

Secure Outsourcing of Circuit Manufacturing

Giuseppe Ateniese^{*}, Aggelos Kiayias[†], Bernardo Magri[‡], Yiannis Tselekounis[†], and
Daniele Venturi[§]

^{*}*Stevens Institute of Technology*

[†]*University of Edinburgh*

[‡]*Friedrich-Alexander-Universität Erlangen-Nürnberg*

[§]*Sapienza University of Rome*

January 3, 2018

Abstract

The fabrication process of integrated circuits (ICs) is complex and requires the use of off-shore foundries to lower the costs and to have access to leading-edge manufacturing facilities. Such an outsourcing trend leaves the possibility of inserting malicious circuitry (a.k.a. hardware Trojans) during the fabrication process, causing serious security concerns. Hardware Trojans are very hard and expensive to detect and can disrupt the entire circuit or covertly leak sensitive information via a subliminal channel.

In this paper, we propose a formal model for assessing the security of ICs whose fabrication has been outsourced to an untrusted off-shore manufacturer. Our model captures that the IC specification and design are trusted but the fabrication facility(ies) may be malicious. Our objective is to investigate security in an *ideal sense* and follows a simulation based approach that ensures that Trojans cannot release any sensitive information to the outside. It follows that the Trojans' impact in the overall IC operation, in case they exist, will be negligible up to simulation.

We then establish that such level of security is in fact achievable for the case of a single and of multiple outsourcing facilities. We present two compilers for ICs for the single outsourcing facility case relying on verifiable computation (VC) schemes, and another two compilers for the multiple outsourcing facilities case, one relying on multi-server VC schemes, and the other relying on secure multiparty computation (MPC) protocols with certain suitable properties that are attainable by existing schemes.

Keywords: Security models, hardware Trojans, circuit compilers, fabless manufacturing, verifiable computation.

Contents

1 Introduction	1	4.2 Compiler based on Input-Private VC	13
1.1 Background	1	4.3 Compiler based on VC without Input Privacy	18
1.2 Previous Work	2		
1.3 Our Contributions	3		
1.4 Related Work	6		
2 Preliminaries	7	5 Case Study II: Multiple Manufacturers	20
2.1 Notation	7	5.1 Compiler based on MPC	20
2.2 Circuits	7	5.2 Compilers based on VC with Multiple Servers	24
3 Secure Circuit Fabrication	8		
3.1 Security	9		
3.2 Undetectability	11	6 Concrete Instantiations	28
4 Case Study I: Single Manufacturer	11	6.1 Case Study I	28
4.1 Prelude: Verifiable Computation	11	6.2 Case Study II	29
		7 Conclusion and Open Problems	31

1 Introduction

1.1 Background

The fabrication process adopted by the semiconductor industry is fundamentally global, involving several parties that may not be trusted. As a result, integrated circuits (ICs) are vulnerable to so-called hardware Trojans that can compromise or disable critical systems, or covertly leak sensitive information [57, 24, 9]. Analogously to a software Trojan, a hardware Trojan is a *back-door* deliberately added to the circuit to disrupt its operation or disable it when certain events occur. A Trojan can be added to the circuit during the design phase, by some malicious designer, or more often during the manufacturing phase, by some malicious off-shore fabrication facility. A hardware Trojan’s objectives may be to modify the functionality of the circuit (e.g., in order to compromise or disable critical systems), modify its specification (e.g., by changing its energy consumption), covertly leak sensitive information (e.g., from a secret memory), or simply disable the entire circuit when instructed to do so [7]. Once the Trojan is inserted into the circuit it can stay active the entire time, or it can be “triggered” by some event such as a special input to the circuit.

Reliably detecting compromised circuit components through testing and reverse engineering appears to be an impossible task given our current technology [15]. Indeed, all non-destructive testing techniques can easily be circumvented by properly obfuscating embedded Trojans. The U.S. military recognized this threat and started two programs, Trust and IRIS, with the intent of developing techniques and metrics to certify ICs going into weapon systems. The main concern is that advanced weapons may appear to work properly but then switch off in combat or when triggered by some special events. Another stated concern is information leakage, where a malicious component is programmed to leak sensitive information [71].

The U.S. military however currently obtains trusted chips through the DOD Trusted Foundry program which is currently managed by the NSA’s Trusted Access Program Office (TAPO). Within this program, a trusted design center and foundry are established through an exclusive partnership with IBM for secure semiconductor fabrication and ASIC services, along with the involvement of several Trusted Suppliers which are accredited by an accreditation authority

(DMEA). The intent of the Trusted Foundry program is to provide national security and defense programs with access to ICs from trusted sources. However, a report by the U.S. Government Accountability Office (GAO) [59], released in April 2015, found that even though the Trusted Foundry program started in 2004, IBM remained the sole-source supplier for leading-edge technologies meeting the criteria put forth by DOD. GAO’s report highlights two main issues: First, it notices that IBM sold its microelectronics fabrication business to a foreign-owned entity (GlobalFoundries). Second, relying on a single source supplier for defense microelectronics hinders competition and thus innovation in this critical area.

1.2 Previous Work

Inspired by the above considerations, in this work we put forward a formal security model for the problem of utilizing off-shore fabrication facilities for IC manufacturing. Our main motivation is that the setting of secure circuit fabrication, while being an extremely important practical problem, almost completely lacks theoretical foundations. We discuss a few remarkable exceptions below.

- Seifert and Bayer [70] introduced a very strong security model for the fabrication of Trojan-resilient circuits, where the produced circuit is required to always have the same output as the original circuit; unfortunately, they show how to achieve their definition only for very limited classes of Trojans (i.e., the adversary is allowed to “corrupt” only a small fraction of the gates in each layer of the IC, and a small fraction of the wires connecting different layers).
- Recently, Wahby *et al.* [72] introduced a new approach to the problem of defeating hardware Trojans in fabless circuit manufacturing. Their model reflects the fact that IC specification and design are trusted but the fabrication facility is not. Rather than testing or reverse engineering the IC hardware received, which only provides limited security, they consider a class of solutions where the IC’s operations are continuously verified. Such an approach makes sense as long as the verification circuitry can be: (i) constructed cheaply, and (ii) run efficiently. These properties are achieved by leveraging a verifiable computation (VC) scheme for the function implemented by the original circuit. Verifiable computation (see, e.g., [41]) is a recent paradigm by which resource-constrained clients can delegate the computation of some function \mathcal{F} , on (possibly private) input X , to an untrusted (and computationally powerful) server, without the server being able to cheat about the outcome of the computation, and with the property that verifying the server’s answer is much more efficient than computing the function from scratch.

In a nutshell, the goal of [72] is to make sure that the produced circuit maintains correctness of the computation, meaning that the output of the circuit is either invalid, or equal to the output of the original circuit. The main drawback is that invalid outputs might be arbitrarily correlated with the secret state of the circuit, which could expose key material in case the produced circuit is a cryptographic circuit. (We will formalize this fact later in the paper.)

- In [35], the authors show how to protect against hardware Trojans using testing-based mechanisms. Their work is based on two existing techniques for Trojan detection, called “input scrambling” and “split manufacturing” [49], for which the authors provide formal models. Hence, they present a generic compiler that transforms any circuit into a new (equivalent) circuit with the following guarantee: Assuming the attacker invokes the circuit q times, and that the device is being tested t times, for $t > q$ uniform on a specific range which is not known to the attacker, the compiled circuit is secure with probability at

least $1 - (q/t)^{\ell/2}$, where ℓ is the number of copies of the sub-circuits whose production is outsourced.

The main limitation is that [35] assumes an a-priori known bound on the number q of interactions between the user and the device; in fact, without such a bound, their construction would require a super-polynomial number of tests. Unfortunately, in many important applications, it is not realistic to assume an upper bound on the value q , and thus it is an important open problem to design a methodology that provides security for an arbitrary polynomial number of interactions between the user/attacker and the device.

- The approach of applying secure distributed computing to defeat hardware Trojans has also been recently explored in [60]. However, this work is more focused on the implementation aspects of this idea, and moreover it assumes that the possibly malicious circuit components run applications that are developed and signed by a trusted software developer.

1.3 Our Contributions

We put forward a formal framework for assessing security of a circuit whose production has been, in part, outsourced to a set of manufacturers that are not trusted. Our security definition implies that using the produced circuit in the wild leaks no information on its secrets. Additionally, the adversarial model we consider does not assume any a-priori bound on the number of executions, and allows the manufacturer(s) to make arbitrary modifications to the outsourced components. In essence, our security model captures any attack in which the backdoored circuit communicates with the user/attacker through the input/output gates of the produced circuit. (This includes digital and analog Trojans, but not hidden antennas as considered in [35].)

With such a framework in hand, we give several design methodologies that achieve our definition with different tradeoffs in terms of security, efficiency, and underlying assumptions. Thus, our work establishes the theoretical feasibility of utilizing off-shore fabrication facilities for IC manufacturing. A more detailed explanation of our main contributions follows below.

Secure circuit fabrication. Let Γ be the original circuit to be produced. Instead of producing Γ directly, we first “compile” it into a different circuit $\hat{\Gamma}$ using an efficient, possibly randomized, procedure Φ that we call an *outsourcing compiler*. The compiler Φ takes as input a description of Γ and returns a description of $\hat{\Gamma}$, together with some auxiliary information specifying how $\hat{\Gamma}$ can be divided into sub-components, and which of these components can be produced off-shore; the remaining components will be instead built in-house. After all components have been produced, the circuit designer re-assembles the circuit $\hat{\Gamma}$ (by combining the outsourced components and the components built in-house), which is then initialized with some initial secret memory M_1 , and used in the wild.

In order to make sense, the above approach needs to satisfy a few important requirements. The first requirement is that Φ needs to be functionality preserving, meaning that the compiled circuit $\hat{\Gamma}$ should compute the same functionality as the original circuit Γ (for all possible initial memories M_1 , and for all possible inputs). The second requirement is that the effort needed to manufacture the trusted sub-components should be (much) less compared to the effort required to manufacture the original circuit Γ . The third requirement is that Φ should be secure, meaning that, under an acceptable assumption about the manufacturers who construct the outsourced components, the produced circuit $\hat{\Gamma}$ can be safely used in real-life applications.

Our security definition follows the simulation paradigm, and is inspired by similar definitions in the setting of tamper-proof circuit compilers [51, 38]. We refer the reader to Section 1.4 for a more detailed comparison between the two approaches. In a nutshell, security of Φ is defined

by requiring that whatever an adversary can learn by interacting with the fabricated circuit $\widehat{\Gamma}$ (produced following the steps outlined above), can be simulated given only black-box access to the original circuit Γ . This essentially means that, no matter how the outsourced components are maliciously modified (e.g., by inserting a hardware Trojan), using circuit $\widehat{\Gamma}$ is as secure as using the original circuit Γ , and thus, in particular, does not leak sensitive information on the secret memory. See Section 3 for a precise definition.

We also consider a weakening of the above definition, in which the simulator is allowed to receive a short advice (or leakage) on the secret memory M_1 . This models a setting where the adversary might be able to learn a short amount of information on the secret memory, but still yields a meaningful security guarantee provided that the original circuit is resilient to such a short leakage. An appealing advantage of this weaker definition is that it might allow for significantly more efficient circuit compilers.

Case study I: Single manufacturer. In Section 4, we show how to construct secure outsourcing compilers that work for *arbitrary* circuits Γ in the setting where *all* outsourcing manufacturers are corrupted. Similarly to [72], our compilers generically leverage a VC scheme for the function \mathcal{F} implemented by Γ . Recent breakthrough research on verifiable computation led to nearly practical schemes that work for any function [66, 26]; some schemes additionally preserve the privacy of the inputs on which the function is being computed on [39]. VC schemes satisfying the latter property are called *input-private*.

The main idea of how to use verifiable computation in order to build secure outsourcing compilers is simple enough to describe it here. The fabrication of the chips that perform the entire bulk of computation will be outsourced to the untrusted fabrication facility, whereas the only circuit components that need to be built in-house are: (i) the component corresponding to the algorithm for encoding the inputs (in case of input-private VC), (ii) the component corresponding to the algorithm run by the client in order to verify correctness of the server’s computation, and (iii) the component used to generate fresh random coins as needed for computing the function (in case of randomized functions). Thanks to the nature of VC, the size of the components in (i) and (ii) is independent of the size of the original circuit computing the function. As for the component in (iii), we can use any existing (and trusted) circuitry for generating true random numbers (RNG). A good example is the Intel on-chip hardware random number generator which can be accessed through the RDRAND instruction available on all modern processors [53, 48].¹

We implement the above idea in two ways, depending on the properties satisfied by the underlying VC scheme, as explained below.

- Our first compiler relies on VC schemes with input-privacy, and achieves our strongest security notion (i.e., no leakage required for the simulation).
- Our second compiler relies on VC schemes without input-privacy, and achieves security provided the original primitive (implemented by the circuit Γ) is resilient against a logarithmic amount of leakage on the private memory.

Remarkably, any public-key encryption or signature scheme is resilient to such an amount of leakage at the price of a polynomial loss in the concrete security, and recently many leakage-resilient schemes [36, 68, 55, 34, 63, 19, 64, 37] have been constructed, where the concrete security does not degrade with the total amount of tolerated leakage.

¹Intel’s generator relies on unpredictable thermal noise to generate bits that are fed to a cryptographic “conditioner” (AES in CBC-MAC mode) which produces a 256-bit seed that is then passed through a NIST SP800-90A-based pseudorandom generator.

The second compiler additionally relies on a special “self-destruct” feature (which is implemented² in one of the components built in-house), meaning that after the first invalid output is ever processed, the entire memory is overwritten. As we show, this is an inherent requirement for this compiler, in that, without such a property, there exist generic attacks that allow to recover the entire private memory. Moreover, such attacks are undetectable by all polynomial-time (black-box) tests. Our definition of undetectability (see Section 3.2) is similar in spirit to analogous definitions in the context of subversion-resilient cryptography [12, 11].

Case study II: Multiple manufacturers. In Section 5, we show how to construct secure outsourcing compilers for arbitrary circuits Γ in the setting where $m \geq 2$ outsourcing manufacturers are available, and a certain unknown subset of them is malicious. This is a strictly stronger assumption compared to the setting of a single manufacturer, nevertheless, as we show, it opens the possibility for more efficient constructions and stronger availability guarantees. We present two different approaches.

- The first approach yields an outsourcing compiler utilizing a general client-server secure multiparty computation (MPC) protocol, i.e., a protocol that, for any function, enables a set of clients to privately communicate their inputs to a set of servers that will perform a computation and return the output to a single designated recipient. We stress that many MPC protocols follow this paradigm (e.g., [30]), while others, as we comment later, can be easily adapted to it.

Given such a protocol, the compiler operates in the following way (see also Section 5.1). For a given circuit Γ it produces the MPC protocol implementing it, isolates the client and recipient computation for manufacturing in-house, and outsources each of the other components (representing a server in the MPC protocol) to the untrusted manufacturers. The key points of this compiler construction are as follows: (i) The client and recipient computation are typically quite lightweight; the client, in many protocols, simply performs an encryption or a secret-sharing operation, and the recipient a secret-reconstruction protocol; in either case, the computation is independent of the circuit that is outsourced. (ii) There are MPC protocols that can tolerate up to $m - 1$ malicious servers, something we can leverage to argue that if at least one of the outsourcing manufacturer is honest the compiled circuit would be safe for use.

Additional properties of the underlying MPC protocol can also be very valuable by our compiler: for instance, if the underlying MPC protocol supports guaranteed output delivery, we can use this guarantee to argue that the final circuit will be resilient to a certain faulty outsourced sub-component. Moreover, if the underlying protocol satisfies the identifiable abort property, cf. [50], we can enable our compiled circuit to switch-off an outsourced sub-component that is discovered to be faulty (or malicious), thus reducing energy consumption.

- The second approach yields outsourcing compilers based on so-called multi-server VC [21, 1]; these are VC schemes where the client outsources the computation of the function to multiple servers. In particular, one of our construction leverages a special type of multi-server VC scheme called Refereed Delegation of Computation (RDoC) [21, 22, 23]. This yields a very simple outsourcing compiler achieving the very strong guarantee that the compiled circuit outputs the same as the (untampered) original circuit with overwhelming probability. Furthermore, in this construction it is also possible to identify the malicious

²The self-destruct feature can be realized in several ways. One option is to have the decoder overwrite the memory with all zeroes after an invalid output is produced. Alternatively, one could let the decoder keep a single bit of state, which is written to 1 as soon as the first invalid output is produced.

components, which can then be switched-off providing additional advantages in terms of efficiency and energy consumption.

