DUPLO: Unifying Cut-and-Choose for Garbled Circuits*

Vladimir Kolesnikov Bell Labs vladimir.kolesnikov@nokia-bell-labs. com Jesper Buus Nielsen Aarhus University jbn@cs.au.dk Mike Rosulek Oregon State University rosulekm@eecs.oregonstate.edu

Roberto Trifiletti Aarhus University roberto@cs.au.dk

ABSTRACT

Cut-and-choose (C&C) is the standard approach to making Yao's garbled circuit two-party computation (2PC) protocol secure against malicious adversaries. Traditional cut-and-choose operates at the level of entire circuits, whereas the LEGO paradigm (Nielsen & Orlandi, TCC 2009) achieves asymptotic improvements by performing cut-and-choose at the level of individual gates. In this work we propose a unified approach called DUPLO that spans the entire continuum between these two extremes. The cut-and-choose step in our protocol operates on the level of arbitrary circuit "components," which can range in size from a single gate to the entire circuit itself.

Ni Trieu

Oregon State University

trieun@eecs.oregonstate.edu

With this entire continuum of parameter values at our disposal, we find that the best way to scale 2PC to computations of realistic size is to use C&C components of intermediate size, and not at the extremes. On computations requiring several millions of gates or more, our more general approach to C&C gives between 4-7x improvement over existing approaches.

In addition to our technical contributions of modifying and optimizing previous protocol techniques to work with general C&C components, we also provide an extension of the recent Frigate circuit compiler (Mood et al, EuroS&P 2016) to effectively express any C-style program in terms of components which can be processed efficiently using our protocol.

CCS CONCEPTS

• Theory of computation → Cryptographic protocols; • Security and privacy → Cryptography; • General and reference → Performance;

CCS'17, , Oct. 30-Nov. 3, 2017, Dallas, TX, USA.

© 2017 Association for Computing Machinery.

ACM ISBN ISBN 978-1-4503-4946-8/17/10...\$15.00

https://doi.org/http://dx.doi.org/10.1145//3133956.3133991

KEYWORDS

Garbled Circuits, Cut-and-Choose, 2PC, UC-secure, Malicious adversary, Implementation, Cryptographic Protocol

1 INTRODUCTION

Garbled Circuits (GC) are currently the most common technique for practical two-party secure computation (2PC). GC has advantages of high performance, low round complexity, low latency, and, importantly, relative engineering simplicity. Both the core garbling technique itself and its application in higher level protocols have been the subject of significant improvement. In the semi-honest model, there have been relatively few asymptotic/qualitative improvements since the original protocols of Yao [59] and Goldreich et al. [18]. The more challenging task of providing security in the presence of malicious parties has seen more striking improvements, such as reducing the number of garbled circuits needed for cutand-choose [32-34, 51], exploring trade-offs between online and offline computation phases [24, 36], and exploring slight weakenings of security [3, 27, 28, 40]. These improvements have brought the malicious security setting to a polished state of affairs, and even small-factor performance improvements are rare.

Cut-and-choose. The focus of this work is to unify two leading approaches for malicious security in GC-based protocols, by viewing them as extreme points on a single continuum. We will find that optimal performance — often significantly better than the state-of-the-art — is generally found somewhere in the middle of the continuum. We start with reviewing the idea of cut-and-choose (C&C) and the two existing approaches which we generalize.

Whole-circuit C&C. Recall, Yao's basic GC protocol is not secure against a cheating GC generator, who can submit a maliciously garbled circuit. Today, C&C is the standard tool in achieving malicious security in secure computation. At the high level, it proceeds in two phases.

C&C phase. The GC generator generates a number of garbled circuits and sends them to GC evaluator, who chooses a subset of them (say, half) at random to be opened (with the help of the generator) and verifies their correctness.

Evaluation phase. If all opened circuits were constructed correctly, the players proceed to securely evaluate the unopened circuits, and take the majority (or other protocol-prescribed) output.

A statistical analysis shows that the probability of the GC generator violating security (by making the evaluator accept an incorrect output) is exponentially small in the number of circuits n.

^{*}The first author was supported by Office of Naval Research (ONR) contract number N00014-14-C-0113. The second author has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement #731583 (SODA). Third and fourth author partially supported by NSF awards #1149647 and #1617197. The fifth author has received funding from the European research Council (ERC) under the European Unions's Horizon 2020 research and innovation programme under grant #669255 (MPCPRO).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Significant progress has been made [7, 23, 24, 32, 36] in reducing the concrete value of *n* needed to achieve a given failure probability. Specifically, if the evaluation phase of the protocol requires a *majority* of unopened circuits to be correct (as in [51]), then $\sim 3s$ circuits are required in total for statistical security 2^{-s} . If the evaluation phase merely requires at least one unopened circuit to be correct (e.g., [7, 32]), then only s circuits are required for the same security. This multiplicative overhead in garbling material due to replication, the replication factor, in the above protocols is 3s and s, respectively. In the amortized setting where parties perform Nindependent evaluations of the same circuit, all evaluations can share a common C&C phase where only a small fraction of circuits needs to be opened. Here, the (amortized) replication factor per evaluation is $O(1) + O(s/\log N)$ for statistical security 2^{-s} [24, 36]. As an example, for N = 1024 and s = 40 the amortized replication factor is around 5.

LEGO. The LEGO paradigm (Large Efficient Garbled-circuit Optimization), introduced by Nielsen & Orlandi [44], works somewhat differently. First, the generator produces many independent *garbled gates* (e.g., NAND gates). Similarly to the whole-circuit C&C, the evaluator chooses a random subset of these gates to be opened and checked. Now, the evaluator randomly assigns the unopened gates into *buckets*. The garbled gates in each bucket are carefully combined in a certain way, so that, as long as a majority of gates in each bucket are correct, the bucket as a whole behaves like a *correct* logical garbled NAND gate. These buckets are then assembled into the final garbled circuit, which is finally evaluated.

The extra step in the LEGO paradigm of randomly assigning unopened gates into buckets improves the protocol's asymptotic replication factor. More precisely, if the evaluated function has N gates, then the LEGO protocol has replication factor $2 + O(s/\log N)$ for security 2^{-s} (compared to s or 3s for conventional whole-circuit C&C). The main disadvantage of the LEGO approach is that there is a nontrivial cost to connect independently generated gates together ("soldering," in LEGO terminology). Since soldering needs to be performed for each wire of the Boolean circuit, LEGO's asymptotic advantages overtake whole-circuit C&C in performance only for circuits of large size. In Section 3 we give more details about the LEGO paradigm.

1.1 DUPLO: building garbled circuits from big pieces

We introduce DUPLO (DUPLO Unifying Procedure for LEGO), a new approach for malicious-secure two-party computation.

As discussed above, the two standard approaches for malicioussecure 2PC perform C&C at the level of entire circuits (whether in the single-execution setting or in the multi-execution setting [37, 49]), or at the level of individual gates (LEGO). DUPLO is a single unifying approach that spans the *entire continuum between these extremes*. The DUPLO approach performs C&C at the level of arbitrary garbled subcircuits (which we refer to as *components*). After the C&C phase has completed, the parties can use the resulting garbled components in any number of 2PC executions, of any (possibly different) circuits that can be built from these components.

What is the value in generalizing C&C in this way? In short, the DUPLO approach unlocks a new degree of freedom in optimizing

practical secure computation. To understand its role, we first review in more detail the costs associated with the C&C techniques (including LEGO).

The most obvious (and often the most significant) cost is the GC replication factor, discussed above. When evaluating a function consisting of *N* components (either entire circuits, gates, or generalized components explored in this work), the replication factor is $O(1) + O(s/\log N)$, for desired security 2^{-s} . Clearly, using smaller components improves the replication factor, since *N* is increased.

The replication factor converges to a lower limit of 2 [63]. As the number of components grows, the benefit of amortization quickly reaches its effective maximum. With practical parameters, there is little improvement to be gained beyond a few million components.

It is when the number of components is "maxed out" that the flexibility of DUPLO starts to have its most pronounced effect. There will be a wide range of different component sizes that all give roughly the same replication factor. Among these choices for component size, it is now best to choose the *largest*, thereby reducing the *cost of soldering*, or connecting the components. This cost is proportional to the number of input/output wires of a component (whole-circuit C&C can be also seen this way, since we have special processing for the inputs and outputs). When a circuit is decomposed into *larger* components, a smaller fraction of wires will cross a boundary between components and therefore require soldering.

In other words, we expect a "sweet spot" for ideal component size, and for computations of realistic size this sweet spot is expected to be between the extremes of gate-level and whole-circuit components. We confirm this analysis by the empirical performance of our prototype implementation. We indeed find such a "sweet spot" between the extremes of component size, as we start considering computations with millions of gates. For these realistic problem sizes, the DUPLO approach improves performance by 4-7x over gate-based and circuit-based C&C. Details are given in Section 7.

Is it realistic to express computations in terms of moderately sized components? We note that the C&C components need to garble identical circuits, i.e. be interchangeable in GC evaluation. Indeed, all NAND gates in LEGO and all circuits in whole-circuit C&C are interchangeable in the sense that they are garblings of the same functionality. One may rightly ask whether it is reasonable to expect realistic computations to be naturally decomposable into interchangeable and non-trivial (i.e. not a single-gate or entirecircuit) subcircuits.

We argue that this is indeed a frequent occurrence in standard (insecure) computation. Standard programming language constructs (standard-size arithmetic operations, subroutine calls, loops, etc.) naturally generate identical subcircuits. Given the recent and growing tendency to automate circuit generation and to build 2PC compilers for higher-level languages [21, 38, 39, 43, 61], it is natural to presume that many practical circuits evaluated by 2PC will incorporate many identical components. Specifically, consider the following scenarios:

• Circuits compiled from higher level languages containing multiple calls to the same subroutine (e.g. algebraic calculations), loops, etc. For example, a boolean circuit for matrix multiplication can be expressed in terms of subcircuits for multiplication and addition.

- Two parties know they will perform many secure computations of a CBC-MAC-based construction (e.g., CMAC) using AES as the block cipher, where one party provides the key and the other provides the message to be authenticated. They can use the AES circuit (or a CBC-composition of several AES circuits) as the main DUPLO component, and use as many components as needed for each evaluation of CMAC. Another example involving AES is to consider the AES round function as the DUPLO component. As this is the same function used internally in AES-128, AES-192 and AES-256 (only the key schedule and number of rounds differ) this preprocessing becomes more independent of the final functionality.
- · Two parties agree on a predetermined low-level instruction set, where for each instruction (represented as a circuit), the parties can produce a large number of preprocessed garbled components without knowing a priori the final programs/functionalities to be computed securely. This CPU/ALU emulation setting has recently been considered in the context of secure computation of MIPS assembly programs [54, 58]. The DUPLO approach elegantly and efficiently provides a way to elevate these results to the malicious setting.

In Section 7 we investigate several of these scenarios in detail, and compare our performance to that of previous work.

1.2 **Related work**

Maliciously secure 2PC using Yao's garbled circuit technique has seen dramatic improvements in recent years, both algorithmic/theoretical The main contribution of the paper is putting forward and techniand implementations. Since the first implementation in [35], tremendous effort has been put into improving concrete efficiency [2, 7, 13, 16, 22-25, 31-34, 36, 37, 41, 44, 46, 47, 49, 51, 52, 55-57, 63] yielding current state-of-the-art prototypes able to securely evaluate an AES-128 computation in 6 ms (multi-execution) or 65 ms (singleexecution). Multi-execution refers to evaluating the same function several times (either in serial or parallel) on distinctly chosen inputs while the more general single-execution setting treats the computation atomically. In addition, some of these protocols allow for dividing the computation into different phases to utilize preprocessing. In the most general case the computation can be split into three consecutively dependent phases. Following the convention of [46] we have:

- Function-independent preprocessing depends only on the statistical and computational security parameters s and k. It typically prepares a given number of gates/components that can be used for later computation.
- Function-dependent preprocessing uses the previously computed raw function-independent material and stitches it together to compute the desired function f.
- Online/Eval phase lastly depends on the parties inputs to the actual computation and is typically much lighter than the previous two phases.

Of notable interest are the protocols of [49] and [55] which represent the current state-of-the-art protocols/prototypes for the multi- and single-execution settings, respectively. Both protocols also support function-dependent preprocessing. With regards to constant-round function-independent preprocessing the works of

[46, 56, 63] are the most efficient, however at this time only the work of [46] provides a public prototype implementation.

The idea of connecting distinct garbled circuits has also previously been studied in [42] by mapping previous output garbled values to garbled input values in a following computation. Their model and approach is different from ours and is mainly motivated by enabling garbled state to be reusable for multiple computations. Finally we point out the recent work of [20] for the semi-honest case of secure 2PC using garbled circuits. [20] likewise considers splitting the function of interest into sub-circuits and processes these independently. As there is no cut-and-choose overhead in the semi-honest setting, their approach is motivated primarily by allowing function-independent preprocessing using the garbled components as building blocks. Although the high-level idea is similar to ours, we apply it in a completely different setting and use different techniques. Further, while malicious security is often significantly more expensive, the efficiency gap in the linking and online phase between [20] and our protocol is surprisingly small. In the application of computing an AES-128 (by preprocessing the required round functions) we see that [20] sends 82 kB in the online phase (link + evaluate) vs. 88 kB using our protocol. For the offline step the gap is larger due to the overhead of C&C in the malicious case. However utilizing amortization this can be reduced significantly and in some cases be as low as 3-5x that of the semi-honest protocols.

Our Contributions and Outline of the Work 1.3

cally and experimentally supporting the idea of generalizing C&C protocols to arbitrary subcircuits. Due to the generality of the approach and the performance benefits we demonstrate, we believe the DUPLO approach will be the standard technique in 2PC compilers. As a lower-level technical contribution, we propose several improvements to garbling and soldering for this setting.

