Pietrzak—Suvradip and Krzysztof have received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (682815 - TOCNeT).

© International Association for Cryptologic Research 2021

K. Nissim and B. Waters (Eds.): TCC 2021, LNCS 13043, pp. 397–428, 2021.

[https://doi.org/10.1007/978-3-030-90453-1\\_14](https://doi.org/10.1007/978-3-030-90453-1_14)

# 1 Hardware Trojans

Preventing attacks on cryptographic hardware that are based on leakage and tampering has been a popular topic both in the theory in the practical research communities [6–10, 15]. Despite being very powerful, the models considered in this area are restricted in the sense that it is typically assumed that a given device has been manufactured correctly, i.e., the adversary is present during the execution of the device, but not when it is produced. As it turns out, this assumption is not always justifiable, and in particular in some cases the adversary may be able to modify the device at the production time. This is because, for economic reasons, private companies and government agencies are often forced to use hardware that they did not produce themselves. The contemporary, highly-specialized digital technology requires components that are produced by many different enterprises, usually operating in different geographic locations. Even a single chip is often manufactured in a production cycle that involves different entities. In a very popular method of hardware production, called the *foundry model*, the product designer is only developing the abstract description of a device. The real hardware fabrication happens in *foundry*. Only few major companies (like Intel) still manufacture chips by themselves [16].

Modifications to the original circuit specification introduced during the manufacturing process (in a way that is hard to detect by inspection and simple testing) are called *hardware Trojans*, and can be viewed as the extreme version of hardware attacks. For more on the practical feasibility of such attacks the reader may consult, e.g., books [13, 16], or popular-science articles [1, 12]. Hardware Trojans can be loosely classified into *digital* and *physical* ones. Physical hardware Trojans can be triggered and/or communicate via a physical side-channel, while digital hardware Trojans only use the regular communication interfaces. In this paper we only consider digital hardware Trojans.

## 1.1 Detecting Digital Hardware Trojans

A simple non-cryptographic countermeasure to detect whether a circuit  $F$  contains a hardware Trojan or follows the specification  $\mathcal{F}$  is testing: one samples inputs  $x_1, \dots, x_T$ , queries  $y_i \leftarrow F(x_i)$  and checks whether  $y_i = \mathcal{F}(x_i)$  for all  $i$ . Two types of digital hardware Trojans discussed in the literature that evade detection by such simple testing are time bombs and cheat codes (see, e.g., [5]). A *time bomb* is a hardware Trojan where the circuit starts deviating after a fixed number of queries. *Cheat codes* refer to hardware Trojans where the circuits deviate on a set of hard-coded inputs. To achieve some robustness against all digital hardware Trojans, solutions using cryptographic tools, in particular verifiable computation (VC) [2, 17] and multiparty computation (MPC) [5] were suggested. In both cases the idea is to take the specification  $\mathcal{F}$  of the desired circuit and replace it with a more sophisticated construction of one or more circuits  $\mathcal{F}^\diamond$ . The circuit(s)  $F^\diamond$  that (presumably) are manufactured according to specification  $\mathcal{F}^\diamond$  are then embedded into a master circuit  $M$  to get a circuit  $M^{F^\diamond}$ .

Method	Specification of circuit $\mathcal{F}^\diamond$ (to be manufactured by untrusted fab) to realize $\mathcal{F}$	Test and (trusted) Master Circuit	Security with Trojans	Functionality when no Trojans present
MPC [5]	<b>Different than <math>\mathcal{F}</math>:</b> Implements functionality $\mathcal{F}$ in terms of a 3-player MPC computation.	<b>Testing: <math>T</math> test queries.</b> <b>Master's computation — non-trivial:</b> Master secret shares inputs and reconstructs outputs. It needs lots of randomness.	Gives $n < T/2$ correct and guaranteed outputs with prob. $1 - \exp(-k)$ . No restriction on $\mathcal{F}$ .	<b>Limited:</b> Behaves like $\mathcal{F}$ for $n = T/2$ queries, then stops.
VC [17,2]	<b>Different than <math>\mathcal{F}</math>:</b> $\mathcal{F}$ plus succinct proof of correct computation	<b>Testing: no test queries.</b> <b>Master's computation — non-trivial:</b> Master verifies succinct proof.	No guarantee on number of outputs, but as long as outputs are provided they are correct. No restriction on $\mathcal{F}$	<b>Ideal:</b> behaves like $\mathcal{F}$ for unbounded number of outputs.
This work	<b>Same as <math>\mathcal{F}</math>:</b> 12 instantiations needed in our provable construction, 2 conjectured to be already sufficient.	<b>Testing: <math>T</math> test queries.</b> <b>Master's computation — very simple:</b> Master only does equality checks and Multiplexer. Needs tiny amount of randomness.	Either the Trojans will be detected with probability $1 - o(1)$ , or at most a $O(1/T)$ fraction of the outputs is wrong. Requires that $\mathcal{F}$ has no evolving state and inputs are independent.	<b>Ideal:</b> (as above)

**Fig. 1.** Comparison of cryptographic solutions with our new construction. We achieve weaker security, but with a *much* simpler construction.

which is proven to follow specification  $\mathcal{F}$  with high probability as long as it produces outputs. The master circuit must be trusted, but hopefully can be much simpler than  $\mathcal{F}$ . We elaborate on these two methods below.

*Using Verifiable Computation.* Here the idea is to let  $\mathcal{F}^\diamond(x)$  output a tuple  $(y, \pi)$  where  $y = \mathcal{F}(x)$  and  $\pi$  is a succinct zero-knowledge proof (see [3]) that  $y$  is the correct output. In the compiled circuit  $M^{\mathcal{F}^\diamond}$  the master  $M$  in input  $x$

invokes  $(y, \pi) \leftarrow F^\diamond$ , then verifies the proof  $\pi$  and only outputs  $y$  if the check passes. If verification fails, the master aborts with a warning. As long as the compiled circuit provides outputs, they are guaranteed to be correct. If there are no Trojans, the number of outputs is unbounded; but if there is a Trojan, they can make the compiled circuit abort already at the first query. See [2, 17] for the details.

*Using Multiparty Computation.* In this case, the idea is to use secure multiparty computation protocols (MPCs, see, e.g., [4]). The compiled circuit  $M^{F^\diamond}$  contains some number of sub-components  $F^\diamond$  that communicate only via the master circuit. In [5], this number is  $3k$  (where  $k$  is a parameter). The sub-components are grouped in triples, each of them executing a 3-party protocol. In order to avoid the “cheat code” attacks, the master secret shares the input between the 3 parties. To get assurance that the sub-components are not misbehaving they are tested before deployment. In order to avoid the “time bomb” attacks, the number of times each sub-component is tested is an independently chosen random number from 1 to  $T$ . The output of each triple is secret-shared between its sub-components. Each of them sends its share to the master circuit, who reconstructs the  $k$  secrets, and outputs the value that is equal to the majority of these secrets. For the details see [5].

*Simple Schemes.* In this work we consider compilers as discussed above, but only particularly simple ones which have the potential of being actually practical. In particular, we require that  $\mathcal{F} \equiv \mathcal{F}^\diamond$ . That is, the specification  $\mathcal{F}^\diamond$  of the functionality given to the untrusted manufacturer is the actual functionality  $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$  we want to implement. Moreover, our master just invokes (a random subset of) the circuits on the input and checks if the outputs are consistent.

This restricted model has very appealing properties. For example, it means one can use our countermeasures with circuits that have already been manufactured. But there are also limitations on what type of security one can achieve. Informally, the security we prove for our construction roughly states that for any constant  $c > 0$  there exists a constant  $c'$  such that no malicious manufacturer can create Trojans which (1) will not be detected with probability at least  $c$ , and (2) if not detected, will output a  $\geq c'/T$  fraction of wrong outputs. Here  $T$  is an upper bound on the number of test queries we can make to the Trojans before they are released.

In particular, we only guarantee that most outputs are correct, and we additionally require that the inputs are iid. Unfortunately, it’s not hard to see that for the simple class of constructions considered these assumptions are not far

from necessary.<sup>1</sup> We will state the security of the VC and MPC solutions using our notion of Trojan-resilience in Sect. 2.7.

It’s fair to ask whether our notion has any interesting applications at all. Two settings in which Trojan resilience might be required are (1) in settings where a computation is performed where false (or at least too many false) outputs would have serious consequences, and (2) cryptographic settings where the circuit holds a key or other secret values that should not leak.

For (1) our compiler would only be provably sufficient if the inputs are iid, and only useful if a small fraction of false outputs can be tolerated. This is certainly a major restriction, but as outlined above, if one doesn’t have the luxury to manufacture circuits that are much more sophisticated than the required functionality, it’s basically the best one can get. Depending on the setting, one can potentially use our compiler in some mode – exploiting redundancy or using repetition – to fix those issues. We sketch some measures in the cryptographic setting below.

For the cryptographic setting (2) our notion seems even less useful: if the adversary can learn outputs of the Trojans, he can use the  $\Theta(1/T)$  fraction of wrong outputs to embed (and thus leak) its secrets. While using the compiler directly might not be a good idea, we see it as a first but major step towards simple and Trojan-resilient constructions in the cryptographic setting. As an example, consider a weak PRF  $\mathcal{F} : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  (a weak PRF is defined like a regular PRF, but the outputs are only pseudorandom if queried on random inputs). While implementing  $\mathcal{F}(k, \cdot)$  using our compiler directly is not a good idea as discussed above,<sup>2</sup> we can compile  $t > 1$  weak PRFs with independent keys and inputs and finally XOR the outputs of the  $t$  master circuits to implement a weak PRF  $\mathcal{F}_3((k_1, \dots, k_t), (x_1, \dots, x_t)) = \bigoplus_{i=1}^t \mathcal{F}(k_i, x_i)$ . Intuitively, the output can only leak significant information about the keys if *all*  $t$  outputs are wrong as otherwise the at least one pseudorandom output will mask everything. If each output is wrong with probability, say  $1/T$  for a modest  $T = 2^{30}$  and we use  $t = 3$ , then for each query we only have a probability of  $1/T^3 = 2^{-90}$  that all  $t = 3$  outputs deviate, which we can safely assume will never happen. Unfortunately, at this point we can’t prove the above intuition and leave this for future work. For one thing, while we know that the XOR will not leak much if at least one of the  $t$  values is correct when the weak PRFs are modelled as ideal ciphers [11], we don’t have a similar result in the computational setting. More importantly, we only

<sup>1</sup> We show that a small fraction of wrong outputs must be allowed in Sect. 2.4. The iid assumption can be somewhat relaxed, but as we don’t have a clean necessary condition we will not discuss this further in this paper. Informally, a sufficient condition seems to just require that there is no (efficiently recognisable) subset of inputs which appear rarely (not more than with probability around  $1/T$ ) but can come in “bursts”, say two such inputs are consecutive with prob.  $\gg 1/T^2$ .

<sup>2</sup> It’s acceptable by our construction if the inputs are iid conditioned on some secret, so the master on input  $x$  and key  $k$  can forward  $(k, x)$  to the circuits. Alternatively the key  $k$  can be hard-coded in the circuit (probably not a good idea if the manufacturer is not trusted in the first place) or, if the circuits have some storage, one can give them  $k$  after receiving the circuits from the manufacturer.

prove that at most a  $1/T$  fraction of the outputs is wrong once a sufficiently large number of queries was made, but to conclude that in the above construction all  $t$  instances fail at the same time with probability at most  $1/T^t$  we need a stronger statement saying that for each individual query the probability of failure is  $1/T$  (we believe that this is indeed true for our construction, but the current proof does not imply this).

Weak PRFs are sufficient for many basic symmetric-key cryptographic tasks like authentication or encryption.<sup>3</sup> Even if a fraction of outputs can be wrong, as long as they don't leak the key (as it seems to be the case for the construction just sketched), this will only affect completeness, but not security. An even more interesting construction, and the original motivation for this work, is a trojan-resilient stream cipher. This could then be used to e.g., generate the high amount of randomness required in side-channel countermeasures like masking schemes. The appealing property of a stream-cipher in this setting is that we don't care about correctness at all, we just want the output to be pseudorandom. It's not difficult to come up with a candidate for such a stream-cipher based on our compiler, but again, a proof will require more ideas. One such construction would start with the weak PRF construction just discussed, and then use two instantiations of it in the leakage-resilient mode from [14].