The main technical ideas behind the compilers based on multi-server VC are similar in spirit to the ones underlying our first compiler based on single-server VC. We refer the reader directly to Section 5.2 for a high-level description of these compilers and their security analysis.

On efficiency. In Section 6, we describe several instantiations of our compilers, based on state-of-the-art construction in the settings of VC and MPC. This yields several concrete constructions with different tradeoffs in terms of security, efficiency, and underlying hardness assumptions. We do stress, however, that the main contributions of this paper are the security models and precise definitions put forward for the problem of fabless ICs manufacturing. The constructions presented in Section 6 are described mainly to evidence the feasibility of our proposed model.

An important remark is that all our approaches require a partitioning and assembly procedure which must be performed in-house. To lower the cost, trusted and untrusted sub-circuits will likely be diced in their own wafers and packaged independently. We note that recent advanced designs and emerging technologies in the electronics assembly and packaging industry make interconnecting different components a reasonably cost-efficient operation.

1.4 Related Work

Hardware Trojans. Prevention of hardware Trojans in ICs is a common practice that might take place during the design, manufacturing, and post-manufacturing stage [69, 58]. However, since it is not always possible to efficiently prevent Trojans insertion, Trojans *detection* has also been vastly explored [15]; once a Trojan is detected, the circuit can be disposed and not used. Common methodologies used to perform Trojans detection vary from invasive ones (that destroy the IC to examine it inside), to non-invasive ones (where the circuit is executed and compared against a trusted copy of the circuit, or against some expected output values). Trojan detection is typically a very expensive and unreliable process, therefore the best practice is usually not to rely on any kind of testing to protect against Trojans.

Explicit *countermeasures* against Trojans also exist, where the objective is to guarantee the functionality or security of the circuit even in the presence of some unknown Trojan. For instance, the so-called “data guards” are designed to prevent a Trojan from being activated and/or to access sensitive data [73]. Another approach is the duplication of logic elements and the division of the sensitive data to independent parts of the circuit [61, 73].

Tamper-proof hardware. The existence of tamper-proof hardware is a physical assumption that found several applications in cryptography, for instance in the setting of secure computation [54]. Here, parties in a protocol can exchange (and rely on) tamper-proof hardware tokens, executing (possibly stateful) programs in a black-box manner, upon inputs of the parties’ choice.

In this vein, our work is related to the challenging scenario where the parties rely on tamper-proof tokens that might have been produced by one (or more), possibly malicious, outsourcing facility. Such a question was studied in [40] for the special case of secure protocols for set intersection, and very recently in [25] for the more general case of arbitrary multi-party functionalities.

Tamper-proof circuits. Our main security definition shares similarities with analogous definitions in the context of protecting circuit implementations against tampering attacks. This line of research received considerable attention in the past few years [51, 38, 27, 56, 28].

The main difference between this setting and the one considered in our paper is that tamper-proof circuit compilers are typically used to protect against fault injection [65] and tampering attacks at run-time; such attacks are usually carried out in an adaptive manner, depending on the outcome of previous attempts. Outsourcing compilers, instead, only protect against (non-adaptive) tampering taking place during the circuit fabrication process. Importantly, the latter restriction allows to obtain security against arbitrary modifications, whereas in circuit tampering one has to consider very restricted attacks (e.g., wire tampering [51] or gate tampering [56]).

Subversion. The above type of non-adaptive tampering is, in fact, reminiscent of the setting of subversion attacks against cryptographic primitives and algorithms. Inspired by the recent revelations of Edward Snowden [67, 5, 46], this line of research recently led to constructing several concrete primitives resisting large classes of subversion attacks [12, 32, 33, 3]. In this light, our work could be interpreted as formalizing the security of circuits that might have been subject to subversion during fabrication.

2 Preliminaries

2.1 Notation

For a string x , we denote its length by $|x|$; if S is a set, $|S|$ represents the number of elements in S ; for a natural number n , $[n]$ denotes the set $\{1, \dots, n\}$. When x is chosen randomly in S , we write $x \leftarrow_s S$. When \mathcal{A} is an algorithm, we write $y \leftarrow \mathcal{A}(x)$ to denote a run of \mathcal{A} on input x and output y ; if \mathcal{A} is randomized, then y is a random variable and $\mathcal{A}(x; r)$ denotes a run of \mathcal{A} on input x and randomness r . An algorithm \mathcal{A} is *probabilistic polynomial-time* (PPT) if \mathcal{A} is randomized and for any input $x, r \in \{0, 1\}^*$ the computation of $\mathcal{A}(x; r)$ terminates in at most $\text{poly}(|x|)$ steps.

We denote with $\lambda \in \mathbb{N}$ the security parameter. A function $\nu : \mathbb{N} \rightarrow [0, 1]$ is negligible in the security parameter (or simply negligible) if it vanishes faster than the inverse of any polynomial in λ , i.e. $\nu(\lambda) = \lambda^{-\omega(1)}$.

The statistical distance between two random variables \mathbf{Z} and \mathbf{Z}' defined over some common set Z is defined as $\Delta(\mathbf{Z}; \mathbf{Z}') = \frac{1}{2} \sum_{z \in Z} |\mathbb{P}[\mathbf{Z} = z] - \mathbb{P}[\mathbf{Z}' = z]|$.

We rely on the following lemma (which follows directly from the definition of statistical distance):

Lemma 1. *Let \mathbf{Z} and \mathbf{Z}' be a pair of random variables, and W be an event defined over the probability space of \mathbf{Z} and \mathbf{Z}' . Then,*

$$\Delta(\mathbf{Z}; \mathbf{Z}') \leq \Delta(\mathbf{Z}; \mathbf{Z}' | \neg W) + \mathbb{P}[W].$$

For two ensembles $\mathbf{Z} := \{Z_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathbf{Z}' := \{Z'_\lambda\}_{\lambda \in \mathbb{N}}$, we write $\mathbf{Z} \equiv \mathbf{Z}'$ to denote that the two ensembles are identically distributed. We also write $\mathbf{Z} \approx_c \mathbf{Z}'$ to denote that the ensembles are computationally indistinguishable, i.e. for all PPT distinguishers \mathcal{D} there exists a negligible function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Delta^{\mathcal{D}}(\mathbf{Z}; \mathbf{Z}') := |\mathbb{P}[\mathcal{D}(z) = 1 : z \leftarrow_s \mathbf{Z}] - \mathbb{P}[\mathcal{D}(z) = 1] : z \leftarrow_s \mathbf{Z}'| \leq \nu(\lambda).$$

2.2 Circuits

A (Boolean) circuit $\Gamma = (V, E)$ is a directed graph. The vertices V are logical gates, and the edges E are wires connecting the gates. For the case of *deterministic* circuits, the gates can be of type AND, XOR and copy, where AND (resp. XOR) have fan-in two and fan-out one, and output

the AND (resp. XOR) operation on the input bits; a copy gate, denoted `copy`, simply forwards the input bit into two output wires. The depth of a circuit is defined as the longest path from an input to an output; the size of a circuit is defined as its total number of gates. Sometimes we explicitly write $\langle \Gamma \rangle$ for the description of the circuit Γ .

A circuit is clocked if it evolves in clock cycles (or rounds). The input and output values of the circuit Γ in clock cycle i are denoted by X_i and Y_i , respectively. A circuit is *probabilistic* if it uses internal randomness as part of its logic. We call such probabilistic logic *randomness gates* and denote them with $\$$. In each clock cycle $\$$ outputs a fresh random bit. Additionally, a circuit may contain memory gates. Memory gates, which have a single incoming edge and any number of outgoing edges, maintain state: at any clock cycle, a memory gate sends its current state down its outgoing edges and updates it according to the value of its incoming edge. Any cycle in the circuit graph must contain at least one memory gate. The state of all memory gates at clock cycle i is denoted by M_i , with M_1 denoting the initial state. When a circuit is run in state M_i on input X_i , the circuit will output Y_i and the memory gates will be in a new state M_{i+1} . We will denote this by $(Y_i, M_{i+1}) \leftarrow \Gamma[M_i](X_i)$.

3 Secure Circuit Fabrication

In this section we put forward a formal model for assessing security of a (cryptographic) circuit whose production is outsourced to one or more untrusted facilities. We start by recalling the standard notion of connected component of a circuit or graph.

Definition 1 (Component). *A circuit $\Gamma' = (V', E')$ is a (connected) component of circuit $\Gamma = (V, E)$ if $V' \subseteq V$, $E' \subseteq E$ and for all $g_1, g_2 \in V'$ we have that $(g_1, g_2) \in E'$ iff $(g_1, g_2) \in E$.*

Next, we introduce the notion of an outsourcing circuit compiler (or simply compiler). In a nutshell, a circuit compiler is an efficient algorithm Φ that takes as input (the description of) a circuit Γ , and outputs (the description of) a compiled circuit $\hat{\Gamma}$. Additionally, Φ returns a list of sub-components $\hat{\Gamma}_i$ of $\hat{\Gamma}$ whose production can be outsourced to one or more external manufacturers, together with the relevant information on how to connect those sub-components with the remaining ones (that need to be built in-house) in order to re-assemble the compiled circuit $\hat{\Gamma}$.

Definition 2 (Outsourcing circuit compiler). *Let Γ be an arbitrary circuit. A (ρ, m) -outsourcing compiler Φ is a PPT algorithm $(\hat{\Gamma}, \text{aux}) \leftarrow \Phi(\Gamma)$, such that the following holds:*

- $\text{aux} := ((\hat{\Gamma}_1, \dots, \hat{\Gamma}_n), \mathcal{M}, (I_1, \dots, I_m))$, with $n \in \mathbb{N}$ and $I_j \subseteq [n]$, for $j \in [m]$, mutually disjoint subsets.
- $(\hat{\Gamma}_1, \dots, \hat{\Gamma}_n)$ are disjoint (connected) components of $\hat{\Gamma}$ such that $V = \bigcup_{i \in [n]} V_i$, where $\Gamma_i = (V_i, E_i)$.
- $\mathcal{M} : V \times V \rightarrow \{0, 1\}$ is a function such that $\mathcal{M}(v, v') = 1$ iff $v, v' \in V_i, V_j$ for some $i \neq j$ and $(v, v') \in E$.

We call $\rho := \frac{\sum_{i \in [n] \setminus I_1 \cup \dots \cup I_m} |\hat{\Gamma}_i|}{|\hat{\Gamma}|}$ the outsourcing ratio of the compiler.

Intuitively, in the above definition, the outsourcing ratio ρ represents the fraction of the compiled circuit (w.r.t. the original circuit) that should be built in-house. Note that the sub-components $(\hat{\Gamma}_i)_{i \in [n]}$ “cover” the entire compiled circuit $\hat{\Gamma}$ (without overlap), and the mapping function \mathcal{M} specifies how to connect the different components in order to reconstruct $\hat{\Gamma}$. The sets of indexes $I_j \subseteq [n]$ represents the sub-components whose production will be outsourced to manufacturer $j \in [m]$. See Fig. 1 for a pictorial representation in a simple toy example.




Figure 1: On the left side we present the description of a (compiled) circuit. On the right side the same circuit is represented as three different components. The mapping function \mathcal{M} establishes the connections between the blue component and the green and red components.

Correctness of an outsourcing compiler demands that the compiled circuit maintains the same functionality of the original circuit.

Definition 3 (Correctness). *We say that an outsourcing compiler Φ is functionality preserving if for all circuits Γ , for all values of the initial memory M_1 , and for any set of public inputs X_1, \dots, X_q , the sequence of outputs Y_1, \dots, Y_q produced by running the original circuit Γ starting with state M_1 is identical to the sequence of outputs produced by running the transformed circuit $\hat{\Gamma}$ starting with state M_1 (with all but negligible probability over the randomness of the compiler and the randomness of the original and compiled circuit).*

3.1 Security

We define security using the simulation paradigm. Our approach is similar in spirit to previous work on tamper-resilient circuit compilers (see, e.g., [51, 38]). In a nutshell, security is defined by comparing two experiments. In the first experiment, also called the real experiment, the circuit designer compiles the circuit and outsources the production of some of the components in the compiled circuit to a set of m untrusted manufacturers. A subset of size t of the manufacturers are malicious, and controlled by a monolithic adversary \mathcal{A} ; of course the circuit designer does not know which manufacturers are malicious and which ones are honest. During production, \mathcal{A} is allowed to completely change the outsourced circuit components under its control, whether by adding, removing or changing gates and/or wires. Later, the designer assembles the circuit by re-combining all the components (the outsourced ones and the ones built in-house). Finally, \mathcal{A} can access the assembled circuit in a *black-box* way, that is, it can observe inputs/outputs produced by running the assembled circuit (with some initial memory M_1).

In the second experiment, also called the ideal experiment, a simulator is given black-box access to the original circuit (initialized with initial memory M_1). The goal of the simulator is to produce an output distribution which is indistinguishable from the one in the real experiment. In its most general form, our definition allows the simulator to obtain a short leakage on the initial memory. This captures the (reasonable) scenario where the adversary, in the real experiment, could learn at most a short amount of information on the private memory.

Real experiment. The distribution $\mathbf{Real}_{\mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1}(\lambda)$ is parameterized by the adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, the set of corrupt manufacturers \mathcal{C} , the compiler Φ , and the original circuit Γ with initial memory M_1 .

1. $(\widehat{\Gamma}, \mathbf{aux}) \leftarrow \Phi(\Gamma)$: In the first step, the description of the original circuit Γ is given as input to the compiler Φ ; the compiler outputs the description of the compiled circuit $\widehat{\Gamma}$ plus the auxiliary information $\mathbf{aux} := ((\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_n), \mathcal{M}, (I_1, \dots, I_m))$ which is used to specify how the compiled circuit is split into sub-components, how the different sub-components are connected (via the mapping function \mathcal{M}), and the subset of sub-components whose production is outsourced to each manufacturer (in the index sets I_j , for $j \in [m]$).
2. $(\{\widehat{\Gamma}'_i\}_{i \in I}, \tau) \leftarrow \mathcal{A}_0(1^\lambda, \{\widehat{\Gamma}_i\}_{i \in I}, \langle \Gamma \rangle, \langle \widehat{\Gamma} \rangle)$: The adversary is given as input the description of the components from the index set $I = \cup_{j \in \mathcal{C}} I_j$, the description of the original circuit Γ , the description of the compiled circuit $\widehat{\Gamma}$, and returns the modified components along with some value τ that may contain some auxiliary state information.
3. $\widehat{\Gamma}' := (\widehat{V}', \widehat{E}')$: The compiled circuit $\widehat{\Gamma}'$ is rebuilt by replacing the components $(\widehat{\Gamma}_i)_{i \in I}$ with the modified components $(\widehat{\Gamma}'_i)_{i \in I}$, and by connecting the different components as specified by the mapping function \mathcal{M} .
4. $\mathcal{A}_1^{\widehat{\Gamma}'[M_1]^{(\cdot)}}(1^\lambda, \tau)$: Adversary \mathcal{A}_1 , with auxiliary information τ , is given oracle access to the rebuilt circuit $\widehat{\Gamma}'$ with compiled private memory M_1 .

Simulation. The distribution $\mathbf{Ideal}_{\mathcal{S}, \mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1, \ell}(\lambda)$ is parameterized by the simulator \mathcal{S} , the adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, the compiler Φ , the set of corrupt manufacturers \mathcal{C} , the original circuit Γ with initial memory M_1 , and some value $\ell \in \mathbb{N}$.

1. $f \leftarrow \mathcal{S}(1^\lambda, \langle \Gamma \rangle, \Phi, \mathcal{A}, \mathcal{C}, \ell)$: Given as input a description of the original circuit, of the compiler and of the adversary, the subset of corrupt manufacturers, and the parameter $\ell \in \mathbb{N}$, the simulator specifies an arbitrary polynomial-time computable function $f : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$.
2. $\mathcal{S}^{\mathcal{A}, \Gamma[M_1]^{(\cdot)}}(1^\lambda, L)$: The simulator takes as input leakage $L = f(M_1)$, and is given oracle access to adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ and to the original circuit Γ with private memory M_1 . We remark that the simulator is restricted to be fully black-box. In particular, \mathcal{S} only accesses the modified sub-components returned by \mathcal{A}_0 in a black-box way (i.e., without knowing their description).

