

Whitewash: Outsourcing Garbled Circuit Generation for Mobile Devices

Henry Carter
Georgia Institute of
Technology
carterh@gatech.edu

Charles Lever
Georgia Institute of
Technology
chazlever@gatech.edu

Patrick Traynor
University of
Florida
traynor@cise.ufl.edu

ABSTRACT

Garbled circuits offer a powerful primitive for computation on a user's personal data while keeping that data private. Despite recent improvements, constructing and evaluating circuits of any useful size remains expensive on the limited hardware resources of a smartphone, the primary computational device available to most users around the world. In this work, we develop a new technique for securely outsourcing the generation of garbled circuits to a Cloud provider. By outsourcing the circuit generation, we are able to eliminate the most costly operations from the mobile device, including oblivious transfers. Our proofs of security show that this technique provides the best security guarantees of any existing garbled circuit outsourcing protocol. We also experimentally demonstrate that our new protocol, on average, decreases execution time by 75% and reduces network costs by 60% compared to previous outsourcing protocols. In so doing, we demonstrate that the use of garbled circuits on mobile devices can be made nearly as practical as it is becoming for server-class machines.

1. INTRODUCTION

Mobile devices have become one of the dominant computing platforms, with approximately 57% market penetration in the United States alone [9]. These devices are capable of gathering and storing all of a user's personal data, from current location and social contacts to banking and electronic payment information. Because of the personal nature of these devices, it is critical that a user's information be protected at all times. Unfortunately, many smartphone applications that require users to send data to application servers make preserving the privacy of this data difficult.

To resolve this issue, a variety of secure multiparty computation techniques exist that could be leveraged to perform computation over encrypted inputs [6, 11, 12, 28]. Currently, the most practically efficient two-party technique is the Yao Garbled Circuit [43]. Despite recent improvements in the efficiency of garbled circuits [27, 41], this technique still requires significant computation and communication resources, rendering it impractical for most smartphones. One possible solution to this imbalance of resources is to blindly outsource the heavy computation to the Cloud.

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.

ACM '14 December 08 - 12 2014, New Orleans, LA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3005-3/14/12..\$15.00

<http://dx.doi.org/10.1145/2664243.2664255>

However, because of the untrusted nature of Cloud providers [42], such a solution fails to provide measurable guarantees for applications requiring high assurance.

In this work, we develop a new protocol for securely outsourcing garbled circuit generation to an untrusted Cloud. We construct a protocol that offloads the role of generating the garbled circuit from the mobile device to the Cloud without exposing any private inputs or outputs. By choosing to outsource this portion of the protocol, we eliminate a significant number of expensive public-key cryptography operations and rounds of communication used in oblivious transfer. The result is a more computationally and bandwidth efficient outsourcing protocol with the strongest security guarantees of any outsourcing technique to date [8, 23].

In this paper, we make the following contributions:

- **Develop a new outsourcing protocol:** We develop the Whitewash¹ outsourcing protocol, which allows a mobile device participating in a two-party secure function evaluation to outsource the generation of the garbled circuit. Our protocol assigns the mobile device the role of circuit generator instead of circuit evaluator, outsourcing a completely different set of operations from previous outsourcing protocols [8, 23]. By reversing the functions of the two players, we fully eliminate the requirement for any oblivious transfers, outsourced or otherwise, to or from the mobile device. This "simple" role reversal requires fundamentally redesigning the outsourcing techniques used in previous work, as well as new security proof formulations.
- **Formal verification and analysis:** We formally prove the security of our outsourcing techniques in the malicious model defined by Kamara et al. [23]. Unlike previous work [8, 23], our protocol provides security when the mobile device is colluding with its Cloud provider against the application server. *This added security guarantee makes our protocol the most secure outsourcing protocol to date.* We then provide an analysis of the reduction in operations between our work and the outsourced oblivious transfer of Carter et al. [8], as well as the Salus framework by Kamara et al. [23]. Specifically, our protocol requires more executions of a pseudorandom number generator in exchange for fewer algebraic group operations and zero-knowledge proofs. Moreover, we significantly reduce the number of rounds of communication required to the mobile device.
- **Implement and evaluate the performance of our protocol:** In our performance evaluation, we demonstrate a maximum improvement of 98% in execution time and 92% improvement in bandwidth overhead compared to Carter et al. [8] (with 75% and 60% average improvement, respectively). For a different

¹A reference to Tom Sawyer, who "outsourced" his chores to his friends without ever revealing the true nature of the work.

test application, when compared to performing computation directly on the mobile device [28], we demonstrated a 96% and 90% improvement in execution time and bandwidth, respectively. These improvements allow for the largest circuits evaluated on any platform to be computed from a mobile device efficiently and with equivalent security parameters to non-mobile protocols.

The rest of this work is organized as follows: Section 2 provides detail on related research; Section 3 describes our threat model and security definition; Section 4 provides a description of the White-wash protocol; Section 5 compares the operations required in our protocol to the protocols by Carter et al. and Kamara et al.; Section 6 describes our empirical performance analysis; and Section 7 provides concluding remarks.

2. RELATED WORK

Fairplay [35] provided the first practically efficient implementation of Yao’s garbled circuit protocol [43], requiring only simple hash and symmetric key operations to securely evaluate an arbitrary function. Since then, a variety of garbled circuit-based secure function evaluation (SFE) protocols have been developed in the semi-honest adversarial model [5, 16, 20, 21, 29, 31, 34, 39]. The latest of these, developed by Huang et al. [16], allows garbled circuits to be evaluated in stages, which makes it the most efficient semi-honest garbled circuit evaluation technique, both in computation and memory requirements. In recent work, several garbled circuit SFE protocols have been developed in the malicious security model, which require significantly more computational resources than semi-honest protocols, but are secure against arbitrary polynomial-time adversaries [18, 26, 28, 30, 33, 37, 40]. The protocol by Shelat and Shen [41] provides a two-party garbled circuit protocol which uses only symmetric-key constructions outside of the oblivious transfer. When combined with Huang’s pipelining approach and the PCF compiler by Kreuter et al. [27], their protocol is among the most efficient maliciously-secure garbled circuit protocols implemented to date. Some efforts have been made to improve the efficiency of these protocols by slightly reducing the adversary model. Many schemes have been developed in the covert adversary model, which allows for some efficiency gains at the cost of security [2, 10, 15, 36]. Huang et al. [17] developed a protocol that leaks only one bit of input to a malicious adversary through dual execution, which was later implemented on GPUs by Husted et al. [19]. In order to further improve the efficiency of garbled circuit protocols, Gordon et al. [13] developed a protocol that combined Oblivious RAM with garbled circuits, allowing sub-linear amortized complexity. However, this protocol only allows this performance gain for functions that can be computed efficiently on a random-access machine. Other work has focused on making circuit compilation and garbling more efficient [4, 38]. These techniques improve the underlying operations found in all garbled circuit execution protocols, providing improvement in all existing techniques.

To further improve the speed of cryptographic protocols on devices with minimal computational resources, the idea of outsourcing cryptographic operations has been explored for many years in the field of Server-assisted cryptography [3]. More recently, Green et al. [14] developed a technique for outsourcing the costly decryption of attribute-based encryption schemes to the cloud without revealing the contents of the ciphertext. Atallah and Frikken [1] developed a set of special-purpose protocols for securely outsourcing Linear Algebra computations to a single cloud server. For data mining applications, Kerschbaum recently developed an outsourced set intersection protocol using homomorphic encryption techniques [24]. While all of these applications provide significant

performance gains for specific cryptographic applications, none of them address outsourcing of general secure computation.

