

Secure Two-Party Computation is Practical

Benny Pinkas¹, Thomas Schneider², Nigel P. Smart³, and Stephen C. Williams³

¹ Dept. of Computer Science,
University of Haifa
Haifa 31905, Israel
`benny@pinkas.net`

² Horst Görtz Institute for IT-Security,
Ruhr-University Bochum,
D-44780 Bochum, Germany,
`thomas.schneider@trust.rub.de`

³ Dept. Computer Science,
University of Bristol,
Woodland Road,
Bristol, BS8 1UB, United Kingdom,
`{nigel,williams}@cs.bris.ac.uk`

Abstract. Secure multi-party computation has been considered by the cryptographic community for a number of years. Until recently it has been a purely theoretical area, with few implementations with which to test various ideas. This has led to a number of optimisations being proposed which are quite restricted in their application. In this paper we describe an implementation of the two-party case, using Yao's garbled circuits, and present various algorithmic protocol improvements. These optimisations are analysed both theoretically and empirically, using experiments of various adversarial situations. Our experimental data is provided for reasonably large circuits, including one which performs an AES encryption, a problem which we discuss in the context of various possible applications.

1 Introduction

That secure multi-party computation can be executed at all is considered one of the main results of the theory of cryptography. Starting with Yao's seminal work [30] many authors have looked at various optimisations and extensions to the basic concept, for both the two-party and the multi-party settings, see for example [7, 10, 11, 18, 20, 23, 29]. Until recently all work on secure multi-party computation has been essentially of a theoretical nature, focusing on feasibility results. However in the last few years a number of practical implementations have appeared [5, 6, 22, 24, 3].

There are many different protocols for secure multi-party computation. Our work focuses on implementation of secure computation and therefore we only mention protocols which have been previously implemented. Secure multi-party computation essentially comes in two flavours. The first approach is typically

based upon secret sharing and operates on an arithmetic circuit representation of the computed function, such as in the BGW (Ben-Or, Goldwasser and Wigderson) or CCD (Chaum, Crepeau and Damgård) protocols [4, 8]. This approach is usually applied when there is an honest majority among the participants (which can only exist if more than two parties participate in the protocol). An alternative approach represents the function as a binary circuit. This approach was used in the original two-party garbled circuit construction of Yao [30], and in the GMW (Goldreich, Micali and Wigderson) multi-party protocol [11].

The arithmetic circuit method is better at representing addition and multiplication operations, where parties have additive shares of secret values, but cannot be used to compute comparisons unless the shares are converted to shares of the binary representation of the values. This approach has been used to great effect in the SIMAP project [6], which has resulted in a “real-life” application of secure multi-party computation to the Danish sugar beet industry [5].

The binary circuit approach handles arithmetic operations, especially multiplications, less efficiently, but can easily compute binary operations such as comparisons. This second approach, which forms the basis of Yao’s construction for the two party case, has been implemented by Malkhi et al. in the Fairplay system [24]. That system also provides a method to compile a given functionality from a representation in a high-level language into a circuit, which is then interpreted by a run-time environment that performs the secure evaluation of this functionality. FairplayMP, an extension of Fairplay to the case of more than two parties using a modified version of the protocol of Beaver et al. [2] has recently been released [3]. All these implementations provide security against semi-honest adversaries only. A major advantage of the binary circuit based systems (Fairplay and FairplayMP) is that they run in a constant number of communication rounds, whereas the SIMAP system has the advantage of being able to process arithmetic operations very efficiently.

Efficient extensions of Yao’s construction to more relevant adversarial models have been a topic of research interest in the last few years. There are several constructions which aim to secure the protocol against malicious adversaries without using generic zero-knowledge protocols. We will focus on the construction of Lindell and Pinkas [20] which is efficient and provides fully simulatable security according to the definition of Canetti [7]⁴. A definition of a weaker class of corruption, “covert adversaries”, and a protocol secure against this type of behavior, was provided by Aumann and Lindell [1]. In [22] an implementation of the basic Lindell–Pinkas protocol was reported upon and experimental data in various security models was provided.

⁴ This construction may be preferable over other two-party protocols with security against malicious adversaries. The construction of Mohassel and Franklin [23] only protects privacy and is not fully simulatable. The construction of Jarecki and Shmatikov [18] requires the use of public-key operations, rather than symmetric key operations, for any gate of the circuit. The construction of Nielsen and Orlandi [26], too, uses public key operations, or rather public-key based commitments, for each key of every wire of the circuit. A precise practical comparison between the different approaches is beyond the scope of the current paper

In this paper we improve on the implementation of [22] in a number of ways. The resulting set of quantitative improvements results in qualitative conclusions: (1) We demonstrate that two-party computation, secure against malicious adversaries, is truly practical, and we experimentally identify the performance bottlenecks which remain after our optimisations. This result should direct further research to the issues which have the largest effect on performance. (2) We experiment with a secure computation of the AES standard, and show that it is indeed feasible, even with security against malicious adversaries. There are a number of applications of such an implementation, some of which we describe below. (3) We provide the first implementation of a protocol with security against covert adversaries and we compare the performance of all 3 types of protocols: malicious, covert and semi-honest.

A more detailed summary of our main results is as follows:

- We improve the communication cost for transmitting the circuits between the parties. In the case when we model the underlying key derivation functions (KDFs) as correlation robust (see discussion below), using the technique of [19] we are able to transmit no information for the XOR gates within the circuit. In this situation we are also able to reduce the data which needs to be sent by 25% for the other gates. When we are not willing to model the KDFs as correlation robust, and we only assume they are pseudo-random functions, we are unable to perform the free XOR optimisation. However we are able to reduce the communication cost for all gates by 50%. Unlike other methods used to improve communication, like [13], our improvement makes a marginal impact on computational costs. We will return to this in a later section.
- In addition to the theoretical analysis we provide experimental data for evaluating “real life” circuits, in both the honest-but-curious, covert and malicious adversary cases; also for the two different methods in the literature that construct the auxiliary circuits in the covert and malicious cases (see [22] and the full version). The implementation for the malicious setting is based on the construction of Lindell and Pinkas [20] which provides security in the sense of full simulatability. Therefore the resulting construction can be used as a black-box primitive in more complex applications. The use of our optimisations results in a considerable performance boost compared to previous experimental results published in [22].
Our optimisations change the performance bottleneck to a different part of the computation; namely, the verification of garbled circuits generated by the circuit constructor. This observation is important for focusing future research on the issues that affect the overhead the most.
- We experiment with secure evaluation of a circuit which computes an AES encryption of a single block. The secure computation of AES involves one party which knows the key, and a different party which has an input block. The second party learns the encryption of the block, while the first party learns nothing. We demonstrate the feasibility of computing this function in the semi-honest, covert and malicious settings.