Definition 4 (Security). *We say that a (ρ, m) -outsourcing circuit compiler Φ is (ℓ, t) -secure if the following conditions are met.*

- (i) Non-triviality: $\rho < 1$, for sufficiently large values of $\lambda \in \mathbb{N}$.
- (ii) Simulatability: For all $\mathcal{C} \subseteq [m]$ of size at most t and for all PPT adversaries \mathcal{A} , for all circuits Γ , and for all initial values of the memory $M_1 \in \{0, 1\}^*$, there exists a simulator \mathcal{S} with running time $\text{poly}(|\mathcal{A}|, |\Gamma[M_1]|)$ such that,

$$\{\mathbf{Real}_{\mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1}(\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Ideal}_{\mathcal{S}, \mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1, \ell}(\lambda)\}_{\lambda \in \mathbb{N}}.$$

We observe that the above definition is only interesting for small values of ℓ (as, e.g., it becomes trivial in case $\ell = |M_1|$). Also notice that the non-triviality condition demands that the ratio between the size of the sub-components of the compiled circuit built in-house, and the size of the original circuit, should be less than one. This is necessary, as otherwise a manufacturer could simply produce the entire circuit by itself, without the help of any off-shore facility. Clearly, the smaller ρ is, the better, as this means that a large fraction of the original circuit production can be outsourced.

3.2 Undetectability

We formally define what it means for an adversarial strategy to be undetectable by all black-box polynomial-time tests. Informally, this means that it is hard to distinguish the output of the original circuit from the output of the compiled circuit (after the outsourced sub-components have been maliciously modified). Importantly, the latter has to hold even if the testing strategy knows the initial content of the private memory and the description of all sub-components in the compiled circuit.

For simplicity, we give the definition in the case of a single manufacturer (i.e., $m = 1$); a generalization to the case $m \geq 2$ is immediate. The formalization borrows ideas from similar definitions in the setting of subversion and algorithm-substitution attacks [12, 11].

Definition 5 (Undetectability). *Let Φ be an outsourcing circuit compiler, and Γ be a circuit. We say that an adversary \mathcal{A} is undetectable for Γ w.r.t. Φ if for all PPT algorithms Test there exists a negligible function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that, for all initial values of the memory M_1 , we have that $\mathbb{P}[\text{Test wins}] \leq 1/2 + \nu(\lambda)$ in the following game:*

1. *The challenger samples $b \leftarrow_{\$} \{0, 1\}$, runs $(\widehat{\Gamma}, (\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_n), \mathcal{M}, I) \leftarrow \Phi(\Gamma)$, and returns $(M_1, \langle \Gamma \rangle, \langle \widehat{\Gamma} \rangle, (\langle \widehat{\Gamma}_1 \rangle, \dots, \langle \widehat{\Gamma}_n \rangle), \mathcal{M}, I)$ to Test .*
2. *Let $\widehat{\Gamma}'$ be the circuit implicitly defined by the sub-components $\{\widehat{\Gamma}_i\}_{i \in [n] \setminus I} \cup \{\widehat{\Gamma}'_i\}_{i \in I}$ together with the mapping function \mathcal{M} , where $\{\widehat{\Gamma}'_i\}_{i \in I} \leftarrow \mathcal{A}(1^\lambda, \{\langle \widehat{\Gamma}_i \rangle\}_{i \in I}, \langle \Gamma \rangle, \langle \widehat{\Gamma} \rangle)$.*
3. *Algorithm Test , can ask polynomially many queries of the type X_i . Upon input such query, the answer from the challenger depends on the value of the bit b :*
 - *In case $b = 0$, the output is $(Y_i, M_{i+1}) \leftarrow_{\$} \Gamma[M_i](X_i)$.*
 - *In case $b = 1$, the output is $(Y_i, M_{i+1}) \leftarrow_{\$} \widehat{\Gamma}'[M_i](X_i)$.*
4. *Algorithm Test outputs a bit b' , and wins iff $b' = b$.*

4 Case Study I: Single Manufacturer

In this section we build secure outsourcing compilers that work for any circuit, in the presence of a *single* malicious manufacturer. The compilers are based on any verifiable computation (VC) scheme (satisfying certain properties) for the function computed by the underlying circuit. We start by recalling the basic definitions for VC schemes in Section 4.1. In Section 4.2, we describe our first compiler, which requires a VC scheme satisfying input privacy. In Section 4.3, we describe our second compiler, which can be instantiated with non-input-private VC schemes; our second compiler requires that once the first invalid output is produced, the compiled circuit overwrites its entire memory with the all-zero string and “self-destructs.” As we will show, this restriction is necessary. We refer the reader to Section 6 for concrete instantiations of our compilers based on state-of-the-art research on verifiable computation.

4.1 Prelude: Verifiable Computation

A verifiable computation scheme allows for a client to outsource the computation of a function to a (untrusted) server; the server produces a proof of correctness along with the output of the function. The client checks the correctness proof to decide whether the output provided by the server is accepted or not.

Definition 6 (Verifiable Computation). *Let \mathcal{F} be a function. A VC scheme $\mathcal{VC} = (\text{KeyGen}, \text{ProbGen}, \text{Compute}, \text{Verify})$ for function \mathcal{F} consists of the algorithms described below.*

- $(SK, PK) \leftarrow \text{KeyGen}(\mathcal{F}, \lambda)$: The (randomized) key generation algorithm takes as input the function \mathcal{F} and the security parameter λ , and outputs a public key PK and a secret key SK .
- $(\Sigma_X, VK_X) \leftarrow \text{ProbGen}_{SK}(X)$: The (randomized) problem generation algorithm takes as input the value X and uses the secret key SK to compute an encoding Σ_X of X and a secret verification key VK_X .
- $\Sigma_Y \leftarrow \text{Compute}_{PK}(\Sigma_X)$: The (deterministic) compute algorithm takes as input the encoded value Σ_X and uses the public key PK to compute an encoding of $Y = \mathcal{F}(X)$.
- $Y \leftarrow \text{Verify}_{SK}(VK_X, \Sigma_Y)$: The (deterministic) verify algorithm takes as input the verification key VK_X and the value Σ_Y ; it uses the secret key SK and VK_X to compute a value $Y \in \{0, 1\}^* \cup \{\perp\}$, where symbol \perp denotes that the algorithm rejects the value Σ_Y .

A typical VC scheme needs to satisfy some properties that we formalize below.

Correctness. A VC scheme is correct if the ProbGen algorithm produces problem instances that allow for a honest server to successfully compute values Σ_Y such that $Y = \mathcal{F}(X)$.

Definition 7 (Correctness for VC schemes). *Let \mathcal{VC} be a VC scheme for some function \mathcal{F} . We say that \mathcal{VC} is correct if for all values X , the following holds:*

$$\mathbb{P} \left[Y = \mathcal{F}(X) : \begin{array}{l} (SK, PK) \leftarrow \text{KeyGen}(\mathcal{F}, \lambda); (\Sigma_X, VK_X) \leftarrow \text{ProbGen}_{SK}(X) \\ \Sigma_Y \leftarrow \text{Compute}_{PK}(\Sigma_X); Y \leftarrow \text{Verify}_{SK}(VK_X, \Sigma_Y) \end{array} \right] = 1.$$

Soundness. A VC scheme is sound if no malicious server can “trick” a client into accepting an incorrect output, i.e, some value Y such that $Y \neq \mathcal{F}(X)$. We require this to hold even in the presence of so-called verification queries [39].

Definition 8 (Soundness for VC schemes). *Let \mathcal{VC} be a VC scheme for some function \mathcal{F} . We say that \mathcal{VC} is sound if for all PPT adversaries \mathcal{A} there exists some negligible function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that $\mathbb{P}[\mathcal{A} \text{ wins}] \leq \nu(\lambda)$ in the following game.*

1. The challenger runs $(SK, PK) \leftarrow \text{KeyGen}(\mathcal{F}, \lambda)$ to obtain the secret key SK and the public key PK , and sends PK to \mathcal{A} .
2. Adversary \mathcal{A} can make the following two types of queries to the challenger, that can be carried out polynomially many times in any order and in an adaptive way.
 - (i) Adversary \mathcal{A} can specify an input X_i ; the challenger computes $(VK_i, \Sigma_i) \leftarrow \text{ProbGen}_{SK}(X_i)$ and sends Σ_i to \mathcal{A} .
 - (ii) Adversary \mathcal{A} can specify verification queries $(i, \hat{\Sigma})$; the challenger computes $Y = \text{Verify}_{SK}(VK_i, \hat{\Sigma})$ and returns 1 if $Y \neq \perp$, otherwise returns 0.
3. Eventually, adversary \mathcal{A} will output a pair (i^*, Σ^*) ; we say that \mathcal{A} wins iff $Y^* \neq \mathcal{F}(X_{i^*})$ and $Y^* \neq \perp$, such that $Y^* = \text{Verify}_{SK}(VK_{i^*}, \Sigma^*)$.

Input privacy. A VC scheme is input-private if no server can learn the input value X that the function is being computed on.

Definition 9 (Input-privacy for VC schemes). *Let \mathcal{VC} be a VC scheme for some function \mathcal{F} . We say that \mathcal{VC} is input private if for all PPT adversaries \mathcal{A} there exists some negligible function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that $\mathbb{P}[\mathcal{A} \text{ wins}] - \frac{1}{2} \leq \nu(\lambda)$ in the following game.*

1. The challenger runs $(SK, PK) \leftarrow \text{KeyGen}(\mathcal{F}, \lambda)$ to obtain the secret key SK and the public key PK , and sends PK to \mathcal{A} .
2. Adversary \mathcal{A} can make the following two types of queries to the challenger, that can be carried out polynomially many times in any order and in an adaptive way.
 - (i) Adversary \mathcal{A} can specify an input X_i ; the challenger computes $(VK_i, \Sigma_i) \leftarrow \text{ProbGen}_{SK}(X_i)$ and sends Σ_i to \mathcal{A} .
 - (ii) Adversary \mathcal{A} can specify verification queries $(i, \hat{\Sigma})$; the challenger computes $Y = \text{Verify}_{SK}(VK_i, \hat{\Sigma})$ and returns 1 if $Y \neq \perp$, otherwise returns 0.
3. Adversary \mathcal{A} chooses two values X_0 and X_1 and sends them to the challenger.
4. The challenger samples a random bit $b \leftarrow_{\$} \{0, 1\}$ and computes $(VK^*, \Sigma^*) \leftarrow \text{ProbGen}_{SK}(X_b)$ forwarding Σ^* to \mathcal{A} .
5. Adversary \mathcal{A} can still specify to the challenger the queries described above, including special verification queries for the verification key VK^* .
6. Finally, \mathcal{A} outputs a bit b' ; we say that \mathcal{A} wins if and only if $b = b'$.

Outsourceability. A VC scheme is outsourceable if the time to encode the input plus the time to run a verification is smaller than the time to compute the function itself.

Definition 10 (Outsourceability for VC schemes). *A VC scheme can be outsourced if it allows efficient generation and efficient verification. This means that for any X and any Σ_Y the time (or circuit size) required for $\text{ProbGen}_{SK}(X)$ plus the time (or circuit size) required for $\text{Verify}_{SK}(VK, \Sigma_Y)$ is $o(T)$, where T is the time (or circuit size) required to compute $\mathcal{F}(X)$.*

VC without input-privacy. The above definitions can be adapted to cast VC schemes without input-privacy, i.e schemes where the server is allowed to learn the input value X used by the client. For such VC schemes, algorithm ProbGen returns the value X in the clear along with the secret verification key VK_X ; the correctness, soundness and outsourceability definitions can easily be adapted to this setting.

4.2 Compiler based on Input-Private VC

In this section we construct an outsourcing circuit compiler by using a VC scheme that satisfies the properties of correctness, soundness, input-privacy and outsourceability. Let Γ be a circuit; the idea is to invoke a VC scheme for the function \mathcal{F} corresponding to the functionality computed by Γ . The compiled circuit will consist of four main components $\hat{\Gamma}_{\text{ProbGen}}$, $\hat{\Gamma}_{\text{Compute}}$, $\hat{\Gamma}_{\text{Verify}}$, and $\hat{\Gamma}_{\mathcal{S}}$. The first three components are the circuit representations of the algorithms ProbGen , Compute and Verify corresponding to the underlying VC scheme; such components hard-wire keys (SK, PK) generated using algorithm KeyGen . The fourth component samples the random coins R_i to be used during each invocation of the circuit.

The production of component $\hat{\Gamma}_{\text{Compute}}$ will then be outsourced to a single untrusted facility, whereas all other components are built in-house (as their implementation needs to be trusted). Notice that the implementation of algorithm KeyGen can be thought of as a pre-processing stage that runs only once (and could be carried out in software).

An important observation is that the size of circuit $\hat{\Gamma}_{\text{Verify}}$ and $\hat{\Gamma}_{\text{ProbGen}}$ is independent, and much smaller, than the size of circuit $\hat{\Gamma}_{\text{Compute}}$. As discussed in the introduction, the size of $\hat{\Gamma}_{\mathcal{S}}$ can also be considered to be constant (consisting only of a few gates). We describe our first compiler below in more details.

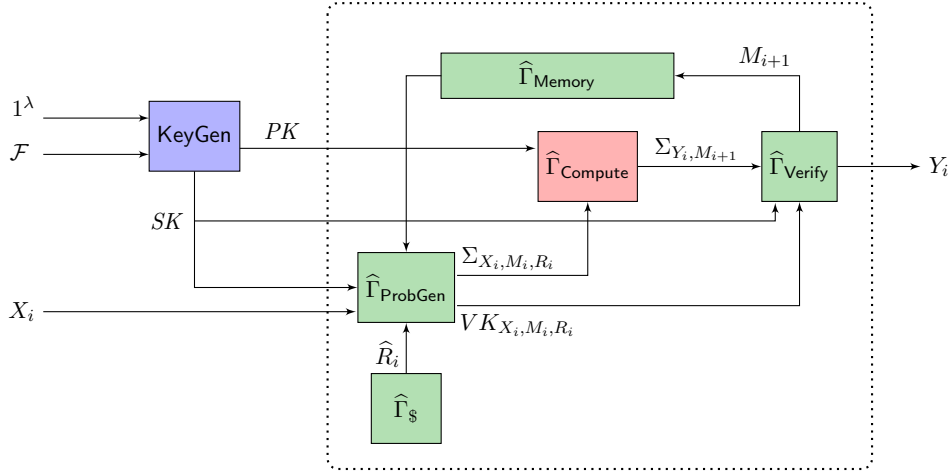


Figure 2: The description of compilers $\Phi_{\mathcal{VC}}^1$. The green components (i.e., $\widehat{\Gamma}_{\text{ProbGen}}$, $\widehat{\Gamma}_{\text{Verify}}$, and $\widehat{\Gamma}_{\S}$) need to be built in-house, while the production of the red component (i.e., $\widehat{\Gamma}_{\text{Compute}}$) can be outsourced; the blue component (i.e., KeyGen) is built only once (not necessarily in hardware). The dotted line depicts the circuit boundaries. The dotted line depicts the circuit boundaries.

The compiler $\Phi_{\mathcal{VC}}^1$. Let Γ be a circuit, and $\mathcal{VC} = (\text{KeyGen}, \text{ProbGen}, \text{Compute}, \text{Verify})$ be a VC scheme for the function \mathcal{F} implemented by Γ . Our first compiler is depicted in Fig. 2, and can be described as follows.