We implemented our solution and integrated it with the state-ofthe-art compiler framework Frigate [43]. Experimentally, we report of a 4-7x improvement in total running time compared to [55] for certain circuits. For the multi-execution setting we also improve the performance of [49] by up to $5 \times$ in total running time. We accomplish the above while at the same time retaining the desirable preprocessing and reactive capabilities of LEGO.

PRELIMINARIES 2

Our DUPLO protocol is a protocol for 2PC that is secure in the presence of malicious adversaries. We define security for 2PC using the framework of Universal Composition (UC), due to Canetti [8]. This framework is demanding, as it guarantees security when such protocols are executed concurrently, in arbitrary environments like the Internet.

A detailed treatment of UC security is beyond the scope of this work. At the high level, security is defined in the real-ideal paradigm. We imagine an ideal interaction, in which parties give their inputs to a trusted third party who computes the desired function f and announces the result. In this interaction, the only thing a malicious party can do is select its input to f. In the real interaction, honest parties interact following the prescribed protocol, while malicious

parties may arbitrarily deviate from the protocol. We say that the protocol *securely realizes f* if the real world is "as secure as" the ideal world. More formally, for every adversary attacking the real protocol, there is an adversary (called "simulator") "attacking" the ideal interaction achieving the same effect.

We assume some familiarity with modern garbled circuit constructions, in particular, the *Free-XOR* optimization of Kolesnikov & Schneider [30]. This is reviewed in Section 3. Free-XOR garbled circuits are secure under a circular correlation-robust hash assumption [11].

3 OVERVIEW OF THE LEGO PARADIGM

We now give more details about the mechanics of the LEGO paradigm. Here we describe the MiniLEGO approach of [13]. We chose MiniLEGO as it is the simplest LEGO protocol to present. At the same time, it contains and conveys all relevant aspects of the paradigm.

3.1 Soldering via XOR-Homomorphic Commitments

The sender generates many individual garbled NAND gates. Each garbled gate g is associated with wire $labels L_g^0, L_g^1$ for the left input wire, labels R_g^0, R_g^1 for the right input wire, and labels O_g^0, O_g^1 for the output wire. Here the superscript of each label indicates the truth value that it represents. In MiniLEGO, all gates are garbled using the Free-XOR optimization of Kolesnikov & Schneider [30]. Therefore, there is a global (secret) value Δ so that $L_g^1 = L_g^0 \oplus \Delta$ and $R_g^1 = R_g^0 \oplus \Delta$ and $O_g^1 = O_g^0 \oplus \Delta$. More generally, a wire label K_g^b can be written as $K_g^b = K_g^0 \oplus b \cdot \Delta$. Importantly, the same Δ is used for *all* garbled gates.

The garbled gate consists of the garbled table itself (i.e., for a single NAND gate, the garbled table consists of two ciphertexts when using the scheme of [62]) along with **XOR-homomorphic commitments** to the "zero" wire labels L_g^0 , R_g^0 , and O_g^0 . A global homomorphic commitment to Δ is also generated and shared among all gates.

To assemble assorted garbled gates into a circuit, the LEGO paradigm uses a technique called **soldering**. Imagine two wires (attached to two unrelated garbled gates) whose zero-keys are A^0 and B^0 , respectively. The sender can "solder" these wires together by decommiting to $S = A^0 \oplus B^0$. We require that such a decommitment can be performed given separate commitments to A^0 and B^0 , and that the decommitment reveals no more than *S*. Importantly, *S* is enough information to allow the receiver to transfer a garbled truth value from the first wire to the second (and vice-versa). For example, if the receiver holds wire label A^b (for unknown *b*), he can compute

$$A^b \oplus S = (A^0 \oplus b \cdot \Delta) \oplus S = B^0 \oplus b \cdot \Delta = B^b$$

which is the garbled encoding of the same truth value, but on the other wire.

Gates are assigned to buckets by the receiver, where each bucket, while possibly containing malicious gates, will be assembled to correctly implement the NAND gate. For the gates inside a bucket, the sender therefore solders all their left wires together, all their right wires together, and all their output wires together with the effect that the bucket can operate on a single set of input labels and produce a single set of output labels. For β gates in a bucket, this gives β ways to evaluate the first gate (use solder values to transfer its garbled inputs to the *i*th bucket gate, evaluate it, then transfer the result back to the first gate). In the most basic form of LEGO, the cut-and-choose ensures that the majority of gates within the bucket are good. Hence the evaluator can evaluate the bucket in β ways and take the majority output wire label. Each bucket therefore logically behaves like a *correct* garbled gate.

The buckets are then assembled into a complete garbled circuit by soldering output wires of one bucket to the input wires of another.

3.2 Recent LEGO Improvements

In recent years several improvements to the LEGO approach has been proposed in the literature. The TinyLEGO protocol [14] provide several concrete optimizations to the above MiniLEGO protocol, most notably a more efficient bucketing technique. The subsequent implementation [46] further optimized the protocol and showed that, combined with the XOR-homomorphic commitment scheme of [10, 15], the LEGO paradigm is competitive with previous state-of-the-art protocols for malicious 2PC, in particular in scenarios where preprocessing is applicable.

In addition to the above works, the protocol of [63] also explores optimizations of LEGO using a different soldering primitive, dubbed XOR-Homomorphic Interactive Hash (XOR-HIH). This technique has a number of advantages over commitments as they allow for a better probability than MiniLEGO and TinyLEGO of catching cheating in the C&C phase. XOR-HIH also yields buckets only requiring a single "correct" gate, whereas MiniLEGO requires a majority and TinyLEGO requires a mixed majority of gates and wire authenticator gadgets. However, due to the communication complexity of the proposed XOR-HIH instantiation being larger than that of the [10, 15] commitment schemes, the overall communication complexity of [63] is currently larger than that of TinyLEGO.

4 OVERVIEW OF OUR CONSTRUCTION

DUPLO protocol big picture. At the high level, our idea is to extend the LEGO paradigm to support components of arbitrary size and distinct functionalities, rather than just a single kind of component that is either a single gate or the entire circuit. The approach is similar in many ways to the LEGO protocol and is broken up into three phases.

In the *function-independent phase*, since some subroutines can be known before fixing the final computed function, the garbler generates many independent garblings of each kind of component, along with related commitments required for soldering. For each kind of component, the parties perform a cut-and-choose over all garbled components. The receiver asks the garbler to open some fraction of these components, which are checked for correctness. The remaining components are assembled randomly into *buckets*. The soldering required to connect components into a bucket is done at this step.

In the *function-dependent phase*, the parties agree on circuits that can be assembled from the available components. The parties perform soldering that connects different buckets together, forming the desired circuits.

In the *online phase*, the parties have chosen their inputs for an evaluation of one of the assembled circuits. They perform oblivious transfers for the evaluator to receive its garbled input, and the garbler also releases its own garbled inputs. The evaluator then evaluates the DUPLO garbled circuit and receives the result.

Challenges and New Techniques. The seemingly simple high-level idea described above encounters several significant technical challenges in its realization. We address the issues in detail in Section 5. Here we mention that the main challenge is that the LEGO paradigm uses the same Free-XOR offset Δ for all garbled components, and its soldering technique crucially relies on this fact. This is not problematic when components are single gates, but turns out to lead to scalability issues for larger components. As a result, we must change the fundamental garbling procedure, and therefore change the soldering approach.

The TinyLEGO approach uses an *input recovery* technique inspired by [32]. The idea is that if the garbler cheats in some components, then the resulting garbled circuit will either give the correct garbled output, or else it will leak the garbler's entire input! In the latter case, the evaluator can simply evaluate the function in the clear. As above, the TinyLEGO approach to this input recovery technique relies subtly on the fact that the components are small, and as a result it does not scale for large components. We introduce an elegant new technique that works for components of any size, and improves the concrete cost of the input recovery mechanism.

Implementation, Evaluation, Integration. We implemented a highperformance prototype of our protocol to explore the effect of varying component sizes in the C&C paradigm. We study a variety of scenarios and parameter choices and find that our generalizations of C&C can lead to significant performance improvement. Details are given in Section 7.

We have adapted the Frigate circuit compiler of Mood et al. [43], which compiles a variant of C into circuits suitable for garbled circuit 2PC applications. We modified Frigate so that subroutines are treated as DUPLO components. As an example, a CBC-MAC algorithm that makes calls to an AES subroutine will be compiled into an "outer circuit" built from atomic AES components, as well as an "inner circuit" that implements the AES component from boolean gates. In our implementation, the inner circuits are then garbled as DUPLO components, and the outer circuits are used to assemble the components into high-level functionalities.

5 DUPLO PROTOCOL DETAILS

We now give more details about the challenges in generalizing the LEGO paradigm, and our techniques to overcome them.

5.1 Different Δ 's

The most efficient garbling schemes use the Free-XOR optimization of [30]. MiniLEGO/TinyLEGO are compatible with Free-XOR, and in fact they enforce that all garbled gates use the same global Free-XOR difference Δ . However, having a common Δ does lead to some drawbacks. In particular, consider the part of the cut-andchoose step in which the receiver chooses some garbled gates to be opened/checked. If we fully open a garbled gate, both wire labels are revealed for each wire. In MiniLEGO, this would reveal Δ and compromise the security of the *unopened* gates, which share the same Δ . To avoid this, the MiniLEGO approach is to make the sender reveal only *one out of the four* possible input combinations to each opened gate (by homomorphically decommitting to the input wire labels). Note that the receiver may now have only a 1/4 probability of detecting an incorrectly garbled gate (the technique of [63] improves this probability to 1/2). The cut-and-choose analysis must account for this probability.

This approach of only partially opening garbled gates does not scale well for large components. If a component has *n* input wires, then the receiver will detect bad components with probability $1/2^n$ in the worst case. In the DUPLO protocol, we garble each component *c* with a separate Free-XOR offset Δ_c (so each gate inside the garbled component uses Δ_c , but other garbled components use different offset). Hence, DUPLO components can be *fully opened* in the cut-and-choose phase, while XOR gates are still free inside each component.

As a result:

- Bad components are detected with probability 1, so the statistical analysis for DUPLO cut-and-choose is better than Mini/TinyLEGO by a constant factor.
- We can use a variant of the optimization suggested in [19] to save bandwidth for cut-and-choose. Initially the sender only sends a short hash of each garbled component. Then to open a component, the sender decommits to the input and output keys as well as the Δ_c used for garbling the component. Hence, communication for the opened components is minimal.

Adapting soldering. It remains to describe how to adapt the soldering procedure to solder wires with different Free-XOR offsets (the MiniLEGO approach relies on the offsets being the same). Here we adapt a technique of [1] for soldering wires. Using the pointand-permute technique for garbled circuits [4], the two wire labels for each wire have random and opposite least-significant bits. We refer to this bit as the *color bit* for a wire label. The evaluator sees the color bit of a wire, but not the truth value of a wire.

In MiniLEGO, the garbler commits to the "zero-key" for each wire, which is the wire label encoding *truth value false*. In DUPLO, we have the garbler generate homomorphic commitments to the following:

- For each wire, commit to the wire label with color bit zero. In this section we therefore use notation K^b to denote a wire label with color bit (not necessarily truth value) b.
- For each wire, commit to an indicator bit σ for each wire that denotes the color bit of the *false* wire label. Hence, wire label K^b has truth value $b \oplus \sigma$.
- For each component *c*, commit to its Free-XOR offset Δ_c .

Consider a wire *i* with labels $(K_i^0, K_i^1 = K_i^0 \oplus \Delta_i)$ and indicator bit σ_i , and another wire *j* in a different component with labels $(K_j^0, K_j^1 = K_j^0 \oplus \Delta_j)$ and indicator bit σ_j . To solder these wires together, the garbler will give homomorphic decommitments to the following solder values:

$$s^{\sigma} = \sigma_i \oplus \sigma_j; \quad S^K = K^0_i \oplus K^0_j \oplus s^{\sigma} \cdot \Delta_j; \quad S^{\Delta} = \Delta_i \oplus \Delta_j$$

Note that the decommitment to S^{Δ} can be reused for all wires soldered between these two components. Now when the evaluator learns wire label K_i^b (with color bit *b* visible), he can compute:

$$\begin{split} K_i^b \oplus S^K \oplus b \cdot S^\Delta &= K_i^b \oplus (K_i^0 \oplus K_j^0 \oplus s^\sigma \cdot \Delta_j) \oplus b \cdot (\Delta_i \oplus \Delta_j) \\ &= b \cdot \Delta_i \oplus (K_j^0 \oplus s^\sigma \cdot \Delta_j) \oplus b \cdot \Delta_i \oplus b \cdot \Delta_j \\ &= K_j^0 \oplus (s^\sigma \oplus b) \cdot \Delta_j = K_j^{s^\sigma \oplus b} \end{split}$$

Also note that a common truth value has opposite color bits on wires *i* & *j* if and only if $s^{\sigma} = \sigma_i \oplus \sigma_j = 1$. Hence, the receiver obtains the wire label $K_j^{s^{\sigma} \oplus b}$ which encodes the same truth value as K_i^b .