Secure evaluation of AES has an impact in a number of scenarios which we will discuss in short here and elaborate on in the full version. The fact that a secure computation of AES is feasible, and can run in a matter of seconds, is quite surprising.

Application 1, OPRF: A secure computation of a pseudo-random function, denoted OPRF for “oblivious prf”, has been defined in [9] for the purpose of secure keyword based searches, and was subsequently used in different applications. The OPRF protocol in [9] is based on the Naor-Reingold prf, which is a number theoretic construction. Our construction has different advantages over the NR based construction, which we detail in the full version.

Application 2, Side Channel Protection: In [12] the authors introduce “one-time programs”, which are programs that can only be executed once and then “self-destruct”. An important advantage of this construction is that the execution of the program reveals no side-channel information. Most of the computation in that construction is essentially done using a garbled Yao circuit.

One of the main applications of smart cards is to compute symmetric encryptions, and therefore the ability to compute AES encryptions by Yao circuits has immediate application in the above scenario. It enables smart cards to perform a one-time computation, secure against side-channel attacks, of AES. This is particularly interesting since in that setting the circuit evaluation need only be secure against semi-honest adversaries, while we show below that semi-honest computation of AES can be run very efficiently, taking only a few seconds.

Application 3, Blind MACs and Blind Encryption: One can think of the operation of obtaining the AES encryption of a message, under the other party’s secret key, as a blind MAC or a blind symmetric encryption. These operations have different applications in secure computation.

Application 4, Third Party Operations on Encrypted Data: We essentially show that encryption and decryption can be implemented using circuits. This enables secure computation of homomorphic operations on encrypted data. This operation is done by a circuit which receives two ciphertexts from one party and a key from the other party, decrypts the ciphertexts, applies some arbitrary mathematical operation to the plaintexts, and then encrypts the result.

2 Yao’s Garbled Circuit Construction

Two-party secure function evaluation makes use of the famous garbled circuit construction of Yao [30] which we briefly overview in this section. The basic idea is to encode the function to be computed via a binary circuit and then to securely evaluate the circuit on the players’ inputs.

2.1 Garbled Circuits

We consider two parties, denoted as P_1 and P_2 , who wish to compute a function securely which is represented as a simple binary circuit. First assume the circuit

consists of only a single gate with two input wires and one output wire. We denote the input wires by w_1 and w_2 , and the output wire by w_3 . The input to w_1 is denoted by b_1 and is known to P_1 , similarly P_2 knows the input to w_2 and this is given by b_2 . Each gate has a unique identifier Gid ; this enables a circuit fan out of greater than one, i.e., it enables the output wire of one gate to be used in more than one other gate. We require that P_2 evaluates the gate on the two inputs, without P_1 learning anything, and without P_2 determining the value b_1 , bar what it can deduce from the output of the gate and its own input. We define the output of the gate by the function $G(b_1, b_2) \in \{0, 1\}$.

The construction of Yao works as follows. P_1 encodes, or garbles, each wire w_i by selecting two different cryptographic keys k_i^0 and k_i^1 of length t . Here t is a computational security parameter which suffices for the length of a symmetric encryption scheme. A random permutation π_i of $\{0, 1\}$ is associated to each wire. The garbled value of wire w_i is then represented by $k_i^{b_i} \| c_i$, where $c_i = \pi_i(b_i)$. We call the value c_i the ‘‘external value’’ of the wire, note that this value is completely independent of the actual value of the wire b_i .

An encryption function $E_{k_1, k_2}^s(m)$ is selected which has as input two keys of length t , a message m , and some additional information s . The additional information s must be unique per invocation of the encryption function, i.e., it is used only once for any choice of keys. The gate itself is then replaced by a four entry table indexed by the values of c_1 and c_2 , and given by

$$c_1, c_2 : E_{k_1^{b_1}, k_2^{b_2}}^{\text{Gid} \| c_1 \| c_2} \left(k_3^{G(b_1, b_2)} \| c_3 \right),$$

where $c_1 = \pi_1(b_1)$, $c_2 = \pi_2(b_2)$, and $c_3 = \pi_3(G(b_1, b_2))$. Each entry in the table corresponds to a combination of the values of the input wires and contains the encryption of the corresponding garbled output value. The resulting look up table, or set of look up tables in general, is called the ‘‘garbled circuit’’.

Player P_1 then sends to P_2 the garbled circuit, the key corresponding to its input value $k_1^{b_1}$, the value $c_1 = \pi_1(b_1)$, and the permutation π_3 . The parties engage in an oblivious transfer (OT) protocol so that P_2 learns the value of $k_2^{b_2} \| c_2$, where $c_2 = \pi_2(b_2)$. Player P_2 can then decrypt the entry in the look up table indexed by (c_1, c_2) using $k_1^{b_1}$ and $k_2^{b_2}$; revealing the value of $k_3^{G(b_1, b_2)} \| c_3$. P_2 determines the value of $G(b_1, b_2)$ by using the mapping π_3^{-1} from c_3 to $\{0, 1\}$.

In the general case the circuit consists of multiple gates. Player P_1 chooses random garbled values for all wires and uses them for constructing tables for all gates. It sends these tables, i.e., the garbled circuit, to P_2 and in addition provides P_2 with the garbled values and the c values of P_1 's inputs, and with the permutations π used to encode the *output* wires of the circuit. Player P_2 uses invocations of oblivious transfer to learn the garbled values and c values of its own inputs to the circuit. Given these values, P_2 can evaluate the gates in the first level of the circuit, compute the garbled values and the c values of their output wires. Player P_2 can then continue with this process and compute the garbled values of all wires in the circuit. Finally P_2 uses the π permutations of the output wires of the circuit to compute the real output values of the circuit.