1. First run $(SK, PK) \leftarrow \text{KeyGen}(\mathcal{F}, \lambda)$ once, obtaining the pair of keys (SK, PK) .
2. Let $\widehat{\Gamma}_{\text{Memory}}$ be a circuit component consisting only of memory gates, as needed by the original circuit Γ , storing the initial value of the private memory M_1 .
3. Let $\widehat{\Gamma}_{\S}$ be a circuit outputting random coins \widehat{R}_i (as needed in each invocation of the compiled circuit).
4. Define a component for each function ProbGen , Compute and Verify of the \mathcal{VC} scheme as explained below.
 - $\widehat{\Gamma}_{\text{ProbGen}}$: This component embeds the secret key SK , and it takes three inputs; the input X_i , the (current) private memory M_i , and random coins $\widehat{R}_i := R_i || R'_i$. It implements function $\text{ProbGen}_{SK}(X_i || M_i || R_i; R'_i)$, that produces two outputs: an encoding Σ_{X_i, M_i, R_i} , and a verification key VK_{X_i, M_i, R_i} .
 - $\widehat{\Gamma}_{\text{Compute}}$: This component embeds the public key PK , and it takes as input the encoding Σ_{X_i, M_i, R_i} . It implements function $\text{Compute}_{PK}(\Sigma_{X_i, M_i, R_i})$, that produces the encoding $\Sigma_{Y_i, M_{i+1}}$ of $(Y_i, M_{i+1}) = \mathcal{F}(X_i, M_i; R_i)$ as output.
 - $\widehat{\Gamma}_{\text{Verify}}$: This component embeds the secret key SK , and it takes two inputs; the encoding $\Sigma_{Y_i, M_{i+1}}$ and the verification key VK_{X_i, M_i, R_i} . It implements function $\text{Verify}_{SK}(VK_{X_i, M_i, R_i}, \Sigma_{Y_i, M_{i+1}})$, to produce the output $Y_i \in \{0, 1\}^* \cup \{\perp\}$, and eventually update the circuit private memory to M_{i+1} .
5. The output of $\Phi_{\mathcal{VC}}^1$ is defined as follows. The first output is a (description of the) compiled circuit $\widehat{\Gamma}$ as depicted in Fig. 2. The auxiliary information aux consists of the components $\widehat{\Gamma}_{\text{ProbGen}}$, $\widehat{\Gamma}_{\text{Compute}}$, $\widehat{\Gamma}_{\text{Verify}}$, $\widehat{\Gamma}_{\text{Memory}}$, and $\widehat{\Gamma}_{\S}$, the mapping function \mathcal{M} that describes the physical connections between such components (i.e., the arrows in Fig. 2), and the index set $I = \{2\}$ specifying the component $\widehat{\Gamma}_{\text{Compute}}$ as a candidate for outsourcing.

Remark 1 (On outsourcing memory gates.). *In the compiler depicted in Fig. 2, $\widehat{\Gamma}_{\text{Memory}}$ is being built in-house. In order to outsource private memory to a potentially malicious manufacturer we*

can modify the above compiler as follows: instead of storing in $\widehat{\Gamma}_{\text{Memory}}$ the value M_i in plaintext, we store $C \leftarrow \text{AE}_{SK'}(M_i)$, where C is the encryption of M_i using a symmetric, semantically secure authenticated encryption scheme, with secret key SK' . Moreover, $\widehat{\Gamma}_{\text{ProbGen}}$ is modified such that when receiving the private memory value C , it first decrypts it using SK' and then executes the original circuit $\widehat{\Gamma}_{\text{ProbGen}}$ on the resulting plaintext. We also substitute $\widehat{\Gamma}_{\text{Verify}}$ so that it outputs the encryption of M_{i+1} , under SK' . This modification enables the simulator to execute the circuit using the all-zeros bit-string as the initial memory value, and security follows by the semantic security of the encryption scheme. Finally, whenever the decryption of C gives \perp the circuit output is \perp .

The theorem below states that the compiler from Fig. 2 satisfies our strongest security notion (i.e., Definition 4 with $\ell = 0$), provided that the underlying VC scheme is correct, sound, input-private, and outsourceable.

Theorem 1. *Let Γ be an arbitrary circuit and let \mathcal{VC} be a verifiable computation scheme for the function \mathcal{F} computed by Γ , satisfying the properties of correctness, soundness, input-privacy and outsourceability. Then the compiler $\Phi_{\mathcal{VC}}^1$ is a correct, $(0, 1)$ -secure $(o(1), 1)$ -outsourcing circuit compiler.*

Proof. We start by showing that the outsourcing ratio parameter ρ of the compiler $\Phi_{\mathcal{VC}}^1$ is always smaller than 1, for sufficiently large values of the security parameter λ , thus meeting the non-triviality condition.

Claim 1. $\rho = o(1)$.

Proof. The non-triviality requirement from Definition 4 states that

$$\rho = \frac{|\widehat{\Gamma}_{\text{ProbGen}}| + |\widehat{\Gamma}_{\text{Verify}}| + |\widehat{\Gamma}_{\S}| + |\widehat{\Gamma}_{\text{Memory}}|}{|\Gamma|} < 1.$$

The claim is that $\rho = o(1)$. By the outsourceability property of the \mathcal{VC} scheme we know that $|\widehat{\Gamma}_{\text{ProbGen}}| + |\widehat{\Gamma}_{\text{Verify}}| = o(|\Gamma|)$, i.e., $\lim_{\lambda \rightarrow \infty} \frac{|\widehat{\Gamma}_{\text{ProbGen}}| + |\widehat{\Gamma}_{\text{Verify}}|}{|\Gamma|} = 0$. By substituting the expression for ρ , and by using the fact that the size of $\widehat{\Gamma}_{\S}$ can be assumed to be constant and that the original and the compiled circuit contain the same number of memory gates, we obtain:

$$\lim_{\lambda \rightarrow \infty} \frac{(|\widehat{\Gamma}_{\text{ProbGen}}| + |\widehat{\Gamma}_{\text{Verify}}| + |\widehat{\Gamma}_{\S}| + |\widehat{\Gamma}_{\text{Memory}}|)/|\Gamma|}{1} = 0.$$

This shows that function ρ converges to 0. As the size of each component is monotonously increasing with the security parameter, for sufficiently large λ , the outsourcing ratio ρ will always be smaller than 1, as desired. \square

Claim 2. *The compiler $\Phi_{\mathcal{VC}}^1$ satisfies correctness.*

Proof. The correctness of the compiler $\Phi_{\mathcal{VC}}^1$ follows immediately from the correctness property of the underlying VC scheme. \square

We proceed to prove security of $\Phi_{\mathcal{VC}}^1$. We need to build a simulator \mathcal{S} that is able to “fake” experiment **Real** for all adversaries \mathcal{A} , for all circuits Γ , and for all initial memory values M_1 . A description of the simulator follows.

- Run the compiler $\Phi_{\mathcal{VC}}^1(\Gamma)$ supplying (a description of) the original circuit Γ ; the output is (a description of) the compiled circuit $\widehat{\Gamma}$, and the auxiliary information $\text{aux} := ((\widehat{\Gamma}_{\text{ProbGen}}, \widehat{\Gamma}_{\text{Compute}}, \widehat{\Gamma}_{\text{Verify}}, \widehat{\Gamma}_{\text{Memory}}, \widehat{\Gamma}_{\S}), \mathcal{M}, \{2\})$.

- The description of component $\widehat{\Gamma}_{\text{Compute}}$ is sent to adversary \mathcal{A} together with the descriptions of Γ and $\widehat{\Gamma}$.
- Adversary \mathcal{A} produces the component $\widehat{\Gamma}'_{\text{Compute}}$ (that may be malicious) and sends it to the simulator; the circuit $\widehat{\Gamma}'$ is assembled using the components $(\widehat{\Gamma}_{\text{ProbGen}}, \widehat{\Gamma}'_{\text{Compute}}, \widehat{\Gamma}_{\text{Verify}}, \widehat{\Gamma}_{\text{Memory}}, \widehat{\Gamma}_{\mathfrak{S}})$, via the mapping function \mathcal{M} .
- Upon input X_i from \mathcal{A} , run $(Y_i, \widetilde{M}_{i+1}) \leftarrow \widehat{\Gamma}'[0^\mu](X_i)$, where $\mu := |M_1|$; if $Y_i = \perp$ then forward \perp to \mathcal{A} , otherwise query X_i to oracle $\Gamma[M_1](\cdot)$ and forward the output to \mathcal{A} .

Let us write \mathbf{R} and \mathbf{S} for the distribution of the random variables in experiment **Real** and **Ideal** of Definition 4. Recall that these variables are parametrized by adversary \mathcal{A} , simulator \mathcal{S} , initial memory M_1 , compiler $\Phi_{\mathcal{VC}}^1$, and circuit Γ , but we omit explicitly writing all these parameters to simplify the exposition.³ We consider a new experiment \mathbf{R}' that is exactly the same as \mathbf{R} , except that for all queries X_i such that $(Y'_i, M_{i+1}) \leftarrow \widehat{\Gamma}'[M_i](X_i)$ with $Y'_i \neq \perp$, the experiment computes $(Y_i, M_{i+1}) \leftarrow \mathcal{F}(X_i, M_i)$ and outputs Y_i .

The claim below shows that experiment \mathbf{R} and \mathbf{R}' are computationally indistinguishable.

Claim 3. *For all PPT adversaries \mathcal{A} , and for all PPT distinguishers \mathcal{D} , there exists a negligible function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that $\Delta^{\mathcal{D}}(\mathbf{R}; \mathbf{R}') \leq \nu(\lambda)$.*

Proof. Let $q \in \text{poly}(\lambda)$ be the number of input queries asked by \mathcal{A} to its oracle, and let us define an event W that becomes true whenever there is some index $i^* \in [q]$ such that in experiment \mathbf{R}' we have $Y_{i^*} \neq Y'_{i^*}$. Then, by Lemma 1, we have that $\Delta^{\mathcal{D}}(\mathbf{R}; \mathbf{R}') \leq \Delta(\mathbf{R}; \mathbf{R}' | \neg W) + \mathbb{P}[W]$. Note that the distributions \mathbf{R} and \mathbf{R}' , conditioned on event W not happening are exactly the same. We proceed to show that the probability of event W is negligible.

Let us assume that there exists some circuit Γ , some initial memory M_1 and a PPT adversary \mathcal{A} provoking event W with non-negligible probability. We build a PPT adversary \mathcal{A}' that uses \mathcal{A} in order to break the soundness property of the underlying VC scheme (cf. Definition 8).

Adversary $\mathcal{A}'(M_1, \mathcal{F}, \Gamma, \Phi)$:

1. Receive PK from the challenger, where $(SK, PK) \leftarrow \text{KeyGen}(\mathcal{F}, \lambda)$.
2. Run $(\widehat{\Gamma}, M_1, \text{aux}) \leftarrow \Phi(\Gamma, M_1)$, and send (the description of) circuit components $(\widehat{\Gamma}_{\text{Compute}}, \Gamma, \widehat{\Gamma})$ to \mathcal{A} . Adversary \mathcal{A} produces the modified component $\widehat{\Gamma}'_{\text{Compute}}$ and sends it to \mathcal{A}' .
3. For all $i \in [q]$, upon input query X_i from \mathcal{A} , do the following.
 - Sample random coins R_i , create the string $Z_i := X_i || M_i || R_i$ and send it to the challenger as an encoding query. In response, the challenger computes $(VK_{Z_i}, \Sigma_{Z_i}) \leftarrow \text{ProbGen}_{SK}(Z_i)$ and sends back the encoding Σ_{Z_i} to \mathcal{A}' .
 - Run $\Sigma_{Y'_i, M_{i+1}} = \widehat{\Gamma}'_{\text{Compute}}(PK, \Sigma_{Z_i})$ and send a query $(i, \Sigma_{Y'_i, M_{i+1}})$ to the challenger; the answer from the challenger is a verification bit d . If $d = 0$, then reply with \perp , otherwise compute $(Y_i, M_{i+1}) = \mathcal{F}(M_i, X_i; R_i)$ and reply with Y_i .
4. Finally, pick a random $i^* \in [q]$ and output the pair $(i^*, \Sigma_{Y'_{i^*}})$.

For the analysis, we note that the above simulation is perfect. In particular the outputs seen by \mathcal{A} retain exactly the same distribution as in experiment \mathbf{R}' . Now \mathcal{A} provokes event W with non-negligible probability, so there exists some index $i^* \in [q]$ such that Y_{i^*} is different from Y'_{i^*} where $(Y'_{i^*}, M_{i^*+1}) \leftarrow \widehat{\Gamma}'[M_{i^*}](X_{i^*})$. Since \mathcal{A}' guesses the right index i^* with probability

³We also omit to mention the set of corrupt manufacturers \mathcal{C} , as in this case there is a single malicious manufacturer.

$1/q$, we obtain that \mathcal{A}' wins the soundness game with non-negligible probability $\frac{1}{q} \cdot \mathbb{P}[W]$. This contradicts the soundness property of the underlying VC scheme, and thus concludes the proof of the claim. \square

Claim 4. *For all PPT adversaries \mathcal{A} , and for all PPT distinguishers \mathcal{D} , there exists a negligible function $\nu' : \mathbb{N} \rightarrow [0, 1]$ such that $\Delta^{\mathcal{D}}(\mathbf{R}'; \mathbf{S}) \leq \nu'(\lambda)$.*

Proof. Let $q \in \text{poly}(\lambda)$ be the number of input queries asked by \mathcal{A} to its oracle. For an index $i \in [q]$ consider the hybrid experiment \mathbf{H}_i that answers the first i queries as in \mathbf{R}' and all the subsequent queries as in \mathbf{S} . We note that experiments \mathbf{H}_{i-1} and \mathbf{H}_i only differ in how the output is computed in position i , and that $\mathbf{H}_0 \equiv \mathbf{S}$ and $\mathbf{H}_q \equiv \mathbf{R}'$.

We now show that for all circuits Γ , all initial memories M_1 , and all PPT adversaries \mathcal{A} , each pair of adjacent hybrids \mathbf{H}_{i-1} and \mathbf{H}_i are computationally indistinguishable. Fix some $i \in [0, q]$, and assume that there exists some circuit Γ , some initial memory M_1 , a PPT adversary \mathcal{A} , and a PPT distinguisher \mathcal{D} that distinguishes between the pair of hybrids \mathbf{H}_{i-1} and \mathbf{H}_i . We build an adversary \mathcal{A}' that uses $(\mathcal{A}, \mathcal{D})$ in order to break the input-privacy property of the underlying VC scheme (cf. Definition 9).

Adversary $\mathcal{A}'(M_1, \mathcal{F}, \Gamma, \Phi, i)$:

1. Receive PK from the challenger, where $(SK, PK) \leftarrow \text{KeyGen}(\mathcal{F}, \lambda)$.
2. Run $(\widehat{\Gamma}, \text{aux}) \leftarrow \Phi(\Gamma)$, and send (the description of) circuit components $(\widehat{\Gamma}_{\text{Compute}}, \Gamma, \widehat{\Gamma})$ to \mathcal{A} . Adversary \mathcal{A} produces the modified component $\widehat{\Gamma}'_{\text{Compute}}$ and sends it to \mathcal{A}' .
3. Upon input query X_j from \mathcal{D} , such that $j \neq i$, answer as follows:
 - Sample random coins R_j . If $j < i$ create the string $Z_j := X_j || M_j || R_j$, else create the string $Z_j := X_j || 0^\mu || R_j$.
 - Send Z_j to the challenger as an encoding query. In response the challenger computes $(VK_{Z_j}, \Sigma_{Z_j}) \leftarrow \text{ProbGen}_{SK}(Z_j)$ and sends Σ_{Z_j} to \mathcal{A}' .
 - Run $\Sigma_{Y_j, M_{j+1}} = \widehat{\Gamma}'_{\text{Compute}}(PK, \Sigma_{Z_j})$, and send a verification query $(j, \Sigma_{Y_j, M_{j+1}})$ to the challenger.
 - The challenger replies with a decision bit d ; if $d = 0$ then return \perp , otherwise compute $(Y_j, M_{j+1}) = \mathcal{F}(M_j, X_j; R_j)$ and return Y_j .
4. Upon input query X_i from \mathcal{A} , answer as follows:
 - Sample random coins R_i , and create two strings $Z_0^* := X_i || M_i || R_i$ and $Z_1^* := X_i || 0^\mu || R_i$.
 - Send Z_0^* and Z_1^* to the challenger as challenge inputs for the input-privacy game; the challenger replies with Σ_{Z^*} (which corresponds to the encoding of either Z_0^* or Z_1^*).
 - Run $\Sigma_{Y^*} = \widehat{\Gamma}'_{\text{Compute}}(PK, \Sigma_{Z^*})$ and send a verification query (i, Σ_{Y^*}) to the challenger; the challenger replies with a decision bit d . If $d = 0$ then return \perp , otherwise compute $(Y_i, M_{i+1}) = \mathcal{F}(M_i, X_i; R_i)$ and return Y_i .
5. Finally, output whatever \mathcal{D} outputs.

