



Research Article

Breaking and Fixing Garbled Circuits When a Gate has Duplicate Input Wires

Raine Nieminen · Thomas Schneider

Technical University of Darmstadt, Darmstadt, Germany
nieminen@crypto.cs.tu-darmstadt.de
schneider@crypto.cs.tu-darmstadt.de

Communicated by David Pointcheval and Nigel Smart

Received 22 September 2022 / Revised 6 June 2023 / Accepted 7 June 2023

Online publication 3 August 2023

Abstract. Garbled circuits are a fundamental cryptographic primitive that allows two or more parties to securely evaluate an arbitrary Boolean circuit without revealing any information beyond the output using a constant number of communication rounds. Garbled circuits have been introduced by Yao (FOCS'86) and generalized to the multi-party setting by Beaver, Micali and Rogaway (STOC'90). Since then, several works have improved their efficiency by providing different garbling schemes and several implementations exist. Starting with the seminal Fairplay compiler (USENIX Security'04), several implementation frameworks decoupled the task of compiling the function to be evaluated into a Boolean circuit from the engine that securely evaluates that circuit, e.g., using a secure two-party computation protocol based on garbled circuits. In this paper, we show that this decoupling of circuit generation and evaluation allows a subtle attack on several prominent garbling schemes. It occurs when violating the implicit assumption on the circuit that gates have different input wires which is most often not explicitly specified in the respective papers. The affected garbling schemes use separate calls to a deterministic encryption function for the left and right input wire of a gate to derive pseudo-random encryption pads that are XORed together. When a circuit contains a gate where the left and right input wire are the same, these two per-wire encryption pads cancel out and we demonstrate that this can result in a complete break of privacy. We show how the vulnerable garbling schemes can be fixed easily.

Keywords. Secure multi-party computation, Garbled circuits, Garbling schemes, Circuits, Attack, Vulnerability.

1. Introduction

Secure Multi-Party Computation (MPC), also known as Secure Function Evaluation (SFE), enables two or more parties to jointly evaluate a function securely over their private inputs. In 1986, Andrew Yao introduced Garbled Circuits (GCs) [39], the first general technique for two-party SFE, which was later generalized to the multi-party setting by Beaver-Micali-Rogaway (BMR) [2]. GCs allow to construct constant-round MPC

protocols and hence are well-suited for high latency network settings, where multi-round protocols such as the Goldreich-Micali-Wigderson (GMW) protocol [12] are often less efficient [37]. Research on GCs has been very active and many different constructions have been proposed that improve the concrete communication and computation complexities, e.g., Point-and-Permute [2], 3-Row Reduction [31], Free-XOR [18], Garbling via AES [20], Fixed-key AES garbling [3], HalfGates [41], and Three-Halves Garbling [36].

Garbled circuits allow to securely evaluate a function that is represented as a Boolean circuit, e.g., consisting of AND and XOR gates. Several frameworks have been built to compile a high-level specification of the function to be evaluated securely (e.g., in Verilog, ANSI-C, or a domain specific language) into a Boolean circuit that is then evaluated using a GC engine. Some examples for such frameworks are Fairplay [28], FairplayMP [5], TASTY [14], FastGC [15], VMCrypt [27], PAL [30], PCF [19], CBMC-GC [11], TinyGarble [38], Frigate [29], and HyCC [8].

The basic idea of a GC is to assign to each wire w_i of the circuit two random-looking *wire labels* $\tilde{w}_i^0, \tilde{w}_i^1$. The length of each of these labels is the computational security parameter κ (often set to $\kappa = 128$ in implementations) such that wire label \tilde{w}_i^0 corresponds to plaintext value 0 and wire label \tilde{w}_i^1 corresponds to plaintext value 1. As the wire labels look random, the evaluator (who obtains exactly one label per wire) learns no information about the underlying plaintext value. For each gate in the circuit, a *garbled table* is generated by encrypting for all $2^2 = 4$ possible input combinations the corresponding output label using the corresponding input labels. To not leak any information from the position in the garbled table, its entries are permuted randomly. To preserve privacy, a crucial property required for privacy of GCs is that given a combination of two input labels only the corresponding output label can be decrypted from the garbled table.

Since garbling (and evaluation of the garbled circuit) requires calls to an encryption function, this is a natural point for optimizations and several encryption functions have been proposed and used in the literature (we give a summary in Sect. 3 and Table 1). In this paper, we show an attack and fixes on garbling schemes that use separate calls to a deterministic encryption function for the left and right input wire to derive per-wire encryption pads that are XORed together. Our attack becomes particularly devastating when Free-XOR garbling [18] is used, as in this case we reconstruct the global offset $R = \tilde{w}^0 \oplus \tilde{w}^1$ which allows to decrypt the entire GC.