DUPLO bucketing. In Section 3.1 we described how [13] used a bucket size that guaranteed a majority of correct AND gates in each bucket. In this work we use the original bucketing technique of [44] that only requires a single correct component in each bucket, but requires a majority bucket of wire authenticator (WA) gadgets on each output wire. The purpose of a WA is to accept or reject a wire label as "valid" without revealing the semantic value on the wire, and as such a simple construction can be based on a hash function and C&C. A WA consists of a "soldering point" (homomorphic commitments to a Δ and a zero-key), along with an unordered pair $\{\mathcal{H}(K_i^0), \mathcal{H}(K_i^0 \oplus \Delta)\}$. A wire label *K* can be authenticated checking for membership $\mathcal{H}(K) \in \{\mathcal{H}(K_i^0), \mathcal{H}(K_i^0 \oplus \Delta)\}$. In order to defeat cheating a C&C step is carried out on the WAs to ensure that a majority of any WA bucket only accepts committed wire labels. The choice of using WAs in this work is motivated by the fact that DUPLO components can be of arbitrary size and are often much larger than a single gate. By requiring fewer such components in total, we therefore achieve much better overall performance as WAs are significantly cheaper to produce in comparison to garbled components.

Avoiding commitments to single bits. We also point out that the separate commitments to the zero-label K_i^0 and the indicator bit σ_i can be combined into a single commitment. The main idea is that the least significant bit of K_i^0 is always zero (being the wire label with color bit zero). Similarly, when using Free-XOR, the offset Δ must always have least significant bit 1. Hence in the solder values S and S^{Δ} , the evaluator knows *a priori* what the least significant bit will be. We can instead use the least significant bit will be. We can instead use the least significant bits of the K_i^0 commitments to store the indicator bit σ_i so that homomorphic openings convey $\sigma_i \oplus \sigma_j$. This approach saves *s* bits of communication per wire commitment over the naive approach of instantiating the bit-commitments using [15] using a bit-repetition code with length *s*.

In the online evaluation phase, the garbler decommits to the indicator bits of the evaluators designated input and output. In this case, the garbler does not want to decommit the entire wire label as this would potentially let the evaluator learn the global difference Δ (if the evaluator learned the opposite label through the OTs or evaluation). To avoid this, we have the garbler generate many commitments to values of the special form $R \parallel 0$ for random $R \in \{0, 1\}^{\kappa-1}$. Using the homomorphic properties of these commitments, this can be done efficiently by having the garbler decommit *s* random linear combinations of these commitments to ensure that

all of them have the desired form with probability $1 - 2^{-s}$. Then when the garbler wants to decommit to a wire label's indicator bit only, it gives a homomorphic decommitment to the wire label XOR a mask R||0, which hides everything but the indicator bit.

5.2 Improved Techniques for Circuit Inputs

We also present a new, more efficient technique for input recovery. The idea of input recovery [32] is that if the sender in a 2PC protocol cheats, the receiver will learn the sender's input (and can hence compute the function output).

Within each DUPLO bucket, the cut-and-choose guarantees at least one correctly garbled component and a majority of correct output-wire authenticators. As such, the evaluator is guaranteed to learn, for each output wire of a component, either 1 or 2 *valid* garbled outputs. If only one garbled output is obtained, then it is guaranteed to be the correct one. Otherwise, the receiver learns both wire labels and hence the Free-XOR offset Δ_c for that component. The receiver can then use the solder values to iteratively learn *both* wire labels on *all* wires in the circuit (at least all the wires in the connected component in which the sender cheated).

However, knowing both wire labels does not necessarily guarantee that the receiver learns their corresponding *truth values*. We need a mechanism so that the receiver learns the truth value for the sender's garbled inputs.

Our approach is to consider special input-components. These consist of an *empty garbled circuit* but homomorphic commitments to a zero-wire-label K and a Free-XOR offset Δ that serve as soldering points. Suppose for every input to the circuit, we use such an input component that is soldered to other components. The sender gives his initial garbled input by homomorphically decommitting to either the zero wire-label K or $K \oplus \Delta$. If the sender cheats within the computation, the receiver will learn Δ . The key novelty in our approach is to use **self-authenticating wire labels**. In an input-gadget, the false wire label must be $H(\Delta) \oplus \Delta$ (the sender will still commit to whichever has color bit zero). Then when the sender cheats, the receiver learns Δ , and can determine whether the sender initially opened $H(\Delta)$ (false) or $H(\Delta) \oplus \Delta$ (true).

This special form of wire labels can be checked in the cut-andchoose for input components. In the final circuit, we assemble input-components into buckets to guarantee that a majority within each bucket is correct. Then the receiver can extract a cheating sender's input according to the majority of input-components in a bucket.

5.3 Formal Description, Security

Our protocol implements secure reactive two-party computation [45], *i.e.*, the computation has several rounds of secret inputs and secret outputs, and future inputs and as well as the specification of future computations might depend on previous outputs.

To be more precise, let \mathcal{F} denote the ideal functionality $\mathcal{F}_{R2PC}^{\mathbb{L},\Phi}$ in Fig. 9 on page 1040 in [45]. Recall that this functionality allows to specify a reactive computation by dynamically specifying the functionality of sub-circuits and how they are linked together. The command (FUNC, t, f) specifies that the sub-circuit identified by thas circuit f. The command (INPUT, t, i, x) gives input x to wire i on sub-circuit *t*. Only one party supplies *x*, the other party inputs (INPUT, *t*, *i*, ?) to instruct \mathcal{F} that the other party is allowed to give an input to the specified wire. The command defines the wire to have value *x*. The command (LINK, t_1, i_1, t_2, i_2) specifies that output wire i_1 of sub-circuit t_1 should be soldered on input wire i_2 of sub-circuit t_2 . When an output value becomes defined to some *x*, this in turn defines the linked input wire to also have value *x*. The command (GARBLE, *t*, *f*) evaluates the sub-circuit *t*. It assumes that all the input wires have already been defined. It runs *f* on these values and defines the output wires to the outputs of *f*. There are also output commands that allow to output the value of a wire to a given party. They may be called only on wires that had their value defined.

The set \mathbb{L} allows to restrict the set of legal sequences of calls to the functionality. We need the restriction that all (FUNC, t, f) commands are given before any other command. This allows us to compute how many times each f is used and do our preprocessing. The function Φ allows to specify how much information about the inputs and outputs of \mathcal{F} is allowed to leak to the adversary. We need the standard setting that we leak the entire sequence of inputs and outputs to the adversary, except that when an honest party has input (INPUT, t, i, x), then we only leak (INPUT, t, i, ?) and when an honest party has output (OUTPUT, t, i, y), then we only leak (OUTPUT, t, i, ?).

With many components, many buckets, and many 2PC executions, the formal description of our protocol is rather involved. It is therefore deferred to Appendix A while we in the full version [29] prove the following theorem.

THEOREM 5.1. Our protocol implements \mathcal{F} in the UC framework against a static, poly-time adversary.

6 SYSTEM FRAMEWORK

In this section we give an overview of the DUPLO framework and our extension to the Frigate compiler that allows to transform a high-level C-style program into a set of boolean circuit components that can be fed to the DUPLO system for secure computation. We base our protocol on the recent TinyLEGO protocol [14], but adapted for supporting larger and distinct components. Our protocol has the the following high-level interface:

- **Setup** A one-time setup phase that initializes the XOR-homomorphic commitment protocol.
- **PreprocessComponent**(n, f) produces *n* garbled representations F_i of *f* that can be securely evaluated.
- **PrepareComponents**(i) produces i input authenticators that can be used to securely transfer input keys from garbler G to evaluator E. In addition, for all F_j previously constructed using PreprocessComponent, this call constructs and attaches all required output authenticators. These gadgets ensure that only a single valid key will flow on each wire of all garbled components (otherwise the evaluator learns the generator's private input).
- **Build**(*C*) Takes a program *C* as input, represented as a DAG where nodes consist of the input/output wires of a set of (possibly distinct) components $\{f_i\}$ and edges consist of links from output wires to input wires for all of these f_i 's. The Build call then looks up all previously constructed F_j for each f_i

and stitches these together using the XOR-homomorphic commitments so that they together securely compute the computation specified by C. This call also precomputes the required oblivious transfers (OTs) for transferring E's input securely.

- **Evaluate**(x, y) Given the plaintext input x of garbler G and y of evaluator E, the parties can now compute a garbled output Z, representing the output of the f(x, y). The system allows both parties to learn the full output, but also distinct output, *e.g.* G can learn the first half of f(x, y) and E learn the second half.
- **Decode** Finally the system allows the parties to decode their designated output. The reason why we have a dedicated decode procedure is to allow partial output decoding. Based on the decoded values the parties can then start a new secure computation on the remaining non-decoded output, potentially adding fresh input as well. The input provided and the new functionality to compute can thus depend on the partially decoded output. This essentially allows branching within the secure computation.

Following the terminology introduced in [46] we have that the Setup, PreprocessComponent, and PrepareComponents calls can be done independently of the final functionality C. These procedures can therefore be used for function-independent preprocessing by restricting the functionality C to be expressible from a predetermined set of instructions. The Build procedure clearly depends on C, but not on the inputs of the final computation, so this phase can implement function-dependent preprocessing. Finally the Evaluate and Decode procedures implement the online phase of the system and depend on the previous two phases to run.

For a detailed pseudocode description of the system as well as a proof of its security we refer the reader to Appendix A and the full version [29], respectively.

6.1 Implementation optimizations

As part of our work we developed a prototype implementation in C++ using the latest advances in secure computation engineering. As the basis for our protocol we start from the libOTe library for efficient oblivious transfer extension [48]. As we in this work require UC XOR-homomorphic commitments to the input and output wires of all components we instantiate our protocol with the efficient construction of [15] and use the implementation of [50] in our prototype.

As already mentioned, our protocol is described in detail in Appendix A. However, for reasons related to efficiency our actual software implementation deviates from the high-level description in several aspects

- In the homomorphic commitment scheme of [15], commitments to random values (chosen by the protocol itself) are cheaper than commitments to values chosen by the sender. Hence, whenever applicable we let the committed key-values be defined in this way. This optimization saves a significant amount of communication since the majority of commitments are to random values.
- Along the same lines we heavily utilize the batch-opening mechanism described in [15]. The optimization allows a

sender to decommit to *n* values with total bandwidth $n\kappa + O(s)$ as opposed to the naive approach which requires $O(n\kappa s)$.

• In the PrepareComponents step we construct all output-wire key authenticators using a single global difference Δ_{ka} . This saves a factor 2x in terms of the required number of commitments and solderings, at the cost of an incorrect authenticator only getting caught with probability 1/2 (as opposed to prob. 1 using distinct differences). However as the number of required key authenticators depends on the total number of output wires of all garbled components the effect of this difference in catching probability does not affect performance significantly when considerings many components.

In addition to the above optimizations, our implementation takes full advantage of modern processors' multi-core capabilities and instruction sets. We also highlight that our code leaves a substantially lighter memory footprint than the implementation of [46] which stores all garbled circuits and commitments in RAM. In addition to bringing down the required number of commitments on the protocol level, our implementation also makes use of disk storage in-between batches of preprocessed component types. This has the downside of requiring disk reads of the garbled components during the online phase, but we advocate that the added flexibility and possibility of streaming preprocessing is well worth this trade-off in performance.

6.2 Frigate Extension

The introduction of Fairplay [39], the first compiler targeted for secure computation (SC), has stimulated significant interest from the research community. Since then, a series of new compilers with enhanced performance and functionality have been proposed, such as CBMC [21], Obliv-C [61], and ObliVM [38]. Importantly, the state-of-the-art compiler, Frigate [43], features a modular and extensible design that simplifies the circuit generation in secure computation. Relying on its rich language features, we provide an extension to the original Frigate framework, in which we divide the **specific** input program into distinct functions. We can then generate a circuit representation for each function which is fully independent from the circuit representation of other functions. Due to this independence we can easily garble each distinct function separately using the DUPLO framework and afterwards solder these back together such that they compute the original source program. As an additional improvement, which is tangential to the main thrust of this work, we construct an AES module that optimizes the number of uneven gates (all even gates can be garbled and evaluated without communication using e.g. [62]).

In the following, we describe the details of our compiler extension. Similar to the Frigate output format, our circuit output contains a set of input and output calls, gate operations, and function calls. The input and output calls consist of wires, which we enumerate and manage. We also use wires to represent declared variables in the source program. Each wire (or, rather its numeric id) is placed in a pool, and is ready for use whenever a new variable is introduced. Our function representation however differs from that of Frigate. In that work, each function reserves a specific set of wire values which requires no overlap among the functions' wires. As a result, Frigate's function representation is dependent on that of other functions. We remove this dependency by creating and managing separate wire pools for each function. In particular, every time a variable is introduced, our compiler searches for the free wires with the smallest indices in the pool of the current working function. Similarly to the original Frigate, our compiler will free the wires it can after each operation or variable assignment. Hence, our function is represented independently of other functions.

We now describe our strategy for constructing our optimized AES circuit. A key component of AES is the Rijndael S-Box [12] which is a fixed non-linear substitution table used in the byte substitution transformation and the key expansion routine. The circuit optimization in our AES-128 source program is described in the context of this S-Box. We note that if we generate the S-Box dynamically using the Frigate compiler, this will not optimize the number of uneven gates substantially. Hence, we create an AES-128 source program that embed a highly optimized S-Box circuit statically. To the best of our knowledge, [6] presents one of the most efficient S-Box circuit representation which contains only 32 uneven gates in a total of 115 gates. Therefore, we integrate this S-Box into our AES-128 source program, which allows our Frigate extension to optimize the number of uneven gates. For the key-expanded AES-128 circuit, which takes a 128-bit plaintext and ten 128-bit round keys as input and outputs a 128-bit ciphertext, this results in 5,120 uneven gates. This is almost a 2x reduction compared the AES-128 circuit originally reported in Frigate. Furthermore, our AES-128 circuit has 640 fewer uneven gates than the circuit reported in Tiny-Garble [54] which is the current best compiler written in Verilog. For completeness we note that for the non-expanded version of AES-128, our compiled circuit results in 6,400 uneven gates.