For the analysis, we note that the above simulation is perfect. In particular, depending on Z^* being either an encoding of Z_0^* or Z_1^* the view of $(\mathcal{A}, \mathcal{D})$ is identical to the view in either experiment \mathbf{H}_{i-1} or \mathbf{H}_i . Hence, \mathcal{A}' retains the same advantage as $(\mathcal{A}, \mathcal{D})$ which contradicts input-privacy of the underlying VC scheme. We conclude that there exist negligible functions

$\nu', \nu'' : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Delta^{\mathcal{D}}(\mathbf{R}'; \mathbf{S}) \leq \sum_{i=1}^q \Delta^{\mathcal{D}}(\mathbf{H}_{i-1}; \mathbf{H}_i) \leq q \cdot \nu'(\lambda) \leq \nu''(\lambda),$$

as desired. \square

The statement now follows by Claim 3 and Claim 4, and by the triangle inequality, as

$$\Delta^{\mathcal{D}}(\mathbf{R}; \mathbf{S}) \leq \Delta^{\mathcal{D}}(\mathbf{R}; \mathbf{R}') + \Delta^{\mathcal{D}}(\mathbf{R}'; \mathbf{S}) \leq \nu(\lambda) + \nu'(\lambda).$$

This finishes the proof. \square

4.3 Compiler based on VC without Input Privacy

In this section, we construct an outsourcing circuit compiler by using any VC scheme that satisfies the properties of correctness, soundness and outsourceability. The construction follows the same ideas of compiler $\Phi_{\mathcal{VC}}^1$ (cf. Section 4.2), with two main differences. First, as we rely on a VC scheme without input-privacy, the component $\widehat{\Gamma}_{\text{ProbGen}}$ now outputs the values X_i, M_i, R_i in the clear. Second, the component $\widehat{\Gamma}_{\text{Verify}}$ needs to implement a special “self-destruct” feature: The first time the component returns the special symbol \perp , the private memory is overwritten with the all-zero string.

As we argue later in this section, the self-destruct feature is *necessary*, in that, without such a feature, generic attacks against our compiler are possible, possibly exposing the entire private memory in an undetectable manner.

The compiler $\Phi_{\mathcal{VC}}^2$. Let Γ be a circuit, and $\mathcal{VC} = (\text{KeyGen}, \text{ProbGen}, \text{Compute}, \text{Verify})$ be a VC scheme for the function \mathcal{F} implemented by Γ . The description of the compiler $\Phi_{\mathcal{VC}}^2(\Gamma)$ can be found in Fig. 3. The case of memory outsourcing is identical to that of the first compiler and we refer the reader to Remark 1. The theorem below establishes that such a compiler is secure, provided that the original circuit to be produced is resilient to a logarithmic (in the security parameter) amount of leakage on its private memory.

Theorem 2. *Let Γ be an arbitrary circuit and let \mathcal{VC} be a verifiable computation scheme for the function \mathcal{F} computed by Γ , satisfying the properties of correctness, soundness and outsourceability. Then the compiler $\Phi_{\mathcal{VC}}^2$ is a correct $(\log(q) + 1, 1)$ -secure $(o(1), 1)$ -outsourcing circuit compiler, where q is the number of oracle queries asked by adversary \mathcal{A} in Definition 4.*

Proof idea. We give an intuition for the security proof. As for our first compiler, correctness of the compiler and the fact that $\rho = o(1)$ follow readily from the correctness and the outsourceability properties of the underlying VC scheme. As for security, we need to build a simulator \mathcal{S} that is able to “fake” the real experiment for all adversaries \mathcal{A} , for all circuits Γ , and for all initial memory values M_1 . The simulator \mathcal{S} is allowed to define an arbitrary polynomial-time computable function $f : \{0, 1\}^* \rightarrow \{0, 1\}^{\log(q)+1}$ that leaks a logarithmic amount of information regarding the initial private memory M_1 . Intuitively, the function f hard-wires the randomness for running the entire real experiment with adversary \mathcal{A} ; note that this randomness includes the randomness for the compiler, and the random coin tosses of \mathcal{A} . The output of the function consists of: (i) a predicate `abort` which tells whether there was some round where the output of circuit $\widehat{\Gamma}'$ differs from the output of the original circuit Γ ; (ii) the index q^* corresponding to the round where self-destruct was triggered.

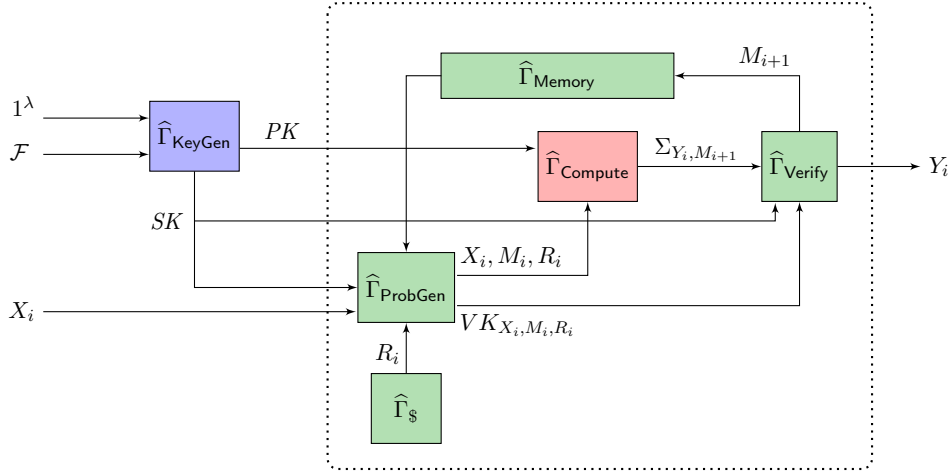


Figure 3: The description of compiler $\Phi_{\mathcal{VC}}^2$. Notice that component $\widehat{\Gamma}_{\text{ProbGen}}$ does not need to hide its input, and that component $\widehat{\Gamma}_{\text{Verify}}$ implements the self-destruct feature. The dotted line depicts the circuit boundaries.

Let Γ be any circuit with input size of n bits, and consider the compiled circuit $\widehat{\Gamma}$ produced by running $\Phi_{\mathcal{VC}}^2(\Gamma)$. Define the following adversarial strategy, aimed at building a modified circuit component $\widehat{\Gamma}'_{\text{Compute}}$ that leaks the entire initial private memory M_1 when given some secret trapdoor information.

$\mathcal{A}^*(X_i, M_i, R_i)$:

1. Choose a random value τ , such that $|\tau| = n$, and store it in a memory location.
2. Upon the first run, duplicate and store the content of memory M_1 into another memory location, and initialize a counter $j := 0$.
3. Upon input a tuple (X_i, M_i, R_i) such that $X_i \neq \tau$, output $\Sigma_{Y_i} = \text{Compute}_{PK}(X_i, M_i, R_i)$. Otherwise, behave as follows:
 - If $M_1[j] = 0$, let Σ_{Y_i} be the all-zero string.
 - Else, compute $\Sigma_{Y_i} = \text{Compute}(X_i, M_i, R_i)$.
 - Update counter $j \leftarrow (j + 1) \bmod |M_1|$ and output Σ_{Y_i} .

Figure 4: Undetectable attack against compiler $\Phi_{\mathcal{VC}}^2$ (without the self-destruct capability).

In case $\text{abort} = 1$ the simulator \mathcal{S} simply gives up and aborts. Otherwise, it emulates the real experiment with the same randomness hard-wired in the function f , answering all the queries from \mathcal{A} up until round $q^* - 1$ using black-box access to the original circuit, whereas all further queries are answered with \perp . On the one hand, this is a perfect simulation as long as $\text{abort} = 0$. On the other hand, the soundness property of the underlying VC scheme ensures that $\text{abort} = 0$ with all but a negligible probability. The proof follows.

Necessity of self-destruct. We show that the self-destruct feature is *necessary* for the security of compiler $\Phi_{\mathcal{VC}}^2$ by presenting an *undetectable* attack (as per Definition 5) against the circuit $\widehat{\Gamma}'$ produced by compiler $\Phi_{\mathcal{VC}}^2$. The attack, which is described in details in Fig. 4, works for a large class of circuits, and leaks the entire initial private memory M_1 embedded in the compiled circuit.

Theorem 3. Let Γ be any circuit with input size $n = \omega(\log \lambda)$ and let Φ_{VC}^2 be the compiler from Fig. 3 without the self-destruct capability. Then, the attack \mathcal{A}^* described in Fig. 4 is undetectable for Γ w.r.t. Φ_{VC}^2 , and leaks the entire initial private memory M_1 .

The proof is given in the full version of the paper.

Proof. The second part of the statement follows directly by observing that knowledge of the trapdoor information τ allows to learn the value M_1 with overwhelming probability.

We proceed to show undetectability. Let \mathbf{G} be the undetectability game described in Definition 5, where the adversary \mathcal{A} is chosen to be adversary \mathcal{A}^* from Fig. 4. Consider the game \mathbf{G}_0 , an identical copy of game \mathbf{G} when $b = 0$, and consider the game \mathbf{G}_1 an identical copy of game \mathbf{G} when $b = 1$. Abusing notation, let us write \mathbf{G}_0 and \mathbf{G}_1 for the distribution of the random variables corresponding to algorithm `Test`'s view in games \mathbf{G}_0 and game \mathbf{G}_1 respectively. For an index $i \in [q]$ consider the hybrid game \mathbf{H}_i that answers the first i queries as in \mathbf{G}_0 and all the subsequent queries as in \mathbf{G}_1 . We note that game \mathbf{H}_{i-1} and \mathbf{H}_i only differ in position i , and that $\mathbf{H}_0 \equiv \mathbf{G}_1$ and $\mathbf{H}_q \equiv \mathbf{G}_0$.

We claim that, for all $i \in [q]$, it holds $\mathbf{H}_{i-1} \approx_c \mathbf{H}_i$. Fix some index i , and define the event W that the i -th query X_i happens to be equal to the secret value τ embedded in the modified component $\widehat{\Gamma}_{\text{Compute}}$ (as described in Fig. 4). By Lemma 1, $\Delta(\mathbf{G}_0; \mathbf{G}_1) \leq \Delta(\mathbf{G}_0; \mathbf{G}_1 | \neg W) + \mathbb{P}[W]$. Clearly, conditioned on event W not happening, the distributions of game \mathbf{G}_0 and \mathbf{G}_1 are identical; this is because in such a case the modified component $\widehat{\Gamma}'_{\text{Compute}}$ behaves exactly like the original component $\widehat{\Gamma}_{\text{Compute}}$.

On the other hand, if $|\tau| = \omega(\log \lambda)$, the probability of event W is negligible. We conclude that for all PPT distinguishers \mathcal{D} there exists a negligible function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Delta^{\mathcal{D}}(\mathbf{G}_0; \mathbf{G}_1) \leq \sum_{i=1}^q \Delta^{\mathcal{D}}(\mathbf{H}_{i-1}; \mathbf{H}_i) \leq \nu(\lambda).$$

□

5 Case Study II: Multiple Manufacturers

In this section we focus on outsourcing compilers in the presence of multiple manufacturers, aiming to improve the efficiency of the resulting circuit at the expense of achieving security in the weaker model where there are $m \geq 2$ manufacturers, a t -fraction of which is malicious (for some threshold $t \leq m - 1$). Our first solution, described in Section 5.1, is based on client-server multi-party computation protocols; our second solution, described in Section 5.2, is based on multi-server VC protocols.

5.1 Compiler based on MPC

Before presenting the compiler, we first revisit the core ideas of MPC, and then we give a generic definition for MPC protocols in the *client-server* model, along the lines of [8].

5.1.1 MPC in the Client-Server Model

In MPC we consider p parties, where each party P_i , for $1 \leq i \leq p$, possesses an input X_i and they all wish to jointly compute the tuple $(Y_1, \dots, Y_p) = \mathcal{F}(X_1, \dots, X_p)$, where P_i receives Y_i . In the client-server model, the parties are divided into two categories: the parties that provide inputs and wish to receive the output of the computation (clients), and those performing the

computation (servers). A t -private MPC protocol ensures that any adversary who controls up to t servers cannot leak any information related to the private inputs of the clients, besides the information that can be inferred by inspecting the output of the computation, and regardless of the number of corrupted clients. In our compiler, the circuit corresponding to the code executed by the servers will be outsourced to a number of possibly malicious manufacturers, that may apply arbitrary modifications against the circuit components. Thus, we require MPC protocols that are secure against active (malicious) attackers.

The general idea behind the compiler is the following. Let Γ be a circuit implementing some functionality \mathcal{F} , and let $\Pi_{\mathcal{F}}$ be a t -private MPC protocol realizing the function \mathcal{F} . Then, assuming the number of malicious manufacturers is at most $t < m$, the circuit $\widehat{\Gamma}$ will implement the code of $\Pi_{\mathcal{F}}$, and each $\widehat{\Gamma}_i$ will implement the code of the i -th server. Below, we define the protocol framework that we are going to use for the rest of this section. The idea is to describe any MPC protocol using its next message function, denoted as `Next`.

Definition 11 (r -round protocols). *Let C, S be sets of probabilistic interactive Turing machines, with cardinalities p, m , respectively. An r -round protocol Π for p clients and m servers is a tuple $(C, S, \text{Enc}, \text{Dec}, \text{Next})$, where $\text{Next} = (\text{Next}_1, \dots, \text{Next}_m)$, described as follows.*

- Setup: Each client computes $(X_i^1, \dots, X_i^m) \leftarrow \text{Enc}(X_i)$, and sends X_i^j to the server indexed by j . Let $\mathbf{in}^j := (X_1^j, \dots, X_p^j)$, and $\tau_j := 0$ (we assume that the network is fully connected, still the properties of the communication channel depend on the instantiation).
- Computation: For $i \in [r]$:
 - If $i \neq r$, for $j \in [m]$ execute $(o_1^j, \dots, o_m^j, \tau_j') \leftarrow \text{Next}_j(\mathbf{in}^j, \tau_j)$, send o_k^j , $k \neq j$, to the server with index k . Set $\mathbf{in}^j = (o_1^j, \dots, o_m^j)$, and $\tau_j = \tau_j'$.
 - If $i = r$, for $j \in [m]$ execute $o^j \leftarrow \text{Next}_j(\mathbf{in}^j, \tau_j)$, and send o_j to `Dec`.
- Output: Execute $(Y_1, \dots, Y_p) \leftarrow \text{Dec}(o_1, \dots, o_m)$, and send Y_j to the client with index j .

For any function \mathcal{F} , the protocol computing \mathcal{F} will be denoted by $\Pi_{\mathcal{F}}$.

Informally, in the first step of the protocol execution, the clients encode their inputs, as it is prescribed by `Enc`, and then the main computation begins. The code executed by the servers at each round is defined by the function `Next` (the next message function). Hence, in the i -th round, server S_j computes `Nextj` upon the outputs and the state information τ produced by the other servers in round $i - 1$. One can also consider deterministic next message functions, assuming the randomness is given as input in each round. Below, we formally define *correctness* and *privacy* for MPC protocols.

Definition 12 (Correctness). *Let \mathcal{F} be a p -party functionality. We say that Π realizes \mathcal{F} with perfect (resp., statistical) correctness if for any input (X_1, \dots, X_p) , the probability that the output delivered to the i -th client during the protocol execution, is different from Y_i , is 0 (resp., negligible in λ), where $(Y_1, \dots, Y_p) = \mathcal{F}(X_1, \dots, X_p)$ for $i \in [p]$.*

Definition 13 ((t, m) -privacy). *Let λ be the security parameter, p be the number of parties (clients) and m be the number of servers, and let \mathcal{A} be an adversary that may corrupt any set of parties $I_c \subseteq [p]$, and servers $I_s \subseteq [m]$, where $|I_s| \leq t$. We say that the protocol Π realizes \mathcal{F} with (t, m) -privacy if there exists a PPT algorithm \mathcal{S} such that for all sufficiently large $\lambda \in \mathbb{N}$,*

$$\text{View}_{I_s, I_c}(\lambda, X_1, \dots, X_p) \approx_c \mathcal{S}(1^\lambda, I_c, I_s, (X_i, Y_i)_{i \in I_c})$$

where $\text{View}_{I_s, I_c}(\lambda, X_1, \dots, X_p)$ denotes the joint view of the servers and clients in I_s and I_c , respectively, within an execution of the protocol upon inputs X_1, \dots, X_p , and $(Y_1, \dots, Y_p) = \mathcal{F}(X_1, \dots, X_p)$.