2. Our Attack

In this section, we show our attack on garbling schemes that use separate calls to a deterministic encryption function to derive per-wire encryption pads that are XORed together. Later in Sect. 3, we show a set of existing constructions and summarize which ones are affected by our attack.

For a gate in the garbled circuit, let \tilde{w}_ℓ be the left input wire label, \tilde{w}_r the right input wire label, and \tilde{w}_o the output wire label. The GC encryption function with per-wire encryption pads is defined as

$$E(\tilde{w}_\ell, \tilde{w}_r; \tilde{w}_o) = \tilde{w}_o \oplus F(\tilde{w}_\ell; T) \oplus F(\tilde{w}_r; T), \quad (1)$$

where $F(k; T)$ is a Pseudo-Random Function (PRF) with key k and tweak T , often instantiated based on AES. The tweak T is set such that it is unique per gate and invocation of the GC encryption function (for the 4 possible input wire label combinations), so it is often set to the gate number concatenated with the position in the garbled table. However, the affected works omit to mention (and implement) that the tweak must also be different for the left and right call of the PRF, presumably because they make the implicit assumption that the two input wires are different, i.e., $\ell \neq r$, which however is never specified in the paper.

Our main observation is that when the tweak is the same for the left and right invocation of the PRF and the input wire labels are the same ($\ell = r$), then the encryption pads cancel out and leave both output wire labels *unencrypted* in the garbled table. Let \tilde{w}^0 be the wire label corresponding to plaintext bit 0 and \tilde{w}^1 be the wire label corresponding to 1. If the input wires for an AND gate are the same, then $\ell = r$ and the garbled table looks as follows (before the permutation of the entries):

$$E(\tilde{w}_\ell^0, \tilde{w}_r^0; \tilde{w}_o^0) = \tilde{w}_o^0 \oplus F(\tilde{w}_r^0; T_{00}) \oplus F(\tilde{w}_r^0; T_{00}) = \tilde{w}_o^0; \quad (2)$$

$$E(\tilde{w}_\ell^0, \tilde{w}_r^1; \tilde{w}_o^0) = \tilde{w}_o^0 \oplus F(\tilde{w}_r^0; T_{01}) \oplus F(\tilde{w}_r^1; T_{01}); \quad (3)$$

$$E(\tilde{w}_\ell^1, \tilde{w}_r^0; \tilde{w}_o^0) = \tilde{w}_o^0 \oplus F(\tilde{w}_r^1; T_{10}) \oplus F(\tilde{w}_r^0; T_{10}); \quad (4)$$

$$E(\tilde{w}_\ell^1, \tilde{w}_r^1; \tilde{w}_o^0) = \tilde{w}_o^0 \oplus F(\tilde{w}_r^1; T_{11}) \oplus F(\tilde{w}_r^1; T_{11}) = \tilde{w}_o^1. \quad (5)$$

Hence, the per-wire encryption pads cancel out for the first and last row of the truth table and leak both output wire labels \tilde{w}_o^0 and \tilde{w}_o^1 . Note that the second and third garbled table entries differ, because the tweaks T_{01} and T_{10} usually contain (or are explicitly concatenated with) the position in the garbled table and hence are different.

2.1. Breaking GCs with Per-Wire Encryption Pads

The implications of our attack depend on the specific garbling scheme.

General Attack. When the input wires are equal ($\ell = r$), the random pads cancel out and the garbled table of the AND gate contains *both wire labels unencrypted*: the zero output wire label \tilde{w}_o^0 (see Eq. (2)) and the one output wire label \tilde{w}_o^1 (see Eq. (5)). Hence, the GC evaluator obtains *both* output wire labels and can evaluate the remaining circuit following the AND gate not only for the output wire label corresponding to the current evaluation, but also for the other wire label. Depending on the circuit, this can leak crucial information as described next.