7 PERFORMANCE

In order to evaluate the performance of our prototype we run a number of experiments on a single server with simulated network bandwidth and latency. The server has two 36-core Intel(R) Xeon(R) E5-2699 v3 2.30 GHz CPUs and 256 GB of RAM. That is, 36 cores and 128 GB of RAM per party. As both parties are run on the same host machine we simulate a LAN and WAN connection using the Linux *tc* command: a LAN setting with 0.02 ms round-trip latency, 1 Gbps network bandwidth; a WAN setting with 96 ms round-trip latency, 200 Mbps network bandwidth.

For both settings, the code was compiled using GCC-5.4. Throughout this section, we performed experiments with a statistical security parameter s = 40 and computational security parameter k = 128. The running times recorded are an average over 10 trials.

We demonstrate the scalability of our implementation by evaluating the following circuits:

AES-128 circuit consisting of 6,400 AND gates. The circuit takes a 128-bit key from one party and a 128-bit block from another party and outputs the 128-bit ciphertext to both. (Note that this functionality is somewhat artificial for secure computation as the AES function allows decryption with the same key; thus the player holding the AES key can obtain the plaintext block. We chose to include the ciphertext output to the keyholder to measure and demonstrate the performance for the case where both parties receive output.)

- **CBC-MAC** circuit with different number of blocks $m \in \{16, 32, 64, 128, 256, 1024\}$ using AES-128 as the block cipher. The circuit therefore consists of 6, 400*m* AND gates. The circuit takes a 128-bit key from one party and *m* 128-bit blocks from another party and outputs a 128-bit block to both.
- **Mat-Mul** circuit consisting of around 4.2 million AND gates. The circuit takes one 16×16 matrix of 32-bit integers from each party as input and outputs the 16×16 matrix product to both.
- **Random** circuit consisting of 2^n AND gates for various *n* where topology of the circuit is chosen at random. The circuit takes 128-bit input from each party and outputs a 128-bit value to both.

7.1 Effect of Decomposition

In this section we show how DUPLO scales for the above-mentioned circuits, when considering subcomponents of varying size. As discussed in Section 1.1, we expect the performance of our protocol to be optimal for a subcomponent size somewhere inbetween the extremes of whole-circuit and gate-level C&C. We empirically validate this hypothesis by running two kinds of experiments, one for the randomly generated circuits and one for the real-world AES-128, CBC-MAC-16 and Mat-Mul circuits. The purpose of the random circuit experiment is to explore the trade-offs in overall performance between different decomposition strategies. For the latter experiment we aim to find their optimal decomposition strategy, both to see how this aligns to the random circuit experiment, but also for use in our later performance comparison in Section 7.2.

Random Circuits. In order to build a random circuit consisting of 2^n AND gates that is easily divisible into different subcomponent sizes we initially generate a number of smaller random circuit containing 2^t AND gates with 256 input wires and 128 output wires. This is done by randomly generating non-connected XOR and AND until exactly 2^t AND gates have been generated. Then for each of these generated gates *i* we assign their two input wires at random from the set of gates with index smaller than *i* (the gate id *i* is also the gate's output wire). Finally we solder 2^{n-t} copies of these components together into a final circuit *C*, thus consisting of 2^n AND gates overall. We consider $n \in \{10, 12, 14, 16, 18, 19, 20\}$ in this experiment, and for each of these we build a circuit of size 2^n using several values of *t*.

As we are only considering relative performance between different strategies in these experiments we evaluate random circuit using a single thread for each party on the previously mentioned LAN setup.¹ We summarize our findings in Figure 1. We initially consider very few, but large subcomponents and then gradually decrease the size. Thus, the x-axis of the figure represents the continuum from whole-circuit C&C (t = n) towards gate-level C&C (t = 0). The overall trend of our experiments is strikingly clear, initially as the number of subcomponents increases (t decreases) the running time goes down as well due to our protocol taking advantage of the amortization benefits offered by the LEGO paradigm. However for all circuit sizes considered it is also apparent that at some point this benefit is outweighed by the overhead of soldering





Figure 1: DUPLO performance for random circuits consisting of 2^n AND gates divided into 2^{n-t} subcomponents.

and committing to the increasing number of input/output wires between the components. It is at exactly this point (the vertex of each graph), in the sweet spot between substantial LEGO amortization and low soldering overhead, that DUPLO has it's optimal performance. We thus conclude that for an ideally decomposable circuit such as the ones generated in this experiment the viability of the DUPLO approach is apparent.

Real-world circuits. The experiments for the random circuits show that the DUPLO approach for C&C does have merit for circuits that can be divided into multiple identical subcomponents. Clearly, this is a very narrow class of functions so in addition we also evaluate our prototype on the previously mentioned real-world circuits in order to investigate their optimal decomposition strategy. We first describe our approach of dividing these circuits into subcomponents.

AES-128 We consider the following three strategies:

- Five kinds of subcomponents: each computing one of the functions of the AES algorithm, that is 1x Key Expansions (1,280 AND gates), 11x AddRoundKey, 10x SubBytes (512 AND gates), 10x ShiftRows, and 9x MixColumns.
- Three kinds of subcomponents: 1x Key Expansions and Initial Round (1,280 AND gates); 9x AES Round Functions (each 512 AND gates); 1x AES Final Round (512 AND gates).
- A single component consisting of the entire AES-128 circuit (6,400 AND gates), *i.e.* whole-circuit C&C.
- **CBC-MAC-16** We consider decomposing this circuit into a single subcomponent of varying size. In each case, the component contains $i \in \{16, 8, 4, 2, 1\}$ AES-128 blocks, meaning each of these consists of 6, 400*i* AND gates.
- **Mat-Mul** In order to multiply two matrices *A*, *B* use the blockmatrix algorithm: We divide *A*, *B* into $m \times m$ 32-bit submatrices $A_{i, j}, B_{i, j}$ for $i, j \in [1, 16/m]$. To compute *AB*, the block



Figure 2: DUPLO performance for N parallel executions of CBC-MAC-16 using different decomposition strategies.

entries $A_{i,k}$ are first multiplied by the block entries $B_{k,j}$ for $k \in [1, m]$, while summing the results over k. It is therefore the case that the experiment contains two different kinds of components, $m \times m$ 32-bit matrix product and $m \times m$ 32-bit matrix addition. In our experiment we consider block matrix sizes $m \in \{16, 8, 4, 2\}$ and the concrete number of AND gates for each kind of component are reported in Table 1.

When performing N = 1, 32, 128 executions of AES-128 in parallel, we observe that our protocol performs best when considering the entire circuit as a single component. For example of performing N = 128, evaluating AES-128 with single component takes 28.18 milliseconds while splitting the AES-128 into three or five relatively small subcomponents, it costs 35.2 and 32.21 milliseconds respectively. This is in contrast to what we observed in the random circuit experiment, but can be explained by the non-uniformity of the considered decomposition strategies. The fact that we split the AES-128 into small subcomponents, some of which are only used once, has a very negative influence on DUPLO performance as there is some overhead associated with preparing each component type while at the same time no LEGO-style amortization can be exploited when preparing only a single copy.

For the CBC-MAC-16 circuit however whole-circuit C&C is not the optimal approach and we summarize the observed performance for the different decompositions in Figure 2. Here we see that the best strategy is to decompose the circuit into many *identical* subcomponents. The trend observed is similar to the random circuit experiments where initially it is best to optimize for many identical subcomponents. In particular for a single execution of CBC-MAC-16 it is best to decompose into 16 copies of the AES-128 circuit yielding around 5x performance increase over the whole-circuit approach. For the parallel executions (which contain overall many more AES-128 circuits) we can see that it is best to consider subcomponents consisting of 4xAES-128 circuits each. The lower relative performance difference between the strategies for the parallel executions is due to there being a minimum of *N* circuits for utilizing LEGO amortization, even for the whole-circuit approach. However



Figure 3: DUPLO performance for N parallel executions of Mat-Mul using different decomposition strategies.

as the number of total subcomponents grow it can be seen that there are savings to be had by grouping executions together.

Block	Componer	nt Size	Number Executions N				
Size	Mult	Add	1	32	128		
2x2	8,192	124	11,160	7,815	7,554		
4x4	65,536	496	14,847	7,539	6,622		
8x8	524,288	1,984	52,334	9,615	7,324		
16x16	4,194,304	0	351,002	11,338	9298		

Table 1: Component sizes and amortized running time per execution for Mat-Mul (ms). Best performance marked in bold.

Finally for the Mat-Mul circuit we see a similar overall trend as in the CBC-MAC-16 experiment (Table 1 and Figure 3). Most notably is the performance increase for a single execution yielding around 31x by considering blocks of size 2×2 instead of a single wholecircuit 16×16 . This experiment indeed highlights the performance potential of the DUPLO approach for large computations that can naturally be decomposed into distinct repeating subcomponents, in this case matrix product and matrix addition. This is in contrast to the previous AES-128 example where this approach was penalized. The difference however is that in the Mat-Mul experiment each subcomponent is repeated several times and therefore all benefit from LEGO amortization.

Experiment Discussion. The above real-world examples show that the DUPLO approach has merit, but the exact performance gains depend significantly on the circuit in question. As a general rule of thumb DUPLO performs best when the circuit can be decomposed into many identical subcomponents as can be seen from the CBC-MAC-16 and Mat-Mul experiments (the more the better). As there is no immediate way of decomposing the AES-128 circuit in this way, we see that performance suffers when the circuit cannot be decomposed into distinct *repeating* parts. However the Mat-Mul experiments show that decomposing the circuit into distinct circuits can certainly have merit, however it is crucial that

Drotocol	Satting	ing N	AES			CBC-MAC-16			Mat-Mul		
	Setting		Ind.Prep	Offline	Online	Ind.Prep	Offline	Online	Ind.Prep	Offline	Online
WMK[55] LAI	LAN	1	X	(X)	125	X	(X)	1,177	X	(X)	43,930
	WAN	1	X	(X)	2,112	X	(X)	11,443	X	(X)	368,190
		1	X	198	8.83	X	3,495	35.52	X	120,200	913
	LAN	32	×	67.77	3.60	×	1,296	20.71	X	36,437	1,247
	LAIN	128	×	40.70	2.86	×	863	18.38	-	-	-
DD [40]		1024	×	24.90	3.06	×	471	17.48	-	-	-
KK [49]	-	1	X	941	527	×	12039	565	X	467,711	1,550
	WAN	32	×	311	472	×	4202	471	X	157,928	1,677
,	WAIN	128	×	192	557	×	2743	573	X	-	-
		1024	X	115	577	X	1762	597	X	-	-
		1	1,506	22.34	2.54	2,594	230	18.14	-	-	-
	LAN	32	119	2.42	0.22	965	38.63	1.34	-	-	-
	LAIN	128	75.64	2.08	0.16	922	37.90	0.87	-	-	-
NST[46]		1024	60.48	1.85	0.14	-	-	-	-	-	-
1831[40]	WAN	1	9,325	223	195	13,812	699	219	-	-	-
		32	599	15.17	6.71	4,158	151	8.54	-	-	-
		128	341	12.75	6.24	3,810	148	7.15	-	-	-
		1024	256	11.81	5.56	-	-	-	-	-	-
	LAN	1	X	371	8.62	799	29.83	41.94	10,268	569	118
DUPLO –		32	×	47.03	0.65	303	10.18	3.72	7,124	331	83.73
		128	×	27.77	0.41	213	11.54	2.49	6,260	303	58.65
		1024	×	17.58	0.30	175	13.30	1.61	-	-	-
	WAN	1	X	7,391	585	8,970	1,370	620	50,856	1,775	744
		32	X	347	19.37	1,477	49.21	23.62	32,098	517	135
		128	X	148	5.55	990	22.23	8.85	27,613	388	101
		1024	×	74.03	1.53	733	15.93	3.83	-	-	-

Table 2: All timings are ms per circuit. Best results are marked in bold. Cells with " λ " denote setting not supported. Cells with "-" denote program out of memory.

Protocol	Setting	CBC-N Offline	IAC-32 Online	CBC-N Offline	IAC-64 Online	CBC-M Offline	AC-128 Online	CBC-M Offline	AC-256 Online	CBC-M. Offline	AC-1024 Online
WMK[55]	LAN	(X)	2,298	(X)	4,539	(X)	9,029	(X)	18,003	(X)	71,787
DUPLO		1,211	68.29	1,877	104	2,991	196	5,072	274	14,167	1,003
WMK[55]	WAN	(X)	21,460	(X)	41,114	(X)	79,157	(X)	155,995	(X)	606,329
DUPLO		12,269	656	15,039	698	19,089	793	27,093	910	81,883	1,733

Table 3: Comparison for CBC-MAC-XX. All timings are ms per circuit. Best results are marked in bold. "(X)" denotes the setting is supported, but we only ran the "everything online" version of WMK.

each subcomponent is repeated a minimal number of times or the non-repeating part of the computation is relatively small.²

7.2 Comparison with Related Work

We also compared our prototype to three related high-performance open-source implementations of malicious-secure 2PC. All experiments use the same hardware configuration described at the beginning of this section. For all experiments we have tried tuning the calling parameters of each implementation to obtain the best performance.