If P_1 additionally requires some output from the circuit then this can be dealt with by standard mechanisms, as described in the full version.

One could use more general gates than 2-to-1 gates, such as n -to- m gates with 2^n entries. However the optimisations we shall present in this paper are most effective when applied to 2-to-1 gates. While we found that more general gates can improve the performance of a naive Yao circuit protocol, they actually decrease the performance of the optimisations. Hence the rest of this paper is restricted to 2-to-1 gates.

2.2 Required Implementation Details

Having described the basic theoretical description of Yao’s protocol and its extensions, we now present a number of implementation details which are needed to understand some of our optimisations. The basic implementation choice of the underlying encryption scheme to be used is the same as the implementation described in [22].

Oblivious transfer: Unlike [22] we do not use the OT scheme of Hazay and Lindell (HL) [15]. Instead we use the OT scheme of Peikert et al. (PVW)[27]. This scheme is UC-secure and hence requires the setup of a Common Reference String (CRS) of a few hundred bits. For our experiments we assume that this is given to the parties. (Alternatively, the parties can run a coin-tossing protocol to generate the CRS, which is possible due to the nature of the CRS used in the PVW scheme.) The batched method of PVW is more efficient per OT than the batched method of HL, especially on the receiver’s side. In particular the CRS can be used for any number of invocations of the OT, whereas the method in HL requires the maximum number of OT’s being executed to be known before the setup is performed. (The setup in HL also requires two ZK-proofs as opposed to a CRS being created in PVW.) The OT stage is not our computational bottleneck, and is unlikely to be, unless one is in the rare situation of having a circuit with a large number of inputs for P_2 and yet a relatively small number of gates. Thus we do not consider optimisations of OT schemes which are secure against only semi-honest or covert adversaries, since the fully secure OT is efficient enough.

Encryption scheme: The only implementation detail we will need from [22] is that the encryption scheme is implemented via

$$E_{k_1, k_2}^s(m) = m \oplus \text{KDF}^{|m|}(k_1, k_2, s)$$

where KDF is a key derivation function, whose $|m|$ bits of output are independent of the two input keys in isolation, and which depends on the value of s . We will instantiate this function as follows⁵

$$\text{KDF}^\ell(k_1, k_2, s) = H(k_1 \| s)_{1\dots\ell} \oplus H(k_2 \| s)_{1\dots\ell}.$$

⁵ In [22] two instantiations were presented, depending on whether we are working in the random oracle model (ROM) or standard model, via truncating, or extending, the output of a suitable hash function H in the standard way as follows

$$\text{KDF}^\ell(k_1, k_2, s) = \begin{cases} H(k_1 \| k_2 \| s)_{1\dots\ell} & H \text{ is modeled as an RO,} \\ H(k_1 \| s)_{1\dots\ell} \oplus H(k_2 \| s)_{1\dots\ell} & H \text{ is modeled as a PRF.} \end{cases}$$

Even if H is a Merkle–Damgård type hash function this will be secure (with the associated issues of length extension), since we are only applying the function to fixed length inputs. Indeed, in our experiments we implement H using SHA-256.

Modeling the hash function, and correlation robustness: In this paper we need to model the underlying hash function H in two ways. In the first we make the usual assumption that it behaves as a pseudo-random function, namely that $H(k\|s)$ is an invocation of a pseudo-random function keyed by k , with the input s . However one of our optimisations requires that we make a stronger assumption on the hash function, namely that it is **correlation robust**. This later property can be stated formally as follows:

Definition 1 (Correlation robustness [16]). *An efficiently computable function $H : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$ is correlation robust if the following distribution is pseudo-random: $(t_1, \dots, t_m, H(t_1 \oplus r), \dots, H(t_m \oplus r))$, where t_1, \dots, t_m and r are chosen at random, and m is polynomial in the security parameter.*

This can also be stated by saying that the function $f_r(x) = H(x \oplus r)$ is a weak pseudo-random function. The definition also implies that the distribution of $(H(t_1), \dots, H(t_m), H(t_1 \oplus r), \dots, H(t_m \oplus r))$ is pseudo-random.

The correlation-robustness assumption is satisfied by a random oracle (or rather by a very weak form of it: a non-programmable, non-extractable random oracle). However, assuming correlation robustness seems as a much weaker requirement than assuming the existence of random oracles. This assumption has been introduced in [16] and was used there for providing security against malicious adversaries for a method of extending oblivious transfer. The correlation robustness assumption has been recently used in the context of oblivious transfer [14, 17] and in the context of secure computation [19, 26].

For our construction, as we deal with circuits with arbitrary fan out, we require a slightly modified definition. Namely that for any set $S = \{s_1, \dots, s_{|S|}\}$ of size which is of the same order as the number gates, the distribution of $(t_1, \dots, t_m, \langle H((t_1 \oplus r)\|s_1), \dots, H((t_m \oplus r)\|s_1) \rangle, \langle H((t_1 \oplus r)\|s_2), \dots, H((t_m \oplus r)\|s_2) \rangle, \dots, \langle H((t_1 \oplus r)\|s_{|S|}), \dots, H((t_m \oplus r)\|s_{|S|}) \rangle)$ is pseudo-random, where t_1, \dots, t_m and r are chosen at random. In other words, all the pads that are used for encrypting table entries are pseudo-random. If one is willing to assume this then our optimisations provide highly efficient protocols. We also provide optimisations for when the user is unwilling to make such an assumption.

3 Structural Optimisations of the Circuit

Yao’s protocol operates on functions which are described as a boolean circuit, and its overhead depends on the size of the circuit. A convenient way of generating

The difference is that the security analysis in the ROM works even if we feed related keys to different invocations of the function. Namely, it is possible to compute, say, $H(k_1\|k_2)$, $H(k_1\|k'_2)$, $H(k'_1\|k_2)$ and $H(k'_1\|k'_2)$ and claim that knowledge of k_1, k_2 does not disclose information about any of the values except $H(k_1\|k_2)$. This is impossible in the standard model. Therefore if $H()$ is modeled as a prf it must be invoked separately with each key.

a representation of a function in this form is to use a compiler which translates a description of a function in a high-level language to a description as a binary circuit. The Fairplay system provides a compiler for this task which operates on functions described in a high-level language called Secure Function Description Language (SFDL) [24, 3]. We use that compiler as the basis of our experiments, but use our own run-time environment to execute the protocol.