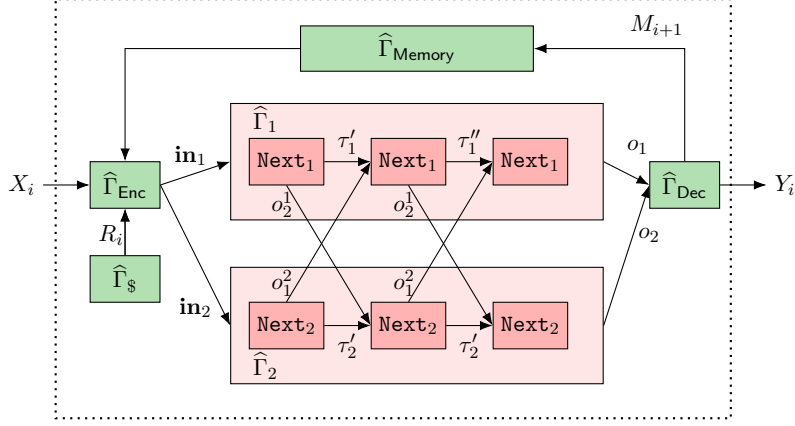


Figure 5: The MPC compiler for the case of $m = 2$ outsourcing facilities (from which *one* can be malicious). The components $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$ can be outsourced, while the connectivity between them and the remaining components are built in-house. The dotted line depicts the circuit boundaries.

The main idea behind the above definition is that the view of the attacker during the protocol execution can be computed based on its own input and output only.

5.1.2 The MPC Compiler

Let Γ be a circuit implementing the function $\mathcal{F}(M_1, \cdot)$, where for any X and $i \in \mathbb{N}$, we have $(Y, M_{i+1}) = \mathcal{F}(M_i, X)$. Let $\Pi_{\mathcal{F}} = (C, S, \text{Enc}, \text{Dec}, \text{Next})$ be an r -round protocol realizing the function \mathcal{F} , over a set of m servers with a single client. The compiler produces $(\widehat{\Gamma}, \text{aux}) \leftarrow \Phi_{\Pi_{\mathcal{F}}}(\Gamma)$, where

- $\widehat{\Gamma}$ is the circuit that implements $\Pi_{\mathcal{F}}$ (depicted in Figure 5 for the case $m = 2$ and $p = 1$), having as a sub-circuit $\widehat{\Gamma}_{\text{Memory}}$, which is a circuit consisting only of memory gates, as needed by the original circuit Γ . During initialization, $\widehat{\Gamma}_{\text{Memory}}$ stores the initial private memory value, M_1 .
- $\text{aux} = ((\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_{m+2}), \mathcal{M}, (I_1, \dots, I_m))$, where
 - $\widehat{\Gamma}_{m+1} = \widehat{\Gamma}_{\text{Enc}}$ and $\widehat{\Gamma}_{m+2} = \widehat{\Gamma}_{\text{Dec}}$, i.e., the circuits $\widehat{\Gamma}_{m+1}$ and $\widehat{\Gamma}_{m+2}$ implement the encoder, Enc , and the decoder Dec , of $\Pi_{\mathcal{F}}$, respectively.
 - For $i \in [m]$, $\widehat{\Gamma}_i$ is the circuit that implements the code of the i -th server, for the entire execution of $\Pi_{\mathcal{F}}$ (r -rounds). Those circuits can be implemented in a straightforward way using the next message function Next_i (cf. the sub-components $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$ in Figure 5).
 - The mapping function \mathcal{M} describes the physical connections between the circuits described above, and I_j , for $j \in [m]$, specifies the components that will be outsourced to the manufacturer with index j . In our case $I_j = \{j\}$.
 - In case the original circuit is randomized, in addition to the components described above, Φ also outputs a circuit $\widehat{\Gamma}_{\text{s}}$ producing random coins R_i (as needed in each invocation of the circuit).

Our construction must be non-trivial (cf. Definition 4), thus the underlying protocol Π must satisfy the following outsourceability property.

Definition 14 (Outsourceability of protocols). *A protocol $\Pi = (C, S, \text{Enc}, \text{Dec}, \text{Next})$ that realizes the function \mathcal{F} can be outsourced if it satisfies the following condition: The circuit computing*

the encoding and decoding procedures (Enc, Dec) must be smaller than the circuit computing the function \mathcal{F} .

We prove the following result.

Theorem 4. *Let \mathcal{F} be any function, and let $\Pi_{\mathcal{F}}$ be a (t, m) -private MPC protocol for \mathcal{F} , satisfying the correctness and outsourceability properties. Then, the compiler $\Phi_{\Pi_{\mathcal{F}}}$ is a correct, $(0, t)$ -secure, $(o(1), m)$ -outsourcing circuit compiler.*

Proof. The correctness property of $\Phi_{\Pi_{\mathcal{F}}}$ follows directly by the correctness property of $\Pi_{\mathcal{F}}$. By the outsourceability property of the $\Pi_{\mathcal{F}}$, we have that $|\text{Enc}| + |\text{Dec}| = o(|\Gamma|)$, which implies that $\lim_{\lambda \rightarrow \infty} \frac{|\text{Enc}| + |\text{Dec}|}{|\Gamma|} = 0$. Having this in mind we can prove that the ratio ρ converges to 0, as λ goes to infinity, using the expression for ρ , and assuming $\widehat{\Gamma}_{\S}$ is of constant size:

$$\lim_{\lambda \rightarrow \infty} \frac{(|\text{Enc}| + |\text{Dec}| + |\widehat{\Gamma}_{\S}|)/|\Gamma|}{1} = 0.$$

Thus, for sufficiently large λ the outsourcing ration ρ is smaller than 1.

Let \mathcal{F} be any functionality and let Γ be the circuit implementing \mathcal{F} . Assuming that $\Pi_{\mathcal{F}}$ is a (t, m) -private MPC protocol for \mathcal{F} , we will prove that $\Phi_{\Pi_{\mathcal{F}}}$ is a $(0, t)$ -secure, circuit compiler. Concretely (cf. Definition 4), we need to prove that for all $\mathcal{C} \subseteq [m]$ of size at most t , all PPT adversaries \mathcal{A} , all circuits Γ , and for all initial values of the memory $M_1 \in \{0, 1\}^*$, there exists a simulator \mathcal{S} with running time $\text{poly}(|\mathcal{A}|, |\Gamma[M_1]|)$ such that

$$\{\mathbf{Real}_{\mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1}(\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Ideal}_{\mathcal{S}, \mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1, \ell}(\lambda)\}_{\lambda \in \mathbb{N}}, \quad (1)$$

for all sufficiently large values of λ . Let \mathcal{A} be an attacker $\Phi_{\Pi_{\mathcal{F}}}$. The idea behind the proof is to relate the interaction between \mathcal{A} and the circuits produced by $\Phi_{\Pi_{\mathcal{F}}}$, with the interaction between an attacker \mathcal{A}' corrupting up to t , with a protocol $\Pi_{\mathcal{F}}$. Then, we will use the simulator \mathcal{S}' that is given by the (t, m) -privacy of $\Pi_{\mathcal{F}}$ to construct a simulator \mathcal{S} , satisfying eq. (1). In what follows, and for the sake of simplicity, we prove the needed assuming \mathcal{A} is a single round attacker, and then we discuss how the proof easily extends to the setting in which we have multiple executions.

By the compiler definition, the protocol $\Pi_{\mathcal{F}}$ that $\Phi_{\Pi_{\mathcal{F}}}$ is based on, consists of two clients, C_1, C_2 , where C_1 is the corrupted client that provides the public input to the circuit, X , and C_2 supplies the circuit with private input, M_i , and m servers. Let Γ be the circuit implementing \mathcal{F} . Given the adversary \mathcal{A} for $\Phi_{\Pi_{\mathcal{F}}}$ we define the adversary $\mathcal{A}' = (\mathcal{A}'_0, \mathcal{A}'_1)$ against $\Pi_{\mathcal{F}}$ as follows:

- **(server corruption)** \mathcal{A}'_0 : run $(\widehat{\Gamma}, \text{aux}) \leftarrow \Phi(\Gamma)$, where $\text{aux} := ((\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_n), \mathcal{M}, (I_1, \dots, I_m))$, and sample $(\{\widehat{\Gamma}'_i\}_{i \in I}, \tau) \leftarrow \mathcal{A}_0(1^\lambda, \{\widehat{\Gamma}_i\}_{i \in I}, \langle \Gamma, \widehat{\Gamma} \rangle)$. Then corrupt the server S_i , for $i \in I$, so that S_i will execute the possibly modified circuit $\widehat{\Gamma}'_i$.
- **(protocol execution)** \mathcal{A}'_1 : participate in the protocol $\Pi_{\mathcal{F}}$ choosing the input for client C_1 (the corrupted client), according to the input value chosen by \mathcal{A}_1 . Concretely, execute the following steps: sample $X \leftarrow \mathcal{A}_1(1^\lambda, \tau)$, define the input of client C_1 to equal to X , receive the output of $\Pi_{\mathcal{F}}$ for client C_1 , Y , for inputs (X, M) , and forward Y to \mathcal{A}_1 .

We define the random variable View_{I_s, I_c} , $I_s = \mathcal{C}$, $I_c = \{1\}$, to be the view of \mathcal{A} while indirectly interacting with $\Pi_{\mathcal{F}}$ through \mathcal{A}'_1 . Clearly, by the definition of \mathcal{A}' , the view of \mathcal{A} while being executed by \mathcal{A}' , matches its view while executing the real world experiment of Definition 4, thus we have

$$\text{View}_{I_s, I_c}(\lambda, X, M) = \mathbf{Real}_{\mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1}(\lambda). \quad (2)$$

Assuming $\Pi_{\mathcal{F}}$ is (t, m) -private against \mathcal{A}' , there exists a simulator \mathcal{S}'' that simulates the view of \mathcal{A}' during the protocol execution. Let \mathcal{S}' be code of \mathcal{S}'' that only outputs the view of \mathcal{A} . Then we have that for all sufficiently large $\lambda \in \mathbb{N}$,

$$\text{View}_{I_s, I_c}(\lambda, X, M) \approx_c \mathcal{S}'(1^\lambda, I_c, I_s, (X, Y)_{i \in I_c}). \quad (3)$$

Now we define the simulator \mathcal{S} for \mathcal{A} against $\Phi_{\Pi_{\mathcal{F}}}$. \mathcal{S} on input $(1^\lambda, \langle \Gamma \rangle, \Phi, \mathcal{A}, \mathcal{C}, 0)$ executes the following steps:

- executes \mathcal{A}_1 with oracle access to $\Gamma[M_1](\cdot)$, and constructs the pair (X, Y) , i.e., it constructs the valid output of \mathcal{F} on input X , chosen by \mathcal{A}_1 .
- executes $o \leftarrow \mathcal{S}'(\lambda, I_c, I_s, (X, Y)_{i \in I_c})$, where $I_s = \mathcal{C}$ and $I_c = \{1\}$, and outputs o .

Clearly, from eq. (3) we have that \mathcal{S} produces outputs which is computationally indistinguishable from $\text{View}_{I_s, I_c}(\lambda, X, M)$, and then using eq. (2) we receive,

$$\mathbf{Real}_{\mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1}(\lambda) \approx_c \mathbf{Ideal}_{\mathcal{S}, \mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1, \ell}(\lambda),$$

and this concludes the proof for attackers executing the protocol only once.

For multi-round attackers against the circuit compiler, we need to have multiple, sequential executions, of the same protocol, as a single execution computes a single circuit output. Moreover, the attacker is non-adaptive, and corrupts the servers only before the first protocol execution. By the composition theorem of [20], we have that any secure MPC protocol is also secure against sequential composition, even for adaptive adversaries. Using a standard hybrid argument, this gives rise to a simulator, \mathcal{S}' , that simulates the view of the attacker for all protocol executions, and the proof idea is identical to one given above: we relate the attacker against the compiler to an attacker against the protocol, and we use \mathcal{S}' to construct a simulator \mathcal{S} for the circuit compiler. \square

The case of memory outsourcing is identical to that of the first compiler and we refer the reader to Remark 1.

5.2 Compilers based on VC with Multiple Servers

Multi-Server VC is a generalization of the VC setting described in Section 4.1, where a client P_0 with private input X outsources the computation of some function \mathcal{F} to a set of servers P_1, \dots, P_m ; a t -fraction of the servers might be corrupted (for $t \leq m - 1$) and thus controlled by a monolithic adversary. In general, the communication pattern of a multi-server VC scheme could be arbitrary, but known protocols typically follow very restricted and simple patterns (see below for some examples); sometimes there is also a setup phase which is run only once, and it is independent of the input X and of the function \mathcal{F} being computed.

5.2.1 Refereed Delegation of Computation

First introduced by Canetti, Riva and Rothblum [21], a Refereed Delegation of Computation (RDoC) scheme is a special case of a multi-server VC scheme with the following simple communication pattern: The client P_0 forwards the input X to each of the servers, and then each server P_i (for $i \in [m]$) sends a single message back to the client P_0 . Note that the servers are not allowed to communicate with each other. In case the servers claim different results for the computation of $\mathcal{F}(X)$, the client (acting as a referee) attempts to determine the correct answer.

Definition 15 (Refereed Delegation of Computation). *Let P_0, P_1, \dots, P_m be probabilistic interactive Turing machines. An RDoC scheme for computing a function \mathcal{F} is an interactive protocol described as follows.*




Figure 6: The Refereed Delegation of Computation compiler for the case of $m = 2$ outsourcing facilities (from which *one* is malicious). The components $\hat{\Gamma}_1$ and $\hat{\Gamma}_2$ are outsourced, while the connectivity between them and the remaining components are built in-house.

- Setup: The client P_0 forwards some input X to the servers P_1, \dots, P_m .
- Computation: Each server P_i , for $i \in [m]$, sends a single message Y^i to the client P_0 .
- Refereeing and output: The client P_0 , eventually after interrogating the servers, outputs some value Y .

We say that an RDoC is correct if for all inputs X , assuming all the servers are honest, at the end of the protocol P_0 outputs the value $\mathcal{F}(X)$.

Soundness of an RDoC scheme demands that, at the end of the referee phase, the client P_0 always outputs the correct value $\mathcal{F}(X)$ with high probability, even if up to $m - 1$ servers are malicious and behave arbitrarily. Note that here, in contrast to the soundness guarantee for single-server VC, the client P_0 never outputs \perp ; in order to stress this difference we call this property *strong* soundness.

Definition 16 (Strong soundness for RDoC). Let $\mathcal{RDC} = (P_0, P_1, \dots, P_m)$ be an RDoC scheme. We say that \mathcal{RDC} satisfies strong soundness if for any input X , any $i \in [m]$, any set of malicious servers $P_1^*, \dots, P_{i-1}^*, P_{i+1}^*, \dots, P_m^*$, there exists a negligible⁴ function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that, at the end of the protocol, the output of client P_0 equals $\mathcal{F}(X)$ with probability at least $1 - \nu(\lambda)$.

Finally, an RDoC scheme should be outsourceable, meaning that the complexity of the referee phase should be independent of the complexity required to evaluate the function itself.

Definition 17 (Outsourceability for RDoC). Let \mathcal{F} be a function with input-size of n bits. An RDoC scheme is outsourceable if the complexity of the client is at most quasi-linear in n , and the complexity of the honest servers is polynomial in the complexity of evaluating \mathcal{F} .

The compiler $\Phi_{\mathcal{RDC}}$. We now describe a natural outsourcing compiler based on any RDoC scheme for a function \mathcal{F} . On a high level, the designer outsources the circuits components corresponding to the servers in the RDoC scheme, while the component corresponding to the client is built in-house.

Let Γ be a circuit implementing the function $\mathcal{F}(M_1, \cdot; \cdot)$, where for any X and $i \in \mathbb{N}$, we have $(Y, M_{i+1}) = \mathcal{F}(M_i, X; R)$. Let $\mathcal{RDC} = (P_0, P_1, \dots, P_m)$ be an RDoC scheme for the function \mathcal{F} . The compiler produces $(\hat{\Gamma}, \text{aux}) \leftarrow \Phi_{\mathcal{RDC}}(\Gamma)$, which we describe next.

⁴As shown in [21], the soundness error can always be made negligible by parallel repetition.