In the general attack, the evaluator obtains from the garbled AND gate with duplicate input wires four candidates for potential wire labels (taking each entry in the garbled table), two of which are the valid wire labels. He already knows one valid wire label from the evaluation of the GC, so the other wire label can be one of the remaining three candidates. Assume that the subsequent circuit contains a “normal” AND gate (i.e., with two different input wires), where w.l.o.g. the evaluator knows one valid wire label \tilde{w}_ℓ^* from the evaluation of the GC, and three candidate wire labels $\tilde{w}_\ell^1, \tilde{w}_\ell^2, \tilde{w}_\ell^3$ for the *left* input. He evaluates the garbled gate on the valid wire label and obtains a valid output wire label \tilde{w}_o^* . Now, he evaluates the garbled gate for each of the candidate wire labels \tilde{w}_ℓ^i ,

$i \in \{1, 2, 3\}$. If the corresponding output wire label is equal to \tilde{w}_o^* , then \tilde{w}_ℓ^i is a valid input wire label, the plaintext output is 0 (because an AND gate evaluates to 0 in three cases), and the right plaintext input is 0 (because the AND gate evaluated to 0 for left inputs 0 and 1). Otherwise, if all output wire labels are different to \tilde{w}_o^* , then the right plaintext input must be 1. Hence, the evaluator has successfully decrypted the plaintext value of the right input wire.

3-Row Reduction [31]. The idea of garbled 3-Row Reduction [31] is to fix one of the four entries in the table to 0^κ which then no longer needs to be sent to the evaluator. Here, we can apply the general attack described above by setting the omitted entry to 0^κ , as we explain in the following.

For a garbled 3-Row Reduction AND gate with duplicate input wires ($\ell = r$), we have two possible cases. In the first case, the first entry as in Eq. (2) (resp. the last entry as in Eq. (5)) of the garbled table is fixed to 0^κ , and therefore the actual wire label becomes all zeros, namely $\tilde{w}_o^0 = 0^\kappa$ (resp. $\tilde{w}_o^1 = 0^\kappa$). The other wire label is contained unencrypted in the last garbled table entry as in Eq. (5) (resp. the first garbled table entry as in Eq. (2)). In the second case, the second entry as in Eq. (3) (resp. third entry as in Eq. (4)) of the garbled table is fixed to 0^κ . In this case, the first entry (as in Eq. (2)) and the last entry (as in Eq. (5)) of the garbled table contain the unencrypted wire labels \tilde{w}_o^0 and \tilde{w}_o^1 .

In both cases, the evaluator now has four candidate wire labels: three taken from the three garbled table entries and the fourth being the all zeros wire label 0^κ . Exactly two out of these four candidates are valid wire labels, which successfully decrypt the two output wire labels corresponding to the gate's plaintext inputs $(0, v_r)$ and $(1, v_r)$, exactly as in the general attack.

Free-XOR [18]. The idea of Free-XOR garbling [18] is that all wire labels are correlated with a global offset R that is supposed to remain hidden from the evaluator, i.e., $R = \tilde{w}_i^0 \oplus \tilde{w}_i^1$ for all wires i . This allows to securely evaluate XOR gates by XORING the corresponding wire labels, and therefore requires no communication for XOR gates. When Free-XOR garbling is used, our attack becomes devastating, since the garbled table of the AND gate where both input wires are the same ($\ell=r$) contains both \tilde{w}_o^0 and \tilde{w}_o^1 , and the global offset $R = \tilde{w}_o^0 \oplus \tilde{w}_o^1$ is leaked. Using this, the evaluator knows *all* wire labels for *all* wires in the circuit and can evaluate it on arbitrary input combinations from which substantial information about the other party's inputs can be derived.

The attack works as follows: As for the general attack and the 3-Row Reduction [31] described above, we obtain four candidate wire labels from an AND gate with duplicate input wires. The evaluator already knows one valid wire label from the GC evaluation, thus it can compute three candidate values for the global offset R , of which exactly one is the correct one. The evaluator can now determine which of the candidate global offsets is the correct one by evaluating *any* other AND gate (with two different input wires) in the circuit for different combinations of input wire labels as follows: An AND gate evaluates to 0 in three cases and to 1 in one case. Let R^* be the candidate global offset to check, and the evaluator knows the left input wire label \tilde{w}_ℓ and the right input wire label \tilde{w}_r . Now, he evaluates the AND gate for three combinations of input wire labels, e.g., $(\tilde{w}_\ell, \tilde{w}_r)$, $(\tilde{w}_\ell \oplus R^*, \tilde{w}_r)$, and $(\tilde{w}_\ell, \tilde{w}_r \oplus R^*)$. If at least two evaluations yield the same output wire label (corresponding to plaintext value 0), then R^* is the correct global offset R .

2.2. AND Gates with Duplicate Input Wires

Our attack assumes an AND gate where both input wires are the same ($\ell = r$). As $\text{AND}(X, X) = X$, this is a trivial identity gate which just outputs the input value. One would assume that such a gate would not occur in practice. We give three examples where such gates can occur naturally.