## 2 Definition and Security of Simple Schemes

For  $m \in \mathbb{N}$ , an  $m$ -redundant simple construction  $\Pi_m = (\mathsf{T}^*, \mathsf{M}^*)$  is specified by a master circuit  $\mathsf{M}^* : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\text{abort}\}$  and a test setup  $\mathsf{T}^* : \mathbb{N} \rightarrow \{\text{fail}, \text{pass}\}$ .<sup>4</sup> The  $*$  indicates that they expect access to some “oracles”. The following oracles will be used: (a)  $\mathsf{F}_1, \dots, \mathsf{F}_m$ —the Trojan circuits that presumably implement the functionality  $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$ , (b)  $\mathsf{F}$ —a trusted implementation of  $\mathcal{F}$  (only available in the test phase), and (c)  $\mathcal{S}$ —a source of random bits (sometimes we will provide the randomness as input instead),

### 2.1 Test and Deployment

The construction  $\Pi_m$  which implements  $\mathcal{F}$  in a Trojan-resilient way using the untrusted  $\mathsf{F}_1, \dots, \mathsf{F}_m$  is tested and deployed as follows.

Lab Phase (test): In this first phase we execute  $\{\text{pass}, \text{fail}\} \leftarrow \mathsf{T}^{\mathsf{F}_1, \dots, \mathsf{F}_m, \mathsf{F}, \mathcal{S}}(T)$

The input  $T$  specifies that each  $\mathsf{F}_i$  may be queried at most  $T$  times. If the output is fail, a Trojan was detected. Otherwise (i.e. the output is pass) we move to the next phase.

Wild Phase (deployment): If the test outputs pass, the  $\mathsf{F}_i$ 's are embedded into the master to get a circuit  $\mathsf{M}^{\mathsf{F}_1, \dots, \mathsf{F}_m, \mathcal{S}} : \mathcal{X} \rightarrow \mathcal{Y} \cup \text{abort}$ .

<sup>3</sup> To encrypt  $m$  sample a random  $r$  and compute the ciphertext  $(r, \mathcal{F}(k, r) \oplus m)$ .

<sup>4</sup> We consider much stronger  $\mathsf{M}^*, \mathsf{T}^*$  for the lower bounds compared to what we require in the constructions as discussed in Sect. 2.5.

## 2.2 Completeness

The completeness requirement states that if every  $F_i$  correctly implements  $\mathcal{F}$ , then the test phase outputs **pass** with probability 1 and the master truthfully implements the functionality  $\mathcal{F}$ . That is, for every sequence  $x_1, x_2, \dots, x_q$  (of arbitrary length and potentially with repetitions) we have

$$\Pr[\forall i \in [q] : y_i = \mathcal{F}(x_i)] = 1 \text{ where for } i = 1 \text{ to } q : y_i := M^{F_1, \dots, F_m, \mathcal{S}}(x_i)$$

The reason we define completeness this way and not simply for all  $x$  we have  $\Pr[M^{F_1, \dots, F_m, \mathcal{S}}(x) = \mathcal{F}(x)]$  is that the Trojan  $F_i$  can be stateful, so the order in which queries are made does matter.

## 2.3 Security of Simple Schemes

We consider a security game  $\text{TrojanGame}(\Pi, T, Q)$  where, for some  $T, Q \in \mathbb{Z}$ , an adversary  $\text{Adv}$  can choose the functionality  $\mathcal{F}$  and the Trojan circuits  $F_1, \dots, F_m$ . We first run the test phase  $\tau \leftarrow T^{F_1, \dots, F_m, \mathcal{F}, \mathcal{S}}(T)$ . We then run the wild phase by querying the master on  $Q$  iid inputs  $x_1, \dots, x_Q$ .

$$\text{for } i = 1, \dots, Q : y_i \leftarrow M^{F_1, \dots, F_m, \mathcal{S}}(x_i).$$

The goal of the adversary is two-fold:

1. They do not want to be caught, if either  $\tau = \text{fail}$  or  $y_i = \text{abort}$  for some  $i \in [Q]$  we say the adversary was detected and define the predicate

$$\text{detect} = \text{false} \iff (\tau = \text{pass}) \wedge (\forall i \in [Q] : y_i \neq \text{abort})$$

2. They want the master to output as many wrong outputs as possible. We denote the number of wrong outputs by  $Y \stackrel{\text{def}}{=} |\{i : y_i \neq \mathcal{F}(x_i)\}|$ .

Informally, we call a compiler (like our simple schemes)  $(\text{win}, \text{wrng})$ -Trojan resilient, or simply  $(\text{win}, \text{wrng})$ -secure, if for every Trojan, the probability that it causes the master to output  $\geq \text{wrng}$  fraction of wrong outputs without being detected is at most  $\text{win}$ . In the formal definition  $\text{win}$  and  $\text{wrng}$  are allowed to be a function of the number of test queries  $T$ .

**Definition 1 ((win, wrng)-Trojan resilience).** *And adversary (win, wrng)-wins in  $\text{TrojanGame}(\Pi, T, Q)$  if the master outputs more than a wrng fraction of wrong values without the Trojans being detected with probability greater than win, i.e.,*

$$\Pr_{\text{TrojanGame}(\Pi, T, Q)} [(\text{detect} = \text{false}) \wedge (Y/Q \geq \text{wrng})] \geq \text{win}$$

For  $\text{win} : \mathbb{N} \rightarrow [0, 1]$ ,  $\text{wrng} : \mathbb{N} \rightarrow [0, 1]$ ,  $q : \mathbb{N} \rightarrow \mathbb{N}$ , we say that  $\Pi$  is  $(\text{win}(T), \text{wrng}(T), q(T))$ -Trojan-resilient (or simply “secure”) if there exists a constant  $T_0$ , such that for all  $T \geq T_0$  and  $Q \geq q(T)$  no adversary  $(\text{win}(T), \text{wrng}(T))$ -wins in  $\text{TrojanGame}(\Pi, T, Q)$ .

We say  $\Pi$  is  $(\text{win}(T), \text{wrng}(T))$  Trojan-resilient if it is  $(\text{win}(T), \text{wrng}(T), q(T))$ -Trojan-resilient for some (sufficiently large) polynomial  $q(T) \in \text{poly}(T)$ .

In all our simple constructions the test and master only use the outputs of the  $F_i$  (and for the test also  $F$ ) oracles to check for equivalence. This fact will allow us to consider somewhat restricted adversaries in the security proof.

**Definition 2 (Generic Simple Scheme).** *A generic simple scheme  $T^*, M^*$  treats the outputs of the  $F_i$  (and for  $T^*$  additionally  $F$ ) oracles like variables. Concretely, two or more oracles can be queried on the same input, and then one checks if the outputs are identical. Moreover the master can use the output of an  $F_i$  as its own output.*

By the following lemma, to prove security of generic simple schemes, it will be sufficient to consider restricted adversaries that always choose to attack the trivial functionality  $\mathcal{F}(x) = 0$  and where the output range of the Trojans is a bit.

**Lemma 1.** *For any generic simple scheme  $\Pi_m$ , assume an adversary  $\text{Adv}$  exists that (win, wrng)-wins in  $\text{TrojanGame}(\Pi_m, T, Q)$  and let  $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$ ,  $F_1, \dots, F_m : \mathcal{X} \rightarrow \mathcal{Y}$  denote its choices for the attack. Then there exists an adversary  $\text{Adv}'$  who also (win, wrng)-wins in  $\text{TrojanGame}(\Pi_m, T, Q)$  and chooses  $\mathcal{F}' : \mathcal{X} \rightarrow \{0, 1\}$ ,  $F'_1, \dots, F'_m : \mathcal{X} \rightarrow \{0, 1\}$  where moreover  $\forall x \in \mathcal{X} : \mathcal{F}'(x) = 0$ .*

*Proof.*  $\text{Adv}'$  firstly runs  $\text{Adv}$  to learn (i) the functionality  $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$  which it wants to attack and (ii) its Trojans  $F_1, \dots, F_m$ . It then outputs (as its choice of function to attack) an  $\mathcal{F}'$  where  $\forall x \in \mathcal{X} : \mathcal{F}'(x) = 0$  and, for every  $i \in [m]$ , it chooses the Trojan  $F'_i$  to output 0 if  $F_i$  would output the correct value, and 1 otherwise. More formally,  $F'_i(x)$  invokes the original Trojan  $y \leftarrow F_i(x)$  and outputs 0 if  $\mathcal{F}(x) = y$  and 1 otherwise.

By construction, whenever one of the  $F'_i$ 's deviates (i.e., outputs 1), also the original  $F_i$  would have deviated. And whenever the test or master detect an inconsistency in the new construction, they would also have detected an inconsistency with the original  $\mathcal{F}$  and  $F_i$ .<sup>5</sup> □

### 2.4 Lower Bounds

By definition, (win, wrng)-security implies (win', wrng')-security for any  $\text{win}' \geq \text{win}$ ,  $\text{wrng}' \geq \text{wrng}$ . The completeness property implies that no scheme is (1, 0)-secure (as by behaving honestly an adversary can (1, 0)-win). And also no scheme is (0, 1)-secure (as  $\Pr[E] \geq 0$  holds for every event  $E$ ). Thus our (win, wrng)-security notion is only interesting if both, win and wrng are  $> 0$ . We will prove the following lower bound:

**Lemma 2 (Lower bound for simple schemes).** *For any  $c > 0$  and  $m \in \mathbb{N}$  there exists a constant  $c' = c'(c, m) > 0$  such that no  $m$ -redundant simple scheme  $\Pi_m$  is  $(c, \frac{c'}{m})$ -Trojan-resilient.*

<sup>5</sup> Let us mention that the opposite is not true (it's possible that for some  $i \neq j$  we have  $F'_i(x) = F'_j(x) = 1$ , while  $F_i(x) \neq F_j(x)$ ). This just captures the observation that an adversary who wants to deviate as often as possible without being detected can wlog. always deviate to the same value.



*Proof.* Adv chooses the constant functionality  $\mathcal{F}(x) = 0$  with a sufficiently large input domain  $|\mathcal{X}| \gg (m \cdot T)^2$  (so that sampling  $m \cdot T$  elements at random from  $\mathcal{X}$  with or without repetition is basically the same). Now Adv samples a random subset  $\mathcal{X}' \subset \mathcal{X}$ ,  $|\mathcal{X}'|/|\mathcal{X}| = \frac{1.1 \cdot c'}{T}$  (for  $c'$  to be determined) and then defines Trojans which deviate on inputs from  $\mathcal{X}'$

$$\forall i \in [m] : F_i(x) = \begin{cases} 1 & \text{if } x \in \mathcal{X}' \text{ (deviate)} \\ 0 & \text{if } x \notin \mathcal{X}' \text{ (correct)} \end{cases}$$

Should the test pass, the master will deviate on each input with probability  $1.1 \cdot c'/T$ , if we set the number of queries  $Q$  large enough, the fraction of wrong outputs will be close to its expectation  $1.1 \cdot c'/T$ , and thus almost certainly larger than  $c'/T$ .

It remains to prove that the test passes with probability  $\geq c$ . By correctness, the testing procedure  $\mathsf{T}^{F_1, \dots, F_m, \mathcal{F}, \mathcal{S}}$  must output **pass** unless one of the total  $\leq m \cdot T$  queries it made to the  $F_i$ s falls into the random subset  $\mathcal{X}'$ . The probability that no such query is made is at least

$$\left(1 - \frac{1.1 \cdot c'}{T}\right)^{m \cdot T}$$

and this expression goes to 1 as  $c'$  goes to 0. We now choose  $c' > 0$  sufficiently small so the expression becomes  $> c$ . To get a quantitative bound one can use the well known inequality  $\lim_{T \rightarrow \infty} (1 - 1/T)^T = 1/e \approx 0.367879$ .  $\square$

The (proof of) the previous lemma also implies the following.

**Corollary 1.** *If a simple scheme  $\Pi_m$  is  $(\text{win}(T), \text{wrng}(T))$  secure with*

1.  $\text{win}(T) \in 1 - o(1)$  then  $\text{wrng}(T) \in o(1/T)$ .
2.  $\text{wrng}(T) \in \omega(1/T)$  then  $\text{win}(T) \in o(1)$ .

The first item means that if Adv wants to make sure the Trojan is only detected with sub-constant probability, then he can only force the master to output a  $o(1/T)$  fraction of wrong outputs during deployment. The second item means that if Adv wants to deviate on a asymptotically larger than  $1/T$  fraction of outputs, it will be detected with a probability going to 1.

*Not Interesting Security for 1-Redundant Schemes.* For  $m = 1$  redundant circuits a much stronger lower bound compared to Lemma 2 holds. The following Lemma implies that no 1-redundant scheme is  $(\epsilon(T), \delta(T))$ -Trojan-resilient for any  $\epsilon(T) > 0$  and  $\delta(T) = 1/\text{poly}(T)$  (say  $\epsilon(T) = 2^{-T}$ ,  $\delta(T) = T^{-100}$ ).