- $\widehat{\Gamma}$ is the circuit that implements \mathcal{RDC} (depicted in Fig. 6 for the case $m = 2$).
- $\text{aux} = ((\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_m, \widehat{\Gamma}_{\text{Enc}}, \widehat{\Gamma}_{\text{Dec}}, \widehat{\Gamma}_{\S}), \mathcal{M}, (I_1, \dots, I_m))$, where:
 - In case the original circuit is randomized, $\widehat{\Gamma}_{\S}$ produces random coins R_i (as needed in each invocation of the circuit).
 - $\widehat{\Gamma}_{\text{Enc}}$ simply forwards the input values X_i, R_i to the m components $\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_m$.
 - $\widehat{\Gamma}_{\text{Dec}}$ takes as input X_i, R_i and the outputs Y_i^1, \dots, Y_i^m from the components $\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_m$, respectively, and implements the code of the client P_0 in the RDoC scheme.
 - For $i \in [m]$, the component $\widehat{\Gamma}_i$ is the circuit that implements the code of the i -th server P_i in the RDoC scheme; notice that each component $\widehat{\Gamma}_i$ is initialized with the initial private memory M_1 of the original circuit.
 - The mapping function \mathcal{M} describes the physical connections between the circuits described above, and I_j , for $j \in [m]$, specifies the components that will be outsourced to the manufacturer with index j . In our case $I_j = \{j\}$.

The theorem below, states that if the underlying RDoC scheme satisfies strong soundness, the above compiler is a secure outsourcing compiler. The main idea behind the proof is that the simulator can simply answer each query from the adversary by forwarding it to its own target oracle that computes the original circuit (with the initial memory M_1 embedded); this is fine because strong soundness guarantees that, with overwhelming probability, the output of the compiled circuit should be the same as that of the original circuit.

Theorem 5. *Let Γ be an arbitrary circuit and let $\mathcal{RDC} = (P_0, P_1, \dots, P_m)$ be an RDoC scheme for the function \mathcal{F} computed by Γ satisfying strong soundness. Then the compiler $\Phi_{\mathcal{RDC}}$ is a correct, $(0, m - 1)$ -secure $(o(1), m)$ -outsourcing circuit compiler.*

Proof. The correctness property of $\Phi_{\mathcal{RDC}}$ follows directly from the correctness of the \mathcal{RDC} protocol. The proof of non-triviality uses the outsourcability property of the RDoC scheme; since the argument is very similar to that used in the proof of Claim 1, we omit it here.

Now we proceed to prove that $\Phi_{\mathcal{RDC}}$ is a $(0, m - 1)$ -secure circuit compiler. We need to show that for all $\mathcal{C} \subseteq [m]$ of size at most $m - 1$, all PPT adversaries \mathcal{A} , all circuits Γ , and for all initial values of the memory $M_1 \in \{0, 1\}^*$, there exists a simulator \mathcal{S} with running time $\text{poly}(|\mathcal{A}|, |\Gamma[M_1]|)$ such that $\{\mathbf{Real}_{\mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1}(\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Ideal}_{\mathcal{S}, \mathcal{A}, \Phi, \mathcal{C}, \Gamma, M_1, \ell}(\lambda)\}_{\lambda \in \mathbb{N}}$, for all sufficiently large values of λ . A description of the simulator \mathcal{S} follows.

- Run the compiler $\Phi_{\mathcal{RDC}}(\Gamma)$ supplying (a description of) the original circuit Γ ; the output is (a description of) the compiled circuit $\widehat{\Gamma}$, and the auxiliary information $\text{aux} := ((\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_m, \widehat{\Gamma}_{\text{Enc}}, \widehat{\Gamma}_{\text{Dec}}, \widehat{\Gamma}_{\S}), \mathcal{M}, (\{1\}, \dots, \{m\}))$.
- Forward the components $\{\widehat{\Gamma}_j\}_{j \in \mathcal{C}}$ to adversary \mathcal{A} , together with the descriptions of Γ and $\widehat{\Gamma}$.
- Adversary \mathcal{A} produces the components $\{\widehat{\Gamma}_j\}_{j \in \mathcal{C}}$ (that may be malicious) and sends it to the simulator; at this point, the circuit $\widehat{\Gamma}'$ could be assembled via the mapping function \mathcal{M} , but this step is not necessary in this proof.
- Upon input X_i from \mathcal{A} , query X_i to oracle $\Gamma[M_1](\cdot)$ and forward the output to \mathcal{A} .

Let us write \mathbf{R} and \mathbf{S} for the distribution of the random variables in experiment \mathbf{Real} and \mathbf{Ideal} of Definition 4. The claim below shows that experiment \mathbf{R} and \mathbf{S} are computationally indistinguishable, and thus concludes the proof of the theorem.

Claim 5. *For all PPT adversaries \mathcal{A} , and for all PPT distinguishers \mathcal{D} , there exists a negligible function $\nu : \mathbb{N} \rightarrow [0, 1]$ such that $\Delta^{\mathcal{D}}(\mathbf{R}; \mathbf{S}) \leq \nu(\lambda)$.*

Proof. Let $q \in \text{poly}(\lambda)$ be the number of input queries asked by \mathcal{A} to its oracle, and let us define an event W that becomes true whenever there is some index $i^* \in [q]$ such that the value Y_{i^*} output in experiment \mathbf{R} is different from the value $\mathcal{F}(M_{i^*}, X_{i^*}; R_{i^*})$. Then, by Lemma 1, we have that $\Delta^{\mathcal{D}}(\mathbf{R}; \mathbf{S}) \leq \Delta(\mathbf{R}; \mathbf{S} | \neg W) + \mathbb{P}[W]$. Note that the distributions \mathbf{R} and \mathbf{S} are identical conditioned on event W *not* happening; this is because the simulator always forwards the inputs from the adversary to its target oracle which evaluates the original circuit (with initial private memory M_1). Hence, it suffices to show that the probability of event W is negligible.

Let us assume that there exists some circuit Γ , some initial memory M_1 , some subset $\mathcal{C} \subset [m]$ (for the corrupted players), and a PPT adversary \mathcal{A} provoking event W with non-negligible probability. We build a PPT adversary \mathcal{A}' that uses \mathcal{A} in order to break the strong soundness property of the RDoC scheme.

Adversary $\mathcal{A}'(M_1, i^*, \mathcal{C}, \mathcal{F}, \Gamma, \Phi_{\text{RDoC}})$:

1. Run the compiler $\Phi_{\text{RDoC}}(\Gamma)$ supplying (a description of) the original circuit Γ ; the output is (a description of) the compiled circuit $\widehat{\Gamma}$, and the auxiliary information $\text{aux} := ((\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_m, \widehat{\Gamma}_{\text{Enc}}, \widehat{\Gamma}_{\text{Dec}}, \widehat{\Gamma}_{\text{S}}), \mathcal{M}, (\{1\}, \dots, \{m\}))$.
2. Forward the components $\{\widehat{\Gamma}_j\}_{j \in \mathcal{C}}$ to adversary \mathcal{A} , together with the descriptions of Γ and $\widehat{\Gamma}$. Adversary \mathcal{A} produces the components $\{\widehat{\Gamma}'_j\}_{j \in \mathcal{C}}$ and sends them to \mathcal{A}' ; the circuit $\widehat{\Gamma}'$ is assembled via the mapping function \mathcal{M} .
3. Upon input each query X_i from \mathcal{A} , proceed as follows:
 - If $i < i^*$, compute $(Y_i, M_{i+1}) \leftarrow \widehat{\Gamma}'[M_i](X_i; R_i)$ and reply with Y_i .
 - If $i = i^*$, forward to the client P_0 in the RDoC game the values (Y^1, \dots, Y^m) , where Y^j is equal to the first message output by the component $\widehat{\Gamma}'_j[M_{i^*}](X_{i^*}, R_{i^*})$ for all $j \in \mathcal{C}$, whereas Y^j is equal to the first message output by $\mathcal{F}(M_{i^*}, X_{i^*}; R_{i^*})$ for all $j \in [m] \setminus \mathcal{C}$. During the referee phase done by P_0 , answer each query by either running the components $\widehat{\Gamma}'_j$ (for all $j \in \mathcal{C}$), or $\widehat{\Gamma}_j$ (for all $j \in [m] \setminus \mathcal{C}$).
 - Stop processing further queries from \mathcal{A} and terminate.

For the analysis, we note that the above simulation is perfect. In particular the outputs seen by \mathcal{A} retain exactly the same distribution as in experiment \mathbf{R} . Also, the messages sent by \mathcal{A}' have the same distribution as in an execution of the RDoC strong soundness game with input $(M_{i^*}, X_{i^*}, R_{i^*})$; this follows by the fact that the components $\{\widehat{\Gamma}'_j\}_{j \in \mathcal{C}}$ and $\{\widehat{\Gamma}_j\}_{j \in [m] \setminus \mathcal{C}}$ run, respectively, the code of the corrupted servers and the honest servers of the underlying RDoC scheme. Now \mathcal{A} provokes event W with non-negligible probability, so that the client P_0 in the RDoC game will output a value different from $\mathcal{F}(M_{i^*}, X_{i^*}; R_{i^*})$ with non-negligible probability. This contradicts strong soundness of the RDoC scheme, and thus concludes the proof of the claim. □

□

5.2.2 Input-Private Multi-Server VC

RDoC schemes satisfy a very strong form of soundness, namely the honest client always outputs the correct value $\mathcal{F}(X)$; in particular, the corrupted servers cannot cause the client to receive an invalid output. A natural relaxation of security is to require that the honest client either outputs $\mathcal{F}(X)$, or it outputs \perp ; however, as we have shown in Section 4.3, when using a non-input-private

VC scheme in an outsourcing compiler, invalid outputs can leak a lot of information about the circuit private memory (in an undetectable manner).⁵

A similar problem appears in the case of multiple servers, and thus a possible solution is to rely on a multi-server VC scheme satisfying correctness, soundness, and input-privacy. Assuming such a VC scheme, one obtains a secure outsourcing compiler for the case of multiple manufacturers (without self-destruct); we omit the formal details, as the construction and its analysis are very similar to that of our compiler from Section 4.2.

Input-private multi-server VC schemes were constructed in [1]. The communication pattern is circular, meaning that P_0 sends a message to P_1 , which then sends a message to P_2 , and so on, until P_m sends a last message back to the client P_0 . Interestingly, for the case of multiple servers, input privacy can be obtained without relying on expensive tools such as fully-homomorphic encryption. We refer the reader to Section 6 for a more in-depth discussion about efficiency/security related to these instantiations.

6 Concrete Instantiations

In this section we propose several instantiations for the compilers analyzed in the previous sections, highlighting several possible tradeoffs between security, efficiency, and underlying hardness assumptions.

6.1 Case Study I

The area of verifiable computation has a long history in the cryptographic literature [4, 62, 44, 41]. We refer the reader to the excellent survey by Walfish and Blumberg [74] for a thorough introduction. By now, several schemes and models for the problem of outsourcing computation

⁵For standard VC schemes with only one server, this attack can be avoided by relying on the self-destruct capability. We do not consider this option here, as it is trivial to obtain a secure outsourcing compiler in the case of multiple manufacturers if one relies on self-destruct.

Compiler	Reference	t	s_{in}	s_{out}	No Self-Destruct	Leakage	Assumption
§ 4.2	[39]	–	$O(n + v)$	$O(s \log s)$	✓	0	KoE + FHE
§ 4.3	[42, 66]	–	$O(n + v)$	$O(s \log s)$	✗	$O(\log q)$	KoE
§ 5.2	[21]	$m - 1$	$O(n \log n)$	$O(s + \lambda \log s)$	✓	0	CRH
	[1]	$m - 1$	$O(\lambda n)$	$O(\lambda s)$	✓	0	DDH
	[1]	$O(m)$	$O(\lambda n)$	$O(\lambda s)$	✓	0	OWF
§ 5.1	[31]	$m - 1$	$O(dm)$	$O(sm + m^3)$	✓	0	SHE
	[75]	$m - 1$	$O(dm^2)$	$O(sm^2 \cdot \lambda / \log s)$	✓	0	OT

Table 1: Comparing our compilers in terms of security, efficiency, and hardness assumptions. We write s, n, v for the size, number of inputs and number of outputs of the original circuit Γ , respectively; as usual m denotes the number of servers of which up to t might be corrupted (note that $t = m$ corresponds to the case of a single manufacturer), and q is the number of oracle queries asked by the adversary in Definition 4. The values $s_{\text{in}}, s_{\text{out}}$ and $\#s_{\text{out}}$ denote, respectively, for the sizes of the components built in house, the size of the outsourced components, and the quantity of outsourced components; d denotes the number of multiplications in Γ . KoE stands for “Knowledge of Exponent” assumptions, DDH for “Decisional Diffie-Hellman”, OWF for “One-Way Functions”, CRH for “Collision-Resistant Hashing”, FHE for “Fully-Homomorphic Encryption”, OT for “Oblivious Transfer” and SHE for “Somewhat Homomorphic Encryption”. The first two (colored) rows represent the compilers with a single outsourcing facility ($m = 1$), while the remaining rows represent the compilers with multiple outsourcing facilities ($m \geq 2$).

are known (see, among others, [2, 10, 16, 13, 52, 1]) Below, we focus only on single server VC schemes suitable for the compilers described in Section 4.

Input privacy. For the compiler of Section 4.2, we need a VC scheme satisfying both soundness and input-privacy (in the presence of verification queries). The only known schemes meeting these requirements are the ones constructed by Fiore, Gennaro, and Pastro [39]. The asymptotic complexity stays the same as the one for non-input-private VC, but with much bigger constants hidden inside the big-oh notation.

No input privacy. For the compiler of Section 4.3, we need a VC scheme satisfying soundness (in the presence of verification queries), but no input-privacy is required. Therefore, we can instantiate this compiler using very efficient schemes based on SNARKs [62, 45, 16, 29, 18, 17, 13, 42, 66, 14, 26, 47]. If s is the size of the original circuit, the complexity of the outsourced components is in $O(s \log s)$, while the size of the components built in-house is in $O(n+v)$, where n, v denote, respectively, the input and output size of the original circuit.

6.2 Case Study II

We describe below a few possible instantiations for the multiple manufacturers compilers of Section 5.

RDoC. The work of [21] gives a very efficient construction of an RDoC scheme for arbitrary polynomial-time computable functions, with computational security based on any collision-resistant hash functions family. The idea is to simply have each server compute the original function. Then, in case the outputs of the servers are not all equal, the client starts a binary search phase to check the computation performed by each of the servers. Such an interrogation phase essentially consists of checking the correctness of intermediate steps of the computation.⁶

The size of the outsourced components $\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_m$ is dominated by the complexity of function \mathcal{F} , plus the overhead relative to the binary search phase; this yields a complexity of roughly $O(s + \lambda \log s)$. The component $\widehat{\Gamma}_{\text{Enc}}$ that is built in-house consists of just a few gates since it simply forwards the current input to the m sub-components. The component $\widehat{\Gamma}_{\text{Dec}}$, that is also built in-house, implements the code of the client in the RDoC scheme, and by definition its size is quasi-linear in the size of the input. Furthermore, as long as the output of all components $\widehat{\Gamma}_1, \dots, \widehat{\Gamma}_m$ agrees, the running time of $\widehat{\Gamma}$ is roughly the same as Γ (assuming all the m components are able to run in parallel); when the output of the components differ, the RDoC protocol is executed, resulting in some overhead.

A nice feature of this construction is that as soon as the components disagree on an output, the malicious components can be identified and switched-off; in the long run, this makes all the subsequent executions of $\widehat{\Gamma}$ almost as efficient as that of Γ . Fig. 7 exemplifies this feature for the case $m = 2$; notice that disabling the malicious component requires one bit of state that needs to be shared between the encoder and the decoder.

Input-private multi-server VC. Multi-server VC schemes satisfying input privacy were constructed in [1]; interestingly, in the multi-server case we can obtain input privacy *without* relying on FHE. An additional feature of these protocols is that the pre-processing stage is independent of the input and of the function to be computed.

⁶The original protocol in [21] is for Turing machines, but we can always transform a Turing machine in an equivalent circuit of related complexity [76, Theorem 2.3].