In their Salus protocol, Kamara et al. [22, 23] developed two protocols for securely outsourcing the computation of arbitrary functions to the cloud. Following Salus, Carter et al. [8] developed an outsourcing protocol based on the maliciously secure garbled circuit protocol by Kreuter et al. [28]. Carter’s protocol outsources the evaluation of garbled circuits by adding in an Outsourced Oblivious Transfer primitive. Their participant configuration is the same configuration found in Kamara’s maliciously secure protocol, where the cloud is made responsible for evaluating the garbled circuit. In this work, we choose to build on Shelat and Shen’s latest protocol [41] since the symmetric execution environment of Huang et al. [18] does not lend itself to outsourcing, and the bootstrapping technique used by Lindell [30] has not been implemented or evaluated in practice. Unlike previous work, we choose to fundamentally rearrange the roles of the participants, outsourcing the generation of the garbled circuits as in Kamara’s covertly secure protocol.

3. OVERVIEW AND DEFINITIONS

3.1 Protocol Goals and Summary

The primary reason for developing an outsourcing protocol for secure function evaluation is to allow two parties of asymmetrical computing ability to securely compute some result. Current two-party computation protocols assume both parties are equipped with equivalent computing resources and so require both parties to perform comparable operations. However, when a mobile device is taking part in computation with an application server, some technique is necessary to reduce the complexity of the operation on the mobile device. Ideally, we can make the mobile device perform some small number of operations that is independent of the size of the circuit being evaluated.

In constructing such a protocol, there are four goals that we would like to satisfy. The first of these goals is correctness. It is necessary that an outsourcing protocol must produce correct output even in the face of malicious players attempting to corrupt the computation. The second desirable guarantee is security. SFE protocols frequently use a simulation-based approach to defining and proving security, which we outline in detail below. Essentially, the goal is to show that each party can learn the output of the computed function and nothing else. Third, an ideal protocol would provide some guarantee of fair release. This guarantee ensures that either both parties receive their outputs from the computation, or neither party receives their output. Our protocol achieves this in all but one corruption scenario by treating the Cloud as an arbiter, who will simultaneously and fairly release the outputs of the protocol using one-time pads. In the scenario where the mobile device and Cloud are colluding, it is possible for the Cloud to terminate the protocol after the mobile device receives output but before the application server receives output. However, this is inherently possible in most two-party garbled circuit protocols. The fourth goal of our protocol is efficiency. Our outsourcing protocol balances a minimal set of operations on the mobile device with efficient multiparty computation operations on the application server and Cloud.

Given these goals, we build our protocol in the two-party server-aided multiparty computation setting. This setting is composed of two parties, the mobile device and an application server, who wish to run a two-party secure computation while keeping both of their inputs private. To assist the mobile device, the server-aided setting adds a third party Cloud provider, which is independent and non-colluding with the application server (more on non-collusion in the following section). The Cloud performs cryptographic operations

for the mobile device, but is not allowed to see any party’s input or output from the computation.

To achieve our goals in this setting, we first select the most efficient two-party garbled circuit computation protocol to date that provides guarantees of correctness and security in the malicious model. We assign the mobile device the role of circuit generator in this protocol, and the application server is assigned the role of circuit evaluator. To outsource the circuit generation operations from the mobile device, we allow the device to generate short random seeds and pass these values to a Cloud computation provider, which then generates the garbled circuits using these seeds to generate randomness. Thus, the mobile device’s work is essentially reduced to (1) generating random strings on the order of a statistical security parameter, and (2) garbling and sending its input values to the evaluating party. In this way, we develop a secure computation protocol where the mobile device performs work that is independent of the size of the function being evaluated. We provide a formal proof of security for our protocol in our technical report [7].

3.2 Non-collusion

To maintain security, previous outsourcing protocols assume that neither party colludes with the Cloud [8, 23]. The theoretical intuition for this constraint, outlined by Kamara et al. [23], is that the existence of an outsourcing protocol where parties can arbitrarily collude would imply a two-party secure multiparty computation protocol where one party performs sub-linear work with respect to the size of the circuit. While this has been shown to be possible using fully homomorphic encryption and, in some cases, oblivious RAM [13], it is not clear that these techniques can be efficiently applied to a garbled circuit outsourcing scheme. Because of this, previous work has left the more complex security model for future study. However, while previous protocols restrict collusion between the Cloud and any party, the sub-linear work implication only applies to cases when the Cloud is generating circuits and colludes with the evaluating party, or vice versa. In the Whitewash protocol, we prove security when the mobile device colludes with the Cloud against the evaluating web application. While this collusion scenario removes the fair release guarantee of our protocol, it in no way compromises the security guarantees of confidentiality of participant’s inputs and outputs. Essentially, it reduces to the two-party computation scenario that the underlying protocol is proven to be secure in. Since the mobile device is paying the Cloud for computation services, we believe it is a more realistic assumption to assume that a Cloud provider could collude maliciously with the paying customer, and note that our protocol is the first outsourcing protocol to provide any security guarantees in the face of collusion with the Cloud.

3.3 Security Constructions

In the two-party computation protocol underlying our work, shelat and Shen implement a number of new and efficient cryptographic checks to ensure that none of the parties participating in the computation behave maliciously. We provide an overview of these security checks in the following section. We refer the reader to shelat and Shen’s work for more formal definitions and proofs [41].

3.3.1 k -probe-resistant input encoding

When working with garbled circuit protocols in the malicious model, the generator has the ability to learn information about the evaluator’s input by corrupting the wire labels sent during the oblivious transfer. This attack, known as selective failure, was first proposed by Mohassel and Franklin [37] as well as Kiraz and Schoenmakers [25]. In the server-aided setting, it is possible that the mo-

bile device and the Cloud could collude and carry out this attack to recover the application server’s input. To prevent this attack, shelat and Shen [41] implement an improved version of the k -probe-resistant input encoding mechanism originally proposed by Lindell and Pinkas [32]. In their protocol, the evaluator does not input her real input y to the computation, but chooses her input \bar{y} such that $M \cdot \bar{y} = y$ for a k -probe resistant matrix M . Intuitively, the idea is that the generator would have to probe the evaluator’s input approximately 2^k times before learning anything about her input y .

3.3.2 2-Universal Hash Function

A second concern with garbled circuits in the malicious model is that the generator may send different input values for each of the evaluated circuits from the cut-&-choose. As in the two-party setting, it is possible for the mobile device to submit inconsistent inputs to the application server in the server-aided setting. To ensure that the generator’s inputs are consistent across evaluation circuits, shelat and Shen implement an efficient witness-indistinguishable proof, which computes a randomized, 2-universal hash of the input value using only arithmetic operations on matrices. Because of the regularity guarantees of a 2-universal hash, the outputs of these hash operations can be seen by the evaluator without revealing any information about the generator’s inputs. However, if any of the hashed input values is inconsistent across evaluation circuits, the evaluator can infer that the generator provided inconsistent inputs, and can terminate the protocol.

3.3.3 Output proof of consistency

When a function being evaluated using garbled circuits has separate, private outputs for the generating and evaluating parties, it is necessary to ensure that the evaluating party does not tamper with the generating party’s output. Since the output must be decoded from the garbled output wires for the majority check at the end of the protocol, if the output is only blinded with a one-time pad, this allows the evaluator the opportunity to change bits of the generator’s output. Our setting faces the same potential for attack from the application server, who is responsible for evaluating the circuit and distributing the blinded output. Several techniques for preventing this kind of tampering have been proposed, but shelat and Shen’s latest protocol [41] implements a witness-indistinguishable proof that uses only symmetric key cryptographic operations. After the evaluator sends the blinded output of computation to the generator, the proof guarantees to the generator that the output value he received was actually generated by one of the garbled circuits he generated. However, it keeps the index of the circuit that produced the output hidden, as this could leak information to the generator.

3.4 Security Model and Definition

Our definition of security is based on the definition proposed by Kamara et al. [23], which we specify for the two-party setting as in Carter et al. [8]. We provide a brief description of the real/ideal world model here and direct readers to the previous work in this space for a more formal definition.