There are a number of general circuit simplifications which can be performed to the output of the Fairplay compiler. We have implemented a number of these, based on two basic ideas: (1) identifying component circuits which can be replaced by simpler combinations of gates, and (2) identifying complicated components whose output must always be zero, or one; this allows for the component to be removed and other subsequent components to be further simplified. A combination of these techniques is surprisingly effective, and allows us to produce circuits which are often 60 percent more efficient than the circuit produced by the Fairplay compiler.

Many of the techniques used are ad-hoc, but the following technique is particularly effective. First, by a technique akin to common sub-expression elimination, we identify sets of gates which can be replaced by a single 3-to-1 gate, and then replace the 3-to-1 gate with a set of 2-to-1 gates which was chosen to minimize the number of non-XOR gates. This is particularly effective when combined with our later technique of Section 4, in the case of correlation robust KDFs, to remove the cost of any XOR gates; however the technique is also successful in the more general case as well. We call a gate *even* if its truth table has an even number of ‘1’ entries (for example, a XOR gate is even), otherwise it is called *odd* (an OR gate, for example, is odd). We show in the full version that it is possible to replace any 3-to-1 even gate with at most a single 2-to-1 non-XOR gate and at most three XOR gates. The optimal transformation rules, which we found by exhaustive search, are listed in the full version.

4 Optimisations with Free XORs, when the KDF is Correlation Robust

In [19] Kolesnikov and Schneider present an optimisation based on the correlation robustness assumption, which allows XOR gates to be evaluated for free, thus doing away with the need to evaluate or transmit the garbled tables for such gates. The optimisation requires that there is a global random value R of bit length t , known only to P_1 , such that for all garbled wires w_i it holds that $k_i^1 = k_i^0 \oplus R$. In other words, the garbling of the 1 value of a wire, is determined purely from XOR-ing the garbled 0 value with the value R . Note that a similar property holds for the external values of the wire: $\pi_i(1) = \pi_i(0) \oplus 1$. With this convention we have that a XOR gate can be implemented by simply XOR-ing together the two garbled input values, and the two external values. Namely, for a XOR gate mapping wires w_1 and w_2 to wire w_3 , it holds that $k_3 = k_1 \oplus k_2$ and $c_3 = c_1 \oplus c_2$. For a full proof of this optimisation see [19]. Note that [19] states

the proof in the random oracle model, but it can be easily seen, as noted in [19], that the proof can be based on the correlation robustness assumption.

Garbled Row Reduction – GRR: The above solution is ideal for XOR gates, but in addition we would like to reduce the size of the tables of the non-XOR gates as well. The following simple optimisation (which was pointed out in [25]) provides a 25 percent reduction in the sizes of the tables needed to represent two-input gates. We can do this in a way which still allows the use of the above trick for free XOR gates. (In general, this method provides a $1/2^n$ reduction in the size of n -to-1 gates, but we will only describe it in detail for the two input case.)

The observation is that instead of defining the two garbled values of the output wires randomly, we can define one of them as a function of garbled values of the two input wires which result in this output value. In other words, we choose an input pair $(b_1, b_2) \in \{0, 1\}^2$, and define the garbled output value of $G(b_1, b_2)$ to be a function of the garbled values of b_1 and b_2 . The gate table therefore need not store an entry for the input combination (b_1, b_2) . In the evaluation phase, if the evaluator has the garbled values of the pair (b_1, b_2) it can compute the corresponding garbled output directly, without consulting the gate table.

Suppose the gate maps wire w_1 and wire w_2 to wire w_3 . As before we let k_i^0 and k_i^1 denote the garbled wire values, $G(b_1, b_2)$ denote the function being implemented by the gate, and we set the external value of the wire to be $c_i = \pi_i(b_i)$. We then define the garbled output value corresponding to the output resulting from the external input values $(c_0, c_1) = (0, 0)$ as

$$k_3^{G(\pi_1^{-1}(0), \pi_2^{-1}(0))} \| c_3 = \text{KDF}^{t+1} \left(k_1^{\pi_1^{-1}(0)}, k_2^{\pi_2^{-1}(0)}, \text{Gid} \| 0 \| 0 \right).$$

In other words, the garbled value is exactly equal to the pseudo-random mask that was used to hide it in the basic protocol. Note that this operation also defines the external value c_3 of this output value. We therefore define π_3 such that $c_3 = \pi_3(G(\pi_1^{-1}(0), \pi_2^{-1}(0)))$. The other garbled value of the output wire, $k_3^{1-G(\pi_1^{-1}(0), \pi_2^{-1}(0))}$ is then chosen as in the free XOR method above, to enable the evaluation of XOR gates for free. The table is then constructed in the standard way except that we do not store, or transmit, its first entry.

On evaluating the garbled gate the evaluator proceeds as in the standard algorithm except when it wishes to access the first entry of the table, i.e., when the external values of both input wires are 0, namely $c_1 = c_2 = 0$. In that case it possesses the garbled values $k_1^{b_1}$ and $k_2^{b_2}$, where $b_1 = \pi_1^{-1}(0)$ and $b_2 = \pi_2^{-1}(0)$. It uses them to compute $k_3^{G(b_1, b_2)}$ and $c_3 = \pi_3(G(b_1, b_2))$, by computing $\text{KDF}^{t+1} \left(k_1^{b_1}, k_2^{b_2}, 0 \| 0 \| \text{Gid} \right)$ as defined in the equation above.

We will denote this optimisation as Garbled Row Reduction, GRR for short, in our future discussions.

Security: We sketch why the above optimisation maintains security. Recall that the proof of security for Yao’s protocol given in [21] shows security against a corrupt P_2 based on a hybrid argument, and on a claim that for each gate it is

infeasible to distinguish between a correct garbled table of this gate and a table which encrypts the same value in all four entries. In order for this argument to apply to the GRR optimisation, it is required to show that it is infeasible to find out if the garbled value assigned to the first table entry, $k_3^{G(\pi_1^{-1}(0), \pi_2^{-1}(0))} \parallel c_3$ is equal to the values encrypted in the other entries. However this value is equal to the mask that is used to encrypt the first entry in Yao’s original protocol, and we know that if a polynomial adversary is given only a single pair of garbled input values then the masks that are used for encrypting the other entries of the table are pseudo-random. Therefore the claim follows.