**Lemma 3 (Lower bound for  $m = 1$ ).** *For any 1-redundant scheme  $\Pi_1$  and any polynomial  $p(T) > 0$ , there is an adversary that  $(1, 1 - 1/p(T))$ -wins in the  $\text{TrojanGame}(\Pi_1, T, Q)$  game for  $Q \geq p(T) \cdot T$ .*

*Proof.* Consider an adversary who chooses a “time bomb” Trojan  $F_1$  which correctly outputs  $\mathcal{F}(x)$  for the first  $T$  queries and also stores those queries, so it can

output the correct value if one of those queries is repeated in the future. From query  $T + 1$  the Trojan outputs wrong values unless it is given one of the first  $T$  queries as input, in which case it outputs the correct value. This Trojan will pass any test making at most  $T$  invocations, while the master will deviate on almost all queries, i.e., all except the first  $T$ .

To see why we store the first  $T$  queries and do not deviate on them when they repeat in the future, consider a master which stores the outputs it observes on the first  $T$  queries so it can later detect inconsistencies.  $\square$

## 2.5 Efficiency of Lower Bound vs. Constructions

For the lower bounds in the previous section, the *only* restriction on the test  $\mathsf{T}^{\mathcal{F}_1, \dots, \mathcal{F}_m, \mathcal{F}, \mathcal{S}}(T)$  is that each  $\mathcal{F}_i$  can only be queried at most  $T$  times. There are no restrictions on the master  $\mathsf{M}^{\mathcal{F}_1, \dots, \mathcal{F}_m, \mathcal{S}}(\cdot)$  at all. In particular, it can be stateful, computationally unbounded, use an arbitrary amount of randomness, and query the  $\mathcal{F}_i$ s on an unbounded number of inputs (as the Trojan  $\mathcal{F}_i$ s can be stateful this is not the same as learning the function table of the  $\mathcal{F}_i$ 's).

While the lack of any restrictions makes the lower bound stronger, we want our upper bounds, i.e., the actual constructions, to be as efficient (in terms of computational, query and randomness complexity) and simple as possible, and they will indeed be very simple.

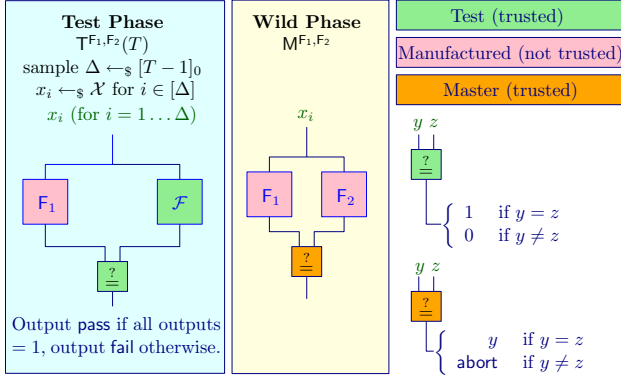
Let us stress that one thing the definition does not allow is the test to pass a message to the master. If we would allow a message of unbounded length to be passed this way no non-trivial lower bound would hold as  $\mathsf{T}$  could send the entire function table of  $\mathcal{F}$  to  $\mathsf{M}$ , which then could perfectly implement  $\mathcal{F}$ . Of course such a “construction” would get against the entire motivation for simple schemes where  $\mathsf{M}^*$  should be much simpler and independent of  $\mathcal{F}$ . Still, constructions where the test phase sends a short message to the master (say, a few correct input/output pairs of  $\mathcal{F}$  which the master could later use to “audit” the Trojans) could be an interesting relaxation to be considered.

## 2.6 Our Results and Conjectures

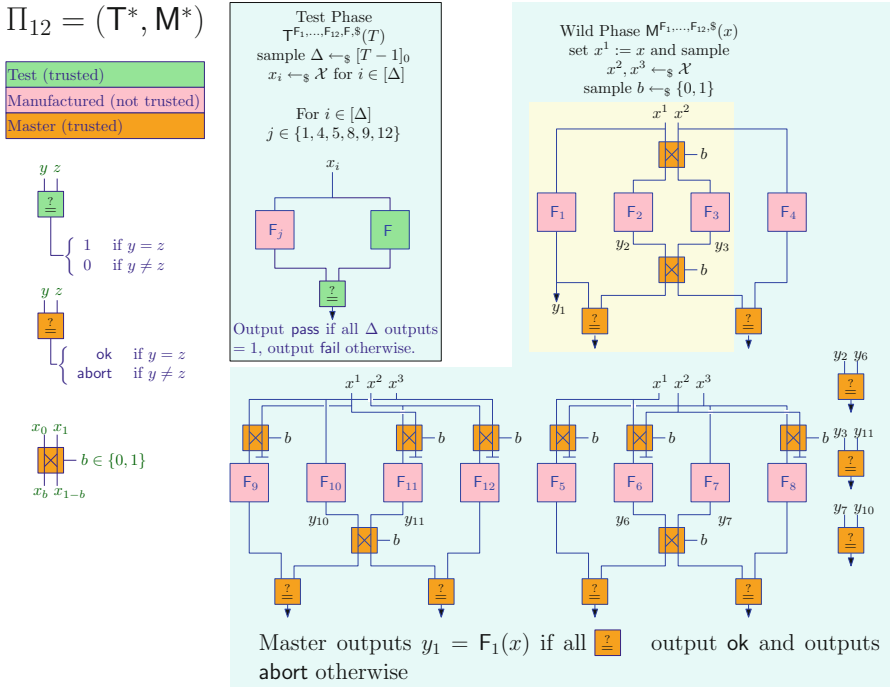
Our main technical result is a construction of a simple scheme which basically matches the lower bound from Lemma 2. Of course for any constant  $c > 0$ , the constant  $c'$  in the theorem below must be larger than in Lemma 2 so there's no contradiction.

**Theorem 1 (Main, optimal security of  $\Pi_{12}$ ).** *For any constant  $c > 0$  there is a constant  $c'$  such that the simple construction  $\Pi_{12}$  from Fig. 3 is  $(c, \frac{c'}{T})$ -Trojan resilient.*

While  $(c, \frac{c'}{T})$ -Trojan-resilience matches our lower bound, the construction is  $m = 12$ -redundant (recall this means we need 12 instantiations of  $\mathcal{F}$  manufactured to instantiate the scheme). While for  $m = 1$  redundancy is not sufficient to get any interesting security, as we showed in Lemma 3, we conjecture that  $m = 2$  is sufficient to match the lower bound, and give a candidate construction.



**Fig. 2.** Construction  $\Pi_2^*$  (discussed in Sect. 2.9), which is  $(c, \frac{c'}{T})$  secure for history-independent Trojans.



**Fig. 3.** Construction  $\Pi_{12}$  for which we prove optimal Trojan-resilience as stated in Theorem 1. Very informally, the security proof is by contradiction: via a sequence of hybrids an attack against  $\Pi_{12}$  is shown to imply an attack where the yellow part basically corresponds to  $\Pi_2^*$  with two history independent circuits. This attack contradicts the security of  $\Pi_2^*$  as stated in Theorem 2. (Color figure online)

*Conjecture 1 (Optimal security of  $\Pi_2^\phi$ ).* For any  $0 < \phi < 1$  and any constant  $c > 0$  there is a constant  $c' = c'(c, \phi)$  such that the simple construction  $\Pi_2^\phi$  from Sect. 3 is  $(c, \frac{c'}{T})$ -Trojan resilient.

The parameter  $\phi$  in this construction basically specifies that the master will query both oracles  $F_1$  and  $F_2$  on a (random)  $T^{-\phi}$  fraction of the input, and check consistency in this case. While the conjecture is wrong for  $\phi = 0$  and  $\phi \geq 1$ , the  $\phi = 0$  case (i.e., when we always query both,  $F_1$  and  $F_2$ ) will be of interest to us as security of the  $\Pi_2^* \stackrel{\text{def}}{=} \Pi_2^0$  construction against a limited adversary (termed history-independent and discussed in Sect. 2.9 below) will be a crucial step towards our proof of our main theorem.

## 2.7 Comparison with VC and MPC

Let us shortly compare the security we achieve with the more costly solutions based on verifiable computation (VC) [2, 17] and multiparty computation (MPC) [5] discussed in the introduction. We can consider (win, wrng)-security as in Definition 1 also for the VC and MPC construction, here one would need change the  $\text{TrojanGame}(\Pi, T, Q)$  from Sect. 2.3 to allow the trojans  $F_i$  to implement a different functionality than the target  $\mathcal{F}$  (for VC one needs to compute an extra succinct proof, for MPC the trojans implement the players in an MPC computation). For VC there's no test (so  $T = 0$ ) and only one  $m = 1$  Trojan, and for MPC and VC we can drop the requirement that the inputs are iid.

In the VC construction the master will catch *every* wrong output (except with negligible probability), so for any polynomial *poly* there is a negligible function *negl* (in the security parameter of the underlying succinct proof system), such that the scheme is  $(1/\text{poly}, \text{negl}, Q)$  secure for any polynomial  $Q$ .

For the MPC construction the master will provide  $Q < c_0 T$  outputs with probability  $c_1^m$  (where  $c_0 \in [0, 1/2]$  and  $c_1 \in [0, 1]$  are some constants), but while outputs are provided they are most likely correct, so for any polynomial  $Q, T$  we have  $(1 - c_1^m, \text{negl}, Q)$  security.

## 2.8 Stateless Trojans

In our security definition we put no restriction on the Trojans  $F_i$  provided by the adversary (other than being *digital* hardware Trojans as discussed in the introduction), in particular, the  $F_i$ 's can have arbitrary complex evolving state while honestly manufactured circuits could be stateless. We can consider a variant of our security definition (Definition 1) where the adversary is only allowed to choose *stateless* Trojan circuits  $F_i$ . Note that the lower bound from Lemma 2 still holds as in its proof we did only consider stateless  $F_i$ 's. There's an extremely simple 1-redundant construction that matches the lower bound when the adversary is only allowed to chose stateless Trojans.

Consider a construction  $\Pi_1 = (\mathsf{T}^*, \mathsf{M}^*)$  where the master is the simplest imaginable: it just forwards inputs/outputs to/from its oracle, if  $F_1$  is stateless

this simply means  $M^{F_1}(\cdot) = F_1(\cdot)$ . The test  $T^{F_1, F, \$}(T)$  queries  $F_1$  and the trusted  $F$  on  $T$  random inputs and outputs fail iff there is a mismatch.

**Proposition 1 (Optimal security for 1-redundant scheme for *stateless Trojans*).** *For any constant  $c > 0$  there is a constant  $c' > 0$  such that  $\Pi_1$  is  $(c, \frac{c'}{T})$ -Trojan resilient if the adversary is additionally restricted to choose a stateless Trojan.*

*Proof.* If  $\text{wrng}'$  denotes the fraction of inputs on which the Trojan  $F_1$  differs from the specification  $\mathcal{F}$  (both chosen by an adversary  $\text{Adv}$ , note that  $\text{wrng}'$  is only well defined here as  $F_1$  is stateless), then  $\text{wrng}'$  must satisfy  $c > (1 - \text{wrng}')^T$  if the adversary wants to  $(c, \text{wrng})$ -win for any  $\text{wrng}$ , as otherwise already the test catches the Trojan with probability  $(1 - \text{wrng}')^T > c$ . For  $c > (1 - \text{wrng}')^T$  to hold  $\text{wrng}' \in \Omega(1/T)$ , in particular,  $\text{wrng}' \geq c'/T$  for some  $c' > 0$ .  $\square$

## 2.9 History-Independent Trojans

A notion of in-between general (stateful) Trojans and stateless Trojans will play a central role in our security proof. We say a trojan  $F_i$  is *history-independent* if its only state is a counter which is incremented by one on every invocation, so it's answer to the  $i$ 'th query can depend on the current index  $i$ , but not on any inputs it saw in the past.

We observe that Lemma 3 stating that no 1-redundant simple scheme can be secure still holds if we restrict the choice of the adversary to history-independent Trojans as the “time-bomb” Trojan used in the proof is history-independent. We will show a 2-redundant construction  $\Pi_2^*$  that achieves optimal security against history-independent Trojans.