When reporting performance of our DUPLO protocol, we split the offline part of the computation into an independent preprocessing (Setup + PreprocessComponent + PrepareComponents) whenever our analysis shows that dividing the computation into subcomponents is optimal - i.e., when evaluating AES-128 we do not have any function-independent preprocessing since the optimal configuration is to let the component consist of the entire circuit. We summarize our measured timings for all the different protocols in Table 2, and now go into more detail:

Better Amortization by Subdivision. The protocol of Rindal & Rosulek (RR in our tables) [49] is currently the fastest malicious-secure 2PC protocol in the multi-execution setting. The protocol of Nielsen et al. (NST) [46] is the fastest that allow for function-independent preprocessing, using the LEGO paradigm.³

The general trend in Table 2 is that as the total complexity (combined cost of all computations) grows, the efficiency of the DUPLO approach becomes more and more apparent. For example, DUPLO is 1.5x times faster (counting total offline+online time) than RR whole-circuit C&C for 1024 AES-128 LAN. For the larger CBC-MAC-16 scenario, the difference 2.5x. For the even larger case of 32 Mat-Mul executions, the difference is 5x. Our experiments clearly

 $^{^2{\}rm This}$ is not the case for the AES-128 circuit as the non-repeating part consists of around 40% of the entire computation.

³The recent 2PC protocol of [56] appears to surpass NST in terms of performance in this setting, but as this implementation is not publicly available at the time of writing we do not consider it for these experiments. However, it can be indirectly compared timing to WMK via RR.

confirm that DUPLO scales significantly better than state-of-the-art amortizing protocols.

When comparing to the LEGO C&C protocol of NST things are harder to compare as they use a much slower BaseOT implementation than we do (1200 ms vs. 200 ms) which especially matters for lower complexity computations. However even when accounting for this difference, in total time, our approach has 2-3x better total performance for AES-128. We note that if Ind. Prep. is applicable for an application then DUPLO cannot compete with NST for small computations, but as demonstrated from our CBC-MAC-16 experiments, once the computation reaches a certain size and we can decompose the target circuit into smaller subcomponents, DUPLO overtakes NST in performance by a factor 5x.

It is interesting to note that the online time of NST is vastly superior to RR and DUPLO, especially for small circuits (2-4x). This is due to the difference between whole-circuit C&C and gate-level C&C where the NST bucket size is relatively small (and thus online computation) even for a single circuit, whereas it needs to be 5-10x larger for the whole-circuit approach. As the number of executions increase we however see that this gap decreases significantly. We believe the reason why NST still outperforms DUPLO in all online running times is that the NST implementation is RAM only, whereas DUPLO writes components and solderings to disk in the offline phases and reads them in the online phase as needed. For RR we notice some anomalies for their online times that we cannot fully explain. We conclude that the throughput measured and reported in our experiments might not be completely fair towards the RR protocol, but might be explained by implementation decisions that work poorly for our particular scenarios. In any case, we do expect DUPLO to perform as fast or faster than RR in the online phase due to less online rounds and data transfer.

Amortized-grade Performance for Single-Execution. The current fastest protocol for single-execution 2PC is due to Wang et al. (WMK) [55]. When comparing to their protocol, we ran all experiments using the "everything online" version of their code since this typically gives the best overall running time. We stress however that the protocol also supports a function-dependent preprocessing phase, but since this is not the primary goal of that work we omit it here.

Unsurprisingly, the protocols designed for the multi-execution settings (including DUPLO) are significantly faster than WMK when considering several executions. However, even in the singleexecution setting, we see that DUPLO scales better and eventually catches up to the performance of WMK for large computations. WMK is 3x faster than DUPLO when the subcomponent is an entire AES-128 circuit. Then, already for CBC-MAC-16 the ability to decompose this into 16 independent AES-128 circuits yields around 1.4x factor improvement over WMK. We further explore this comparison in Table 3, by evaluating even larger circuits in the single-execution setting. For larger CBC-MAC circuits, DUPLO is around 4.7x faster on LAN and 7.4x on WAN.

7.2.1 Bandwidth Comparison. As a final comparison we also consider the bandwidth requirements of the different protocols. In addition to the previous three protocols we here also include the recent work of Wang et al. (WRK) [56]. To directly compare we report on the data required to transfer from constructor to receiver

Protocol	#Execs	Ind. Prep	Dep. Prep	Online
WMK[55]	1	×	X	9.66 MB
RR [49]	32	×	3.75 MB	25.76 kB
	128	×	2.5 MB	21.31 kB
	1024	×	1.56 MB	16.95 kB
NST [46]	1	14.94 MB	226.86 kB	16.13 kB
	32	8.74 MB	226.86 kB	16.13 kB
	128	7.22 MB	226.86 kB	16.13 kB
	1024	6.42 MB	226.86 kB	16.13 kB
WRK [56]	1	2.86 MB	570 kB	4.86 kB
	32	2.64 MB	570 kB	4.86 kB
	128	2.0 MB	570 kB	4.86 kB
	1024	2.0 MB	570 kB	4.86 kB
DUPLO	1 32 128 1024	× × ×	12.94 MB 2.60 MB 1.96 MB 1.59 MB	19.36 kB 18.97 kB 18.96 kB 18.96 kB

Table 4: Comparison of the data sent from constructor to evaluator AES-128 with k = 128 and s = 40. All numbers are per AES-128. Best results marked in **bold**.

in Table 4 for different number of AES-128 executions. We stress that these numbers are all from the same AES-128 circuit [53] and not from our optimized Frigate version. As already established for AES-128, DUPLO performs best by treating the entire circuit as a single component, hence we do not distinguish between Ind. Prep and Dep. Prep in the table. However we do stress that DUPLO only requires solderings from the input-wires to the output-wires of potentially large components, so for applicable settings we expect the Dep. Prep of DUPLO to be much lower than that of NST and WRK as they require solderings for each gate. It can be seen that for a single AES-128 component DUPLO cannot compare with the protocol of WRK in terms of overall bandwidth. This is natural as the replication factor is much lower for gate-level C&C in this case. However as the number of circuits grows we see that DUPLO's bandwidth requirement decreases significantly per AES-128 to a point where it is actually better than WRK by a factor 1.6x at 1024 executions. For the online phase it is clear that WRK's bandwidth is better than our protocol as we require decommitting the garbled input keys for the evaluator which induces some overhead. However we note that our implementation is not optimal in terms of online bandwidth in that we have chosen flexibility over minimizing rounds and bandwidth. For a dedicated application DUPLO's online bandwidth can be reduced by around 2x by combining the evaluate and decode phases and running batch-decommit of the evaluator input wires along with the output indicator bits.

REFERENCES

- Arash Afshar, Zhangxiang Hu, Payman Mohassel, and Mike Rosulek. 2015. How to Efficiently Evaluate RAM Programs with Malicious Security. In EURO-CRYPT 2015, Part I (LNCS), Elisabeth Oswald and Marc Fischlin (Eds.), Vol. 9056. Springer, 702–729. https://doi.org/10.1007/978-3-662-46800-5_27
- [2] Arash Afshar, Payman Mohassel, Benny Pinkas, and Ben Riva. 2014. Non-Interactive Secure Computation Based on Cut-and-Choose. In EUROCRYPT 2014 (LNCS), Phong Q. Nguyen and Elisabeth Oswald (Eds.), Vol. 8441. Springer, 387– 404. https://doi.org/10.1007/978-3-642-55220-5_22
- [3] Gilad Asharov and Claudio Orlandi. 2012. Calling Out Cheaters: Covert Security with Public Verifiability. In ASIACRYPT 2012 (LNCS), Xiaoyun Wang and Kazue Sako (Eds.), Vol. 7658. Springer, 681–698. https://doi.org/10.1007/

978-3-642-34961-4_41

- [4] Donald Beaver, Silvio Micali, and Phillip Rogaway. 1990. The Round Complexity of Secure Protocols (Extended Abstract). In STOC 1990. ACM Press, 503–513.
- [5] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. 2012. Foundations of garbled circuits, See [60], 784–796.
- [6] Joan Boyar and Rene Peralta. 2009. New logic minimization techniques with applications to cryptology. Cryptology ePrint Archive, Report 2009/191. (2009). http://eprint.iacr.org/2009/191
- [7] Luís T. A. N. Brandão. 2013. Secure Two-Party Computation with Reusable Bit-Commitments, via a Cut-and-Choose with Forge-and-Lose Technique - (Extended Abstract). In ASIACRYPT 2013, Part II (LNCS), Kazue Sako and Palash Sarkar (Eds.), Vol. 8270. Springer, 441–463. https://doi.org/10.1007/978-3-642-42045-0_23
- [8] Ran Canetti. 2001. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In FOCS 2001. IEEE Computer Society Press, 136–145.
- [9] Ran Canetti and Juan A. Garay (Eds.). 2013. CRYPTO 2013, Part II. LNCS, Vol. 8043. Springer.

[10] Ignacio Cascudo, Ivan Damgård, Bernardo David, Nico Döttling, and Jesper Buus Nielsen. 2016. Rate-1, Linear Time and Additively Homomorphic UC Commitments. In CRYPTO 2016, Part III (LNCS), Matthew Robshaw and Jonathan Katz (Eds.), Vol. 9816. Springer, 179–207. https://doi.org/10.1007/978-3-662-53015-3_7

- [11] Seung Geol Choi, Jonathan Katz, Ranjit Kumaresan, and Hong-Sheng Zhou. 2012. On the Security of the "Free-XOR" Technique. In *TCC 2012 (LNCS)*, Ronald Cramer (Ed.), Vol. 7194. Springer, 39–53.
- [12] Joan Daemen and Vincent Rijmen. 2002. The Design of Rijndael: AES The Advanced Encryption Standard. Springer. https://doi.org/10.1007/978-3-662-04722-4
- [13] Tore Kasper Frederiksen, Thomas Pelle Jakobsen, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. 2013. MiniLEGO: Efficient Secure Two-Party Computation from General Assumptions. In EUROCRYPT 2013 (LNCS), Thomas Johansson and Phong Q. Nguyen (Eds.), Vol. 7881. Springer, 537–556. https://doi.org/10.1007/978-3-642-38348-9_32
- [14] Tore Kasper Frederiksen, Thomas P. Jakobsen, Jesper Buus Nielsen, and Roberto Trifiletti. 2015. TinyLEGO: An Interactive Garbling Scheme for Maliciously Secure Two-Party Computation. Cryptology ePrint Archive, Report 2015/309. (2015). http://eprint.iacr.org/2015/309
- [15] Tore Kasper Frederiksen, Thomas P. Jakobsen, Jesper Buus Nielsen, and Roberto Trifiletti. 2016. On the Complexity of Additively Homomorphic UC Commitments. In TCC 2016-A, Part I (LNCS), Eyal Kushilevitz and Tal Malkin (Eds.), Vol. 9562. Springer, 542–565. https://doi.org/10.1007/978-3-662-49096-9_23
- [16] Tore Kasper Frederiksen and Jesper Buus Nielsen. 2013. Fast and Maliciously Secure Two-Party Computation Using the GPU. In ACNS 2013 (LNCS), Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini (Eds.), Vol. 7954. Springer, 339–356. https://doi.org/10.1007/978-3-642-38980-1_21
- [17] Juan A. Garay and Rosario Gennaro (Eds.). 2014. CRYPTO 2014, Part II. LNCS, Vol. 8617. Springer.
- [18] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In STOC 1987, Alfred Aho (Ed.). ACM Press, 218–229.
- [19] Vipul Goyal, Payman Mohassel, and Adam Smith. 2008. Efficient Two Party and Multi Party Computation Against Covert Adversaries. In EUROCRYPT 2008 (LNCS), Nigel P. Smart (Ed.), Vol. 4965. Springer, 289–306.
- [20] Adam Groce, Alex Ledger, Alex J. Malozemoff, and Arkady Yerukhimovich. 2016. CompGC: Efficient Offline/Online Semi-honest Two-party Computation. Cryptology ePrint Archive, Report 2016/458. (2016). http://eprint.iacr.org/2016/458
- [21] Andreas Holzer, Martin Franz, Stefan Katzenbeisser, and Helmut Veith. 2012. Secure two-party computations in ANSI C, See [60], 772–783.
- [22] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. 2011. Faster Secure Two-Party Computation Using Garbled Circuits. In USENIX Security 2011. USENIX Association.
- [23] Yan Huang, Jonathan Katz, and David Evans. 2013. Efficient Secure Two-Party Computation Using Symmetric Cut-and-Choose, See [9], 18–35. https://doi.org/ 10.1007/978-3-642-40084-1_2
- [24] Yan Huang, Jonathan Katz, Vladimir Kolesnikov, Ranjit Kumaresan, and Alex J. Malozemoff. 2014. Amortizing Garbled Circuits, See [17], 458–475. https://doi. org/10.1007/978-3-662-44381-1_26
- [25] Nathaniel Husted, Steven Myers, abhi shelat, and Paul Grubbs. 2013. GPU and CPU parallelization of honest-but-curious secure two-party computation. In ACSAC 2013, Charles N. Payne Jr. (Ed.). ACM, 169–178. https://doi.org/10.1145/ 2523649.2523681
- [26] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2015. Actively Secure OT Extension with Optimal Overhead. In CRYPTO 2015, Part I (LNCS), Rosario Gennaro and Matthew J. B. Robshaw (Eds.), Vol. 9215. Springer, 724–741. https: //doi.org/10.1007/978-3-662-47989-6_35
- [27] Vladimir Kolesnikov and Alex J. Malozemoff. 2015. Public Verifiability in the Covert Model (Almost) for Free. In ASIACRYPT 2015, Part II (LNCS), Tetsu Iwata and Jung Hee Cheon (Eds.), Vol. 9453. Springer, 210–235. https://doi.org/10.1007/ 978-3-662-48800-3 9
- [28] Vladimir Kolesnikov, Payman Mohassel, Ben Riva, and Mike Rosulek. 2015. Richer Efficiency/Security Trade-offs in 2PC. In TCC 2015, Part I (LNCS), Yevgeniy Dodis

and Jesper Buus Nielsen (Eds.), Vol. 9014. Springer, 229–259. https://doi.org/10.1007/978-3-662-46494-6_11