In the real world, both parties participating in the computation (the mobile device and application server) provide an input to the computation and an auxiliary input of random coins, while the single party designated as the outsourcing party (the Cloud) provides only random auxiliary input. The party evaluating the circuit in this computation is assumed to be non-colluding with the outsourcing party, as defined by Kamara et al. Some subset of these parties $A = (A_1, A_2, A_3)$ are corrupted and can deviate arbitrarily from the protocol. For the i^{th} honest party, let OUT_i be its output, and for the i^{th} corrupted party, let OUT_i be its view of the protocol

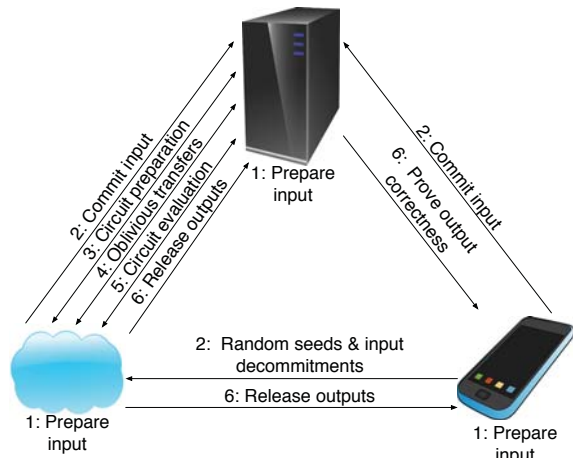


Figure 1: The complete Whitewash protocol. Note that MOBILE performs very little work compared to SERVER and CLOUD.

execution. Then the i^{th} partial output of a real protocol execution with input x is defined as:

$$REAL^{(i)}(k, x; r) = \{OUT_j : j \in H\} \cup OUT_i$$

Where H is the set of honest parties, r is all random coins of all participants, and k is the security parameter.

In the ideal world, each party provides the same inputs as in the real world, however, they are sent to a trusted oracle which performs the secure computation. Once the trusted oracle completes the computation, it returns the output to the participating parties and no output to the outsourcing party. If any party aborts early or sends no input to the oracle, the oracle aborts and does not send the output to any party. For the i^{th} honest party, let OUT_i be its output to the oracle, and for the i^{th} corrupted party, let OUT_i be an arbitrary output value produced by the party. Then the i^{th} partial output of an ideal protocol execution in the presence of independent malicious simulators $S = (S_1, S_2, S_3)$ is defined as:

$$IDEAL^{(i)}(k, x; r) = \{OUT_j : j \in H\} \cup OUT_i$$

Where H , r , and k are defined as before. Given this model, security is formally defined as:

DEFINITION 1. *An outsourcing protocol securely computes the function f if there exists a set of probabilistic polynomial-time (PPT) simulators $\{Sim_i\}_{i \in [3]}$ such that for all PPT adversaries (A_1, A_2, A_3) , inputs x , auxiliary inputs z , and for all $i \in [3]$:*

$$\{REAL^{(i)}(k, x; r)\}_{k \in N} \stackrel{c}{\approx} \{IDEAL^{(i)}(k, x; r)\}_{k \in N}$$

Where $S = (S_1, S_2, S_3)$, $S_i = Sim_i(A_i)$, and r is uniformly random.

4. PROTOCOL

4.1 Participants

Given a mobile device and a web or application server who wish to jointly compute a function, there are three participating parties in the Whitewash protocol:

- **SERVER:** “SERVER” refers to the web or application server participating in the joint computation. She is assumed to have large computational resources and is responsible for evaluating the garbled circuits.

- **MOBILE:** “MOBILE” refers to the mobile device participating in the joint computation. He is assumed to have limited processing power, memory, and communication bandwidth. MOBILE is tasked with garbling the circuit to be evaluated by SERVER.
- **CLOUD:** The outsourcing party “CLOUD” is responsible for relieving MOBILE of the majority of his computational load, but is not trusted with knowing either party’s input to or output from the joint computation.

4.2 Protocol

Common Inputs: Security parameters k (key length) and σ (the number of circuits generated for the cut-&-choose); a commitment scheme $com(x; c)$ with committed value x and commitment key c ; and a function $f(x, y)$.

Private Inputs: MOBILE inputs x and SERVER inputs y .

Outputs: Two outputs f_s, f_m for SERVER and MOBILE, respectively.

Phase 1: Pre-computation

1. **Preparing inputs:** MOBILE randomly generates $r \in \{0, 1\}^{2k + \log(k)}$ as his input to the 2-universal circuit. He also generates $e \in \{0, 1\}^{|f_m|}$ as a one-time pad for his output. SERVER computes her k -probe-resistant matrix M and \bar{y} such that $M \cdot \bar{y} = y$. MOBILE’S input to the circuit will be $\bar{x} = x \| e \| r$ and SERVER’S input will be \bar{y} . We denote the set of indices $[m_s] = \{1, \dots, |\bar{y}|\}$ and $[m_m] = \{1, \dots, |\bar{x}|\}$.
2. **Preparing circuit randomness:** MOBILE generates random seeds $\{\rho^{(j)}\}_{j \in [\sigma]}$ for generating the circuits and sends them to CLOUD.

Phase 2: Input commitments

1. **Committing to MOBILE’S inputs:** For each circuit $j \in [\sigma]$, input bit $i \in [m_m]$, and $b \in \{0, 1\}$ MOBILE uses $\rho^{(j)}$ to generate commitment keys $\theta_{i,b}^{(j)}$. Using the same random seeds, these keys will later be generated by CLOUD to commit to the input wire labels corresponding to MOBILE’S input. MOBILE then commits to his own inputs as $\{\Gamma^{(j)}\}_{j \in [\sigma]}$ as:

$$\Gamma^{(j)} = \{com(\theta_{i,\bar{x}_i}^{(j)}; \gamma_i^{(j)})\}_{i \in [m_m]}$$

using independently generated random commitment keys $\gamma_i^{(j)}$. MOBILE sends $\{\Gamma^{(j)}\}_{j \in [\sigma]}$ to SERVER and the commitment keys $\{\gamma_i^{(j)}\}_{i \in [m_m], j \in [\sigma]}$ to CLOUD.

2. **Committing to CLOUD’S inputs:** To allow for a fair release of the outputs, CLOUD inputs one-time pads to blind both parties’ outputs. CLOUD randomly generates $p_s \in \{0, 1\}^{|f_s|}$ and $p_m \in \{0, 1\}^{|f_m|}$, as well as $r_c \in \{0, 1\}^{2k + \log(k)}$ as its input to the 2-universal circuit. We denote CLOUD’S input as $z = p_s \| p_m \| r_c$, and the indices of CLOUD’S input wires as $[m_c] = \{1, \dots, |z|\}$.

For each circuit $j \in [\sigma]$ and input bit $i \in [m_c]$, CLOUD uses $\{\rho^{(j)}\}_{j \in [\sigma]}$ to generate the garbled input wire keys $(K_{i,0}^{(j)}, K_{i,1}^{(j)}, \pi_i^{(j)})$, where $K_{i,0}^{(j)}, K_{i,1}^{(j)} \in \{0, 1\}^k$ and the permutation bit $\pi_i^{(j)} \in \{0, 1\}$. To locate the correct key for bit b on input wire w_i of circuit j , we designate the label $W_{i,b}^{(j)} = (K_{i,b}^{(j)}, b \oplus \pi_i^{(j)})$.

Let $\{w_{m_s+i}\}_{i \in [m_c]}$ be the input wires for CLOUD. CLOUD then commits to the label pairs for its input wires as $\{\Psi^{(j)}\}_{j \in [\sigma]}$,

where

$$\Psi^{(j)} = \{com(W_{m_s+i,0\oplus\pi_i}^{(j)}; \psi_{i,0\oplus\pi_i}^{(j)}), \\ com(W_{m_s+i,1\oplus\pi_i}^{(j)}; \psi_{i,1\oplus\pi_i}^{(j)})\}_{i \in [m_c]}$$

using commitment keys $\psi_{i,b}^{(j)}$ generated with the random seed $\rho^{(j)}$. CLOUD then commits to its inputs as $\{\Xi^{(j)}\}_{j \in [\sigma]}$ as:

$$\Xi^{(j)} = \{com(\psi_{i,z_i}^{(j)}; \xi_i^{(j)})\}_{i \in [m_c]}$$

using independently generated random commitment keys. CLOUD sends $\{\Psi^{(j)}\}_{j \in [\sigma]}$ and $\{\Xi^{(j)}\}_{j \in [\sigma]}$ to SERVER.