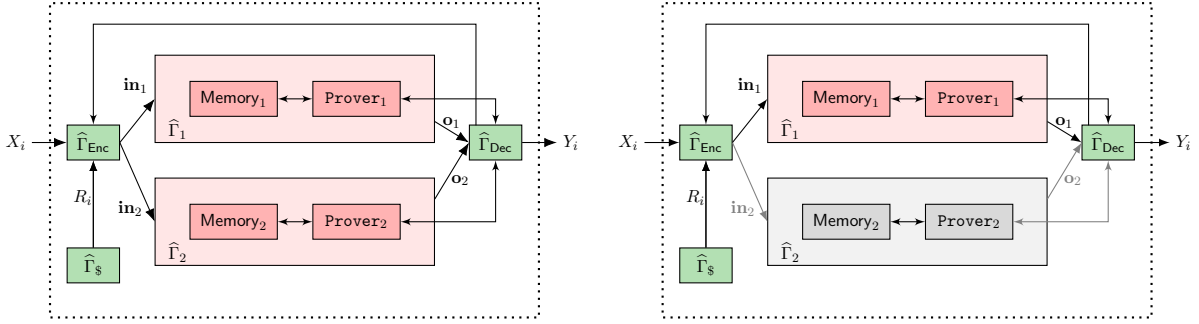


Figure 7: The Refereed Delegation of Computation compiler using the protocol of [21] for the case of $m = 2$ outsourcing facilities (out of which at most one is malicious). The components $\hat{\Gamma}_1$ and $\hat{\Gamma}_2$ are outsourced, while the connectivity between them and the remaining components are built in-house. On the left side both components $\hat{\Gamma}_1$ and $\hat{\Gamma}_2$ are enabled, while on the right side the component $\hat{\Gamma}_2$ was caught trying to cheat and was therefore disabled by $\hat{\Gamma}_{\text{Dec}}$. To keep track of the disabled components, $\hat{\Gamma}_{\text{Enc}}$ is required to store one bit of state that can be written by $\hat{\Gamma}_{\text{Dec}}$.

Based on [1], we obtain the following instantiations: (i) An outsourcing compiler for the special case $m = 2$, with computational security based on the sole assumption that OWF exist; the complexity of the client is in $O(\lambda \cdot n + \lambda \cdot v)$, while each server needs to generate and evaluate a garbled circuit. (ii) An outsourcing compiler tolerating up to $t \leq m - 1$ corrupted facilities, with computational security based on the DDH assumption; the complexity of the client is in $O(\lambda \cdot n)$, while the servers rely on re-randomizable garbled circuits, NIZK proofs and a DDH-based proxy re-encryption scheme. (iii) An outsourcing compiler tolerating a constant fraction of corrupted facilities, with computational security based on the sole assumption that OWF exist.

Client-server MPC. Many MPC protocols satisfy the outsourceability property, as the values that feed the main computation, i.e., the output of the encoder, are independent of the function that is being evaluated, and mostly depend on the number of parties, as in the case of [43] (where the same holds for decoding). An explicit (t, m) -private protocol is given in [30], for $t < m/2$, in which there is a pre-processing phase that can be implemented by the encoder, with running time independent of the function that is being evaluated. The construction uses secure point-to-point and broadcast channels, that can be implemented directly between the components, and besides privacy it also guarantees output delivery.

We can also easily adapt the SPDZ protocol [31] to the client-server setting. The SPDZ protocol requires a pre-processing phase that is performed by the parties, and that will feed the encoder circuit who will perform the actual encoding (which is only a linear operation). The complete protocol requires a linear number of public-key operations in the circuit size s , with the encoder requiring only a linear number of operations in m and the number of multiplications of the original circuit. The efficiency of the pre-processing stage can be further improved [6]. This construction does not guarantee output delivery, but it is secure against adversaries that corrupt up to $m - 1$ sub-components. Finally, the construction of [50] can also be adapted to the client-server scenario, and additionally provides security with *identifiable aborts*, meaning that whenever the protocol terminates due to an abort message, it is guaranteed that at least one malicious party will be caught. As explained above, such a feature might be useful for deactivating adversarial circuit components, improving efficiency and energy consumption even further.

7 Conclusion and Open Problems

We put forward a simulation-based security definition for assessing security of ICs whose fabrication has (partially) been outsourced to an untrusted off-shore manufacturer. Importantly, the size of the components built in-house must be much smaller than the size of the original IC. Our definition implies a strong guarantee, essentially saying that no matter how the manufacturer modified the outsourced IC sub-components, using the re-assembled IC in the wild cannot leak sensitive information about the private memory.

With such a framework in hand, we proposed several compilers that meet our security definition and work for any circuit, both for the case where there is a single outsourcing facility and the case of multiple outsourcing facilities. Our constructions generically leverage VC schemes or MPC protocols for the function underlying the produced IC, with certain properties.

There are several interesting open questions related to our work. First, it might be interesting to explore variations of our model, for example by considering the case where there are several (non-colluding) manufacturers involved in the fabrication process. In such a case, it might be possible to obtain significant efficiency improvements. Second, one could try to instantiate our compilers with specialized VC schemes that are tailored for specific functionalities. Although there are already some schemes with this feature—e.g., [39] constructs VC schemes tailored for (multi-variate) polynomials and linear combinations—to the best of our knowledge, there is no concrete VC scheme for verifying the computation of a specific cryptographic functionality (such as AES). Third, it would be interesting to explore different approaches in order to build compilers meeting our security definition without relying on verifiable computing or multiparty computation.

References

- [1] Prabhanjan Ananth, Nishanth Chandran, Vipul Goyal, Bhavana Kanukurthi, and Rafail Ostrovsky. Achieving privacy in verifiable computation with multiple servers - without FHE and without pre-processing. In *PKC*, pages 149–166, 2014.
- [2] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. From secrecy to soundness: Efficient verification via secure computation. In *ICALP*, pages 152–163, 2010.
- [3] Giuseppe Ateniese, Bernardo Magri, and Daniele Venturi. Subversion-resilient signature schemes. In *ACM CCS*, pages 364–375, 2015.
- [4] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *ACM STOC*, pages 21–31, 1991.
- [5] James Ball, Julian Borger, and Glenn Greenwald. Revealed: how US and UK spy agencies defeat internet privacy and security. *Guardian Weekly*, September 2013.
- [6] Carsten Baum, Ivan Damgård, Tomas Toft, and Rasmus Winther Zakarias. Better preprocessing for secure multiparty computation. In *ACNS*, pages 327–345, 2016.
- [7] Mark Beaumont, Bradley Hopkins, and Tristan Newby. Hardware trojans — prevention, detection, countermeasures (a literature review). Technical report, Australian Government Department of Defence, 07 2011.
- [8] Donald Beaver. Commodity-based cryptography (extended abstract). In *ACM STOC*, pages 446–455, 1997.

- [9] Georg T. Becker, Francesco Regazzoni, Christof Paar, and Wayne P. Burleson. Stealthy dopant-level hardware trojans: extended version. *J. Cryptographic Engineering*, 4(1):19–31, 2014.
- [10] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In *ASIACRYPT*, pages 134–153, 2012.
- [11] Mihir Bellare, Joseph Jaeger, and Daniel Kane. Mass-surveillance without the state: Strongly undetectable algorithm-substitution attacks. In *ACM CCS*, pages 1431–1440, 2015.
- [12] Mihir Bellare, Kenneth G. Paterson, and Phillip Rogaway. Security of symmetric encryption against mass surveillance. In *CRYPTO*, pages 1–19, 2014.
- [13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: verifying program executions succinctly and in zero knowledge. In *CRYPTO*, pages 90–108, 2013.
- [14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *USENIX Security Symposium*, pages 781–796, 2014.
- [15] Shivam Bhasin and Francesco Regazzoni. A survey on hardware trojan detection techniques. In *IEEE ISCAS*, pages 2021–2024, 2015.
- [16] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *ICS*, pages 326–349, 2012.
- [17] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKS and proof-carrying data. In *ACM STOC*, pages 111–120, 2013.
- [18] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In *TCC*, pages 315–333, 2013.
- [19] Elette Boyle, Gil Segev, and Daniel Wichs. Fully leakage-resilient signatures. *J. Cryptology*, 26(3):513–558, 2013.
- [20] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 2000.
- [21] Ran Canetti, Ben Riva, and Guy N. Rothblum. Practical delegation of computation using multiple servers. In *ACM CCS*, pages 445–454, 2011.
- [22] Ran Canetti, Ben Riva, and Guy N. Rothblum. Two protocols for delegation of computation. In *ICITS*, pages 37–61, 2012.
- [23] Ran Canetti, Ben Riva, and Guy N. Rothblum. Refereed delegation of computation. *Inf. Comput.*, 226:16–36, 2013.
- [24] Rajat Subhra Chakraborty, Seetharam Narasimhan, and Swarup Bhunia. Hardware trojan: Threats and emerging solutions. In *IEEE HLDVT*, pages 166–171, 2009.

- [25] Nishanth Chandran, Wutichai Chongchitmate, Rafail Ostrovsky, and Ivan Visconti. Universally composable secure two and multi-party computation in the corruptible tamper-proof hardware token model. *IACR Cryptology ePrint Archive*, 2017:1092, 2017.
- [26] Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *IEEE Symposium on Security and Privacy*, pages 253–270, 2015.
- [27] Dana Dachman-Soled and Yael Tauman Kalai. Securing circuits against constant-rate tampering. In *CRYPTO*, pages 533–551, 2012.
- [28] Dana Dachman-Soled and Yael Tauman Kalai. Securing circuits and protocols against $1/\text{poly}(k)$ tampering rate. In *TCC*, pages 540–565, 2014.
- [29] Ivan Damgård, Sebastian Faust, and Carmit Hazay. Secure two-party computation with low communication. In *TCC*, pages 54–74, 2012.
- [30] Ivan Damgård and Yuval Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In *CRYPTO*, pages 378–394, 2005.
- [31] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, pages 643–662, 2012.
- [32] Jean Paul Degabriele, Pooya Farshim, and Bertram Poettering. A more cautious approach to security against mass surveillance. In *FSE*, pages 579–598, 2015.
- [33] Yevgeniy Dodis, Chaya Ganesh, Alexander Golovnev, Ari Juels, and Thomas Ristenpart. A formal treatment of backdoored pseudorandom generators. In *EUROCRYPT*, pages 101–126, 2015.
- [34] Yevgeniy Dodis, Kristiyan Haralambiev, Adriana López-Alt, and Daniel Wichs. Efficient public-key cryptography in the presence of key leakage. In *ASIACRYPT*, pages 613–631, 2010.
- [35] Stefan Dziembowski, Sebastian Faust, and François-Xavier Standaert. Private circuits III: hardware trojan-resilience via testing amplification. In *ACM CCS*, pages 142–153, 2016.
- [36] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302, 2008.
- [37] Antonio Faonio, Jesper Buus Nielsen, and Daniele Venturi. Mind your coins: Fully leakage-resilient signatures with graceful degradation. In *ICALP*, pages 456–468, 2015.
- [38] Sebastian Faust, Krzysztof Pietrzak, and Daniele Venturi. Tamper-proof circuits: How to trade leakage for tamper-resilience. In *ICALP*, pages 391–402, 2011.
- [39] Dario Fiore, Rosario Gennaro, and Valerio Pastro. Efficiently verifiable computation on encrypted data. In *ACM CCS*, pages 844–855, 2014.
- [40] Marc Fischlin, Benny Pinkas, Ahmad-Reza Sadeghi, Thomas Schneider, and Ivan Visconti. Secure set intersection with untrusted hardware tokens. In *CT-RSA*, pages 1–16, 2011.
- [41] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO*, pages 465–482, 2010.

- [42] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *EUROCRYPT*, pages 626–645, 2013.
- [43] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *ACM STOC*, pages 218–229, 1987.
- [44] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In *ACM STOC*, pages 113–122, 2008.
- [45] Shafi Goldwasser, Huijia Lin, and Avi Rubinfeld. Delegation of computation without rejection problem from designated verifier cs-proofs. *IACR Cryptology ePrint Archive*, 2011:456, 2011.
- [46] Glenn Greenwald. No place to hide: Edward Snowden, the NSA, and the U.S. surveillance state. *Metropolitan Books*, May 2014.
- [47] Jens Groth. On the size of pairing-based non-interactive arguments. In *EUROCRYPT*, pages 305–326, 2016.
- [48] Mike Hamburg, Paul Kocher, and Mark Marson. Analysis of Intel’s Ivy Bridge digital random number generator. Technical report, Cryptography Research, Inc., 03 2012.
- [49] Frank Imeson, Ariq Emtenan, Siddharth Garg, and Mahesh V. Tripunitara. Securing computer hardware using 3d integrated circuit (IC) technology and split manufacturing for obfuscation. In *USENIX Security Symposium*, pages 495–510, 2013.
- [50] Yuval Ishai, Rafail Ostrovsky, and Vassilis Zikas. Secure multi-party computation with identifiable abort. In *CRYPTO*, pages 369–386, 2014.
- [51] Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and David Wagner. Private circuits II: keeping secrets in tamperable circuits. In *EUROCRYPT*, pages 308–327, 2006.
- [52] Thomas P. Jakobsen, Jesper Buus Nielsen, and Claudio Orlandi. A framework for outsourcing of secure computation. In *CCSW*, pages 81–92, 2014.
- [53] Benjamin Jun and Paul Kocher. The Intel random number generator. Technical report, Cryptography Research, Inc., 04 1999.
- [54] Jonathan Katz. Universally composable multi-party computation using tamper-proof hardware. In *EUROCRYPT*, pages 115–128, 2007.
- [55] Jonathan Katz and Vinod Vaikuntanathan. Signature schemes with bounded leakage resilience. In *ASIACRYPT*, pages 703–720, 2009.
- [56] Aggelos Kiayias and Yiannis Tselekounis. Tamper resilient circuits: The adversary at the gates. In *ASIACRYPT*, pages 161–180, 2013.
- [57] Lang Lin, Markus Kasper, Tim Güneysu, Christof Paar, and Wayne Burleson. Trojan side-channels: Lightweight hardware trojans through side-channel engineering. In *CHES*, pages 382–395, 2009.
- [58] Eric Love, Yier Jin, and Yiorgos Makris. Enhancing security via provably trustworthy hardware intellectual property. In *IEEE HOST*, pages 12–17, 2011.

- [59] Marie A. Mak. Trusted Defense Microelectronics: Future Access and Capabilities Are Uncertain. Technical report, United States Government Accountability Office, 10 2015.
- [60] Vasilios Mavroudis, Andrea Cerulli, Petr Svenda, Dan Cvrcek, Dusan Klinec, and George Danezis. A touch of evil: High-assurance cryptographic hardware from untrusted components. In *ACM CCS*, pages 1583–1600, 2017.
- [61] David R. McIntyre, Francis G. Wolff, Christos A. Papachristou, and Swarup Bhunia. Dynamic evaluation of hardware trust. In *IEEE HOST*, pages 108–111, 2009.
- [62] Silvio Micali. Computationally sound proofs. *SIAM J. Comput.*, 30(4):1253–1298, 2000.
- [63] Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. *SIAM J. Comput.*, 41(4):772–814, 2012.
- [64] Jesper Buus Nielsen, Daniele Venturi, and Angela Zottarel. Leakage-resilient signatures with graceful degradation. In *PKC*, pages 362–379, 2014.
- [65] Martin Otto. *Fault Attacks and Countermeasures*. PhD thesis, University of Paderborn, Germany, 2006.
- [66] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *IEEE Symposium on Security and Privacy*, pages 238–252, 2013.
- [67] Nicole Perlroth, Jeff Larson, and Scott Shane. N.S.A. able to foil basic safeguards of privacy on web. *The New York Times*, September 2013.
- [68] Krzysztof Pietrzak. A leakage-resilient mode of operation. In *EUROCRYPT*, pages 462–482, 2009.
- [69] Miodrag Potkonjak. Synthesis of trustable ics using untrusted CAD tools. In *DAC*, pages 633–634, 2010.
- [70] Jean-Pierre Seifert and Christoph Bayer. Trojan-resilient circuits. In Al-Sakib Khan Pathan, editor, *Securing Cyber-Physical Systems*, chapter 14, pages 349–370. CRC Press, Boca Raton, London, New York, 2015.
- [71] Brian Sharkey. Trust in Integrated Circuits Program. Technical report, DARPA, 03 2007.
- [72] Riad S. Wahby, Max Howald, Siddharth J. Garg, Abhi Shelat, and Michael Walfish. Verifiable asics. In *IEEE S&P*, pages 759–778, 2016.
- [73] Adam Waksman and Simha Sethumadhavan. Silencing hardware backdoors. In *IEEE Symposium on Security and Privacy*, pages 49–63, 2011.
- [74] Michael Walfish and Andrew J. Blumberg. Verifying computations without reexecuting them. *Commun. ACM*, 58(2):74–84, 2015.
- [75] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In *ACM CCS*, 2017.
- [76] Ingo Wegener. *The complexity of Boolean functions*. Wiley-Teubner series in computer science. B. G. Teubner, 1987.