- [29] Vladimir Kolesnikov, Jesper Buus Nielsen, Mike Rosulek, Ni Trieu, and Roberto Trifiletti. 2017. DUPLO: Unifying Cut-and-Choose for Garbled Circuits. Cryptology ePrint Archive, Report 2017/344. (2017). http://eprint.iacr.org/2017/344.
- [30] Vladimir Kolesnikov and Thomas Schneider. 2008. Improved Garbled Circuit: Free XOR Gates and Applications. In ICALP 2008, Part II (LNCS), Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz (Eds.), Vol. 5126. Springer, 486–498.
- [31] Benjamin Kreuter, abhi shelat, and Chih-Hao Shen. 2012. Billion-Gate Secure Computation with Malicious Adversaries. In USENIX Security 2012. USENIX Association.
- [32] Yehuda Lindell. 2013. Fast Cut-and-Choose Based Protocols for Malicious and Covert Adversaries, See [9], 1–17. https://doi.org/10.1007/978-3-642-40084-1_1
- [33] Yehuda Lindell and Benny Pinkas. 2007. An Efficient Protocol for Secure Two-Party Computation in the Presence of Malicious Adversaries. In EUROCRYPT 2007 (LNCS), Moni Naor (Ed.), Vol. 4515. Springer, 52–78.
- [34] Yehuda Lindell and Benny Pinkas. 2011. Secure Two-Party Computation via Cutand-Choose Oblivious Transfer. In TCC 2011 (LNCS), Yuval Ishai (Ed.), Vol. 6597. Springer, 329–346.
- [35] Yehuda Lindell, Benny Pinkas, and Nigel P. Smart. 2008. Implementing Two-Party Computation Efficiently with Security Against Malicious Adversaries. In SCN 2008 (LNCS), Rafail Ostrovsky, Roberto De Prisco, and Ivan Visconti (Eds.), Vol. 5229. Springer, 2–20.
- [36] Yehuda Lindell and Ben Riva. 2014. Cut-and-Choose Yao-Based Secure Computation in the Online/Offline and Batch Settings, See [17], 476–494. https: //doi.org/10.1007/978-3-662-44381-1_27
- [37] Yehuda Lindell and Ben Riva. 2015. Blazing Fast 2PC in the Offline/Online Setting with Security for Malicious Adversaries. In ACM CCS 2015, Indrajit Ray, Ninghui Li, and Christopher Kruegel: (Eds.). ACM Press, 579–590.
- [38] C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi. 2015. ObliVM: A Programming Framework for Secure Computation. In 2015 IEEE Symposium on Security and Privacy. 359–376. https://doi.org/10.1109/SP.2015.29
- [39] Dahlia Malkhi, Noam Nisan, Benny Pinkas, and Yaron Sella. 2004. Fairplay a Secure Two-party Computation System. In USENIX Security 2004. USENIX Association.
- [40] Payman Mohassel and Matthew Franklin. 2006. Efficiency Tradeoffs for Malicious Two-Party Computation. In PKC 2006 (LNCS), Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin (Eds.), Vol. 3958. Springer, 458–473.
- [41] Payman Mohassel and Ben Riva. 2013. Garbled Circuits Checking Garbled Circuits: More Efficient and Secure Two-Party Computation, See [9], 36–53. https://doi.org/10.1007/978-3-642-40084-1_3
- [42] Benjamin Mood, Debayan Gupta, Kevin R. B. Butler, and Joan Feigenbaum. 2014. Reuse It Or Lose It: More Efficient Secure Computation Through Reuse of Encrypted Values. In ACM CCS 2014, Gail-Joon Ahn, Moti Yung, and Ninghui Li (Eds.). ACM Press, 582–596.
- [43] B. Mood, D. Gupta, H. Carter, K. Butler, and P. Traynor. 2016. Frigate: A Validated, Extensible, and Efficient Compiler and Interpreter for Secure Computation. In 2016 IEEE European Symposium on Security and Privacy (EuroS&P). 112–127. https://doi.org/10.1109/EuroSP.2016.20
- [44] Jesper Buus Nielsen and Claudio Orlandi. 2009. LEGO for Two-Party Secure Computation. In TCC 2009 (LNCS), Omer Reingold (Ed.), Vol. 5444. Springer, 368–386.
- [45] Jesper Buus Nielsen and Samuel Ranellucci. 2016. Reactive Garbling: Foundation, Instantiation, Application. In ASIACRYPT 2016, Part II (LNCS). Springer, 1022– 1052. https://doi.org/10.1007/978-3-662-53890-6_34
- [46] Jesper Buus Nielsen, Thomas Schneider, and Roberto Trifiletti. 2017. Constant Round Maliciously Secure 2PC with Function-independent Preprocessing using LEGO. In 24. Annual Network and Distributed System Security Symposium (NDSS'17). The Internet Society.
- [47] Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. 2009. Secure Two-Party Computation Is Practical. In ASIACRYPT 2009 (LNCS), Mitsuru Matsui (Ed.), Vol. 5912. Springer, 250–267.
- [48] Peter Rindal. 2017. libOTe: an efficient, portable, and easy to use Oblivious Transfer Library. https://github.com/osu-crypto/libOTe. (2017).
- [49] Peter Rindal and Mike Rosulek. 2016. Faster Malicious 2-Party Secure Computation with Online/Offline Dual Execution. In USENIX Security 2016. USENIX Association.
- [50] Peter Rindal and Roberto Trifiletti. 2017. SplitCommit: Implementing and Analyzing Homomorphic UC Commitments. Cryptology ePrint Archive, Report 2017/407. (2017). http://eprint.iacr.org/2017/407
- [51] abhi shelat and Chih-Hao Shen. 2011. Two-Output Secure Computation with Malicious Adversaries. In EUROCRYPT 2011 (LNCS), Kenneth G. Paterson (Ed.), Vol. 6632. Springer, 386–405.
- [52] abhi shelat and Chih-Hao Shen. 2013. Fast two-party secure computation with minimal assumptions. In ACM CCS 2013, Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung (Eds.). ACM Press, 523–534.

- [53] Nigel Smart and Stefan Tillich. 2017. Circuits of Basic Functions Suitable For MPC and FHE. (2017). http://www.cs.bris.ac.uk/Research/CryptographySecurity/MPC/
- [54] Ebrahim M. Songhori, Siam U. Hussain, Ahmad-Reza Sadeghi, Thomas Schneider, and Farinaz Koushanfar. 2015. TinyGarble: Highly Compressed and Scalable Sequential Garbled Circuits. In 2015 IEEE Symposium on Security and Privacy. IEEE Computer Society Press, 411–428. https://doi.org/10.1109/SP.2015.32
- [55] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. 2017. Faster Secure Two-Party Computation in the Single-Execution Setting. In EUROCRYPT 2017 (LNCS), Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.), Vol. 10212. 399–424. https: //doi.org/10.1007/978-3-319-56617-7_14
- [56] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. 2017. Authenticated Garbling and Efficient Maliciously Secure Two-Party Computation. Cryptology ePrint Archive, Report 2017/030. (2017). http://eprint.iacr.org/2017/030.
- [57] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. 2017. Global-Scale Secure Multiparty Computation. Cryptology ePrint Archive, Report 2017/189. (2017). http://eprint.iacr.org/2017/189.
- [58] Xiao Shaun Wang, S. Dov Gordon, Allen McIntosh, and Jonathan Katz. 2016. Secure Computation of MIPS Machine Code. In ESORICS 2016, Part II (LNCS). Springer, 99–117. https://doi.org/10.1007/978-3-319-45741-3_6
- [59] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets (Extended Abstract). In FOCS 1986. IEEE Computer Society Press, 162–167.
- [60] Ting Yu, George Danezis, and Virgil D. Gligor (Eds.). 2012. ACM CCS 2012. ACM Press.
- [61] Samee Zahur and David Evans. 2015. Obliv-C: A Language for Extensible Data-Oblivious Computation. Cryptology ePrint Archive, Report 2015/1153. (2015). http://eprint.iacr.org/2015/1153
- [62] Samee Zahur, Mike Rosulek, and David Evans. 2015. Two Halves Make a Whole -Reducing Data Transfer in Garbled Circuits Using Half Gates. In EUROCRYPT 2015, Part II (LNCS), Elisabeth Oswald and Marc Fischlin (Eds.), Vol. 9057. Springer, 220–250. https://doi.org/10.1007/978-3-662-46803-6_8
- [63] Ruiyu Zhu and Yan Huang. 2017. Faster LEGO-based Secure Computation without Homomorphic Commitments. Cryptology ePrint Archive, Report 2017/226. (2017). http://eprint.iacr.org/2017/226

A PROTOCOL DETAILS

We describe and analyse the protocol in the UC framework. We will here give an abstract description that lends itself to a security analysis. In Section 6.1 we describe some of the optimisations that were done in the implementation and why they do not affect the security analysis. We describe the protocol for two parties, the garbler G and the evaluator E. We will describe the protocol in the hybrid model with ideal functionalities \mathcal{F}_{HCOM} and \mathcal{F}_{OT} for xorhomomorphic commitment and one-out-of-two oblivious transfer. The precise description of the ideal functionalities are standard by now and can be found in [15] and [26]. Here we will denote the use of the functionalities by some pseudo-code conventions. When using \mathcal{F}_{HCOM} it is G that is the committer and E that is the receiver. When G executes COMMIT(cid, x) for cid $\in \{0, 1\}^*$ and $x \in \{0, 1\}^{\kappa}$, then \mathcal{F}_{HCOM} stores (cid, x_{cid}) (where $x_{cid} = x$) and outputs cid to E. When G executes $OPEN(cid_1, \ldots, cid_c)$, where each cid_i was output to E at some point, then \mathcal{F}_{HCOM} outputs $(\operatorname{cid}_1, \ldots, \operatorname{cid}_c, \oplus_{i=1}^c x_{\operatorname{cid}_i})$ to E. When G executes an open command, then the commitment identifies $(\operatorname{cid}_1, \ldots, \operatorname{cid}_c)$ are always already known by E. If \mathcal{F}_{HCOM} outputs($\operatorname{cid}_{1}^{\prime}, \ldots, \operatorname{cid}_{c}^{\prime}, \bigoplus_{i=1}^{c} x_{\operatorname{cid}_{i}}$) where some $\operatorname{cid}_{i}^{\prime} \neq \operatorname{cid}_{i}$ then E always tacitly aborts the protocol. Similarly the cid used in the commit command is always known and E aborts if G uses a wrong one. When using \mathcal{F}_{OT} it is G that is the sender and E that is the receiver. We assume that we have access to a special OT which has a special internal state $\Delta \in \{0,1\}^{\kappa}$, which is chosen by G once and for all at the initialisation of the ideal functionality by executing OTINIT(Δ). After that, when G executes OTSEND(id, x_0) for id $\in \{0, 1\}^*$ and $x_0 \in \{0, 1\}^{\kappa}$ and E executes OTRECEIVE(id, b) for $b \in \{0, 1\}$, then \mathcal{F}_{OT} outputs (id, x_h) to E, where $x_1 = x_0 \oplus \Delta$. If the protocol specifies that G is to execute OTRECEIVE(id, b) and it does not or uses a wrong id, then E will always detect this and will tacitly abort.

When we instruct a party to send a value, we tacitly assume the receiver stores it under the same name when it is received.

When we instruct a party to check a condition, we mean that the party will abort if the condition is false.

When a variable like K_{id} is created in our pseudo-code, it can be accessed by another routine at the same party using the same identifier. Sometimes we use the **store** and **retrieve** key-words to explicitly do this. To save on notation, it will sometimes be done more implicitly, when it cannot lead to ambiguity. In general, if an uninitialised variable like K_{id} is used in a protocol, then there is an implicit "**retrieve** K_{id} " in the line before.

We assume that we have a free-xor garbling scheme (Gb, Ev) which has correctness, obliviousness and authenticity. We recall these notions now. The key length is some κ . The input to Gb is a poly-sized circuit *C* computing a function $C : \{0, 1\}^n \to \{0, 1\}^m$ along with $(K_1^0, \ldots, K_n^0, \Delta) \in (\{0, 1\}^{\kappa})^{n+1}$, where $lsb(\Delta) = 1$. The output is $(L_1^0, \ldots, L_m^0) \in (\{0, 1\}^{\kappa})^m$ and a garbled circuit *F*. Here *F* is the garbled version of *C*. Define $K_i^1 = K_i^0 \oplus \Delta$. For $x \in \{0, 1\}^n$ define $K^x = (K_1^{x_1}, \dots, K_n^{x_n})$. This is the garbled input, *i.e.*, the garbled version of x. Define $L_i^1 = L_i^0 \oplus \Delta$. For $y \in \{0, 1\}^m$ define $L^y = (L_1^{y_1}, \dots, L_m^{y_m})$. This is the garbled output. The input to Ev is a garbled circuit F and a garbled input $(K_1, \ldots, K_n) \in (\{0, 1\}^{\kappa})^n$. The output is \perp or a garbled output $(L_1, \ldots, L_m) \in (\{0, 1\}^{\kappa})^m$. The scheme being free-xor means the inputs and outputs are of the above form. Correctness says that if you do garbled evaluation, you get the correct output. Obliviousness says that if you are given F but not given $(K_1^0, \ldots, K_n^0, \Delta)$, then the garbled input leaks no information on the plaintext input (or output). Authenticity says that if you are given only a garbled circuit for *C* and a garbled input for *x*, then you cannot compute the garbled output for any other value than the correct value C(x). These notions have been formalized in [5]. We will in fact require extended versions of these notions as we use a reactive garbling scheme in the sense of [45]. In a reactive garbling scheme one can make several independent garblings and then later solder an output wire id with keys (K_{id}^0, K_{id}^1) onto an input wire id' with keys $(K_{id'}^0, K_{id'}^1)$ in another circuit. This involves releasing some information to the evaluator which allows the evaluator later to compute $K_{id'}^b$ from K_{id}^b for either b = 0 or b = 1. The notion of reactive garbling scheme is given in [45]. We will use the reactive garbling scheme from [1]. We will later describe how to solder in [1] and we then recall the notion of reactive garbling scheme from [45] to the detail that we need in our proofs.