5 Optimisations without Free Xors, when the KDF is not Correlation Robust

One may not want to assume the KDF is correlation robust, or perhaps the proportion of XOR gates in the circuit is so low that making this assumption is not as effective. In these situations, too, we would like to reduce the overhead required by the Yao circuit. This section describes an optimisation which reduces the size of every two-input gate by 50%, but which, unfortunately, cannot be combined with the free XOR method of Section 4.

The underlying idea is that if we are not using the free XOR trick then the two values of the output wire can be chosen independently.⁶ The 50% reduction in the size of the gate tables is based on Shamir secret sharing [28]. It makes use of a finite field \mathbb{F}_{2^t} . Recall that t is the bit length of the keys used to represent the garbled values of the wires. We can therefore interpret keys as elements of \mathbb{F}_{2^t} and vice versa. We also interpret small integers such as 1, 2, 3 etc. as elements in \mathbb{F}_{2^t} . For example if we think of \mathbb{F}_{2^t} as $\mathbb{F}_2[X]/(f(X))$, for some polynomial of degree t , then the integer 3 can be interpreted as $x + 1$.

As before we assume a garbled table indexed by the external values, c_1 and c_2 , and each entry corresponds to the value being output, on input of the values $k_1^{b_1}$ and $k_2^{b_2}$ where $b_i = \pi_i^{-1}(c_i)$. We set the rows of the gate table to be numbered $1, \dots, 4$, and therefore set $r = 2c_1 + c_2 + 1$ to be the row number of table entry (c_1, c_2) . We define the value used to mask this entry as

$$K_r \parallel M_r = \text{KDF}^{t+1}(k_1^{b_1}, k_2^{b_2}, s) \quad (1)$$

where $s = \text{Gid} \parallel c_1 \parallel c_2$, K_r is a bit string of length t bits and M_r is a single bit used to mask the external value of the output. We use a different method for optimising odd and even gates. The truth table of each gate, and therefore also the information whether the gate is odd or even, is known to the circuit evaluator. Therefore it can compute each gate according to the right method. (The only information hidden from the evaluator is the values passing on intermediate wires of the circuit.)

⁶ This allows for possible extensions of the GRR method, and in the full version we detail another optimisation method, which we call Garbled Table Reduction (GTR), which reduces the size for the garbled tables needed to represent odd 2-to-1 gates by 1/3, and the size of tables of even 2-to-1 gates by 1/2.

5.1 Odd 2-to-1 Gates

Suppose we are implementing an OR-gate, where the external values of $c_1 = 0$ and $c_2 = 0$ correspond to the real input values $(0, 0)$, the other cases will follow immediately from the following. This means that the values $r = 2, 3$ and 4 should evaluate to the same output value k_3^1 , whilst $r = 1$ should evaluate to the output value k_3^0 . We first define over \mathbb{F}_{2^t} a polynomial $P(X)$ of degree two, by interpolating the polynomial which intersects the three points $(2, K_2)$, $(3, K_3)$ and $(4, K_4)$, where each K_r value was defined according to equation (1). (This is the value which in the other constructions was used to mask entry r of the table.) The garbled output value k_3^1 is defined to be $k_3^1 = P(0)$. We also compute $K_5 = P(5)$ and $K_6 = P(6)$. We then define a second polynomial $Q(X)$, also of degree two, by interpolating the polynomial which intersects the three points $(1, K_1)$, $(5, K_5)$ and $(6, K_6)$, where K_1 was defined according to equation (1). The garbled output value k_3^0 is now defined by $k_3^0 = Q(0)$. The garbled table is replaced by the two values (K_5, K_6) . In addition, for each of the four original rows, the external value for the output wire in the r th row is encrypted using the bit M_r , defined in equation (1). The total amount of data sent for the gate is therefore $2t + 4$ bits.

Player P_2 then, given two key values $k_1^{b_1}$ and $k_2^{b_2}$ plus two external values c_1 and c_2 , computes, using equation (1) the value of K_r and M_r for $r = 2c_1 + c_2 + 1$. Recall that the evaluator knows r but not b_1 or b_2 . It then uses the two supplied values of K_5 and K_6 to interpolate the polynomial passing through the points (r, K_r) , $(5, K_5)$ and $(6, K_6)$. The result is either $Q(X)$ or $P(X)$, depending on whether $r = 1$ or not. Player P_2 then recovers the associated secret value $k_3^{b_3}$, by evaluating the polynomial at the point $X = 0$. Using M_r the evaluator can also decrypt the encryption of the external value of the output wire and so obtains c_3 . Hence the evaluator recovers the correct value of the output wire.

5.2 Even 2-to-1 Gates

The only non-trivial even 2-to-1 gates are the XOR and NXOR gate, since all other gates can be replaced by wires. Again let us assume the external input values $c_1 = 0$ and $c_2 = 0$ correspond to the real input values $(0, 0)$, and assume we are dealing with a XOR gate. Then the entries 1 and 4 in the standard garbled table will correspond to the same output key, namely $k_3^{\pi_3^{-1}(0)}$. Any other case will follow from the following description.

Player P_1 first creates a linear polynomial $P(X)$ over \mathbb{F}_{2^t} which interpolates the two points $(1, K_1)$ and $(4, K_4)$. The value of $k_3^{\pi_3^{-1}(0)}$ is defined to be equal to $P(0)$. If the external value of this output value is 0 then we store $P(5)$ into the first row of the new table of this gate, otherwise we store $P(5)$ as the second entry. Then P_1 creates another linear polynomial $Q(X)$ which interpolates the two points $(2, K_2)$ and $(3, K_3)$. The value of $k_3^{\pi_3^{-1}(1)}$ is then defined to be $Q(0)$, and the value $Q(5)$ is stored in the remaining row of our new table. The external values of the output wires are now encrypted and stored, using the M_r values

as before as a separate sub-table of 4 bits in length. Thus, the total amount of data required to represent the gate is $2t + 4$ bits.