Phase 3: Circuit construction

1. **Constructing the objective circuit:** SERVER sends \mathbf{M} to CLOUD, then SERVER and CLOUD run a coin flipping protocol to randomly determine the 2-universal hash matrix $\mathbf{H} \in \{0, 1\}^{k \times m_m}$. These two matrices can be used to generate the new circuit C that computes the function $g : (\bar{x}, \bar{y}) \rightarrow (\perp, (h_m, h_c, c_s, c_m))$, where $h_m = \mathbf{H} \cdot \bar{x}$, $h_c = \mathbf{H} \cdot z$, $g_m = f_s(x, \mathbf{M} \cdot \bar{y})$, $c_m = g_m \oplus e \oplus p_m$, $g_s = f_s(x, \mathbf{M} \cdot \bar{y})$, and $c_s = g_s \oplus p_s$. MOBILE will need the values $h_c || c_m$ to recover his output. We denote the set of indices corresponding to these values as $O_m = \{1, \dots, |h_c| + |c_m|\}$.
2. **Committing to input and output wire label pairs:** Using the same method as in Phase 2, CLOUD uses $\{\rho^{(j)}\}_{j \in [\sigma]}$ to generate the input wire keys for both SERVER and MOBILE'S input as well as the output wire keys for MOBILE'S output (these output keys must be committed for the witness indistinguishable proof of MOBILE'S output correctness). Let $\{w_i\}_{i \in [m_m]}$ be the input wires for MOBILE, $\{w_{m_m+i}\}_{i \in [m_s]}$ be the input wires for SERVER, and $\{w_i\}_{i \in O_m}$. CLOUD then commits to the label pairs in MOBILE'S input, SERVER'S input, and MOBILE'S output as $\{\Theta^{(j)}, \Omega^{(j)}, \Phi^{(j)}\}_{j \in [\sigma]}$, where

$$\Theta^{(j)} = \{com(W_{i,0\oplus\pi_i}^{(j)}; \theta_{i,0\oplus\pi_i}^{(j)}), \\ com(W_{i,1\oplus\pi_i}^{(j)}; \theta_{i,1\oplus\pi_i}^{(j)})\}_{i \in [m_m]}$$

$$\Omega^{(j)} = \{com(W_{m_m+i,0}^{(j)}; \omega_i^{(j)}), com(W_{m_m+i,1}^{(j)}; \omega_i^{(j)})\}_{i \in [m_s]}$$

$$\Phi^{(j)} = \{com(W_{i,0}^{(j)}; \phi_i^{(j)}), com(W_{i,1}^{(j)}; \phi_i^{(j)})\}_{i \in O_m}$$

using commitment keys generated with the random seed $\rho^{(j)}$. CLOUD then sends these commitments to SERVER.

Phase 4: Oblivious transfers (OT)

1. **Oblivious transfers:** CLOUD and SERVER execute m_s input oblivious transfers and σ circuit oblivious transfers as follows:
 - (a) **Input:** For each $i \in [m_s]$, both parties run a $\binom{2}{1}$ -OT where CLOUD inputs

$$\left(\{(W_{m_m+i,0}^{(j)}; \omega_i^{(j)})\}_{j \in [\sigma]}, \{(W_{m_m+i,1}^{(j)}; \omega_i^{(j)})\}_{j \in [\sigma]} \right)$$

while SERVER inputs \bar{y}_i . Once SERVER receives all of her garbled input wire labels, she uses the decommitment keys obtained in the OTs to check the committed wire values in $\{\Omega^{(j)}\}_{j \in [\sigma]}$. If any of the labels received in the OT do not match the committed wire labels, SERVER terminates the protocol.

- (b) **Circuit:** SERVER selects a set of circuits to be evaluated $S \subset [\sigma]$ such that $|S| = \frac{2\sigma}{5}$, as in shelat and Shen's protocol [40]. She represents this set with a bit string $s \in \{0, 1\}^\sigma$ such that the j^{th} bit $s_j = 1$ if $j \in S$ and $s_j = 0$ otherwise. SERVER and CLOUD perform $\sigma \binom{2}{1}$ -OTs where, for every $j \in [\sigma]$, CLOUD inputs

Protocol	SYM	GROUP	OT	CT
CMTB	$ x $	$\frac{2\sigma}{5}(y + 1)$	k	yes
Salus	$\frac{2\sigma}{5}(x + y + f(x, y))$	-	-	yes
WW	$\sigma(x + \frac{2}{5} f_m(x, y))$	-	-	no

Table 1: Operations required on the mobile device by three outsourcing protocols. Here, SYM is the symmetric cryptographic operations, GROUP is the group algebraic operations, OT is the oblivious transfers, and CT is whether the protocol requires a coin toss. Recall that k is the security parameter, σ is the number of circuits generated, x is the mobile device's input, and y is the application server's input.

$(\rho^{(j)}, (\{\gamma_i^{(j)}\}_{i \in [m_m]} || \{\Xi_i^{(j)}\}_{i \in [m_c]}))$, while SERVER inputs s_j . This allows SERVER to learn either the randomness used to generate the check circuits or MOBILE and CLOUD'S inputs for the evaluation circuits without CLOUD knowing which circuits are being checked or evaluated.

Phase 5: Evaluation

1. **Circuit evaluation:** Using $\rho^{(j)}$, CLOUD garbles the objective circuit C as $G(C)^{(j)}$ for all $j \in [\sigma]$ and pipelines these circuits to SERVER using Huang's technique [16]. Depending on whether the circuit is a check circuit or an evaluation circuit, SERVER performs one of two actions:
 - (a) **Check:** For each $j \in [\sigma] \setminus S$, SERVER checks to see if $\rho^{(j)}$ can correctly regenerate the committed wire values $\{\Theta^{(j)}, \Omega^{(j)}, \Phi^{(j)}, \Psi^{(j)}\}$ and the circuit $G(C)^{(j)}$.
 - (b) **Evaluate:** For each $j \in S$, SERVER checks that she can correctly decommit MOBILE'S input by recovering half of $\Theta^{(j)}$ from the keys committed in $\Gamma^{(j)}$. She does the same for CLOUD'S input, recovering half of $\Psi^{(j)}$ from the keys committed in $\Xi^{(j)}$.

If any of the above checks fail, SERVER aborts the protocol. Otherwise, she evaluates the circuits $\{G(C)^{(j)}\}_{j \in [\sigma] \setminus S}$. Each circuit outputs the values $(h_m^{(j)}, h_c^{(j)}, c_s^{(j)}, c_m^{(j)})$ for $j \in [\sigma] \setminus S$.

2. **Majority output selection and consistency check:** Let (h_m, h_c, c_s, c_m) be the output of the majority of the evaluated circuits. If no majority value exists, SERVER aborts the protocol. Otherwise, she checks that $h_m^{(j)} = h_m$ and $h_c^{(j)} = h_c$ for all $j \in [\sigma] \setminus S$. If any of MOBILE or CLOUD'S hashed input values do not match, SERVER aborts the protocol.

Phase 6: Output proof and release

1. **Proof of output authenticity:** SERVER and MOBILE perform the proof of output authenticity from shelat and Shen's protocol [41] using the commitments to MOBILE'S output wires $\{\Phi^{(j)}\}_{j \in [\sigma] \setminus S}$ and the values $h_c || c_m$.
2. **Output release:** CLOUD simultaneously releases the input one-time pads p_s and p_m to SERVER and MOBILE. SERVER and MOBILE then hash the pads and check to see if the hash values output by the circuit $h_c = H \cdot p_s || p_m$. If the hashes do not match, SERVER and MOBILE abort the protocol. Otherwise, SERVER receives $c_s \oplus p_s$ as her output and MOBILE receives $c_m \oplus p_m \oplus e$ as his output.