We finally assume that we have access to a programmable random oracle $H : \{0, 1\}^{\kappa} \rightarrow \{0, 1\}^{\kappa}$. Note that this in particular implies that H is collision resistant.

We assume that we are to securely compute one circuit C which consist of sub-circuits C and solderings between input wires and output wires of these sub-circuits. We call the position in C in which a sub-circuit C is sitting a *slot* and each slot is identified by some identifier id. There is a public mapping from identifiers id to the corresponding sub-circuit C. If $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, then the inputs wires and output wires of the slot are identified by id.in.1,..., id.in.*n* and id.out.1,..., id.out.*m*. Sub-circuits sitting at a slot are called functional sub-circuits. There are also some special sub-circuits:

- E in-gates, with n = 0 and m = 1. These are for letting E input a bit. The output wire is identified by id.out.1.
- G in-gates, also with n = 0 and m = 1. These are for letting G input a bit. The output wire is identified by id.out.1.
- E out-gates, with n = 1 and m = 0. These are the output gates of E. The input wire is identified by id.in.1.
- G out-gates, with n = 1 and m = 0. These are the output gates of G. The input wire is identified by id.in.1.

Besides a set of named sub-circuits, the circuit C also contains a set S of solderings (id_1, id_2) , where id_1 is the name of an output wire of a sub-circuit and id₂ is the name of an input wire of a sub-circuit. We require that all input wires of all sub-circuits are soldered to exactly one output wire of a sub-circuit and that there are no loops. This ensures we can plaintext evaluate the circuit as follows. For each in-gate id assign a bit x_{id} and say that id.out.1 was evaluated. Now iteratively: 1) for each soldering (id1, id2) where id_1 was evaluated, let $x_{id_2} = x_{id_1}$ and say id_2 was evaluated, and 2) for each sub-circuit where all input wires were evaluated, run C on the corresponding bits, assign the result to the output wires and say they are evaluated. This way all out-gates will be assigned a unique bit. The goal of our protocol is to let both parties learn their own output bits without learning any other information. We assume some given evaluation order of the sub-circuits that allows to plaintext evaluate in that order.

We assume that we have two functions $L^1, \alpha^1 : \mathbb{N} \to \mathbb{N}$ for setting the parameters of the cut-and-choose. Consider the following game parametrised by $n \in \mathbb{N}$. First the adversary picks $L = L^{1}(n)$ balls. Let $\alpha = \alpha^{1}(n)$. Some of them are green and some are red. The adversary picks the colours. Then we sample uniformly at random $L - \alpha n$ of the balls. If any of the sampled balls are red, the adversary immediately loses the game. Then we uniformly at random throw the remaining αn balls into *n* buckets of size α . The adversary wins if there is a bucket with only red balls. We assume that L^1 and α^1 have been fixed such that the probability that any adversary wins the game is 2^{-s} , where *s* is the security parameter. Note that the functions depend on s, but we ignore this dependence in the notation. We assume that we have two other functions $L^2, \alpha^2 : \mathbb{N} \to \mathbb{N}$. We consider a game similar to the above, but where the adversary wins if all sampled balls are green and there is a bucket with a majority of red balls. We assume that L^2 and α^2 have been fixed such that the probability that any adversary wins the game is 2^{-s} .

Overview of Notation.

- id: generic identifier, just a bit-string naming an object.
- Δ_{id} : the *difference* with identifier id. Defined to be the value in the commitment with identifier id.dif. It should hold that $lsb(\Delta_{id}) = 1.$
- σ_{id} : the *indicator* bit with identifier id. Defined to be the value in the commitment with identifier id.ind.
- $K_{id}^{\sigma_{id}}$: *base-key* with identifier id. Defined to be the value in the commitment with identifier id.base. It should hold that $lsb(K_{id}^{\sigma_{id}}) = 0.$
- $K_{id}^{1-\sigma_{id}}$: esab-key with identifier id. Defined to be $K_{id}^{\sigma_{id}} \oplus \Delta_{id}$.
- K_{id}^{0} : 0-key with identifier id.

- K_{id}^1 : 1-key with identifier id.
- K_{id} : key held by E. It should hold that $K_{id} \in \{K_{id}^0, K_{id}^1\}$. $L^1(n)$: total number of objects to create when one component should be good per bucket and *n* buckets are needed.
- $\alpha^{1}(n)$: bucket size when one component should be good per bucket and n buckets are needed.
- $L^2(n)$: as above, but majority in each bucket is good.
- $\alpha^2(n)$: as above, but majority in each bucket is good.
- Par(id): mapping from input wire id to the unique parent output wire. This is well-defined given the soldering set *S*.
- rco: A special wire index used in recovering inputs of a corrupted G.
- E: a set of identifiers id for which id is an in-gate.
- *I*: a set of identifiers id with which RECOVER could be called but which are not in \mathcal{E} .

Main Structure. We assume that E knows an input bit x_{id} for each E in-gate id before it is evaluated and that G knows an input bit x_{id} for each G in-gate id before it is evaluated. The inputs are allowed to depend adaptively on previous outputs. At the end of the protocol E knows an output bit y_{id} for each E out-gate id and G knows an output bit y_{id} for each G out-gate id.

During the pre-processing G will commit to key material for all wires. The keys K_{id}^0 and K_{id}^1 will be well defined from these committed values, even if G is corrupt. We then implement the input protocols and the evaluation protocols such that it is guaranteed that for each wire, E will learn $K_{id} \in \{K_{id}^0, K_{id}^1\}$.

We then implement the input protocols such that it is guaranteed that for each G-input gate id the evaluator will learn some $K_{id} \in$ $\{K_{id}^0, K_{id}^1\}$. This holds even if G or E is corrupted. If G is honest, it is guaranteed that $K_{id} = K_{id}^{x_{id}}$. If G is corrupted, then x_{id} is defined by $K_{id} = K_{id}^{x_{id}}$. Furthermore, for each E-input gate E will learn $K_{id} \in \{K_{id}^0, K_{id}^1\}$. This holds even if G or E is corrupted. If E is honest, it is guaranteed that $K_{id} = K_{id}^{x_{id}}$. If E is corrupted, then x_{id} is defined by $K_{id} = K_{id}^{x_{id}}$. This ensures that after the input protocols, an input bit x_{id} is defined for each input wire, called the plaintext value of the wire. This allows us to mentally do a plaintext evaluation of the circuits, which gives us a plaintext bit for each output wire and input wire of all components. We denote the bit defined for wire id by x_{id} . We call this the correct plaintext value of the wire. Note that this value might be known to neither E nor G. However, by security of the input protocols E will learn the correct key $K_{id}^{x_{id}}$ for in-gates. We then implement the evaluation protocol such that E iteratively will also learn the correct keys $K_{id}^{x_{id}}$ for all internal wires id. For the G-output wires, the evaluator E will just send $K_{id}^{x_{id}}$ to G who knows (K_{id}^0, K_{id}^1) and can decode to x_{id} . By authenticity E cannot force an incorrect output. For the E-output wires, the evaluator E will be given the indicator bit for the keys which will allow to compute exactly x_{id} from $K_{id}^{x_{id}}$. That the evaluator learns nothing else will follow from obliviousness of the garbling scheme.

Below we give a full description of the Duplo protocol, decomposed into subprotocols. Security proofs are deferred to the full version.

function GENWIRE(id, K, Δ) ▶ Require: $lsb(\Delta) = 1$ $\mathbf{G}{:}\,K_{\mathsf{id}}^0 \leftarrow K$ ▶ the 0-key $\begin{array}{l} \mathbf{G} \colon \boldsymbol{\Delta}_{\mathsf{id}}^{\mathsf{id}} \leftarrow \boldsymbol{\Delta} \\ \mathbf{G} \colon K_{\mathsf{id}}^1 \leftarrow K_{\mathsf{id}}^0 \oplus \boldsymbol{\Delta}_{\mathsf{id}} \end{array}$ ▶ the difference ▶ the 1-key G: $\sigma_{id} \leftarrow lsb(K_{id}^0)$ ▶ the indicator bit G: Сомміт(id.dif, Δ_{id}) G: Сомміт(id.ind, σ_{id}) G: COMMIT(id.base, $K_{id}^{\sigma_{id}}$) end function function VerWire(id) G: Open(id.dif); E: receive Δ_{id} E: check $lsb(\Delta_{id}) = 1$ G: Open(id.ind); E: receive σ_{id} G: OPEN(id.base); E: receive $K_{id}^{\sigma_{id}}$ E: check $lsb(K_{id}^{\sigma_{id}}) = 0$ E: $K_{id}^{1-\sigma_{id}} \leftarrow K_{id}^{\sigma_{id}} \oplus \Delta_{id}$ E: store $K_{id}^{0}, K_{id}^{1}, \Delta_{id}, \sigma_{id}$ end function **function** GENSOLD(id₁, id₂) G: OPEN(id₁.ind, id₂.ind); E: **receive** σ_{id_1,id_2} if $\sigma_{id_1,id_2} = 0$ then G: Open(id₁.base, id₂.base) else G: OPEN(id₁.base, id₂.base, id₂.dif) end if E: receive K_{id_1, id_2} G: Open(id₁.dif, id₂.dif) $\mathsf{E}: \textbf{receive} \ \Delta_{\mathsf{id}_1,\mathsf{id}_2}$ **check** $lsb(\Delta_{id_1, id_2}) = 0$ **check** $lsb(K_{id_1, id_2}) = \sigma_{id_1, id_2}$ end function **function** EvSolD(id₁, id₂, *K*) E: **return** $K \oplus K_{id_1, id_2} \oplus lsb(K)\Delta_{id_1, id_2}$ end function **function** EvSolD(id₁, id₂) E: retrieve K_{id1} $E: K_{id_2} \leftarrow EvSold(id_1, id_2, K_{id_1})$ end function function GenKeyAuth(id) $\mathbf{G} {:} K_{\mathsf{id}}^0 \leftarrow \{0,1\}^{\kappa}$ $\mathbf{G}: \Delta_{\mathsf{id}} \leftarrow \{0, 1\}^{\kappa - 1} \times \{1\}$ G: GENWIRE(id, $K_{id}^0, \Delta_{id}^0)$ $G: A_{id} \leftarrow \{H(K_{id}^0), H(K_{id}^1)\}$ G: send A_{id} end function function VerKeyAuth(id) VERWIRE(id) E: **check** $A_{id} = \{H(K_{id}^0), H(K_{id}^1)\}$

end function

 $\begin{array}{l} \textbf{function} \ \texttt{PRePROCESSKA} \\ \ell \leftarrow \texttt{\#output} \ \texttt{wires} \ \texttt{of} \ \texttt{all} \ \texttt{functional} \ \texttt{sub-circuits} \end{array}$

Let $L = L^2(\ell)$ ▶ # KAs generated Let $\alpha_{ka} = \alpha^2(\ell)$ bucket size $\forall_{i=1}^{L}$: GENKEYAUTH(preka.i) E: Sample $V \subset [L]$ uniform of size $L - \alpha_{ka}\ell$. E: send V $\forall_{i \in V}$: VerKeyAuth(preka.*i*) for all functional sub-circuits id do **for all** $j = 1, ..., m_{id}$ **do** pick α_{ka} uniform, fresh KAs $i \notin V$ rename them to ids id.ka.1, . . . , id.ka. α_{ka} . $\forall_{i=2}^{\alpha_{ka}}$: GenSold(id.ka.1, id.ka.*i*) end for end for end function function GENINKEYAUTH(id) $\Delta_{\mathsf{id}} \leftarrow \{0,1\}^{\kappa-1} \times \{1\}$ $K_{id}^{0} \leftarrow H(\Delta_{id})$ GENWIRE(id, K_{id}^0, Δ_{id}^0) $A_{\text{id}} \leftarrow \{H(K_{\text{id}}^0), H(K_{\text{id}}^1)\}$ G sends A_{id} end function function VerInKeyAuth(id) VerKeyAuth(id) E: check $K_{id} = H(\Delta_{id})$ end function function PreProcessInKA $\ell \leftarrow \#$ input wires in CLet $L = L^2(\ell)$ ▶ # KAs generated Let $\alpha_{inka} = \alpha^2(\ell)$ ▶ bucket size $\forall_{i=1}^{L}$: GenInKeyAuth(preka.i) E: Sample $V \subset [L]$ uniform of size $L - \alpha_{inka}\ell$. E: send V $\forall_{i \in V}$: VERINKEYAUTH(preka.i) for all input wires id do pick $\alpha_{\sf inka}$ uniform, fresh KAs $i \notin V$ rename them to ids id.ka.1, . . . , id.ka. α_{inka} . $\forall_{i=2}^{\alpha_{\text{inka}}}$: GenSold(id.ka.1, id.ka.*i*) end for end function function PreProcessOTINIT $G: \Delta_{ot} \leftarrow \{0, 1\}^{\kappa}$ G: COMMIT(ot, Δ_{ot}) G: OTINIT(Δ_{ot}) for $i \in [s]$ do G: $R_i \leftarrow \{0, 1\}^{\kappa}$, OTsend (R_i, ot_i) G: COMMIT (ot_i, R_i) $\mathsf{E}: b_i \leftarrow \{0, 1\}, R_{b_i} \leftarrow \mathsf{OTreceive}(b_i, \mathsf{ot}_i)$ E: send (R_{b_i}, b_i) G: receive (\bar{R}_i, \bar{b}_i) ; check $\bar{R}_i = R_{\bar{b}_i}$ if $\bar{b}_i = 0$ then $G: OPEN(ot_i)$ else $G: OPEN(ot_i, ot)$

