



Auditable, Available and Resilient Private Computation on the Blockchain via MPC

Christopher Cordi¹, Michael P. Frank¹, Kasimir Gabert¹, Carollan Helinski¹,
Ryan C. Kao¹, Vladimir Kolesnikov²(✉), Abraham Ladha²,
and Nicholas Pattengale¹

¹ Sandia National Laboratories, Albuquerque, NM, USA

² Georgia Institute of Technology, Atlanta, GA, USA

kolesnikov@gatech.edu

Abstract. Simple but mission-critical internet-based applications that require extremely high reliability, availability, and verifiability (e.g., auditability) could benefit from running on robust public programmable blockchain platforms such as Ethereum. Unfortunately, program code running on such blockchains is normally publicly viewable, rendering these platforms unsuitable for applications requiring strict privacy of application code, data, and results.

In this work, we investigate using MPC techniques to protect the privacy of a blockchain computation. While our main goal is to hide both the data and the computed function itself, we also consider the standard MPC setting where the function is public.

We describe GABLE (Garbled Autonomous Bots Leveraging Ethereum), a blockchain MPC architecture and system. The GABLE architecture specifies the roles and capabilities of the players. GABLE includes two approaches for implementing MPC over blockchain: Garbled Circuits (GC), evaluating universal circuits, and Garbled Finite State Automata (GFSA).

We **formally model and prove** the security of GABLE implemented over garbling schemes, a popular abstraction of GC and GFSA from (Bellare et al., CCS 2012).

We analyze in detail the performance (including Ethereum gas costs) of both approaches and discuss the trade-offs. We implement a simple prototype of GABLE and report on the implementation issues and experience.

This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for NNSA under contract DE-NA0003525. This report describes objective technical results and analysis. Any subjective views or opinions that might be expressed in this report do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Approved for public release SAND2022-3532 C.

© Springer Nature Switzerland AG 2022

S. Dolev et al. (Eds.): CSCML 2022, LNCS 13301, pp. 281–299, 2022.

https://doi.org/10.1007/978-3-031-07689-3_22

1 Introduction

Current programmable blockchain platforms such as Ethereum provide a “global computer,” an always-on public computing resource. They guarantee reliability, availability and auditability for computations implemented as smart contracts, which are posted to the blockchain and subsequently executed. Even in the event of a widespread disaster, any still-functioning part of the Ethereum network renders this computing resource *available*. Organizations can take advantage of such a robustly available computing facility to execute particularly mission-critical computational tasks, if this computation can be done *privately*.

We note that the performance of secure multiparty computation (MPC) has been steadily improving and is practical today even for large functions/inputs.

In this work, we explore the motivation, use cases, architectures, and constructions for securing (i.e. ensuring privacy of) a sensitive computation done on sensitive inputs on a public blockchain network using MPC techniques.

We present GABLE (Garbled Autonomous Bots Leveraging Ethereum), an architecture and a system for running MPC on the Ethereum blockchain. We consider two different base approaches, garbled finite state machines/automata (GFSA), and garbled circuits (GC); we present both in a unified manner as garbling schemes. We show how functions can be computed privately. At a high level, in our architecture there are four types of players, or participant roles:

1. *Garbler* is a player who generates encrypted function and input encodings, publishes the former on Ethereum, and distributes the latter to other players. Garbler may be a single trusted entity or run distributedly, e.g., by an MPC over a private chain.
2. *Input Provider* or *Writer* is a player who contributes (encrypted) inputs to the computation.
3. *Input Unlocker* or just *Unlocker* is a designated player that manages encrypted inputs, preventing input providers from learning unauthorized information about the computation and adapting their input based on it.
4. *Output Recipient* or *Reader* is a player who may obtain a designated output from the computation.

We emphasize that players may play several of these logical roles simultaneously (with some exceptions; cf. trust model Sect. 2.3). For example, an input provider may also be an output recipient.

1.1 Motivation

Existing public programmable smart-contract blockchains such as Ethereum provide a highly robust and accessible computing platform. An entity whose operations may require execution of business or strategic logic that needs to function (using inputs from various sources) with a high degree of assurance of its reliability and availability may implement that functionality as a smart contract running on a blockchain. Another desirable blockchain property is auditability, due to the generally immutable nature of data committed to a consensus chain.