5. COMPARISON WITH PREVIOUS OUTSOURCING PROTOCOLS

In this section, we compare the asymptotic complexity and security guarantees of the Whitewash protocol to two previous outsourcing techniques: the protocol developed by Carter et al. [8], which we call "CMTB" for the remainder of this work, and the Salus framework developed by Kamara et al. [23]. We refer to our

Whitewash protocol as WW.

When examining the complexity of each protocol, recall that one of our main goals is to optimize the efficiency on the mobile device. Thus, we examine the number of operations each protocol requires on the mobile device itself. When compared to the underlying shelat-Shen protocol, Whitewash adds extra input values from the Cloud, but does not add any steps to the computation that increase the complexity of operations performed on the application server or the Cloud. Thus, for a discussion of the application server and Cloud protocol complexity, we refer the reader to the original work by shelat and Shen [41].

5.1 Comparison to CMTB

The underlying two-party computation protocols of Whitewash and CMTB follow similar structures in terms of the security checks that are performed. However, Kreuter, shelat, and Shen’s (KSS) protocol [28], which underlies CMTB, uses a number of algebraic operations to perform input consistency checks and output proofs of consistency. The protocol developed by shelat and Shen [41], which underlies Whitewash, removes these expensive cryptographic primitives in favor of constructions that use only efficient, symmetric-key operations. In addition to the improvements to the underlying protocol, Whitewash outsources the generation side of two-party computation, while CMTB outsources the evaluation side. In CMTB, since neither the mobile device or the Cloud could garble inputs before computation, a specially designed Outsourced Oblivious Transfer (OOT) protocol is necessary to deliver the mobile device’s inputs to the evaluating Cloud in a secure, privacy-preserving manner. By swapping roles in the Whitewash protocol, we allow the mobile device to garble its own inputs, removing the need for an OT protocol to be performed from the mobile device. While Whitewash still requires OTs between the Cloud and the evaluating party, these operations can be parallelized, while the OOT protocol acts as a non-parallelizable bottleneck in computation.

5.1.1 Asymptotic Complexity

Table 1 shows this complexity for both Whitewash and CMTB. Note that for the mobile device, Whitewash requires significantly more symmetric key operations for garbling its own input and verifying the correctness of its output. By contrast, the OOT protocol in CMTB requires very few symmetric key operations, but requires several instantiations of an oblivious transfer. In addition, CMTB requires that the mobile device check the application server’s input consistency and verify the correctness of the output using algebraic operations (e.g., modular exponentiations and homomorphic operations). Considering the fact that modular exponentiation is significantly more costly than symmetric key operations, removing these public key operations from the phone is a significant efficiency improvement for Whitewash. We also note that CMTB requires a two-party fair coin toss at the mobile device, which is not required by Whitewash.

5.1.2 Security Guarantees

The removal of the OOT protocol in Whitewash not only increases its efficiency when compared to CMTB, it also allows for stronger security guarantees. In CMTB, security was only possible if none of the parties collude, since the mobile device possessed information that would allow the Cloud to recover both input wire labels for all of the mobile input wires after the OOT. If the mobile device and Cloud collude in the Whitewash protocol, it simply removes the guarantee of fair release and makes the protocol equivalent to the underlying two-party computation protocol. Thus, the only guarantee lost is that of fair release at the end of the proto-

col, since a colluding mobile device and Cloud may not release the one-time pad used to blind the evaluating party’s output. We believe that this represents a more realistic security setting, since the mobile device is paying for the assistance of the Cloud and may collude.

5.2 Comparison to Salus

When considering the operations performed on the mobile device, the Salus protocol and the Whitewash protocol both make the mobile device responsible for generating circuit randomness and garbling its own inputs. However, the Whitewash protocol requires an added proof of output consistency that is not included in Salus. While this proof adds some complexity to the protocol, it allows Whitewash to handle functions where both parties get different output values, while Salus is designed to handle functions with a single, shared output value. In addition, the Whitewash protocol outsources the generation of the garbled circuit, while the malicious secure Salus protocol outsources the evaluation. By swapping the roles of the outsourced task and adding in consistency checks at the evaluating party, the Whitewash protocol guarantees security in a stronger adversarial model.

5.2.1 Asymptotic Complexity

Both the Whitewash and Salus protocols use only efficient, symmetric key operations, but there is a slight tradeoff in the number of operations required (Table 1). Salus only requires operations for the $\frac{2\sigma}{5}$ evaluated circuits, but requires those operations for each bit of both party’s inputs and the shared output. By contrast, Whitewash requires that the mobile device’s input be committed for all σ circuits generated, but then only requires correctness proof of the output wires on the $\frac{2\sigma}{5}$ evaluated circuits. When the application server’s input is significantly longer than the mobile device’s, this will cause the Salus protocol to be less efficient than Whitewash. However, in the average case where both inputs are approximately the same length, this will mean that Whitewash requires more operations. This small tradeoff in efficiency is justified by the fact that Whitewash provides security in a stronger adversarial model than Salus. We also note that Salus requires a two-party fair coin toss before the protocol begins, which is not required by Whitewash.

5.2.2 Security Guarantees

The Salus protocol provides equivalent security guarantees to CMTB, guaranteeing security when none of the parties are colluding. This is a result of outsourcing the evaluation to the Cloud while allowing the mobile device to generate circuit randomness. If the mobile device colludes with the Cloud, they can trivially recover all of the other party’s inputs. By outsourcing the generation of the garbled circuit and adding in additional consistency checks at the evaluating party, Whitewash guarantees security under this type of collusion. As stated above, the only guarantee lost is that of fair output release, which ultimately reduces Whitewash to the security of the underlying two-party computation protocol.

6. PERFORMANCE EVALUATION

Our protocol significantly expands upon the implementations of the PCF garbled circuit generation technique [27] and shelat and Shen’s garbled circuit evaluation protocol [41]. For experimental comparison to previous protocols, we acquired the code implementation of the outsourcing protocol by Carter et al. [8] directly from the authors, as well as an Android port of the two-party garbled circuit protocol developed by Kreuter, shelat, and Shen [28]. We would like to thank the authors of [8, 27, 28, 41] for making their

Circuit	Input Size (Bits)	Total Gates		Non-XOR Gates	
		KSS	PCF	KSS	PCF
HAM (1600)	1,600	24,379	32,912	17,234	6,375
HAM (16384)	16,384	262,771	376,176	186,326	101,083
MAT (3x3)	288	424,748	92,961	263,511	27,369
MAT (5x5)	800	1,968,452	433,475	1,221,475	127,225
MAT (8x8)	2,048	8,067,458	1,782,656	5,006,656	522,304
MAT (16x16)	8,192	64,570,969	14,308,864	40,076,631	4,186,368
DIJK 10	112/1,040	259,232	530,354	118,357	291,490
DIJK 20	192/2,080	1,653,380	2,171,088	757,197	1,192,704
DIJK 50	432/5,200	22,109,330	13,741,514	10,170,407	7,549,370
RSA-256	256/512	934,092,960	673,105,990	602,006,981	235,925,023

Table 2: Input size and circuit size for all test circuits evaluated.

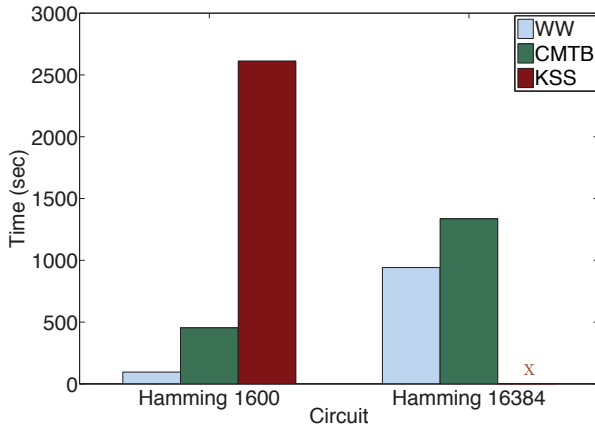


Figure 2: Execution time (ms) for Hamming Distance with input sizes of 1,600 and 16,384 bits for $\sigma = 256$ (note: log scale). Note that without outsourcing, only very small inputs can be computed over. Additionally, even for a large number of input bits, performing OTs on the servers still produces a faster execution time.

code available and for assisting us with our evaluation².