**Theorem 2 (History-Independent Security of  $\Pi_2^*$ ).** *For any constant  $c > 0$  there is a constant  $c' = c'(c) > 0$  such that  $\Pi_2^*$  from Fig. 2 is  $(c, \frac{c'}{T})$ -Trojan resilient if the adversary is additionally restricted to choose a history-independent Trojans.*

The technical Lemma 4 we prove and which implies this theorem, actually implies a stronger statement: for any positive integer  $k$ , the above holds even if we relax the security notion and declare the adversary a winner as long a Trojan is detected by the test or master at most  $k - 1$  times. What this exactly means is explained in Sect. 4.2. Note that this notion coincides with the standard notion for  $k = 1$ .

The  $\Pi_2^*$  scheme is just the  $\Pi_2^\phi$  scheme from Conjecture 1 for  $\phi = 0$ , where we conjecture that  $\Pi_2^\phi$  is (in some sense) optimally secure for  $0 < \phi < 1$ . For  $\phi$  the conjecture is wrong, but somewhat ironically we are only able to *prove* security against history-independent Trojans for  $\phi = 0$ , and this result will be key towards proving the security of  $\Pi_{12}$  as stated in our main Theorem 1.

### 2.10 Proof Outline

The proof of our main Theorem 1 stating that  $\Pi_{12}$  is optimally Trojan-resilient is done in two steps. As just discussed, we first prove security of  $\Pi_2^*$  against *history-independent* Trojans, and in a second step we reduce the security of  $\Pi_{12}$  against general Trojans to the security of  $\Pi_2^*$  against history-independent Trojans. We outline the main ideas of the two parts below.

*Part 1: Security of  $\Pi_2^*$  Against History-Independent Trojans (Theorem 2, Lemma 4).*  $\Pi_2^*$  is a very simple scheme where the test  $T^{F_1, F_2, F, \mathcal{S}}$  just queries  $F_1$  on a random number  $\Delta, 0 \leq \Delta < T$  of inputs and checks if they are correct (the test ignores  $F_2$ ). The master  $M^{F_1, F_2, \mathcal{S}}(x)$  queries  $y \leftarrow F_1(x)$  and  $y' \leftarrow F_2(x)$  on  $x$  and aborts if they disagree.

In the proof of Lemma 4 we define  $p_i$  and  $q_i$  as the probability that  $F_1$  and  $F_2$  outputs a wrong value in the  $i$ th query on a random input, respectively. As  $F_1, F_2$  are history independent, this is well defined as this probability only depends on  $i$  (but not previous queries).

Let the variable  $\Phi_\Delta$  denote the number of times the Trojans will be detected conditioned on the random number of test queries being  $\Delta$ . This value is (below  $Q$  is the number of queries to the master and we use the convention  $q_i = 0$  for  $i < 0$ )

$$\mathbb{E}[\Phi_\Delta] = \sum_{i=1}^{Q+\Delta} |p_i - q_{i-\Delta}| \tag{1}$$

In this sum, the first  $\Delta$  terms account for the test, and the last  $Q$  terms for the wild-phase. Moreover let  $Y_\Delta$  denote the number of times  $F_1$  deviates (and thus the master outputs a wrong value), its expectation is

$$\mathbb{E}[Y_\Delta] = \sum_{i=\Delta+1}^{Q+\Delta} p_i$$

To prove Trojan-resilience of  $\Pi_2^*$  as stated in Lemma 4 boils down to proving that, for most  $\Delta$ , whenever the probability of  $\Phi_\Delta = 0$  (i.e., the Trojan is not detected) is constant, the fraction of wrong outputs  $Y_\Delta/Q$  must be “small” (concretely,  $O(1/T)$ ).

The core technical result establishing this fact is Lemma 5. Unfortunately, this Lemma only establishes this fact for the expectation, i.e., whenever  $\mathbb{E}[\Phi_\Delta]$  is small, also  $\mathbb{E}[Y_\Delta]$  is small. Here is where we use the fact that the  $F_1, F_2$  are history independent: in the history independent case  $\Phi_\Delta$  and  $Y_\Delta$  can be written as the sum of independent boolean variables, so using a Chernoff bound it follows that their actual value will be close to their expectation with high probability.

It is instructive to see why for example setting  $p_i = q_i = \delta$  for some fixed  $\delta > 0$  does not contradict Theorem 2. To contradict it, the fraction of wrong outputs (which here is simply  $\delta$ ) must be  $\omega(1/T)$ . In this case,  $\mathbb{E}[\Phi_\Delta] = \Delta \cdot \delta = \omega(\Delta/T)$ , which to contradict the lemma must be at least constant, which in turn means  $\Delta \in o(T)$  must hold. As  $\Delta, 0 \leq \Delta < T$  is uniform, t’s  $o(T)$  with  $o(1)$  probability, but for a contradiction we also need this probability to be constant.

*Part 2: Reducing the Security of  $\Pi_{12}$  Against General Trojans to the Security of  $\Pi_2^*$  Against History Independent Trojans (Theorem 1).* While the random shift  $\Delta$  makes  $\Pi_2^*$  secure against history-independent attacks (like time-bombs, where a Trojan starts deviating after some fixed number of queries), it succumbs to cheat codes: as the master always queries  $F_1, F_2$  on the same inputs, a Trojan can specify some set of trigger inputs, and after receiving such a trigger the Trojans will deviate forever. By making the fraction of inputs that are triggers sparse, the Trojans will likely not be detected during testing (a  $1/T$  fraction will survive testing with constant probability).

To prevent such a coordination via the inputs, in  $\Pi_{12}$  inputs are somewhat randomly assigned to the different Trojans. In particular, as emphasized in the yellow area in Fig. 3, the  $F_1$  is always queried on the input  $x$ , and then the random bit  $b$  determines whether  $F_2$  or  $F_3$  are queried on  $x$ . If an input  $x$  were to trigger the Trojans to always deviate,  $F_1$  and one of  $F_2$  and  $F_3$  will be triggered, say it's  $F_2$ . But now, as soon as  $F_3$  is queried in a future round the master will abort as  $F_1$  will deviate, but  $F_3$  will not (except if this query also happens to be a trigger, which is unlikely as triggers must be sparse to survive the testing phase).

This just shows why a particular attack does not work on  $\Pi_{12}$ . But we want a proof showing security against all possible Trojans. Our proof proceeds by a sequence of hybrids, where we start with assuming a successful attack on  $\Pi_{12}$ , and then, by carefully switching some circuits and redefining them by hard coding “fake histories”, we arrive at a hybrid game where there is still a successful attack, but now the circuits in the yellow area basically correspond to two the  $\Pi_2^*$  construction instantiated with history-independent Trojans, but such an attack contradicts our security proof for  $\Pi_2^*$  as stated in Lemma 4.

### 3 Conjectured Security of 2-Redundant Schemes

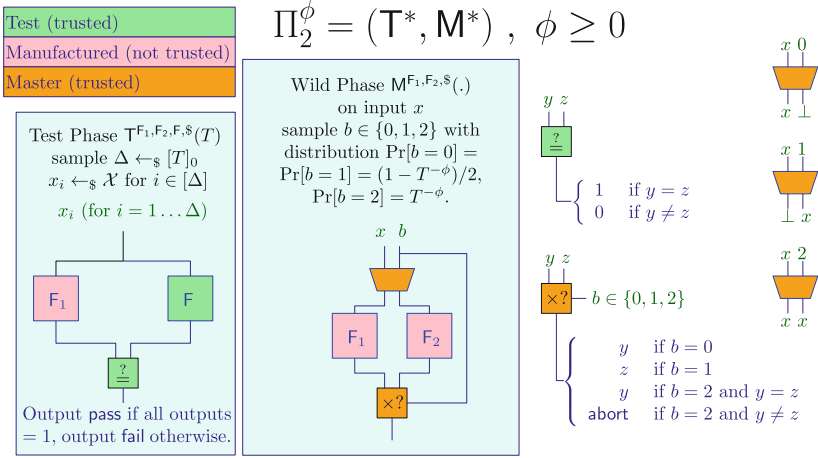
While the main technical result in this paper is a simple scheme  $\Pi_{12}$  that provably achieves optimal security as stated in Theorem 1, this construction is not really practical as it is 12-redundant. Recall that  $k$ -redundant means the master needs  $k$  instantiations of the functionality  $\mathcal{F}$ , so it's in some sense the hardware cost. For this section let us also define a computational cost: the **rate** of a simple construction is the average number of invocations to its  $F_i$  oracles the master  $M^{F_1, \dots, F_m, \mathcal{S}}(\cdot)$  makes with any query.

#### 3.1 A 2-Redundant Scheme $\Pi_2^\phi$

We will now define a scheme  $\Pi_2^\phi$  which in terms of redundancy and rate is as efficient as we possibly could hope for a scheme with non-trivial security: it's 2-redundant and has a rate of slightly above (the trivial lower bound of) 1. The construction  $\Pi_2^\phi = (M^*, T^*)$ , where  $\phi \in \mathbb{R}, \phi \geq 0$  is illustrated in Fig. 4.

**test:** In the test phase,  $T^{F_1, F_2, F}(T)$  picks a random  $\Delta, 0 \leq \Delta \leq T - 1$ , then queries  $F_1$  on  $\Delta$  random inputs and checks if the outputs are correct by comparing with the trusted  $F$ .

**master:** With probability  $1 - T^{-\phi}$  the master  $M^{F_1, F_2, \mathcal{S}}(x)$  picks either  $F_1$  or  $F_2$  at random, queries it on  $x$  and uses the output as its output. Otherwise, with probability  $T^{-\phi}$ , the master queries both oracles and outputs **abort** if their outputs don't match, and forwards the output of  $F_1$  otherwise.



**Fig. 4.** Construction  $\Pi_2^\phi$  from Conjecture 1.

Our Conjecture 1 states that this construction achieves optimal security (optimal in the sense of matching the lower bound from Lemma 2) for any  $0 < \phi < 1$ , i.e.,

For any  $0 < \phi < 1$  and any constant  $c > 0$  there is a constant  $c'$  such that  $\Pi_2^\phi$  is  $(c, \frac{c'}{T})$ -Trojan resilient.

We discuss how  $\Pi_2^\phi$  performs against typical Trojans like time-bombs and cheat codes. Our conjecture only talks about  $(\text{win}(T), \text{wng}(T))$ -security where the winning probability  $\text{win}(T) = c$  is a constant, and here the exact value of  $\phi$  does not seem to matter much as long as it is bounded away from 0 and 1. For  $\text{win}(T) = o(1)$  the parameter  $\phi$  will matter as those attacks will illustrate. (the  $o(1)$  always denotes any value that goes to 0 as  $T \rightarrow \infty$ ).

**Proposition 2 (Time Bomb against  $\Pi_2^\phi$ ).** *For any  $\phi$ , there exists an adversary that  $(\Theta(T^{-\phi}), 1 - o(1))$ -wins in  $\text{TrojanGame}(\Theta(\Pi_2^\phi), T, \omega(T))$*

*Proof (sketch).* Let Adv choose the constant functionality  $\forall x \in \mathcal{X} : \mathcal{F}(x) = 0$ , and a Trojan  $F_1$  which outputs the correct value 0 for the first  $T$  queries, and 1 for all queries  $> T$ , while  $F_2$  always outputs 1.

$F_1$  will always pass the test. The master will **abort** iff one of its first  $T - \Delta$  queries to  $F_1$  (where the output is 0) is a “ $b = 2$ ” query (as then  $F_1(x) = 1 \neq 0 = F_2(x)$ ). With probability  $T^{-\phi}$  we have  $\Delta \geq T - T^\phi$ , and in this case such a bad



event only happens with constant probability (using  $(1-\epsilon)^{1/\epsilon} \approx 1/e = 0.368\dots$ ). So the Trojan will not be detected with probability  $T^{-\phi}/e$ , and in this case also almost all outputs will be wrong.  $\square$

**Proposition 3 (Cheat Code against  $\Pi_2^\phi$ ).** *For any  $\phi$ , there exists an adversary that  $(\Theta(T^{\phi-1}), 1 - o(1))$ -wins in  $\text{TrojanGame}(\Theta(\Pi_2^\phi), T, \omega(T))$*

*Proof (sketch).* Let Adv choose the constant functionality  $\forall x \in \mathcal{X} : \mathcal{F}(x) = 0$ . The Trojans  $F_1, F_2$  output 0 until they get a query from a “trigger set”  $\mathcal{X}' \subset \mathcal{X}$ , after this query they always deviate and output 1.