Circuits composed of NANDs only. In hardware synthesis, especially for ASICs, it is common to represent the function as NAND gates only. Here, inverters are built as $\text{NOT}(X) = \text{NAND}(X, X)$. As our attack described in Sect. 2.1 for AND gates works similarly for NAND gates, it immediately works for NAND-only functions that contain at least one NOT gate.

Agreeing on the circuit to be evaluated. SFE protocols usually assume that both parties “agree” on the circuit to be evaluated securely. When implementing this in practice, one party would suggest a circuit (or a high-level program from which the circuit is compiled) and the other party would agree to jointly evaluate it, e.g., after having tested its correctness for some inputs or inspecting the high-level program. If the party who suggests which circuit to evaluate is the one that evaluates the GC, it could suggest a circuit where one gate has duplicate input wires and then apply our attack. Overall, it is not enough for the other party to test that the circuit computes the intended function correctly.

Garbling with a hardware token [16]. The main bottleneck in protocols using GC is sending the garbled circuit which has size linear in the number of AND gates in the circuit. To circumvent this, the server can send to the client a hardware token (e.g., a smartcard) that generates the GC on his behalf as proposed in [16]. If the token has constant memory (independent of the size of the circuit), the client can send a description of the circuit gate-by-gate to the token. The token locally derives the wire labels corresponding to the gate’s wire indices, generates the garbled table, and sends it to the client. Note that this does not work with garbling schemes such as 3-Row Reduction [31], HalfGates [41], or Three-Halves Garbling [36], where the output wire label is a function of the input wire labels. The protocol adds additional measures that at the end of the protocol allow the server to verify that the client has given the “right” circuit to the token before the outputs of the function can be decrypted (for this the token locally stores a hash of the circuit description seen). Our attack can be applied to this token-based protocol as now the client can simply give an AND gate with duplicate input wires and immediately decrypt intermediate wires, hence circumventing the output decryption mechanism.

3. Affected Garbling Schemes

In this section, we list several prominent garbling schemes and denote if they are affected by our attack presented in Sect. 2. For clarity, we are presenting the encryption function constructions along the schemes using the notation from Sect. 2.

Two Parties. We summarize the garbling schemes in Table 1, where the *permutation bit* (for Point-and-Permute, see [2]) of a wire label \tilde{w} is denoted as $p(\tilde{w})$ and k is a fixed (and public) symmetric key. As can be seen from Table 1, the encryption function of several garbling schemes for the two-party setting are vulnerable to our attack [13, 25, 28, 31, 33].

Interestingly, [4] explicitly states the requirement on circuits that gates must have different input wires ($\ell < r$), so their construction Garble1 with a Dual-Key Cipher

Table 1. Encryption functions in different garbled circuit constructions and whether the scheme is vulnerable to our attack .

Scheme	Encryption Function $E(\tilde{w}_\ell, \tilde{w}_r; \tilde{w}_o)$	Vulnerable
3-Row Reduction [31]		✓
LPS StM [25]	$\tilde{w}_o \oplus F(\tilde{w}_\ell; T) \oplus F(\tilde{w}_r; T)$	✓
PSSW StM [33]		✓
Fairplay [28]		✓
Standard Assumptions [13]		✓
Garble1 w. DKC from PRF [4]	$\tilde{w}_o \oplus F(\tilde{w}_\ell; T \parallel p(\tilde{w}_\ell) \parallel p(\tilde{w}_r)) \oplus F(\tilde{w}_r; T \parallel p(\tilde{w}_\ell) \parallel p(\tilde{w}_r))$	✗
Formally Verified [1]		✓ / ✗
Free-XOR [18]		✗
LPS ROM [25]	$\tilde{w}_o \oplus H(\tilde{w}_\ell \parallel \tilde{w}_r \parallel T)$	✗
PSSW ROM [33]		✗
Proof of Yao [23]	$\tilde{w}_o \oplus F(\tilde{w}_\ell; r_1) \oplus F(\tilde{w}_r; r_2), \quad r_1, r_2 \in_R \{0, 1\}^\kappa$	✗
AES256 [20]	$\tilde{w}_o \oplus F(\tilde{w}_\ell \parallel \tilde{w}_r; T)$	✗
Fixed-key AES [3]	$\tilde{w}_o \oplus F(k; 2\tilde{w}_\ell \oplus 4\tilde{w}_r \oplus T) \oplus 2\tilde{w}_\ell \oplus 4\tilde{w}_r \oplus T$	✗
HalfGates [41]	—	✗
Three-Halves [36]	—	✗

T : tweak, F : pseudo-random function, H : cryptographic hash function modeled as random oracle

(DKC) instantiated from a PRF is not affected by our attack. In fact, [4] mentions in footnote 9 of their full version that [31] “cannot handle a wire being used twice as an input to another gate”, which might refer to the weakness discovered in our work. This construction is used in the formally verified secure two-party implementation of [1], where the EasyCrypt code contains the restriction that input wires are different ($\ell < r$), but this is not mentioned in the paper [1].