6.1 Test Environment

For evaluating our test circuits, we perform our experiments with a single server performing the role of Cloud and Application server, communicating with a mobile device over an 802.11g wireless connection. The server is equipped with 64 cores and 1TB of memory, and we partition the work between cores into parallel processing nodes using MPI. The mobile device used is a Samsung Galaxy Nexus with a 1.2 GHz dual-core ARM Cortex-A9 processor and 1 GB of RAM, running Android 4.0.

The large input sizes examined in the Hamming Distance trials required us to use a different testbed. For inputs as large as 16,384 bits, the phone provided by the above computing facility would overheat and fail to complete computation. Because the gate counts for Hamming Distance are significantly smaller than the other test circuits, we were able to run these experiments on a local testbed. We used two servers with Dual Intel Xeon E5620 processors, each

²We contacted the authors of the Salus protocol [23] in an attempt to acquire their framework to compare the actual performance of their scheme with ours. Because they were unable to release their code, no sound comparison to their work beyond an asymptotic analysis was possible. Our code is available at <http://www.foryourphoneonly.org>.




Figure 3: Execution time (ms) for the Matrix-Multiplication problem with input size varying between 3×3 matrices and 16×16 matrices for $\sigma = 256$ (note: log scale). This figure clearly shows that the oblivious transfers, consistency checks, and larger circuit representations of CMTB add up to a significant overhead as input size and gate count increase. By contrast, Whitewash requires less overhead and increases more slowly in execution time as gate counts and input size grow.

with 4 hyper-threaded cores at 2.4 GHz each for the Cloud and the Application server. Each server is running the Linux kernel version 2.6, and is connected by a VLAN through a 1 Gbps switch. Our mobile device is a Samsung Galaxy Note II with a 1.6 GHz quad-core processor with Arm Cortex A9 cores and 2 GB of RAM, running the Android operating system at version 4.1. The phone connects to the two servers through a Linksys 802.11g wireless router with a maximum data rate of 54 Mbps. While this test environment represents optimistic connection speeds that may not always be available in practice, it allows us to consider the performance of the protocol without interference from variable network conditions, and mirrors the test environments used in previous work [8, 28, 41]. For all experiments except RSA-256, we take the average execution time over ten test runs, with a confidence interval of 95%. For RSA-256, we ran 3 executions.

6.2 Experimental Circuits

To evaluate the performance of our protocol, we run tests over the following functions. We selected the following test circuits because they exercise a range of the two major variables that affect the speed of garbled circuit protocols: input size and gate counts.

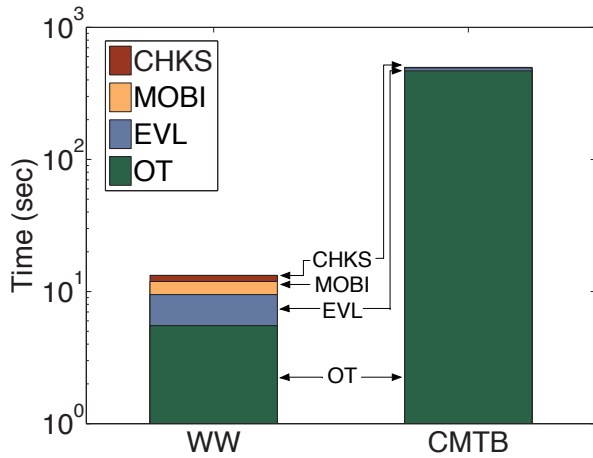


Figure 4: Microbenchmarking execution times (ms) for Whitewash and CMTB over the Matrix-Multiplication problem. We denote the total time spent in computation for Whitewash as “MOBI”. Since the mobile device is linked with “CHKS” and “OT” in CMTB, we do not separate out the mobile time for that protocol. Notice the dominating amount of time required to perform oblivious transfers. Moving these operations off the mobile device removes a significant computation bottleneck.

In addition, these programs are becoming somewhat standard test applications, having been used as benchmarks in a large amount of the related literature [8, 27, 28, 41]. All of the programs are implemented with the algorithms used by Kreuter et al. [27] except for Dijkstra’s algorithm, which matches the implementation used by Carter et al. [8]:

- **Hamming Distance (HAM):** The Hamming Distance circuit accepts binary string inputs from both parties and outputs the number of locations at which those strings differ. This circuit demonstrates performance for a small number of gates over a wide range of input sizes. We consider input strings of length 1,600 bits and 16,384 bits.
- **Matrix Multiplication (MAT):** Matrix multiplication takes an $n \times n$ matrix of 32-bit integer entries from each party and outputs the result of multiplying the matrices together. This circuit demonstrates performance when both input size and gate count vary widely. We consider square matrix inputs where $n = 3, 5, 8, \text{ and } 16$.
- **Dijkstra’s Algorithm (DIJK):** This version of Dijkstra’s algorithm takes an undirected weighted graph with a grid structure and a maximum node degree of four from the first party, and a start and end node from the second party. The circuit outputs the shortest path from the start node to the end node to the second party, and nothing to the first. For an n node graph, the graph description from the first party requires $104n$ input bits, while the start and end node descriptions require $8n + 32$ bits. We consider graphs with $n = 10, 20, \text{ and } 50$ nodes. Due to an error in the PCF compiler, we were unable to compile a program for graphs larger than 50 nodes.
- **RSA Function (RSA):** The RSA function (i.e., modular exponentiation) accepts an RSA message from one party and an RSA public key from the other party and outputs the encryption of the input plaintext under the input public key. Specifically, one party inputs the modulus $n = pq$ for primes p and q , as well as the encryption key $e \in \mathbb{Z}_{\phi(n)}$. The other party inputs a message $x \in \mathbb{Z}_n^*$, and the circuit computes $x^e \pmod n$.

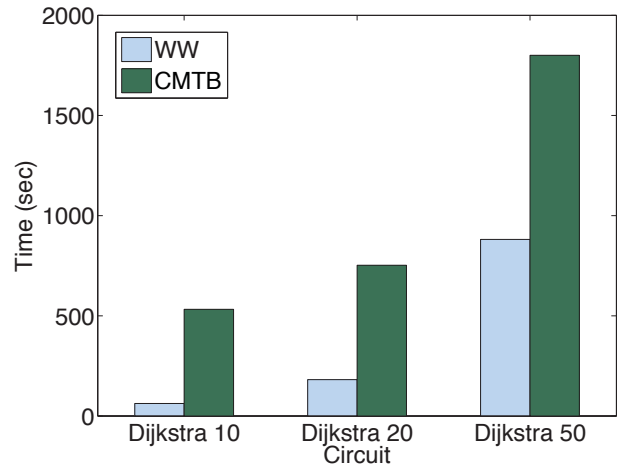


Figure 5: Execution time (s) for Dijkstra’s algorithm with input sizes of 10, 20, and 50 node graphs for $\sigma = 256$. This figure shows that the Whitewash protocol allows for computation that was only feasible to be executed in a close to practically useful time frame.

This circuit demonstrates performance for small input sizes over very large gate counts. We consider the case where the input values $x, n,$ and e are 256 bits each. While these are not secure parameters in practice, the function itself provides a complex circuit that is scalable on input size and useful for benchmarking our protocol.

For each test circuit, we consider the time required to execute and the bandwidth overhead to the mobile device. Table 2 shows the input size and gate counts for each test circuit, showing the exact range of values tested for these two circuit variables.