end if E: receive R_i E: check $\tilde{R}_i = R_{b_i}$ end for end function ▷ $C: \{0,1\}^n \to \{0,1\}^m$ **function** GENSUB(id, *C*) $\mathbf{G}: (K_1, \ldots, K_n) \leftarrow (\{0, 1\}^{\kappa})^n$ $\mathbf{G}: \Delta_{\mathsf{id}} \leftarrow \{0, 1\}^{\kappa - 1} \times \{1\}$ $G: (L_1, \ldots, L_m, F_{id}) \leftarrow Gb(K_1, \ldots, K_n, \Delta_{id})$ G: send F_{id} $\forall_{i=1}^{n}$: GenWire(id.in. i, K_i, Δ_{id}) $\forall_{i=1}^{m}$: GenWire(id.out.*i*, L_i , Δ_{id}) end function function VerSub(id, C) ▷ $C: \{0,1\}^n \to \{0,1\}^m$ $\mathsf{E}: \Delta_{\mathsf{id}} \leftarrow \Delta_{\mathsf{id}.\mathsf{in}.1}$ E: $\forall_{i=2}^{n}$: **check** $\Delta_{\text{id.in.}i} = \Delta_{\text{id}}$ E: $\forall_{i=1}^{m}$: **check** $\Delta_{\text{id.out.}i} = \Delta_{\text{id}}$ $E: \forall_{i=1}^{n} : K_{i} \leftarrow K_{id,in,i}^{0}$ $E: \forall_{i=1}^{m} : L_{i} \leftarrow K_{id,out,i}^{0}$ $E: \mathbf{check} (L_{1}, \dots, L_{m}, F_{id}) = \mathrm{Gb}(K_{1}, \dots, K_{n}, \Delta_{id})$ end function function EvSuB(id) $\mathsf{E}: \forall_{i=1}^n : K_i \leftarrow \mathbf{retrieve} \, K_{\mathsf{id.in.}i}$ $\mathsf{E}: (L_1, \ldots, L_m) \leftarrow \mathsf{Ev}(F_{\mathsf{id}}, K_1, \ldots, K_n)$ $\mathsf{E}: \forall_{i=1}^m : \mathbf{store} \, K_{\mathsf{id.out.}i} \leftarrow L_i$ end function **function** GENSOLDSUB(id₁, id₂) $\begin{array}{l} \forall_{i=1}^{n} : \text{GenSold}(\text{id}_{1}.\text{in}.i,\text{id}_{2}.\text{in}.i) \\ \forall_{i=1}^{m} : \text{GenSold}(\text{id}_{2}.\text{out}.i,\text{id}_{1}.\text{out}.i) \\ \forall_{i=1}^{m} : \text{GenSold}(\text{id}.1.\text{out}.i,\text{id}.1.\text{out}.i.\text{ka}.1) \end{array}$ end function function PreProcessSub for all sub-circuit types C do Let ℓ be the number of times *C* is used. Let $L = L^1(\ell)$ ▶ #circuits generated Let $\alpha_{id} = \alpha^1(\ell)$ ▶ bucket size $\forall_{i=1}^{L}$: GENSUB(C.pre.*i*, C) E: Sample $V \subset [L]$ uniform of size $L - \alpha_{id}\ell$. E: send V $\forall_{i \in V} : \text{VerSub}(C.\text{pre.}i, C)$ for all slots id where C occurs do pick α_{id} uniform, fresh circuits $i \notin V$ rename them to have ids id.1, ..., id. α_{id} . $\forall_{i=2}^{\alpha_{id}}$: GenSoldSub(id.1, id.*i*) end for end for end function function AssembleSubs for all functional sub-circuits id do

 $\forall_{i=1}^{n} : id_{par.i} \leftarrow Par(id.1.in.i)$

```
\forall_{i=1}^{n} : GENSOLD(id<sub>par.i</sub>, id.1.in.i)
       end for
end function
function EvKAs(id, \mathcal{K}_{id})
      \alpha \leftarrow \alpha_{ka}
                                                     ▶ if generated using PreProcessKA
                                                           ▶ if gen. using PreProcessInKA
       \alpha \leftarrow \alpha_{inka}
       \begin{array}{l} \forall_{i=1}^{\alpha} : A_i \leftarrow A_{\mathrm{id.ka.}i} \\ \mathcal{L} \leftarrow \emptyset \end{array} 
                                                                           ▶ get key authenticators
       for K \in \mathcal{K}_{id} do
             K_1 = K
              \forall_{i=2}^{\alpha} : K_i = \text{EvSold}(\text{id.ka.1}, \text{id.ka.}i, K)
              if \#\{i \in \{1, ..., \alpha\} | A_i(K_i) = \top\} > \alpha/2 then
                     \mathcal{L} \leftarrow \mathcal{L} \cup \{K\}
              end if
       end for
       if \mathcal{L} = \{K\} then return K
       else if \mathcal{L} = \{K_0, K_1\} then
             \Delta \leftarrow K_0 \oplus K_1
             Recover(id, \Delta)
       else abort
       end if
end function
function EvSubs(id)
                                                            ▶ Evaluate first circuit in bucket
       \forall_{i=1}^{n} : \mathrm{id}_{\mathrm{par.}j} \leftarrow \mathrm{Par}(\mathrm{id.1.in.}j)
      \forall_{j=1}^{n} : \text{EvSolD}(\text{id}_{\text{par},j}, \text{id}.1.\text{in},j)
\forall_{j=1}^{n} : \text{retrieve } K_{\text{id}.1.\text{in},j}
       EvSub(id.1)
       \forall_{i=1}^m : retrieve K_{id.1.out.j}
       \begin{aligned} &\forall_{j=1}^{m} : \text{EvSold}(\text{id.1.out.}j, \text{id.1.out.}j.\text{ka.1}) \\ &\forall_{j=1}^{m} : \mathcal{K}_{j} \leftarrow \{K_{\text{id.1.out.}j.\text{ka.1}}\} \end{aligned} 
       Ym
                                                                                                       ▶ key sets
                                              ▶ Evaluate remaining circuits in bucket
       for i = 2, ..., \alpha_{id} do E:
             \forall_{i=1}^{n} : EvSold(id.1.in.j, id.i.in.j)
             EvSub(id.i)
              \forall_{i=1}^m : EvSolD(id.i.out.j, id.1.out.j)
              \begin{array}{l} \forall_{j=1}^{m} : \texttt{EvSold}(\texttt{id.1.out.}j, \texttt{id.1.out.}j.\texttt{ka.1}) \\ \forall_{j=1}^{m} : \mathcal{K}_{j} \leftarrow \mathcal{K}_{j} \cup \{K_{\texttt{id.1.out.}j.\texttt{ka.1}}\} \end{array} 
       end for
       \forall_{j=1}^m : K_{\text{id.1.out.}j} \leftarrow \text{EvKAs}(\text{ka}_j, \mathcal{K}_j)
end function
function INPUTG(id)
                                                                      ▶ id is an ID of a G in-gate
       G: retrieve the input bit x_{id} for id
       G: K \leftarrow K_{id.ka.1}^{x_{id}}
G: send K
       E: K_{id} \leftarrow EvKAs(id, \{K\})
       E: store K<sub>id</sub>
end function
function INPUTE(id)
                                                                    ▶ id is an ID of an E in-gate
       G: R_{\text{ot}_{id}} \leftarrow \{0, 1\}^{\kappa}, OTsend(R_{\text{ot}_{id}}, \text{ot}_{id})
       G: COMMIT(ot_{id}, R_{ot_{id}})
```

 $\mathsf{E}: b_{\mathsf{ot}_{\mathsf{id}}} \leftarrow \{0, 1\}, R_{b_{\mathsf{ot}_{\mathsf{id}}}} \leftarrow \mathsf{OTreceive}(b_{\mathsf{ot}_{\mathsf{id}}}, \mathsf{ot}_{\mathsf{id}})$

19

E: **retrieve** the input bit x_{id} for id E: send $f_{id} = x_{id} \oplus b_{ot_{id}}$ G: $e_{id} = f_{id} \oplus \sigma_{id}$ $E: id' \leftarrow id.ka.1$ if $e_{id} = 0$ then G: Open(id', ot_{id}) else G: Open(id', ot_{id}, ot) end if $\mathsf{E} : \mathbf{receive} \ D = K_{\mathrm{id}'} \oplus R_{\mathrm{ot}_{\mathrm{id}}} \oplus e_{\mathrm{id}} \Delta_{\mathrm{ot}}$ G: Open(id'.dif, ot); E : receive $S_{id} = \Delta_{id'} \oplus \Delta_{ot}$ G: Open(id.ind); E : receive σ_{id} $\mathsf{E}: K = D \oplus R_{b_{\mathrm{ot}_{\mathrm{id}}}} \oplus (x_{\mathrm{id}} \oplus \sigma_{\mathrm{id}}) S_{\mathrm{id}}$ $\mathsf{E}: K_{\mathsf{id}'} \leftarrow \mathsf{EvKAs}(\mathsf{id}, \{K\})$ E: **check** $lsb(K_{id'}) = x_{id} \oplus \sigma_{id}$ E: store K_{id'} end function

 $\begin{array}{ll} \textbf{function OUTPUTE(id)} & \triangleright \mbox{ id: ID of an E output gate} \\ E: \textbf{retrieve soldering (id_1, id.out.1) from } C. \\ E: \textbf{retrieve } K_{id_1}. \\ G: OPEN(id_1.ind); E: \textbf{receive } \sigma_{id_1} \\ E: y_{id} \leftarrow lsb(K_{id_1}) \oplus \sigma_{id_1}. \\ end function \end{array}$

 $\begin{array}{l} \textbf{function OUTPUTG(id)} \qquad \triangleright \mbox{ id: ID of an G output gate} \\ E: \textbf{retrieve soldering (id_1, id.out.1) from } C \\ E: \textbf{retrieve } K_{id_1} \\ E: \textbf{send } K_{id_1} \\ G: \textbf{receive } K_{id_1} \\ G: \textbf{check } K_{id_1} \in \{K^0_{id_1}, K^1_{id_1}\} \\ G: y_{id} \leftarrow b \mbox{ where } K^b_{id_1} = K_{id_1}. \end{array}$

function MAIN(C)
 PREPROCESSKA()
 PREPROCESSINKA()
 PREPROCESSOTINIT()
 PREPROCESSOUB()
 AssembleSUBS()
 for all sub-circuit id in evaluation order do
 if id is a G in-gate then INPUTG(id)
 else if id is an E in-gate then INPUTE(id)
 else if id is a G out-gate then OUTPUTG(id)
 else if id is a E out-gate then OUTPUTG(id)
 else EvSUBS(id)
 end if
 end for
end function

function RecoverInputBit(id, Δ) $\triangleright \Delta = \Delta_{id.ka.1}$ id′ ← id.ka $\Delta_1 \leftarrow \Delta$ $\alpha \leftarrow \alpha_{inka}$ $\forall_{j=1}^{\alpha}$ retrieve $K_j \leftarrow K_{id'.j}$; $\forall_{j=2}^{\alpha} \text{ retrieve } \Delta_{\mathrm{id}',1,\mathrm{id}',j}; \Delta_j \leftarrow \Delta_{\mathrm{id}',1,\mathrm{id}',j} \oplus \Delta_1 \\ \mathrm{if}' \# \{j \in \{1,\ldots,\alpha\} \mid H(\Delta_j) = K_j\} > \alpha/2 \text{ then}$ Yα $x_{id} \leftarrow 0$ else $x_{id} \leftarrow 1$ end if end function function PreProcessRecovery $G: \Delta_{\mathsf{rco}} \leftarrow \{0, 1\}^{\kappa}$ G: Commit(rco, Δ_{rco}) G: $\forall id \in I \cup \mathcal{E}$: Open(id.ka.1.dif, rco) E: $\forall id \in \mathcal{I} \cup \mathcal{E}$: receive $\Delta_{id, rco}$ G: $\forall id \in \mathcal{E}$: Open(rco, id.ka.1.dif) E: $\forall id \in \mathcal{E}$: receive $\Delta_{rco, id}$ end function $\triangleright \Delta = \Delta_{\mathsf{id}.\mathsf{ka}.1}$ **function** RecoverINPUTBITS(id, Δ) retrieve $\Delta_{id, eecov}$ $\Delta_{rco} \leftarrow \Delta_{id, rco} \oplus \Delta$ for all $id' \in \mathcal{E}$ do retrieve $\Delta_{\text{rco, id'}}$ $\Delta_{\mathsf{id}'} \leftarrow \Delta_{\mathsf{rco}, \mathsf{id}'} \oplus \Delta_{\mathsf{rco}}$ RecoverInputBit(id', $\Delta_{id'}$) end for end function **function** Recover(id, Δ) $\triangleright \Delta = \Delta_{id,ka,1}$ RecoverInputBits(id, Δ) go to recovery mode end function function OUTPUTG(id) ▶ In recovery mode E: **retrieve** soldering (id₁, id.out.1) from CE: retrieve *x*_{id1} E: retrieve $K_{id_1}^0$ E: retrieve $\Delta_{id_1}^0$ E: send $K_{id_1}^{x_{id_1}}$ end function function OUTPUTE(id) In recovery mode E: **retrieve** soldering (id₁, id.out.1) from *C*. E: retrieve x_{id_1} . $E: y_{id} \leftarrow x_{id_1}$. end function