If we set  $|\mathcal{X}'|/|\mathcal{X}| = 1/T$ , then the test will pass with constant probability  $(1 - 1/T)^\Delta \geq (1 - 1/T)^T \approx 1/e$ . Assuming the Trojans passed the test phase, the master will *not* catch the Trojans if the first trigger query to  $F_1$  and  $F_2$  happen in-between the same  $b = 2$  queries (or in such a query). This happens with probability  $\approx T^{\phi-1}$ .  $\square$

The two propositions above imply that the adversary can always  $(T^{\max\{-\phi, \phi-1\}}, 1 - o(1))$ -win by either using a time-bomb or cheat-code depending on  $\phi$ . The winning probability is minimized if  $-\phi = \phi - 1$  which holds for  $\phi = 0.5$ . We conjecture that the above two attacks are basically all one can do to attack  $\Pi_2^\phi$ .

*Conjecture 2 (Security of  $\Pi_2^{0.5}$  for low winning probabilities).* For  $\text{win}(T) \in \omega(T^{-0.5})$ ,  $\Pi_2^{0.5}$  is  $(\text{win}(T), \text{wrng}(T))$ -Trojan resilient for some  $\text{wrng}(T) \in o(1)$ .

## 4 A Scheme for History-Independent Trojans

In this section we define the simple scheme  $\Pi_2^\star$  and prove its security as claimed in Theorem 2. Recall that a history-independent Trojan circuit is a stateless circuit, except that it maintains a counter. We recall that a trojan is *history-independent* if its state is a counter which is incremented by one on every invocation, so its answer to the  $i$ 'th query can depend on the current index  $i$  and current input  $x_i$ , but not on any inputs it saw in the past.

### 4.1 Notation

For an integer  $n$  we define  $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$  and  $[n]_0 \stackrel{\text{def}}{=} \{0, \dots, n\}$ . We will also use the Chernoff bound.

### 4.2 Security of $\Pi_2^\star$

**Relaxing the Winning Condition.** We can think of the security experiment  $\text{TrojanGame}(\Pi_2^\star, T, Q)$  as proceeding in rounds. First, for a random  $\Delta \in [T]$ , we run the test for  $\Delta$  rounds (in each querying  $F_1$  and  $F$  on a random input and checking equivalence), and then  $Q$  rounds for querying the master (in each round querying  $F_1$  and  $F_2$  and checking for equivalence). The adversary immediately

loses the game if a comparison fails (outputs 0 in the test or **abort** in the master) in any round.

We consider a relaxed notion of (win, wrng)-winning, “relaxed” as we make it easier for the adversary, and thus proving security against this adversary gives a stronger statement. We define (win, wrng)- $k$ -winning like (win, wrng)-winning, but the adversary is allowed to be detected in up to  $k - 1$  rounds, so (win, wrng)-1-winning is (win, wrng)-winning.

This relaxed notion is not of practical interest, as one would immediately abort the moment a Trojan is detected. We consider this notion as we need it for the security proof of our main Theorem 1, where we will only be able to reduce security of  $\Pi_{12}$  to the security of  $\Pi_2^*$  (against history-independent Trojans) if  $\Pi_2^*$  satisfies this stronger notion.

**Proof of Theorem 2.** The following lemma implies Theorem 2 for  $k = 1$ , as discussed after the statement of the theorem the lemma below is somewhat more general as we’ll need the stronger security for any  $k$ .

**Lemma 4.** *For any constant  $c > 0$  and positive integer  $k$ , there exists a constant  $c'$ , and integer  $T_0$  and polynomial  $q(\cdot)$  such that no adversary  $\text{Adv}$  exists that only chooses history-independent Trojans and that for any*

$$T \geq T_0 \quad , \quad Q \geq q(T)$$

*can  $(c, c'/T)$ - $k$ -win  $\text{TrojanGame}(\Pi_2^*, T, Q)$ .*

*Proof.* For a given  $c > 0$  define

$$c'' = \max\{64k, -256 \ln(c/2)\}$$

we then set  $c', q(T)$  and  $T_0$  as

$$c' = c''/c^2 \quad , \quad q(T) = \frac{5 \cdot T^2 c}{c''} + 5T \quad , \quad T_0 = 1 \tag{2}$$

These values are just chosen so that later our inequalities work out nicely, we did not try to optimise them.

By Lemma 1 we can consider the security experiment where an adversary  $\text{Adv}$  chooses the constant functionality  $\mathcal{F} : \mathcal{X} \rightarrow 0$  as target and the two (history-independent) Trojans  $F_1, F_2 : \mathcal{X} \rightarrow \{0, 1\}$  output a bit (so they can either correctly output 0 or deviate by outputting 1). As the  $F_1, F_2$  are history independent, we can think of  $F_1$  as a sequence  $F_1^1, F_1^2, \dots$  of functions where  $F_1^i$  behaves like  $F_1$  on the  $i$ th query. Let  $p_i$  and  $q_i$  denote the probability that  $F_1$  and  $F_2$  deviates on the  $i$ th query, respectively

$$p_i \stackrel{\text{def}}{=} \Pr_{x \leftarrow \mathcal{X}} [F_1^i(x) = 1] \quad , \quad q_i \stackrel{\text{def}}{=} \Pr_{x \leftarrow \mathcal{X}} [F_2^i(x) = 1]$$

In  $\text{TrojanGame}(\Pi_2^*, T, Q)$ , for  $\delta \in [T-1]_0$  let the variable  $Y_\delta$  denote the number of wrong outputs by  $F_1$  conditioned on the number of test queries  $\Delta \leftarrow_{\S} [T-1]_0$  being  $\Delta = \delta$ . The expectation is

$$\mathbb{E}(Y_\delta) = \sum_{i=1}^Q p_{i+\delta} \quad (3)$$

Let the variables  $\Phi_\delta^T$  and  $\Phi_\delta^M$  denote the number of times the test and the master “catch” a Trojan conditioned on  $\Delta = \delta$

$$\mathbb{E}[\Phi_\delta^T] = \sum_{i=1}^{\delta} |p_i| \quad , \quad \mathbb{E}[\Phi_\delta^M] \geq \sum_{i=1}^Q |p_{i+\delta} - q_i|$$

let  $\Phi_\delta \stackrel{\text{def}}{=} \Phi_\delta^T + \Phi_\delta^M$  denote the total number of times the Trojans are detected, and  $\Phi'_\delta$  being the same but we ignore the last  $\delta$  queries. With the convention that  $q_i = 0$  for  $i < 1$

$$\mathbb{E}[\Phi_\delta] \geq \sum_{i=1}^{Q+\delta} |p_i - q_{i-\delta}| \quad , \quad \mathbb{E}[\Phi'_\delta] \geq \sum_{i=1}^Q |p_i - q_{i-\delta}| \quad , \quad \mathbb{E}[\Phi_\delta] \geq \mathbb{E}[\Phi'_\delta] \quad (4)$$

As we consider *history-independent* Trojans the  $Y_\delta, \Phi_\delta$  variables are sums of independent Bernoulli random variables. Using a Chernoff bound we will later be able to use the fact that for such variables are close to their expectation with high probability.

*Claim.* For any  $\delta \in [T-1]_0, \tau \in [T-1-\delta]$  (so  $\delta + \tau \leq T-1$ )

$$\mathbb{E}[\Phi_\delta] + \mathbb{E}[\Phi_{\delta+\tau}] \geq \tau \cdot \frac{\mathbb{E}[Y_\delta] - T}{Q + T} \quad (5)$$

*Proof (of Claim).* Assume for a moment that  $p_1, \dots, p_\tau = 0$  as required to apply Lemma 16, then

$$\mathbb{E}[\Phi_\delta] + \mathbb{E}[\Phi_{\delta+\tau}] \stackrel{(4)}{\geq} \mathbb{E}[\Phi'_\delta] + \mathbb{E}[\Phi'_{\delta+\tau}] \quad (6)$$

$$\geq \sum_{\Delta \in \{\delta, \delta+\tau\}} \sum_{i=1}^Q |p_i - q_{i-\Delta}| \quad (7)$$

$$\stackrel{\text{Lemma 5}}{\geq} \tau \cdot \frac{\sum_{i=1}^Q p_i}{Q} = \tau \cdot \frac{\mathbb{E}[Y_0]}{Q} \quad (8)$$

$$\geq \tau \cdot \frac{\mathbb{E}[Y_\delta] - T}{Q} \quad (9)$$

The last step used  $\mathbb{E}[Y_0] + \delta \geq \mathbb{E}[Y_\delta]$  and  $\delta \leq T$ .

We now justify our assumption  $p_1, \dots, p_\tau = 0$ . For this change the security experiment and replace the Trojans  $F_1, F_2$  chosen by the adversary with Trojans

that first behave correctly for the first  $T$  inputs, and only then start behaving like  $F_1, F_2$  (technically, the new Trojans deviate with probabilities  $p'_i, q'_i$  satisfying  $p'_1, \dots, p'_T = 0, q'_1, \dots, q'_T = 0$  and for  $i > T : : p'_i = p_{i-T}$  and  $q'_i = q_{i-T}$ ). At the same time, we increase  $Q$  to  $Q + T$ . This change leaves  $\mathbb{E}[Y_\delta]$  unaffected, while  $\mathbb{E}[\Phi'_\delta], \mathbb{E}[\Phi'_{\delta+\tau}]$  can only increase. This proves the claim, note that in (5) the denominator is  $Q + T$  not  $Q$  as in (9) to account for this shift.  $\triangle$

*Claim.* For all but at most a  $c/2$  fraction of the  $\delta \in [T]_0$

$$\mathbb{E}[\Phi_\delta] \geq \frac{c \cdot T}{8} \cdot \frac{\mathbb{E}[Y_\delta] - T}{Q + T} \tag{10}$$

*Proof (of Claim).* We use Eq. (5) which can be understood as stating that if  $\mathbb{E}[\Phi_\delta]$  is “small” for some  $\delta$ , then all  $\mathbb{E}[\Phi_{\delta'}]$  with  $|\delta - \delta'|$  large enough can’t be small too. Concretely, consider any  $\delta$  for which (if no such  $\delta$  exists the claim already follows)

$$\mathbb{E}[\Phi_\delta] < \frac{c \cdot T}{8} \cdot \frac{\mathbb{E}[Y_\delta] - T}{Q + T}$$

then for all  $\delta' \in [T]_0$  for with  $|\delta - \delta'| \geq \frac{c \cdot T}{4}$  (note this is at least a  $c/2$  fraction) by Eq. (5)

$$\mathbb{E}[\Phi_\delta] + \mathbb{E}[\Phi_{\delta'}] \geq \frac{c \cdot T}{4} \cdot \frac{\mathbb{E}[Y_\delta] - T}{Q + T}$$

the two equations above now give

$$\mathbb{E}[\Phi_{\delta'}] \geq \frac{c \cdot T}{8} \cdot \frac{\mathbb{E}[Y_\delta] - T}{Q + T}$$

as claimed.  $\triangle$

To prove the lemma we need to show that when  $Q$  is sufficiently large, any adversary attacking at least  $c'/T$  fraction of times, can win at most  $k$  times with probability less than  $c$ . Since the duration of test phase  $\delta$  is chosen randomly from the set  $\{0, \dots, T - 1\}$ , we start with the following equation:

$$\frac{1}{T} \sum_{\delta=0}^{T-1} \Pr[(Y_\delta/Q \geq c'/T) \wedge (\Phi_\delta < k)] < c \tag{11}$$

Let  $c_\delta$  denote the probability the adversary  $k$ -wins conditioned on  $\Delta = \delta$

$$c_\delta \stackrel{\text{def}}{=} \Pr[(Y_\delta/Q \geq c'/T) \wedge (\Phi_\delta < k)] \tag{12}$$

With this notation we need to show

$$\frac{1}{T} \sum_{\delta=0}^{T-1} c_\delta < c$$

which follow from the claim below

*Claim.*  $c_\delta < c/2$  holds for all  $\delta$ , except (the at most  $c/2$  fraction of) the  $\delta \in [T]_0$  for which (10) does not hold

*Proof (of Claim).* Consider any  $\delta$  for which (10) holds. If for this  $\delta$   $\Pr[Y_\delta \geq Q \cdot c'/T] < c/2$  we're done as by (12) also  $c_\delta < c/2$  (using that  $\Pr[a \wedge b] \leq \Pr[a]$  for all events  $a, b$ ). To finish the proof of the claim we need to show that otherwise, i.e., if

$$\Pr[Y_\delta \geq Q \cdot c'/T] \geq c/2 \quad (13)$$

then

$$\Pr[\Phi_\delta < k] < c/2 \quad (14)$$

as this again implies  $c_\delta < c/2$ . Equation (13) (using  $\Pr[V \geq x] \geq p \Rightarrow \mathbb{E}[V] \geq x \cdot p$  which follows from Markov's inequality) gives

$$\mathbb{E}[Y_\delta] \geq Q \cdot c' \cdot c/2T \quad (15)$$

Plugging this into (10), then using our choice (2) of  $c' = c''/c^2$  and in the last step of  $Q \geq q(T) = 5 \cdot T^2 \cdot c/c'' + 5T$  (this bound for  $q(T)$  was just chosen so the last inequality below works out nice).