6.3 Execution Time

In all experiments, the efficiency gains of removing oblivious transfers and public key operations are immediately apparent. To examine how Whitewash compares to generating garbled circuits directly on the mobile device, we considered Hamming Distance as a simple problem (Figure 2). Even with a relatively small gate count, garbling the circuit directly on the mobile device is only possible for the small input size of 1,600 bits. Whitewash is capable of executing this protocol in 96 seconds, while running the computation directly on the mobile device takes 2,613 seconds, representing a 96% performance improvement through our outsourcing scheme. For the very large input size of 16,384 bits, computation directly on the mobile device ceases to be possible. When comparing to CMTB, this circuit further illustrates the cost of oblivious transfers on the mobile device. Even with the significantly reduced number of OTs allowed by the OOT protocol in CMTB (80 OTs), performing 16,384 malicious secure oblivious transfers between two servers in Whitewash still runs 30% faster than CMTB.

The matrix-multiplication circuit provides a good overview of average-case garbled circuit performance, as it represents a large range of both gate counts and size of inputs. For the input size of a 3×3 matrix, the Whitewash protocol runs in an average of 12 seconds, while CMTB requires 493 seconds, representing a 98% improvement (see Figure 3). Upon inspecting the micro benchmarking breakdown of each protocol’s execution in Figure 4, we observe a significant speedup simply by moving oblivious transfers off of the mobile device. Even though the number of OTs required by CMTB is essentially constant based on their application of the

Ishai OT extension, performing standard malicious secure oblivious transfers in parallel between the servers is much more efficient than requiring that the phone perform these costly operations. In addition, if we examine the amount of execution time where the phone participates in Whitewash, we see that the mobile device (“MOBI” in Figure 4), takes around 1 second, and is idle during the majority of computation. By contrast, both the OT and consistency check phases of CMTB require the mobile device to participate in a significant capacity, totaling to almost 8 minutes of the computation. Having the phone perform as little work as possible means that the Whitewash protocol performance is nearly equivalent to running the same computation between two server-class machines.

To examine the performance of Whitewash for a more practical application, we considered the Dijkstra’s algorithm circuit used to implement privacy-preserving navigation by Carter et al. [8]. They point out that this application, which has uses from military convoys to industrial shipping routes, is a significant first step in providing privacy for the growing genre of location-based mobile applications. Unfortunately, the PCF compiler does not optimize the Dijkstra’s circuit as well as the previous experimental programs, which is evident in Table 2. In the 10 and 20 node graphs, the PCF compiler even produces a *larger* circuit than the compiler used by KSS. However, despite evaluating larger circuits, the Whitewash protocol still outperforms CMTB in execution time, running 88%, 76%, and 51% faster in the 10, 20, and 50 node cases respectively (shown in Figure 5). As circuit compilers continue to improve and produce smaller circuits, the performance gains of the Whitewash protocol will be even larger. In this experiment, we also noticed that because Whitewash evaluates and checks circuits simultaneously, it created contention for the network stack in our test server. In a truly distributed environment where each server node has dedicated network resources, the highly parallelizable structure of shelat and Shen’s protocol would allow Whitewash to execute faster. Given that Whitewash can execute Dijkstra’s algorithm obliviously on the order of minutes, it allows computation considered only feasible for previous schemes to be performed in a nearly practical execution time.

The previous experiments clearly show that outsourcing is necessary to run circuits of any practical size. For our final test circuit, we consider an extremely complex problem to demonstrate the ability of outsourcing protocols in the worst-case. The RSA-256 circuit evaluated by Kreuter et al. in [27] and shelat and Shen in [41] represents one of the largest garbled circuits ever evaluated by a malicious secure protocol. For the RSA-256 problem, Whitewash completed the computation in 515 minutes. CMTB was unable to complete one execution of the protocol. A large part of this efficiency improvement results from the underlying protocol of Whitewash, which uses only symmetric-key operations outside of the oblivious transfers between the servers. The reduced non-XOR gate counts and more compact circuit representation of the PCF compiler also contribute to this improvement. Ultimately, because Whitewash ensures that the phone participates minimally in the protocol, it no longer acts as a bottleneck on computation. We essentially reduce performance of our outsourcing protocol to that of the underlying two-party protocol, allowing this technique for outsourcing to benefit as more improvements are made in non-outsourced garbled circuit protocols. In addition, this minimal level of interactivity allows us to run these protocols with 256 circuits, equivalent to a security parameter of approximately 80-bit security, which is agreed by the research community to be an adequate security parameter. Finally, the phone is only active for a few seconds during this large computation, keeping its system resources free for other user applications (or to preserve battery power) while the servers complete the com-

Circuit	Bandwidth (MB)			Reduction Over	
	WW	CMTB	KSS	CMTB	KSS
HAM (1600)	23.56	41.05	240.33	42.62%	90.20%
HAM (16384)	241.02	374.03	x	35.56%	x
MAT (3x3)	4.26	11.50	x	62.97%	x
MAT (5x5)	11.79	23.04	x	48.82%	x
MAT (8x8)	30.15	51.14	x	41.05%	x
MAT (16x16)	120.52	189.52	x	36.41%	x
DIJK 10	1.67	20.21	x	91.73%	x
DIJK 20	2.85	35.28	x	91.93%	x
DIJK 50	6.38	80.49	x	92.08%	x
RSA-256	3.97	x	x	x	x

Table 3: Bandwidth measures for all experiment circuits. Note that there is as much as a 84% reduction in bandwidth when using the Whitewash protocol.

putation. *This shows that Whitewash is capable of evaluating the same circuits as the most efficient desktop-based garbled circuit protocols with a minimal overhead cost.* Exact execution times are shown in our technical report [7].

6.4 Network Bandwidth

The Whitewash protocol not only improves the speed of execution when outsourcing garbled circuit computation, it also significantly reduces the amount of bandwidth required by the mobile device to participate in the computation. Table 3 shows the bandwidth used by the mobile device for each test circuit. In the best case, for Dijkstra’s algorithm over 50 node graphs, we observed a 92% reduction in bandwidth between Whitewash and CMTB. This is a result of the mobile device not performing OTs and only sending relatively small symmetric-key values instead of algebraic elements for consistency checks. For all test circuits, we observed a small decrease in the amount of improvement between the two protocols as the input size increased. This is because the number of commitments sent by the phone in Whitewash increases as the size of the input grows, while CMTB performs a fixed number of OTs as the input size increases. However, the oblivious transfers still require a significant enough amount of bandwidth to make removing them the most efficient option. When comparing to not outsourcing garbled circuit generation, the cost of oblivious transfers and sending several copies of the garbled circuit to the evaluator quickly adds up to a significant bandwidth cost. For the smallest circuit evaluated, outsourcing the circuit garbling reduces the required amount of bandwidth by 90%. The importance of these bandwidth reductions is further highlighted when considering mobile power savings. With data transmission costing roughly 100 times as much power as computation on the same amount of data, any reduction in the bandwidth required by a protocol implies a critical improvement in practicality.

One challenge encountered during the implementation of the Whitewash protocol was the extensive use of hardware-specific functions used to implement commitment schemes in shelat and Shen’s code. Rather than try to port this code over to Android, which would require significant development of hardware-specific libraries, we chose to implement the protocol in an equivalent manner by having the Cloud generate the part of the commitments which requires these functions and send them to the mobile device. The mobile device then finishes generating the commitments that match its input and forwards them to the evaluator. Our proofs of security remain valid even with this small protocol modification. Our preliminary implementation using instructions specific to the ARM architecture has shown that we could further reduce the measured bandwidth values by over 60%. With already significant bandwidth

reductions from previous outsourcing schemes, our protocol will see further improvements as mobile hardware begins to incorporate more machine-specific libraries.

7. CONCLUSION

With the increasingly pervasive and personal nature of mobile computing, garbled circuits provide a solution that preserves both privacy and application functionality. However, to make these computationally expensive protocols usable on mobile devices, secure outsourcing to the cloud is necessary. We develop a new scheme that eliminates the most costly operations, including oblivious transfers, from the mobile device. By requiring that the mobile device instead produce the randomness required for circuit generation, we significantly reduce the number of algebraic group operations and communication rounds for the mobile device. At the same time, we bolster the security guarantees against certain types of collusion, yielding a more secure protocol than any other in this space. Our performance evaluation shows average gains of 75% for execution time and 60% for bandwidth over the previous outsourcing protocol. These improvements allow large circuits representing practical applications to be computed efficiently from a mobile device. As a result, we show that garbled circuit protocols can be made nearly as efficient for mobile devices as they are for server-class machines.