Player P_2 given two key values $k_1^{b_1}$ and $k_2^{b_2}$ plus two external values c_1 and c_2 , computes the value of K_r and M_r . Using M_r it can determine the external value of the output wire. If this external value is zero then using the first entry of our garbled table and the value of K_r , the evaluator recovers $P(X)$ and hence $P(0) = k_3^{\pi_3^{-1}(0)}$. If the external value is one then using the second entry of the table and the value K_r , the evaluator recovers $Q(X)$ and hence $Q(0) = k_3^{\pi_3^{-1}(1)}$.

Security: We sketch why the above optimisations maintain security. Given a pair of garbled values of the input wires, P_2 can compute a garbled output value, but cannot distinguish the other garbled output value from random. This is because that other garbled value is defined using a linear combination with a value which is unknown to P_2 . This fact can be used in a, somewhat modified, security proof in the spirit of the proof of Yao’s protocol in [21].

6 Some Experimental Results

We now present some experimental results. In our results we separate out pre-computation time, i.e., generating the required garbled circuits, from the rest of the computation. This is because it depends on the application whether one should consider this time as part of the computation time or not.

There are two major conclusions of our experiments. Firstly, assuming the KDF is correlation robust then the GRR optimisation produces the most efficient implementation. Secondly we conclude that rather large circuits can be practically evaluated using the methods described. Thus secure two-party computation has become more of a reality than one might previously have thought.

Example 1 – Evaluation a Simple Circuit: First we present results for a simple circuit, where we took the circuit for which each of P_1 and P_2 ’s input is a 32-bit integer. The output for P_2 should be the single bit resulting from the application of the comparison operator on the inputs. The output for P_1 will be a six bit integer resulting from the scalar product of the bits of the two inputs, i.e. the number of ones in the string obtained from forming the bit-wise “and” of the two strings.

Applying the Fairplay compiler to this functionality we obtain a circuit with 689 gates. We produce two circuits from this output; the first, denoted $C_{2,3}$, is to allow comparison with the existing state of the art, namely the methods of [22]. This is a circuit which uses 2-to-1 and 3-to-1 gates and has 245 gates. The second circuit we use, denoted C_{xor} , replaces, via the techniques of Section 3, all complex gates with 2-to-1 gates, and tries to minimise the number of non-XOR gates in the circuit. This circuit has 531 gates, 240 of which are non-XOR gates. An extra six gates are needed in each circuit so as to encode P_1 ’s for transmission back to P_1 , without P_2 learning the value.

The above circuit sizes are purely to implement the functionality, they do not include the extra wires and gates required to transmit P_1 's output back to P_1 (for details of how this is done see the full version), nor do they include the extension of the circuit to cope with P_2 's input in the case of Covert and Malicious adversaries. (We refer to the two methods for encoding P_2 's input as the *independent inputs* and the *random combinations* methods. For the details of these methods see [20] or the full version. These methods add a set of XOR gates to the circuit, which transform P_2 's inputs using a random linear encoding.) The sizes of the extended circuits, and the resulting run-times are given in Table 1, which measures the total elapsed wall times in seconds for the various cases.

The calculations were performed on two machines with Intel Core 2 Duo's running at 3.0 GHz, with 4GB of RAM connected by a 1GB ethernet. The hash function $H()$ used in the protocol was implemented as SHA-256.

Table 1. Experimental Results For Example 1 (Times are in seconds)

Adv.	Input Enc.	Method	No. Gates	% XOR Gates	Precomp Time	Send Time	OT Time	Calc Time	Total Time	Total KBytes
Semi-Honest		Base	251	11	0	0	2	0	2	46
		PRF-SS	537	55	0	0	1	0	1	34
		CoR-GRR	537	55	0	0	1	0	1	22
		ROM-GRR	537	55	0	0	1	0	1	22
Covert	Indep. Inputs	Base	419	38	7	1	4	6	18	1188
		PRF-SS	705	61	8	0	2	7	17	969
		CoR-GRR	705	61	6	1	3	5	15	682
		ROM-GRR	705	61	1	1	2	0	4	629
Covert	Random Comb.	Base	1247	79	9	2	4	7	22	2275
		PRF-SS	1535	82	9	1	3	7	20	1646
		CoR-GRR	1555	82	7	1	3	5	16	682
		ROM-GRR	1555	82	1	1	3	0	5	629
Malic.	Indep. Inputs	Base	1571	83	171	80	47	54	352	180599
		PRF-SS	1857	85	175	79	39	67	360	173942
		CoR-GRR	1857	85	147	78	37	39	301	164323
		ROM-GRR	1857	85	141	71	37	38	287	161741
Malic.	Random Comb.	Base	3029	89	163	75	19	64	321	167276
		PRF-SS	2799	90	161	74	16	69	320	158904
		CoR-GRR	2781	90	117	75	16	39	247	140265
		ROM-GRR	2802	90	117	69	16	37	239	137609

The column of “Total KBytes” contains the total number of kilobytes of data which were transferred during the run of the protocol. The column “Method” details the type of computation used, as follows:

- Base: Denotes the optimisations proposed in [22], extended to the case of Covert and Honest adversaries, which we use for comparison purposes, as our baseline implementation. This uses the $C_{2,3}$ circuit mentioned above,

the KDF which is secure in the standard model, and the OT of Hazay-Lindell [15] as opposed to that of Peikert et al. [27].

- PRF-SS: This denotes using the secret sharing based method of Section 5, to reduce the size of the garbled tables. For this the KDF is assumed to be a PRF, but not correlation robust.
- CoR-GRR: This denotes an implementation which is only secure assuming the KDF is correlation robust. It uses the free XOR trick and the method of Garbled Row Reduction, from Section 4, to reduce the size of the remaining garbled tables.
- ROM-GRR: As above for CoR-GRR but all hash functions used are modelled as random oracles. This means we can implement our KDF via a single hash function call, based on the method described in Footnote 5.

The column denoted “No. of gates” describes the number of gates, and the percentage of XOR gates, in the extended circuit (which transfers P_1 ’s outputs and applies the extension described in the full version, encoding P_2 ’s input).