$$\begin{aligned} \mathbb{E}[\Phi_\delta] &\stackrel{(10)}{\geq} \frac{c \cdot T}{8} \cdot \frac{\mathbb{E}[Y_\delta] - T}{Q + T} \\ &\stackrel{(15)}{\geq} \frac{c \cdot T}{8} \cdot \frac{Q \cdot c' \cdot c - T}{Q + T} \\ &\stackrel{(2)}{\geq} \frac{Q \cdot c'' - 2T^2 \cdot c}{16(Q + T)} \\ &\geq c''/32 \end{aligned}$$

Using the Chernoff bound with  $\epsilon = 1/2$  and  $c'' \geq -256 \ln(c/2)$  (refer to Appendix in the extended technical report for details).

$$\Pr[\Phi_\delta < c''/64] \leq \Pr[\Phi_\delta < \mathbb{E}[\Phi_\delta]/2] \leq e^{-\mathbb{E}[\Phi_\delta]/8} \leq e^{-c''/256} \leq c/2$$

With our choice (2) of  $c'' = \max\{64k, -256 \ln(c/2)\}$  we get the bound  $\Pr[\Phi_\delta < k] \leq c/2$  claimed in (14).  $\triangle$

□

### 4.3 A Technical Lemma

Consider any  $t, z \in \mathbb{N}$ ,  $z > t$ ,  $t = 0 \pmod{2}$  and  $p_1, \dots, p_z \in [0, 1]$ . Denote with  $\bar{p} \stackrel{\text{def}}{=} \frac{\sum_{i=1}^z p_i}{z}$  be the average value.

**Lemma 5.** *For any  $q_1, \dots, q_z \in [0, 1]$ , (defining  $q_i = 0$  for  $i \leq 0$ ) and integers  $\Delta', \tau$  where  $0 \geq \Delta', \tau \geq 0$ , if  $p_1 = p_2 = \dots = p_\tau = 0$  then*

$$\sum_{\Delta \in \{\Delta', \Delta' + \tau\}} \sum_{i=1}^z |p_i - q_{i-\Delta}| \geq \tau \cdot \bar{p} \quad (16)$$

We refer the reader to our technical report for the full proof, but let us observe that for example it implies, that if  $p_1 = p_2 = \dots = p_t = 0$ , then

$$\frac{1}{t} \sum_{\Delta=0}^{t-1} \sum_{i=1}^z |p_i - q_{i-\Delta}| = \frac{1}{t} \sum_{\Delta'=0}^{t/2-1} \underbrace{\sum_{\Delta \in \{\Delta', \Delta'+t/2\}} \sum_{i=1}^z |p_i - q_{i-\Delta}|}_{\geq t \cdot \bar{p}/2 \text{ by (16)}} \geq \frac{t}{4} \cdot \bar{p} \quad (17)$$

Looking ahead, the lhs. of Eq. (17) will denote the expected number of times the master circuit detects an inconsistency in the experiment, while  $\bar{p}$  denotes the fraction of outputs where  $F_1$  diverts. So if the fraction of wrong outputs is larger than  $4/t$ , the master circuit will on average raise an alert once. To get a bound on the security the *expected* number of alerts is not relevant, only in the probability that it's larger than one, as this means that a Trojan was detected. The more fine grained statement Eq. (16) will be more useful to argue this.




Fig. 5. Illustration of the main variables used in the proof of Lemma 5.

## 5 A 12-Redundant Scheme $\Pi_{12}$


In this section we define a scheme  $\Pi_{12}$  and we will show that the lower bound for achievable security for very simple schemes (shown in Lemma 2) is asymptotically tight. Our proof is constructive - the analysis of our  $\Pi_{12}$  construction shows that it is  $(c, \frac{c}{T})$ -Trojan resilient for suitable constants.

Our  $\Pi_{12}$  scheme operates with three independent input streams and one independent bit stream. On each query, every circuit in  $\Pi_{12}$  receives one of the three inputs and produces an output. The master circuit then checks the consistency of the outputs, i.e. verifies if there is no mismatch between any pair of circuits receiving the same input.

As stated in Sect. 1, digital Trojans mainly employ two types of strategies: time bombs (where time is measured in the number of usages) and cheat codes (as a part of the input). To counter these strategies,  $\Pi_{12}$  desynchronizes the

circuits in two ways. First, some of the circuits are tested in the test phase for a randomly chosen time (already employed in the  $\Pi_2^*$  scheme). This effectively makes it difficult for time bomb Trojans to coordinate the time in which they start deviating. In  $\Pi_{12}$ , half of the circuits are tested for  $T$  times where  $T$  is a random variable with uniform distribution on  $[t]$ .

The second method of desynchronization involves using the value of the aforementioned input bit to *alternate* the way inputs are distributed among the circuits. Consequently, cheat code Trojans are rendered ineffective as only a subset of the circuits share the same input. Moreover, at any given point in time a circuit never “knows” which alternating state it is in (i.e. it does not know whether its output would be compared with deviating circuits or not).



**Fig. 6.** In a given group of circuits, depending on the value of  $b$ , the leftmost and rightmost circuits (outer circuits) are paired with the circuits in the middle (inner circuits). Circuits in a pairing are given the same input, and their outputs are checked for equivalence.

The main building block of the  $\Pi_{12}$ -scheme is a group of four circuits: two outer ones and two inner ones (see Fig. 6). On each query, every group of circuits receives two inputs - the first is given to the outer circuit on the left and the second to the outer circuit on the right. Additionally, in every step a fresh decision/alternation bit  $b$  is sampled. According to its value these two inputs are given to the inner circuits.  $\Pi_{12}$  consists of three such groups. Crosschecks are performed whenever two distinct circuits receive the same input (both within a group and among groups).

The proof that the construction  $\Pi_{12}$  is actually Trojan-resilient starts with assuming that it is not secure, goes via a hybrid argument and leads to a contradiction with security of  $\Pi_2^*$  construction. In every hybrid we change the construction slightly by swapping some pairs of circuits, arguing that the advantage of the adversary does not change much between each successive hybrids. In the final hybrid we show that the modified construction contains  $\Pi_2^*$  as a sub-construction. It turns out, that any adversary who wins with reasonable good probability in the final hybrid can be used to build an adversary who breaks the security of  $\Pi_2^*$  which is a contradiction with Theorem 2.

### 5.1 The $\Pi_{12}$ Scheme

We will now define our  $\Pi_{12}$  construction. It is illustrated in Fig. 3. We view our 12-circuit construction as three groups of four circuits each. Group 1 consists of circuits  $F_1, \dots, F_4$ , group 2 consists of  $F_5, \dots, F_8$ , and group 3 consists of  $F_9, \dots, F_{12}$ . At the beginning the three independent and identically distributed sequences of inputs are sampled. Moreover, independent sequence of bits is sampled (it is used to alternate the inputs' distribution in the wild). For every query in the wild, the construction performs two steps: (i) the *querying step*, where the inputs are distributed to all the 12 circuits depending on the value of the corresponding bit (ii) the *cross-checking step*, where the master circuit checks the consistency of the outputs of the circuits who receive the same inputs.

Now we can take a closer look on our construction. There are three pairs of circuits that share the same input throughout the course of the game regardless of the value of the random bit (see Fig. 3). For instance, the circuit pairs  $(F_2, F_6)$ ,  $(F_3, F_{11})$  and  $(F_7, F_{10})$  share the exact same inputs throughout the game. The outer two circuits within each circuit group (circuits  $F_i$  for  $i \equiv 0, 1 \pmod 4$ ) are uniquely paired with exactly one of the middle circuits, i.e. given the same input, depending on the value of the random bit  $b_i$  sampled by the master circuit at each step of the game. For instance, in circuit group 1 if  $b_i = 0$ ,  $F_1$  is paired with  $F_2$  and both circuits given  $x_i^1$  as input, and  $F_4$  is paired with  $F_3$  and both given  $x_i^2$  as input. After the querying phase, the master cross-checks the output of the circuits which share the same input streams. If any of the cross checks in any round fail, then the master aborts and the adversary loses. We now provide a more detailed description of the construction as follows:

**test:** In the test phase,  $T^{F_1, \dots, F_{12}, F}(T)$  picks a random  $\Delta$  such that  $0 \leq \Delta \leq T - 1$ , then queries  $F_1, F_4, F_5, F_8, F_9$  and  $F_{12}$  on  $\Delta$  random and independent inputs  $x_i^1, x_i^4, x_i^5, x_i^8, x_i^9$  and  $x_i^{12}$  respectively and checks if the outputs of the corresponding circuits are correct by comparing them with the trusted  $F$ .

**master:** The master samples three independent input streams  $\vec{x}_1 = (x_1^1, x_2^1, x_3^1, \dots)$ ,  $\vec{x}_2 = (x_1^2, x_2^2, x_3^2, \dots)$ ,  $\vec{x}_3 = (x_1^3, x_2^3, x_3^3, \dots)$  and an independent bit string  $\vec{b} = (b_1, b_2, \dots)$ . The operation of the master circuit is split into two phases: (i) querying phase and (ii) cross-checking phase.

*Querying step.* For all  $i \in [Q]$ , it queries the functions  $F_1, F_2, \dots, F_{12}$  as follows:

1. If  $b_i = 0$ ,
  - The functions  $F_1, F_2, F_5, F_6$  get  $x_i^1$  as input,
  - The functions  $F_3, F_4, F_{11}, F_{12}$  get  $x_i^2$  as input, and
  - The functions  $F_7, F_8, F_9, F_{10}$  get  $x_i^3$  as input.
2. if  $b_i = 1$ ,
  - The functions  $F_1, F_3, F_9, F_{11}$  get  $x_i^1$  as input,
  - The functions  $F_2, F_4, F_6, F_8$  get  $x_i^2$  as input, and
  - The functions  $F_5, F_7, F_{10}, F_{12}$  get  $x_i^3$  as input

*Cross-Checking Step.* For all  $i \in [Q]$ , the master circuit pairwise compares the outputs of the circuits that receive the same inputs (refer to the technical report



for the details of the cross-checking phase). If at any round any of the checks fail, the master outputs **abort** and the adversary loses.

*Output.* If all the checks succeed in the cross-checking phase, the master outputs the output of the circuit  $F_1$ , i.e.,  $\vec{y} = F_1(x^1)$  as the output of  $\Pi_{12}$ .

## 5.2 Security of $\Pi_{12}$

In this section we prove Theorem 1, which states that the construction presented in Sect. 5.1 is  $(c, \frac{c'}{T})$ -secure for appropriate choice of constants  $c$  and  $c'$ . More precisely, we show that the security of the 2-circuit construction from Sect. 2.9 can be reduced to the security of the 12-circuit construction presented above. Before proceeding with the proof, we introduce some useful definitions and notations.

### 5.2.1 History Hardcoded Circuits and Plaits

We observe that the notation  $F(x)$  for *stateful* circuits is ambiguous, since its value depends also on the *history* of queries to  $F$  (which is not provided as a parameter). We can thus assume that each  $F$  is associated with some stream  $\mathbf{x} = (x_1, x_2, \dots)$  and that  $F(x_i) := F(x_i | x_1, x_2, \dots, x_{i-1})$ . This notation uniquely describes the  $i$ -th query to  $F$  given the stream  $\mathbf{x}$ .

In our proof we will however need a slightly different notion called *history-hardcoded circuits*. Given any stateful circuit  $F$  and two arbitrary streams  $\mathbf{x} = (x_1, x_2, x_3, \dots)$  and  $\mathbf{w} = (w_1, w_2, w_3, \dots)$ , we say  $F^{\mathbf{x}}$  is an  $\mathbf{x}$ -history-hardcoded circuit if at the  $i$ -th query it *hardcodes* the stream values  $x_1, \dots, x_{i-1}$  as its history and takes  $w_i$  from the stream  $\mathbf{w}$  as the input to query  $i$ . Thus  $F^{\mathbf{x}}$  on the  $i$ -th query with input  $w_i$  returns the value:  $F^{\mathbf{x},i}(w_i) = F(w_i | x_1, x_2, \dots, x_{i-1})$  and on the  $i + 1$ -th query returns  $F^{\mathbf{x},i+1}(w_{i+1}) = F(w_{i+1} | x_1, x_2, \dots, x_{i-1}, x_i)$ . We call the stream  $\mathbf{x}$  the *hardcoded history stream* and  $\mathbf{w}$  the *input stream*.