Acknowledgments This material is based on research sponsored by DARPA under agreement number FA8750-11-2-0211. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

- [1] M. J. Atallah and K. B. Frikken. Securely outsourcing linear algebra computations. In *Proceedings of the ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2010.
- [2] Y. Aumann. Security Against Covert Adversaries: Efficient Protocols for Realistic Adversaries. *Journal of Cryptology*, 18(3):554–343, 2010.
- [3] D. Beaver. Server-assisted cryptography. In *Proceedings of the workshop on New security paradigms (NSPW)*, 1998.
- [4] M. Bellare, V. T. Hoang, S. Keelveedhi, and P. Rogaway. Efficient garbling from a fixed-key blockcipher. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2013.
- [5] J. Brickell and V. Shmatikov. Privacy-preserving graph algorithms in the semi-honest model. In *Proceedings of the international conference on Theory and Application of Cryptology and Information Security*, 2005.
- [6] H. Carter, C. Amrutkar, I. Dacosta, and P. Traynor. For your phone only: custom protocols for efficient secure function evaluation on mobile devices. *Journal of Security and Communication Networks (SCN)*, 7(7):1165–1176, 2014.
- [7] H. Carter, C. Lever, and P. Traynor. Whitewash: Outsourcing garbled circuit generation for mobile devices. Cryptology ePrint Archive, Report 2014/224, 2014. <http://eprint.iacr.org/>.
- [8] H. Carter, B. Mood, P. Traynor, and K. Butler. Secure Outsourced Garbled Circuit Evaluation for Mobile Devices. In *Proceedings of the USENIX Security Symposium*, 2013.
- [9] comScore. comScore Reports February 2013 U.S. Smartphone Subscriber Market Share. http://www.comscore.com/Insights/Press_Releases/2013/4/comScore_Reports_February_2013_U.S._Smartphone_Subscriber_Market_Share, 2013.
- [10] I. Damgård, M. Geisler, and J. B. Nielsen. From passive to covert security at low cost. In *Proceedings of the 7th international conference on Theory of Cryptography*, 2010.
- [11] I. Damgård, V. Pastro, N. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. In *Proceedings of the Annual International Cryptology Conference on Advances in Cryptology*, 2012.
- [12] C. Gentry, S. Halevi, and N. P. Smart. Homomorphic evaluation of the AES circuit. In *Advances in Cryptology - CRYPTO*, 2012.
- [13] S. D. Gordon, J. Katz, V. Kolesnikov, A.-I. B. Labs, F. Krell, and M. Raykova. Secure Two-Party Computation in Sublinear (Amortized) Time. In *Proceedings of the ACM conference on Computer and communications security (CCS)*, 2012.
- [14] M. Green, S. Hohenberger, and B. Waters. Outsourcing the Decryption of ABE Ciphertexts. In *Proceedings of the USENIX Security Symposium*, 2011.
- [15] C. Hazay and Y. Lindell. Efficient Protocols for Set Intersection and Pattern Matching with Security Against Malicious and Covert Adversaries. *Journal of Cryptology*, 23(3):422–456, 2008.
- [16] Y. Huang, D. Evans, J. Katz, and L. Malka. Faster Secure Two-Party Computation Using Garbled Circuits. In *Proceedings of the USENIX Security Symposium*, 2011.
- [17] Y. Huang, J. Katz, and D. Evans. Quid-pro-quo-tocols: Strengthening semi-honest protocols with dual execution. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2012.
- [18] Y. Huang, J. Katz, and D. Evans. Efficient secure two-party computation using symmetric cut-and-choose. In *Advances in Cryptology—CRYPTO*, 2013.
- [19] N. Husted, S. Myers, abhi shelat, and P. Grubbs. GPU and CPU parallelization of honest-but-curious secure two-party computation. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, 2013.
- [20] A. Iliiev and S. W. Smith. Small, Stupid, and Scalable: Secure Computing with Faerieplay. In *The ACM Workshop on Scalable Trusted Computing*, 2010.
- [21] S. Jha, L. Kruger, and V. Shmatikov. Towards practical privacy for genomic computation. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2008.
- [22] S. Kamara, P. Mohassel, and M. Raykova. Outsourcing multi-party computation. Cryptology ePrint Archive, Report 2011/272, 2011. <http://eprint.iacr.org/>.
- [23] S. Kamara, P. Mohassel, and B. Riva. Salus: A system for server-aided secure function evaluation. In *Proceedings of the ACM conference on Computer and communications security (CCS)*, 2012.
- [24] F. Kerschbaum. Collusion-resistant outsourcing of private set intersection. In *Proceedings of the ACM Symposium on Applied Computing*, 2012.
- [25] M. Kiraz and B. Schoenmakers. A Protocol Issue for The Malicious Case of Yao’s Garbled Circuit Construction. In *Proceedings of the Symposium on Information Theory in the Benelux*, 2006.
- [26] M. S. Kiraz. *Secure and Fair Two-Party Computation*. PhD thesis, Technische Universiteit Eindhoven, 2008.
- [27] B. Kreuter, a. shelat, B. Mood, and K. Butler. PCF: A portable circuit format for scalable two-party secure computation. In *Proceedings of the USENIX Security Symposium*, 2013.
- [28] B. Kreuter, a. shelat, and C. Shen. Billion-Gate Secure Computation with Malicious Adversaries. In *Proceedings of the USENIX Security Symposium*, 2012.
- [29] L. Kruger, S. Jha, E.-J. Goh, and D. Boneh. Secure Function Evaluation with Ordered Binary Decision Diagrams. In *Proceedings of the ACM conference on Computer and communications security (CCS)*, 2006.
- [30] Y. Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. In *Advances in Cryptology—CRYPTO*, 2013.
- [31] Y. Lindell and B. Pinkas. Privacy preserving data mining. In *Proceedings of the Annual International Cryptology Conference on Advances in Cryptology*, 2000.
- [32] Y. Lindell and B. Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In *Proceedings of the annual international conference on Advances in Cryptology*, 2007.
- [33] Y. Lindell and B. Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. In *Proceedings of the conference on Theory of cryptography*, 2011.
- [34] L. Malka. Vmccrypt: modular software architecture for scalable secure computation. In *Proceedings of the 18th ACM conference on Computer and communications security*, 2011.
- [35] D. Malkhi, N. Nisan, B. Pinkas, and Y. Sella. Fairplay—a secure two-party computation system. In *Proceedings of the USENIX Security Symposium*, 2004.
- [36] A. Miyaji and M. S. Rahman. Privacy-preserving data mining in presence of covert adversaries. In *Proceedings of the international conference on Advanced data mining and applications: Part I*, 2010.
- [37] P. Mohassel and M. Franklin. Efficiency tradeoffs for malicious two-party computation. In *Proceedings of the Public Key Cryptography conference*, 2006.
- [38] B. Mood, L. Letaw, and K. Butler. Memory-efficient garbled circuit generation for mobile devices. In *Proceedings of the IFCA International Conference on Financial Cryptography and Data Security (FC)*, 2012.
- [39] N. Nipane, I. Dacosta, and P. Traynor. “Mix-In-Place” anonymous networking using secure function evaluation. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, 2011.
- [40] a. shelat and C.-H. Shen. Two-output secure computation with malicious adversaries. In *Proceedings of the Annual international conference on Theory and applications of cryptographic techniques*, 2011.
- [41] a. shelat and C.-H. Shen. Fast two-party secure computation with minimal assumptions. In *Proceedings of the ACM conference on Computer and communications security (CCS)*, 2013.
- [42] D. Talbot. Security in the ether. <http://www.technologyreview.com/featuredstory/416804/security-in-the-ether/>, 2009.
- [43] A. C. Yao. Protocols for secure computations. In *Proceedings of the Annual Symposium on Foundations of Computer Science*, 1982.