We note that the most recent garbling schemes [36,41] are not vulnerable to our attack as they go beyond encrypting single table entries. However, these schemes rely on a non-standard circularity assumption [9], whereas today’s best garbling scheme under standard assumptions [13] is affected by our attack.

Multiple Parties. In addition to the schemes in Table 1, Ben-Efraim et al. [6] described an encryption function for multi-party garbling based on the approach of Beaver-Micali-Rogaway (BMR) [2] using a *double-key pseudo-random function* as $F^2(\tilde{w}_\ell, \tilde{w}_r; T \parallel j)$, where j is the index of the party. Implementations of this can also be vulnerable to our attack, e.g., if $F^2(\tilde{w}_\ell, \tilde{w}_r; T \parallel j) = F(\tilde{w}_\ell; T \parallel j) \oplus F(\tilde{w}_r; T \parallel j)$ and the full construction for N parties is

$$\bigoplus_{i=1}^N (F(\tilde{w}_\ell; T \parallel j) \oplus F(\tilde{w}_r; T \parallel j)) \oplus \tilde{w}_o .$$

A similar construction was presented in [7] before the authors described a different construction using a key-homomorphic PRF. Later in their paper the authors wanted to “ensure that the same PRF is not queried twice with the same input” (see Definition

Table 2. MPC frameworks, their encryption function construction, and whether the framework is vulnerable to our attack .

MPC Framework	Encryption Function	Vulnerable
Fairplay [28]	Fairplay [28]	✓
TASTY [14]	Similar to [18, 25, 33] (ROM)	✗
FastGC [15]		✗
VMCrypt [27]		✗
ABY [10]		✗
TinyGarble [38]		✗
Obliv-C [40]	HalfGates [41]	✗
OblivM [26]		✗

7 in [7]). However, as in the two-party garbling schemes, the technical details are not explicitly presented and the encryption function remains vulnerable.

The focus in these works seems to be always on the case where two (or more) gates share a common input wire, hence the encryption functions are designed to include some gate identifier, i.e., the tweak T . We emphasize that this is not enough, since the same gate can also have duplicate input wires.

4. Affected MPC Frameworks

In this section, we list several MPC frameworks and denote if they are vulnerable to our attack presented in Sect. 2. We summarize the frameworks in Table 2 alongside with the used encryption function construction (cf. Sect. 3), and if the framework is vulnerable.

Fairplay [28] uses an affected encryption function (see Table 1) and is vulnerable to our attack. We note that their compiler replaces AND gates with duplicate input wires with identity gates (similarly, the Frigate compiler [29] replaces them with a wire), so such gates are not generated automatically. However, one can input a circuit containing an AND gate with duplicate input wires into Fairplay’s GC evaluation engine making our attack work (see Appendix A for details).

The most recent MPC frameworks such as TASTY [14], FastGC [15], VMCrypt [27], ABY [10], TinyGarble [38], Obliv-C [40], and OblivM [26] implement a more recent encryption function for garbling and are not vulnerable to our attack.

5. Fixing Garbled Circuit with Per-Wire Encryption Pads

In this section, we give two countermeasures against our attack on Garbled Circuits (GCs) with per-wire encryption pads and AND gates with duplicate input wires.

1. The first countermeasure is to implement a simple validity check to make sure that the two input wires are indeed distinct, i.e., $\ell \neq r$.
2. The second countermeasure is to modify the encryption function such that the two PRFs are called on differing values even when both inputs are the same.