For the Covert and Malicious cases the “Input Enc.” column denotes whether we use the Independent Inputs technique or the Random Combinations technique for the extended circuit construction. See the full version for details. From the table we can deduce the following conclusions:

- The running time in the semi-honest setting is about 10-20 times faster than in the covert setting, which is in turn about 15-20 times faster than in the malicious setting.
- A lot of the extra data needed to be transmitted in the Malicious case is related to the large number of commitments and decommitments which need to be transmitted. Thus our optimisation techniques are less effective in the Malicious case. This points to a clear direction for future research in optimising the Malicious case.
- If one is not willing to assume that the KDF is correlation robust we see that using our technique based on secret sharing can reduce the amount of data being transmitted, compared to the base scheme, without increasing the computational cost.
- In all cases we see that the correlation robust variant using Garbled-Row-Reduction is the most efficient variant. The extra efficiency comes from the free XOR’s which reduce both the number of encryption/decryptions which need to be performed and also the amount of data needing to be transmitted.
- Note that if we assume the random oracle model, and so could implement our KDF via a single hash function call then for Covert adversaries the protocols run significantly faster. That this does not apply as much to the Malicious case is due to the fact that most of the time in the Malicious case is spent with creating, sending and verifying the various commitments.

We pause to compare our two optimisations with the optimisation in bandwidth suggested in [13]. In our system P_1 , the circuit constructor, sends commitments to all circuits that it constructs and to its own inputs, and a random subset of these committed values are checked by P_2 . In [13] it is suggested that

P_1 commits to a random seed, and uses this to generate the circuit. Then only the commitment to this seed, and eventually its decommitment, need to be transmitted. This means that P_2 needs to compute the circuit given the seed. Whilst this optimisation clearly significantly reduces the consumed bandwidth, it actually leads to a significant increase in the time needed to perform the protocol. To see this consider our Covert experiments in Table 1. The optimisation in [13] would reduce practically to zero, the entry for the “Send Time” column, but P_2 would now need to recompute almost all of the calculations in the “Precomp Time” column. Thus the technique of [13] is only to be compared to ours in the situation where bandwidth is very expensive and CPU time is very cheap.

Before passing onto our larger example we note the following. If we let p denote the proportion of XOR gates within a circuit, and we let N denote the amount of data needed to be sent per circuit in the standard Yao construction, then the average amount of data needed to be sent per circuit gate when using the free XOR gates and GRR methods is $3/4 \cdot (1 - p) \cdot N$. Whereas if we do not use the free XOR gate method and instead use the method based on secret sharing, this value becomes $N/2$. Hence, if we are willing to assume correlation robust KDFs, then the method which uses secret sharing and does not use the free XOR method, will be more efficient as long as the fraction of XOR gates, p , is smaller than $1/3$. However as can be seen from the column entitled “% XOR Gates”, this proportion is generally much larger than $1/3$, especially in the case of Covert and Malicious adversaries where we have had to extend the circuit by a large linear component. This expansion is performed to cope with possible adversarial behaviour related to P_2 's input, see the full version for details. One should note that these theoretical estimates of bandwidth are never achieved fully in practice due to overheads in the underlying data transmission mechanism and the fact that they assume a bit-oriented communication mechanism, whereas practical communication is performed in bytes. Hence the saving we achieve in gate transmission is about 5-10% less than one would predict purely by theory.

Example 2 - Evaluating AES: As our second example we created a circuit which computes an AES encryption of a single 128-bit block with respect to a 128-bit key. Here P_1 's input is the secret key, and P_2 's input is the message block. We require that P_2 learns the encryption of its message under P_1 's secret key, and that P_1 learns nothing. Compiling such a circuit using the Fairplay compiler, and applying various optimisations, resulted in a circuit, which we denote by $C_2^{(1)}$, with 33880 gates, where each gate is a 2-to-1 gate. This circuit was derived in a way to try to minimize the number of non-XOR gates. Again, we stress, the above circuit size purely implements the AES functionality, it does not include the extension of the circuit to cope with P_2 's input in the case of Covert and Malicious adversaries. Note that the key schedule takes up only about 15% of the circuit, hence encrypting a sequence of message blocks as in CBC-Mode encryption will scale almost linearly with respect to our data.

We repeated our experiments from above, but in Table 2 we only present the times for the most efficient choice for the input encoding.

Table 2. Experimental Results for Example 2 (Again times are in seconds)

Adv.	Input Enc.	Method	No. Gates	% XOR Gates	Precomp Time	Send Time	OT Time	Calc Time	Total Time	Total KBytes
Semi-Honest		Base	28216	56	5	2	4	3	14	3162
		PRF-SS	33880	66	5	1	3	3	12	1752
		CoR-GRR	33880	66	2	1	2	2	7	503
		ROM-GRR	33880	66	1	1	3	2	7	503
Covert	Indep. Inputs	Base	28600	56	96	47	18	45	206	51899
		PRF-SS	34264	67	92	36	13	50	191	29380
		CoR-GRR	34264	67	40	21	11	23	95	9078
		ROM-GRR	34264	67	22	21	11	6	60	8942
Malic.	Random Comb.	Base	40253	69	1250	448	39	887	2624	987442
		PRF-SS	45944	75	1184	392	34	829	2439	711729
		CoR-GRR	45960	75	483	270	34	361	1148	406010
		ROM-GRR	45881	75	453	276	35	350	1114	417907

We conclude that performing the Yao protocol is certainly feasible on complicated functionalities such as AES encryption. For the case of honest and covert adversaries we again see that the computation and bandwidth consumed, when we use correlation robust KDFs and the GRR method, greatly reduces in comparison to the base case. If one is not willing to assume correlation robust KDFs (or use the ROM) then our secret sharing based optimisation greatly reduces the bandwidth without affecting the run times. For the malicious case the improvement in the secret sharing based version is less pronounced due to the large number of commitments which need to be transmitted and opened. This clearly points to the place where future optimisation research needs to be performed, namely in reducing the number of commitments needed in the situation of malicious adversaries. However even without such future optimisation we note that performance can be significantly reduced by taking advantage of the inherent parallelism in the algorithm in the Malicious case (in which P_1 generates many commitments and P_2 verifies a subset of them). For web service or cloud computing applications, where server farms are common place, an improvement in computational time by a factor around s_1 could be expected.

We end by noting that many application domains of a secure evaluation of AES, for example the one-time program example from [12], require only security against semi-honest adversaries. Hence, such applications are already within the reach of practical realisation. Furthermore, this application requires no computation of the OT or data to be sent. Thus the party generating the one-time-program will take the time needed in our *Precomp Time* column, and the evaluator (after querying the one-time-memory) will take the time needed in the *Calc Time* column.