For a random variable  $\mathbf{X}$  which takes values from  $\{X_1, X_2, \dots\}$  and a circuit  $F$  we define another random variable  $F^{\mathbf{X}}$  as follows. Its value for  $\mathbf{X} = \mathbf{x}$  is simply  $F^{\mathbf{x}}$ . We will call this random variable an  $\mathbf{X}$ -*history-hardcoded* circuit. Note that as long as  $F^{\mathbf{X}}$  receives inputs from a stream  $\mathbf{W}$  *independent* from  $\mathbf{X}$ , we can say that  $F^{\mathbf{X}}$  is a *history-independent* circuit.

We emphasize that when the hardcoded history stream is equal to the actual input stream, the history-hardcoded circuit returns the same results as the original stateful circuit receiving the same input stream. In other words:

$$F(X_i) = F^{\mathbf{X},i}(X_i), \quad (18)$$

for all  $i \in \mathbb{N}$  with probability 1.

Another idea exploited in our construction is the concept of *alternating* inputs depending on the values of random bits. We will express this idea using the notion of  $\mathbf{b}$ -*plaits*, where  $\mathbf{b}$  is a stream of random bits. A  $\mathbf{b}$ -plait of two streams  $\mathbf{a}^0$  and  $\mathbf{a}^1$  is a new stream  $(\mathbf{a}^0 \mathbf{a}^1)_{\mathbf{b}}$ , where its  $i$ -th value is either  $a_i^0$  from stream  $\mathbf{a}^0$  or  $a_i^1$  from stream  $\mathbf{a}^1$  depending on the  $i$ -th value of the *decision* stream  $\mathbf{b}$ . More precisely:

$$(\mathbf{a}^0 \mathbf{a}^1)_{\mathbf{b}} = (a_1^{b_1}, a_2^{b_2}, a_3^{b_3}, \dots)$$

In our construction, there is only one decision stream used for every plait, therefore the  $\mathbf{b}$  will be omitted for simplicity. Thus to express the plait of two streams  $\mathbf{a}^0, \mathbf{a}^1$  we will simply write  $\mathbf{a}^0\mathbf{a}^1$ . A plait of two identical streams of say  $\mathbf{s}$  will simply be written as  $\mathbf{s}$ , rather than  $\mathbf{ss}$ .

Similarly to  $\mathbf{b}$ -plaits of streams we can define the plaits of history-hardcoded circuits. Let  $G_0^{\mathbf{x}^0}$  be an  $\mathbf{x}^0$ -history-hardcoded circuit and  $G_1^{\mathbf{x}^1}$  be an  $\mathbf{x}^1$ -history-hardcoded circuit. We say  $(G_0^{\mathbf{x}^0}G_1^{\mathbf{x}^1})_{\mathbf{b}}$  is  $\mathbf{b}$ -*plait* for  $G_0^{\mathbf{x}^0}, G_1^{\mathbf{x}^1}$  iff

$$(G_0^{\mathbf{x}^0}G_1^{\mathbf{x}^1})_{\mathbf{b}}^i(\mathbf{x}) = G_{\mathbf{b}_i}^{\mathbf{x}^{\mathbf{b}_i,i}}(\mathbf{x}). \quad (19)$$

Note that the plaited circuit  $(G_0^{\mathbf{x}^0}G_1^{\mathbf{x}^1})_{\mathbf{b}}$  can be expressed as a function of  $G_0, G_1$  and streams  $\mathbf{x}^0, \mathbf{x}^1$ . Looking ahead, this notion of plaited circuits will be crucial in our final reduction of the security of  $\Pi_{12}$  to  $\Pi_2^*$ .

Finally, we define an operation on history-hardcoded circuits in the context of our construction:

**Swap( $F^{\mathbf{x}}, G^{\mathbf{t}}$ ):** Given two history hard coded circuits  $F^{\mathbf{x}}$  and  $G^{\mathbf{t}}$  in our construction, this operation physically exchanges the positions of both circuits. That is, that  $F^{\mathbf{t}}$  physically replaces  $G^{\mathbf{x}}$  and vice versa. Swapped circuits keep their histories, but since they change their place in the construction, they now receive potentially different inputs and are crosschecked with different circuits.

An important notion related to the **Swap** operation which we will exploit in a proof is a *red edge*. We say there is a *red edge* in the  $k$ -th query between two history hard coded circuits  $F^{\mathbf{x}}$  and  $G^{\mathbf{t}}$  iff after performing the **Swap**( $F, G$ ) operation there is a change in either of the outputs of the swapped circuits on the  $k$ -th query compared to the outputs of the circuits if the **Swap** operation was not performed. Looking ahead, the notion of swaps and red edges would be used in our proof to show that modifying the original  $\Pi_{12}$  construction by some **Swap** operations does not change much the security parameters.

Now, given these definitions, we are ready to present an intuition that lies behind our construction. We might (and should) ask the authors “but why 12 circuits?”. The reason is understandable: it is hard to perform any direct proof for history-dependent circuits; things become too complicated. Fortunately, there exist reductions. As long as we have a valid proof of Theorem 2 for history-independent circuits, we can try to find some construction of history-dependent circuits which can be reduced to it. The main goal is to design the crosschecks is such a way, that, informally speaking, making circuit *more history-dependent* make the whole construction *more secure*. It is not hard to believe in such a statement; thanks to the alternating random bit, you never know which of some two circuits will receive a specific input. If these two circuits are very history dependent and have independent histories, there is a high probability, that on the given input they would answer differently. Thank to crosschecks, the master may detect such inconsistency with high probability. To make a practical advantage of this remark, we need to perform many **Swap** operations and analyze the behaviour of various parameters describing our construction. We were able to handle such design and analysis for 12 circuits construction.


Now we will give a more detailed description of the intuition. As written a few lines before, the main idea of the proof is to reduce the construction which consists of (possibly) history-dependent circuits to  $\Pi_2^*$ .  $\Pi_2^*$  consists of 2 history-independent circuits (alternatively speaking - pairs of circuits with different hardcoded histories, independent of the inputs that they receive). The **Swap** operation  $\text{Swap}(F^x, G^t)$  is legit whenever either one of the conditions holds - the circuits  $F$  and  $G$  are engaged in the cross-checking process as pictured in the Fig. 6 (e.g. circuits  $F_1$  and  $F_4$  or circuits  $F_6$  and  $F_7$  in the Fig. 7 ( $\text{Hyb}_0$ ) or the circuits received the same inputs before performing any swaps (e.g. circuits  $F_2$  and  $F_7$  swapped in  $\text{Hyb}_2$  which are placed at the positions of  $F_2$  and  $F_6$  from  $\text{Hyb}_0$  in the Fig. 7). Now, the main idea of the proof is that by performing a series of **Swap** operations on the setting with 3 rows of 4-circuit groups, we are able to end up with a setting  $\text{Hyb}_2$  that contains 2 pairs of history-independent circuits at the place of cross-checked circuit pairs  $(F_1, F_2)$  and  $(F_3, F_4)$ . We need just 1 **Swap** operation in the middle row to have history-independent circuits in the place of  $F_1$  and  $F_4$ , but for  $F_2$  (and  $F_2$ ) we will need an additional input stream that goes with a new row.

We are now ready to proceed to the proof of Theorem 1.

### 5.2.2 Proof of Theorem 1

The proof of Theorem 1 proceeds in two parts. We ultimately want to prove a reduction from the security of  $\Pi_{12}$  to that of  $\Pi_2^*$ . Nevertheless recall in Lemma 4 the security of  $\Pi_2^*$  crucially depends on history independent circuits. Thus the first part of our proof constructs a sequence of three hybrids,  $\text{Hyb}_0, \text{Hyb}_1, \text{Hyb}_2$ , to get a pair of history independent circuits,  $F_4^2$  and  $F_7^3 F_{10}^3$ , in the final hybrid.  $\text{Hyb}_0$  is the original construction. To get from  $\text{Hyb}_0$  to  $\text{Hyb}_1$ , we perform the **Swap** operation on the following pairs of circuits in  $\text{Hyb}_0$ :  $(F_1^1, F_4^2)$ ;  $(F_6^{12}, F_7^3)$ ;  $(F_{10}^3, F_{11}^{21})$ . To get from  $\text{Hyb}_1$  to  $\text{Hyb}_2$ , we perform the **Swap** operation on the following pairs of circuits in  $\text{Hyb}_1$ :  $(F_2^{12}, F_7^3)$ ;  $(F_3^{21}, F_{10}^3)$  (refer to Fig. 7). Note that in the final hybrid  $\text{Hyb}_2$ , it is crucial that  $F_4^2$  and  $F_7^3 F_{10}^3$  are not just history independent, but also take in the same inputs from input stream 1 regardless of the value of the random bit ( $F_7^3 F_{10}^3$  takes inputs from stream 1 due to the definition of plaited circuit in (19)). This will be necessary for the second part of our proof which uses  $F_4^2$  and  $F_7^3 F_{10}^3$  in the final hybrid as the two history independent circuits needed for the  $\Pi_2^*$  construction and uses the  $\Pi_2^*$  construction with these circuits as a subroutine.

*Proof.* For a given  $F_1, \dots, F_{12}$ , we can define some random variables as follows. Let  $\phi_{F_j^A; B}$  be the total number of queries, where  $F_j^A$  gets input from a stream  $B$  and has a mismatch with any other circuit getting input from the same stream in this query. We will refer to random variables related to the  $i$ -th hybrid by adding a superscript  $i$ . For example,  $\phi_{F_1^1; 2}^0 = 0$ , since in  $\text{Hyb}_0$  no crosschecks are made between  $F_1^1$  and the circuits receiving inputs from stream 2. Let  $\Phi$  be the total number of mismatches detected by the master circuit. Recall from Sect. 2.3 that  $Y$  is the total number of mistakes the master circuit makes. The probability



**Fig. 7.** Hybrids with the circuits and their corresponding plaited hardcoded history and input streams (above and below each circuit in black respectively). In Hyb<sub>2</sub>, F<sub>4</sub><sup>2</sup> and the plaited circuit F<sub>7</sub><sup>3</sup>F<sub>10</sub><sup>3</sup> (in red) are history independent. (Color figure online)

space of these random variables is the set of all choices of a stream of random bits **b** and streams of random inputs **1, 2, 3** and a number of tests  $\Delta$ .

We prove our statement by contradiction. To this end, we assume that

$$\exists c > 0 \forall c' > 0, T_0 \in \mathbb{N}, q \in \text{poly} \exists T > T_0, Q > q(T) \exists \text{Adv} \text{ such that} \quad (20)$$

$$\text{Adv} \left( c, \frac{c'}{T} \right) \text{ --wins } \text{TrojanGame}(H_{12}, T, Q)$$

Therefore for some  $c$  and for all  $c'$  there exists an infinite set  $\mathcal{T} \subset \mathbb{N}$  such that for every  $t \in \mathcal{T}$  there exists an infinite set  $\mathcal{Q}_t \subset \mathbb{N}$  such that for every  $t \in \mathcal{T}, z \in \mathcal{Q}_t$  there exists an adversary  $\text{Adv} = \text{Adv}(c, c', z, t)$  such that the following formula is true:

$$\Pr \left[ \Phi^0 = 0 \quad \text{and} \quad Y^0 \geq c' \cdot \left( \frac{z}{t} \right) \right] \geq c. \quad (21)$$

Now we will look what happens to inequality (21) as we move through each hybrid:

Hyb<sub>0</sub>: Hybrid 0 corresponds to the original construction due to equality (18). Hence, the probability that the adversary  $\text{Adv}(c, c', z, t)$  wins in this hybrid is precisely that in Eq. (21).

Hyb<sub>1</sub>: In this hybrid we simply perform three Swap operations on the following pairs of circuits: (F<sub>1</sub><sup>1</sup>, F<sub>4</sub><sup>2</sup>); (F<sub>6</sub><sup>12</sup>, F<sub>7</sub><sup>3</sup>); (F<sub>10</sub><sup>3</sup>, F<sub>11</sub><sup>21</sup>).

*Claim.*  $\Pr \left[ \phi_{F_4^1, 1}^1, \phi_{F_7^3, 12}^3, \phi_{F_{10}^3, 21}^3 \leq k \wedge \Phi^0 = 0 \wedge Y^1 \geq c' \cdot \left( \frac{z}{t} \right) - 3k \right] \geq c - 3 \cdot 2^{-k}.$

Proof of the claim is in the technical report.