Modern garbling schemes such as fixed-key AES garbling [3] already provide this inherently. Otherwise, we can simply append a 0 to the left PRF call and a 1 to the right PRF call. More specifically, we modify Eq. (1) into the following:

$$E(\tilde{w}_\ell, \tilde{w}_r; \tilde{w}_o) = \tilde{w}_o \oplus F(\tilde{w}_\ell; T \parallel 0) \oplus F(\tilde{w}_r; T \parallel 1). \quad (6)$$

So far, including the gate index in the tweak T ensured that PRF calls are unique when the same wire is used as input to multiple AND gates. Moreover, including the permutation bits (resp. position in the garbled table) ensured that the PRF calls are unique for the four garbled table entries as the input wire labels are used to encrypt multiple garbled table entries. By additionally including the left/right bit as in Eq. (6), we ensure that also within one garbled table entry the PRF calls are unique for the left and right input wire, even if this is the same wire.

Since most MPC frameworks decouple the task of compiling the function to be evaluated into a Boolean circuit from the secure evaluation of the Boolean circuit, we recommend to apply *both* countermeasures: The circuit compiler should check for all gates that the two input wires are always distinct, independent of the selection of the underlying encryption function (which could be even unknown at the compilation stage). Additionally, the implementation of the GC generation and evaluation engine should ensure that the protocol is secure even if a gate has duplicate input wires, e.g., by using the unaffected encryption functions from Table 1, or by applying the first or second countermeasure for the vulnerable schemes.

6. Related Work

Our work fits into a series of works that describe subtle issues in the area of garbling schemes, including their usage in higher-level protocols, proving their security, and their implementation.

In [3, §3 “An Insecurity in Prior Works” and Fig. 3], Bellare et al. present a very related attack on garbling schemes with two independent PRF calls. They show that when using Free-XOR [18], a circuit can be built s.t. the input wire labels of the two (different) input wires to an AND gate are the same, and then the per-wire encryption pads cancel out resulting in a break of privacy. This is exactly the same observation as ours, but more general as it even works for circuits where no gate has duplicate input wires. In our paper, we focus on the related issue when an AND gate has duplicate input wires which also applies to garbling schemes without using Free-XOR, show that privacy breaks for a large class of circuits, and survey the affected garbling schemes and implementations.

Kiraz and Schoenmakers [17] describe an attack on Garbled Circuit (GC)-based protocols with stronger security using the Cut-and-Choose technique as it was performed in the initial work of Pinkas [32] and the implementation in Fairplay [28]. In these protocols, multiple GCs are generated of which some are opened and checked for correctness. The problem discovered by Kiraz and Schoenmakers stems from the way OTs are used with which the evaluator obviously obtains the wire labels corresponding to his inputs. They show that classical OT is not enough and propose to use committing OT instead.

The issues pointed out were resolved in subsequent Cut-and-Choose-based protocols such as [22,24,25,33].

Choi et al. [9] show that not only a variant of correlation robustness (as claimed in [18]), but also a form of circular security is required to prove security of the Free-XOR technique [18] using a weaker assumption than a random oracle.

Levi and Hazay [21] give side-channel attacks on implementations of GCs where all wire labels are correlated due to the Free-XOR technique [18].

7. Conclusion

In this work, we presented an attack against garbling schemes that use deterministic encryption functions to derive per-wire encryption pads for the left and right input wire that are XORed together. The attack stems from the implicit, but oftentimes not explicitly mentioned assumption that the input wires to an AND gates must be different. We showed that our attack affects several prominent garbling schemes and can occur in practical systems. We also showed how to easily and efficiently fix the vulnerable constructions to prevent our attack.

In the future, we hope that our findings help to decouple or at least make explicit the required properties about the circuit (or more general the function to be computed) which is the interface between compilation and the underlying secure computation protocols. This is important not only in papers with cryptographic constructions and proofs, but especially in implementation frameworks and in languages and compilers that aim at formally verified secure computation, e.g., [1,34,35].

Acknowledgements

We thank Claudio Orlandi for pointing us to the related attack and observations in [3, §3], cf. 6. This project received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 850990 PSOTI). It was co-funded by the Deutsche Forschungsgemeinschaft (DFG) within SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251805230.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Our Attack on Fairplay

In this section, we show how to apply our attack on the Fairplay MPC framework [28]. We start with the original circuit in Fig. 1a, where Alice is the attacker with two input wires and Bob has one input wire.

First, Alice turns the original circuit into the modified circuit shown in Fig. 1b, where the first input of Alice goes into both inputs of the first AND gate. Note that both circuits compute the same function. Since the compiler of Fairplay either removes or replaces the first AND gate with an identity gate, Alice must manually modify the circuit description in Fairplay’s Secure Hardware Definition Language (SHDL). The corresponding file `Opt.circuit` is given in List. 1.