7 Acknowledgments

All authors would like to thank the EU projects CACE and eCrypt-2 for partially funding the work in this paper. The first author was also partially funded by the ERC project SFEROT. The third author was partially funded by a Royal Society Wolfson Merit Award, and the fourth author was partially funded by EPSRC. We would also like to thank Benny Applebaum, Yehuda Lindell and Ahmad-Reza Sadeghi for various comments and discussions.

References

1. Y. Aumann and Y. Lindell. Security against covert adversaries: Efficient protocols for realistic adversaries. In *Theory of Cryptography Conference – TCC 2007*, Springer-Verlag LNCS 4392, 137–156, 2007.
2. D. Beaver, S. Micali and P. Rogaway. The round complexity of secure protocols. In *22nd STOC*, 503–513, 1990.
3. A. Ben-David, N. Nisan and B. Pinkas. FairplayMP: a system for secure multiparty computation. In *Computer and Communications Security – CCS 2008*, ACM, 257–266, 2008.
4. M. Ben-Or, S. Goldwasser and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *20th STOC*, 1–10, 1988.
5. P. Bogetoft, D.L. Christensen, I. Dãmgard, M. Geisler, T. Jakobsen, M. Krøigaard, J.D. Nielsen, J.B. Nielsen, K. Nielsen, J. Pagter, M. Schwartzbach and T. Toft. Secure multiparty computation goes live, To appear in *Financial Cryptography and Data Security – FC 2009*.
6. P. Bogetoft, I. Dãmgård, T. Jakobsen, K. Nielsen, J. Pagter. and T. Toft. A practical implementation of secure auctions based on multiparty integer computation. In *Financial Cryptography and Data Security – FC 2006*, Springer-Verlag LNCS 4107, 142–147, 2006.
7. R. Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, **13**(1), 143–202, 2000.
8. D. Chaum, C. Crepeau and I Dãmgård. Multiparty unconditionally secure protocols. In *20th STOC*, 11–19, 1988.
9. M.J. Freedman, Y. Ishai, B. Pinkas and O. Reingold. Keyword search and oblivious pseudorandom functions. In *Theory of Cryptography Conference – TCC 2005*, Springer-Verlag LNCS 3378, 303–324, 2005.
10. O. Goldreich. *Foundations of Cryptography: Volume 2 – Basic Applications*. Cambridge Univ. Press, 2004.
11. O. Goldreich, S. Micali and A. Wigderson. How to play any mental game – A completeness theorem for protocols with honest majority. In *19th STOC*, 218–229, 1987.
12. S. Goldwasser, Y.T. Kalai and G.N. Rothblum. One-time programs. In *Advances in Cryptology – Crypto 2008*, Springer-Verlag LNCS 5157, 39–56, 2008.
13. V. Goyal, P. Mohassel and A. Smith. Efficient two party and multi-party computation against covert adversaries. In *Advances in Cryptology – EuroCrypt 2008*, Springer-Verlag LNCS 4965, 289–306, 2008.
14. D. Harnik, Y. Ishai, E. Kushilevitz and J.B. Nielsen. OT-Combiners via secure computation. In *Theory of Cryptography Conference – TCC 2008*, Springer-Verlag LNCS 4948, 393–411, 2008.

15. C. Hazay and Y. Lindell. Oblivious transfer, polynomial evaluation and set intersection. Manuscript, 2008
16. Y. Ishai, J. Kilian, K. Nissim and E. Petrank. Extending oblivious transfers efficiently. In *Advances in Cryptology – Crypto 2003*, Springer-Verlag LNCS 2729, 145–161, 2003.
17. Y. Ishai, M. Prabhakaran and A. Sahai. Founding cryptography on oblivious transfer - efficiently. In *Advances in Cryptology – Crypto 2008*, Springer-Verlag LNCS 5157, 572–591, 2008.
18. S. Jarecki and V. Shmatikov. Efficient two-party secure computation on committed inputs. In *Advances in Cryptology – EuroCrypt 2007*, Springer-Verlag LNCS 4515, 97–114, 2007.
19. V. Kolesnikov and T. Schneider. Improved garbled circuit: Free XOR gates and applications. In *Automata, Languages and Programming – ICALP 2008*, Springer-Verlag LNCS 5126, 486–498, 2008.
20. Y. Lindell and B. Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In *Advances in Cryptology – EuroCrypt 2007*, Springer-Verlag LNCS 4515, 52–78, 2007.
21. Y. Lindell and B. Pinkas. A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology*, **22**, 161–188, 2009
22. Y. Lindell, B. Pinkas, N.P. Smart. Implementing two-party computation efficiently with security against malicious adversaries. In *Security and Cryptography for Networks – SCN 2008*, Springer-Verlag LNCS 5229, 2–20, 2008.
23. P. Mohassel and M.K. Franklin. Efficiency tradeoffs for malicious two-party computation. In *Public Key Cryptography – PKC 2006*, Springer-Verlag LNCS 3958, 458–473, 2006.
24. D. Malkhi, N. Nisan, B. Pinkas and Y. Sella. Fairplay — a secure two-party computation system. In *Proc. of 13th USENIX Security Symposium*, 2004.
25. M. Naor, B. Pinkas and R. Sumner. Privacy Preserving Auctions and Mechanism Design. In Proc. of the 1st ACM conf. on Electronic Commerce, November 1999.
26. J.B. Nielsen and C. Orlandi. LEGO for two party secure computation. In *Theory of Cryptography Conference – TCC 2009*, LNCS 5444, 368-386, 2009.
27. C. Peikert, V. Vaikuntanathan and B. Waters. A framework for efficient and composable oblivious transfer. In *Advances in Cryptology – Crypto 2008*, Springer-Verlag LNCS 5157, 554–571, 2008.
28. A. Shamir. How to share a secret. *Communications of the ACM*, **11**, 612–613, 1979.
29. D. Woodruff. Revisiting the efficiency of malicious two-party computation. In *Adv. in Cryptology – EuroCrypt 2007*, Springer-Verlag LNCS 4515, 79–96, 2007.
30. A. Yao. How to generate and exchange secrets. In *27th FOCS*, 162–167, 1986.