Hyb<sub>2</sub>: In this hybrid we simply perform two Swap operations on the following pairs of circuits: (F<sub>2</sub><sup>12</sup>, F<sub>7</sub><sup>3</sup>); (F<sub>3</sub><sup>21</sup>, F<sub>10</sub><sup>3</sup>).

*Claim.*

$$\Pr \left[ (\phi_{\mathbb{F}_4^2, \mathbf{1}}^2 \leq 3k) \wedge (Y^2 \geq c' \left(\frac{z}{t}\right) - 5k) \right] \geq c - 3 \cdot 2^{-k}. \quad (22)$$

*Proof.* Every **Swap** operation performed in  $\text{Hyb}_2$  changes the value of  $Y^2, \phi_{\mathbb{F}_4^2, \mathbf{1}}^2$  by at most  $k$  (since inequality (refer to our extended technical report for details) holds). The inequality is explicit.  $\triangle$

*Claim.* For every  $k \in \mathbb{N}$  and every adversary  $\text{Adv}$  who  $(c, \frac{c'}{t})$ -wins  $(\Pi_{12}, T, Q) - \text{TrojanGame}$  there exists an adversary  $\text{Adv}'$  who  $(c - 3 \cdot 2^{-k}, \frac{c'}{t} - \frac{5k}{z}) - (3k + 1)$ -wins the game  $\text{TrojanGame}(\Pi_2^*, T, Q)$ .

We want to conclude, that the above statement contradicts Lemma 4. So we want to show, that this incorrect statement is implied by our construction.

$$\begin{aligned} \exists \tilde{c} > 0 \forall \tilde{c}' > 0, T_0 \in \mathbb{N}, q \in \text{poly} \exists T > T_0, Q > q(T) \exists \widetilde{\text{Adv}} \text{ such that} \\ \widetilde{\text{Adv}} \left( \tilde{c}, \frac{\tilde{c}'}{T} \right) \text{-wins } \text{TrojanGame}(\Pi_2^*, T, Q). \end{aligned} \quad (23)$$

Let  $k = 2 + \log(\frac{1}{c})$  and  $\tilde{c} = c - 3 \cdot 2^{-k} = \frac{c}{4} > 0$ . Choose  $\tilde{c}'$  arbitrarily and let  $c' = \cdot \tilde{c}'$ . Let  $\tilde{T} = T$ . Let

$$\widetilde{Q}_t = \left\{ z \in Q_t : z > t \left( \frac{5k}{\tilde{c}'} + 1 \right) \right\}.$$

Obviously  $\widetilde{Q}_t$  is infinite. As a result, for every  $q \in \text{poly}$  there exists  $z \in \widetilde{Q}_t$  such that  $z > q(t)$ .

Now we can construct the adversary  $\widetilde{\text{Adv}}$  which would break the security of  $\Pi_2^*$  which lead us to contradiction. Thanks to the analysis of the hybrids we know, that for  $\text{Adv}$  the inequality (22) holds. Define the circuits  $\widetilde{F}_1, \widetilde{F}_2$  given to  $\widetilde{\text{Adv}}$  in the following way:

$$\widetilde{F}_1 = F_4^2, \quad \widetilde{F}_2 = F_7^3 F_{10}^3,$$

where the latter is a **b**-plait (as defined in Eq. (19)). Actual values of streams **2, 3, b** are sampled uniformly and independently by  $\widetilde{\text{Adv}}$ , and hardcoded in  $\widetilde{F}_1, \widetilde{F}_2$ . Obviously  $\widetilde{F}_1, \widetilde{F}_2$  are history independent, therefore  $\widetilde{\text{Adv}}$  meets the requirements for the  $\Pi_2^*$  scheme.

Now we can bound a random variable  $\widetilde{\Phi}$  - the number of queries in a  $(\Pi_2^*, T, Q) - \text{TrojanGame}$  where the adversary is caught on deviating. If  $\phi_{\mathbb{F}_4^2, \mathbf{1}}^2 \leq 3k$ , then  $\widetilde{\Phi} \leq 3k$ , since if in the  $i$ -th query there was any inconsistency between  $\widetilde{F}_1, \widetilde{F}_2$ , there must had been a mismatch between  $F_4^3$  and any other circuit receiving the same input. Which concludes in:

$$\begin{aligned} \tilde{c} = c - 3 \cdot 2^{-k} &\leq \Pr \left[ (\phi_{\mathbb{F}_4^2, \mathbf{1}}^2 \leq 3k) \wedge (Y^2 \geq c' \left(\frac{z}{t}\right) - 5k) \right] \\ &\leq \Pr \left[ \widetilde{\Phi} \leq 3k \wedge (\widetilde{Y} \geq c' \left(\frac{z}{t}\right) - 5k) \right] \\ &\leq \Pr \left[ \widetilde{\Phi} \leq 3k \wedge (\widetilde{Y} \geq \tilde{c}' \left(\frac{5k}{\tilde{c}'} + 1\right) - 5k = \tilde{c}') \right]. \end{aligned}$$

Note that  $\tilde{Y}$  is the number of mistakes made by the master circuit in the  $\text{TrojanGame}(II_2^*, T, Q)$  and the last inequality comes from  $z > t \left( \frac{5k}{c'} + 1 \right)$ . We conclude, that that there exists  $\tilde{c}$ , such that for every  $\tilde{c}'$  there exists  $\widetilde{\text{Adv}}$  who  $(\tilde{c}, \frac{c'}{t})$ - $(3k + 1)$ -wins  $\text{TrojanGame}(II_2^*, T, Q)$ . It is with contradiction with Lemma 4, which ends the proof.  $\square$

### 5.3 Reapplying the Hybrid Argument

In the previous section, we have used the outputs produced on the input stream **1** in place of  $F_1$  as an output stream of the construction. By symmetry, the argument from the previous section works for the input stream **2** in place of  $F_4$ . Now we will show that in fact the outputs from  $F_5$  or  $F_8$  may be used as an output stream of the construction, which also implies the possibility of using input stream **3** to produce the output stream of the construction.

Now, in  $\text{Hyb}_0$  (Fig. 7) firstly swapping the labels of the input streams **3** and **2** on the input bit 0 and the labels of the input streams **3** and **1** on the input bit 1 and secondly visually swapping the rows 1 and 2 does not change the setting. We achieve the  $\text{Hyb}_0''$  construction as shown in Fig. 8.




Fig. 8.  $\text{Hyb}_0''$  construction.

We can still reapply the hybrid argument from the previous section to the modified  $\text{Hyb}_0''$  construction by applying the following swaps.  $\text{Hyb}_0'' \rightarrow \text{Hyb}_1''$  swaps :  $(F_5^1, F_8^2)$ ;  $(F_2^{12}, F_3^3)$ ;  $(F_{11}^3, F_{10}^{21})$ .  $\text{Hyb}_1'' \rightarrow \text{Hyb}_2''$  swaps :  $(F_3^{12}, F_6^3)$ ;  $(F_7^{21}, F_{11}^3)$ . Finally we conclude that the output streams of either  $F_5$  or  $F_8$  (by symmetry) may be used as an output stream of the construction, what implies that by taking outputs from circuit  $F_8$  on input bit 0 and circuit  $F_5$  on input bit 1 could be used as an output of the construction produced with input stream **3**.

The above argument implies that in each round our construction may output 3 outputs produced by consuming inputs from the same number of 3 input streams.

## 6 Outlook and Open Problems

In this work we introduced countermeasures against hardware Trojans which compared to existing solutions based on multiparty or verifiable computation are extremely simple and efficient, but only achieve limited security guarantees (i.e., we only guarantee that *most* outputs are correct or the Trojan is detected with high probability) in a restricted setting (iid inputs and no evolving state).

Because of this, the scope of application for our schemes is limited, but as discussed in the introduction, we believe they will serve as a first but major step towards solving some of the main application targets like randomness generation. In particular, creating pseudorandomness for “randomness hungry” side-channel countermeasures like masking in a Trojan-resilient way is one of the main motivations for this work. The reason our simple schemes are a promising starting point towards Trojan-resilient pseudorandomness generation is the fact that in most settings (like masking) one does not need that the pseudorandomness is correctly generated, only that it is indistinguishable from uniform, so the relaxed security of our schemes is not a deal breaker. Another reason is the fact that one could use some of the pseudorandomness that is created to implement the master’s randomness source, thus making it deterministic. Fleshing these ideas out will require a better understanding of amplification and circularity issues in this setting.

The main concrete technical question left open problem in this work is to prove the security of the “minimal” and thus really practical scheme  $\Pi_2^\phi$  as stated in Conjecture 1. A positive resolution of the conjecture will need techniques that go beyond our proof via history-independence used in the proof for  $\Pi_{12}$ .

## References

1. Adee, S.: The hunt for the kill switch, spectrum (2008). <https://tinyurl.com/j95zbxma>
2. Ateniese, G., Kiayias, A., Magri, B., Tselekounis, Y., Venturi, D.: Secure outsourcing of cryptographic circuits manufacturing. In: Baek, J., Susilo, W., Kim, J. (eds.) ProvSec 2018. LNCS, vol. 11192, pp. 75–93. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01446-9\\_5](https://doi.org/10.1007/978-3-030-01446-9_5)
3. Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: Goldwasser, S. (ed.) ITCS, pp. 326–349. ACM (2012). <https://doi.org/10.1145/2090236.2090263>
4. Cramer, R., Damgård, I., Nielsen, J.B.: Secure Multiparty Computation and Secret Sharing. Cambridge University Press, Cambridge (2015). <http://www.cambridge.org/de/academic/subjects/computer-science/cryptography-cryptography-and-coding/secure-multiparty-computation-and-secret-sharing?format=HB&isbn=9781107043053>
5. Dziembowski, S., Faust, S., Standaert, F.: Private circuits III: hardware trojan-resilience via testing amplification. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS, pp. 142–153. ACM (2016). <https://doi.org/10.1145/2976749.2978419>

6. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: FOCS, pp. 293–302. IEEE Computer Society (2008). <https://doi.org/10.1109/FOCS.2008.56>
7. Gennaro, R., Lysyanskaya, A., Malkin, T., Micali, S., Rabin, T.: Algorithmic Tamper-Proof (ATP) security: theoretical foundations for security against hardware tampering. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 258–277. Springer, Heidelberg (2004). [https://doi.org/10.1007/978-3-540-24638-1\\_15](https://doi.org/10.1007/978-3-540-24638-1_15)
8. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: securing hardware against probing attacks. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 463–481. Springer, Heidelberg (2003). [https://doi.org/10.1007/978-3-540-45146-4\\_27](https://doi.org/10.1007/978-3-540-45146-4_27)
9. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (1996). [https://doi.org/10.1007/3-540-68697-5\\_9](https://doi.org/10.1007/3-540-68697-5_9)
10. Kocher, P., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (1999). [https://doi.org/10.1007/3-540-48405-1\\_25](https://doi.org/10.1007/3-540-48405-1_25)
11. Maurer, U., Pietrzak, K., Renner, R.: Indistinguishability amplification. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 130–149. Springer, Heidelberg (2007). [https://doi.org/10.1007/978-3-540-74143-5\\_8](https://doi.org/10.1007/978-3-540-74143-5_8)
12. Mitra, S., Wong, H.S.P., Wong, S.: Stopping hardware trojans in their tracks, spectrum (2015). <https://tinyurl.com/5emst8f2>
13. Mukhopadhyay, D., Chakraborty, R.S.: Hardware Security: Design, Threats, and Safeguards, 1st edn. Chapman & Hall/CRC, Boca Raton (2014)
14. Pietrzak, K.: A leakage-resilient mode of operation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 462–482. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-01001-9\\_27](https://doi.org/10.1007/978-3-642-01001-9_27)
15. Standaert, F.-X., Malkin, T.G., Yung, M.: A unified framework for the analysis of side-channel key recovery attacks. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 443–461. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-01001-9\\_26](https://doi.org/10.1007/978-3-642-01001-9_26)
16. Tehranipoor, M., Salmani, H., Zhang, X.: Integrated Circuit Authentication: Hardware Trojans and Counterfeit Detection. Springer, Cham (2013). <https://doi.org/10.1007/978-3-319-00816-5>
17. Wahby, R.S., Howald, M., Garg, S., Shelat, A., Walfish, M.: Verifiable ASICs. In: IEEE SP, pp. 759–778. IEEE Computer Society (2016). <https://doi.org/10.1109/SP.2016.51>