The idea is that now the evaluator Alice gets both input wire labels for the final AND gate and can perform our *General Attack* as explained in Sect. 2.1 from which she can determine the output value of the XOR gate. This value reveals the input value of Bob, since Alice knows the other input value for the XOR circuit (and the output value from the attack). In this example also the output value is revealed, since Alice now knows both input values for the final AND gate.

A similar attack can be applied to more complex circuits in order to reveal Bob’s input values.

Listing 1. Modified circuit for our attack (see Fig. 1(b)) in Fairplay’s Secure Hardware Definition Language (SHDL).

```

0 input //output$input.bob$0
1 input //output$input.alice[1]$0
2 input //output$input.alice[0]$0
3 gate arity 2 table [ 0 0 0 1 ] inputs [ 2 2 ] // duplicate input wires
4 gate arity 2 table [ 0 1 1 0 ] inputs [ 0 1 ]
5 output gate arity 2 table [ 0 0 0 1 ] inputs [ 4 3 ] //output$output.bob$0
    
```

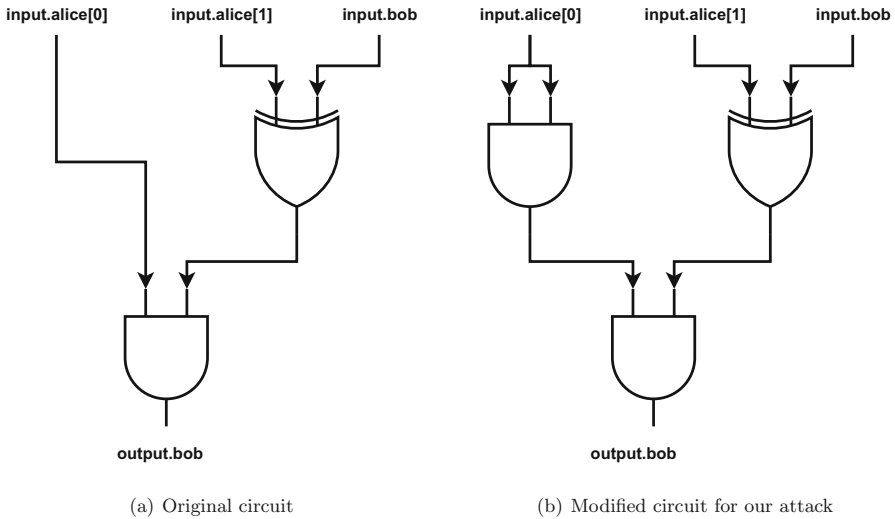


Fig. 1. Example circuit .

References

- [1] J.B. Almeida, M. Barbosa, G. Barthe, F. Dupressoir, B. Grégoire, V. Laporte, V. Pereira, A fast and verified software stack for secure function evaluation, in *CCS* (2017)
- [2] D. Beaver, S. Micali, P. Rogaway, The round complexity of secure protocols, in *STOC* (1990)
- [3] M. Bellare, V.T. Hoang, S. Keelveedhi, P. Rogaway, Efficient garbling from a fixed-key blockcipher, in *S&P* (2013)
- [4] M. Bellare, V.T. Hoang, P. Rogaway, Foundations of garbled circuits, in *CCS* (2012). Full version: <https://ia.cr/2012/265>
- [5] A. Ben-David, N. Nisan, B. Pinkas, FairplayMP: A system for secure multi-party computation, in *CCS* (2008)
- [6] A. Ben-Efraim, Y. Lindell, E. Omri, Optimizing semi-honest secure multiparty computation for the Internet, in *CCS* (2016)
- [7] A. Ben-Efraim, Y. Lindell, E. Omri, Efficient scalable constant-round MPC via garbled circuits, in *ASIACRYPT* (2017)
- [8] N. Büscher, D. Demmler, S. Katzenbeisser, D. Kretzmer, T. Schneider, HyCC: Compilation of hybrid protocols for practical secure computation, in *CCS* (2018)
- [9] S.G. Choi, J. Katz, R. Kumaresan, H.-S. Zhou, On the security of the “Free-XOR” technique, in *TCC* (2012)
- [10] D. Demmler, T. Schneider, M. Zohner, ABY—A framework for efficient mixed-protocol secure two-party computation, in *NDSS* (2015)
- [11] M. Franz, A. Holzer, S. Katzenbeisser, C. Schallhart, H. Veith, CBMC-GC: An ANSI C compiler for secure two-party computations, in *Compiler Construction (CC)* (2014)
- [12] O. Goldreich, S. Micali, A. Wigderson, How to play ANY mental game, in *STOC* (1987)
- [13] S. Gueron, Y. Lindell, A. Nof, B. Pinkas, Fast garbling of circuits under standard assumptions, in *CCS* (2015)
- [14] W. Henecka, S. Kögl, A. Sadeghi, T. Schneider, I. Wehrenberg, TASTY: Tool for automating secure two-party computations, in *CCS* (2010)
- [15] Y. Huang, D. Evans, J. Katz, L. Malka, Faster secure two-party computation using garbled circuits, in *USENIX Security* (2011)
- [16] K. Järvinen, V. Kolesnikov, A. Sadeghi, T. Schneider, Embedded SFE: Offloading server and network using hardware tokens, in *FC* (2010)
- [17] M. Kiraz, B. Schoenmakers, A protocol issue for the malicious case of Yao’s garbled circuit construction, in *Information Theory in the Benelux (SITB)* (2006)
- [18] V. Kolesnikov, T. Schneider, Improved garbled circuit: Free XOR gates and applications, in *ICALP* (2008)
- [19] B. Kreuter, A. Shelat, B. Mood, K. Butler, PCF: A portable circuit format for scalable two-party secure computation, in *USENIX Security* (2013)
- [20] B. Kreuter, A. Shelat, C. Shen, Billion-gate secure computation with malicious adversaries, in *USENIX Security* (2012)
- [21] I. Levi, C. Hazay, Garbled circuits from an SCA perspective: Free XOR can be quite expensive..., in *CHES* (2023)
- [22] Y. Lindell, B. Pinkas, An efficient protocol for secure two-party computation in the presence of malicious adversaries, in *EUROCRYPT* (2007)
- [23] Y. Lindell, B. Pinkas, A proof of security of Yao’s protocol for two-party computation. *J. Cryptol.* (2009)
- [24] Y. Lindell, B. Pinkas, Secure two-party computation via cut-and-choose oblivious transfer. *J. Cryptol.* (2012)
- [25] Y. Lindell, B. Pinkas, N.P. Smart, Implementing two-party computation efficiently with security against malicious adversaries, in *SCN* (2008)
- [26] C. Liu, X.S. Wang, K. Nayak, Y. Huang, E. Shi, OblivM: A programming framework for secure computation, in *S&P* (2015)
- [27] L. Malka, VMCrypt: Modular software architecture for scalable secure computation, in *CCS* (2011)
- [28] D. Malkhi, N. Nisan, B. Pinkas, Y. Sella, Fairplay—Secure two-party computation system, in *USENIX Security* (2004)

- [29] B. Mood, D. Gupta, H. Carter, K. Butler, P. Traynor, Frigate: A validated, extensible, and efficient compiler and interpreter for secure computation, in *EuroS&P* (2016)
- [30] B. Mood, L. Letaw, K. Butler, Memory-efficient garbled circuit generation for mobile devices, in *FC* (2012)
- [31] M. Naor, B. Pinkas, R. Sumner, Privacy preserving auctions and mechanism design, in *Electronic Commerce (EC)* (1999)
- [32] B. Pinkas, Fair secure two-party computation, in *EUROCRYPT* (2003)
- [33] B. Pinkas, T. Schneider, N.P. Smart, S.C. Williams, Secure two-party computation is practical, in *ASIACRYPT* (2009)
- [34] A. Rastogi, M.A. Hammer, M. Hicks, Wysteria: A programming language for generic, mixed-mode multiparty computations, in *S&P* (2014)
- [35] A. Rastogi, N. Swamy, M. Hicks, Wys*: A DSL for verified secure multi-party computations, in *Principles of Security and Trust (POST)* (2019)
- [36] M. Rosulek, L. Roy, Three halves make a whole? Beating the half-gates lower bound for garbled circuits, in *CRYPTO* (2021)
- [37] T. Schneider, M. Zoher, GMW vs. Yao? Efficient secure two-party computation with low depth circuits, in *FC* (2013)
- [38] E.M. Songhori, S.U. Hussain, A. Sadeghi, T. Schneider, F. Koushanfar, TinyGarble: Highly compressed and scalable sequential garbled circuits, in *S&P* (2015)
- [39] A.C. Yao, How to generate and exchange secrets, in *FOCS* (1986)
- [40] S. Zahur, D. Evans, Obliv-C: A language for extensible data-oblivious computation. *Cryptology ePrint Archive 1153* (2015)
- [41] S. Zahur, M. Rosulek, D. Evans, Two halves make a whole—Reducing data transfer in garbled circuits using half gates, in *EUROCRYPT* (2015)