However, the deploying entity may wish to keep private the following computation data and metadata:

- Input values provided to the computation, including their semantics.
- The nature of the computation that is being performed on the inputs. Most generally, we may wish for all information about the structure and function of the computation to be obscured.
- Any internal data within the computation.
- Semantics of any intermediate or final outputs produced by the computation.

Most of these features are implied by standard MPC properties, while hiding the computed function is the goal of Private Function Evaluation (PFE) [25].

Use Cases. We outline two use-case scenarios which illustrate the potential usefulness and practicality of our system. Of course, many others are possible.

Sealed-Bid Auctions: Auditability and Security. Consider a simple sealed-bid auction with simultaneous bidding, where the soundness of the computation may be fully audited after the fact, if needed. This may be useful, e.g., in public-interest auctions, such as auctioning airwaves to cellular providers, or electricity delivery contracts to utilities. Other cases (e.g., auctioning of a privately-held company to a select set of bidders) may demand strict privacy.

While such auctions may be easily run via an MPC, running them on a blockchain offers a much higher level of transparency and trust (and user engagement) than more traditional means of execution. We report some analysis of costs for our methods on simple multi-bidder auctions in Sect. 5.

Transactive Energy. Refers to market-based mechanisms for managing the exchange of energy in a wide-area electrical grid [32]. Participants in such a market, e.g., utility companies, may wish to engage in automatic electronic negotiation without revealing their strategies. A capability such as GABLE could provide a solution. Each participant deploys their own garbled bot expressing their private negotiation strategy; these bots then negotiate with each other on the public blockchain, while hiding their internal negotiation strategies.

Private-Public Chain Integration. The encrypted function may be generated (garbled) by an MPC executed over a private network. This assures trust in the garbling, facilitates future opening of the secure computation (auditing) if needed (cf. Sect. 2.3), and integrates public and private chains.

1.2 Contribution

Blockchain and MPC technologies provide essentially complementary features to a computation. A blockchain assures availability (ability to provide input, compute, and retrieve the output when needed) and reliability (including assurance that only authorized players can submit input and that the output of the computation is correct), even if direct communication channels among the parties

are cut. MPC ensures privacy and correctness, guaranteeing partial reliability (output correctness) but nothing about availability.

Our system, GABLE (Garbled Autonomous Bots Leveraging Ethereum) combines the best of the two technologies. In our security model, we require not only the data, but, optionally, also the computed function to be private (standard MPC evaluates publicly known functions). In this work, we:

- Design a general architecture for available and resilient MPC over the Ethereum blockchain. It allows authorized players to submit inputs at any time, and to retrieve outputs as soon as they are available. Our design specifies the roles and capabilities of the players.
- Design two approaches for implementing MPC on the blockchain: Garbled Circuits (GC), including GC evaluating universal circuits, and Garbled Finite State Automata (GFSA).
- *Formally state and prove* the security of our system against malicious players. Note, we *do not formally define* the properties of auditability, availability and resilience of the computation. They are derived from the corresponding intuitive properties of the underlying Ethereum blockchain. We specifically discuss achieving auditability in Sect. 2.3.
- Analyze in detail the performance characteristics (including gas costs) of both approaches and discuss the trade-offs.
- Implement a simple prototype and several demonstration applications, and report on the implementation issues and experience.¹

1.3 Results and Evaluation

We prototyped GABLE using GFSA and experimented on an Ethereum testnet. See Sect. 5 for details of our experiments and results; here, we summarize our findings. First, a simple provenance-tracking application scenario was modeled using 5 cycles of reactive functionality in a simple 6-state FSA; the cost to run this demo on the Ethereum mainnet would have been less than US\$3.00 on the day of the test. Next, we demonstrated a 5-state machine implementing MPC for the classic 2-party “Millionaire’s Problem” for configurable input lengths based on a bit-serial comparison algorithm; the cost of that demo would have been about \$0.50 per bit of input, and can be further reduced. Finally, we compared costs for GFSA versus a garbled universal circuit (GUC) implementation on a multi-bit auction problem, showing that the cost overhead of GUC grows only modestly (polylogarithmically) with the number of bidders, as expected, and becomes less expensive than GFSA for $B > 7$ bidders, at which point the cost is about \$20 per bit of price data. Note, GC *without* functional privacy would cost far less.

¹ A reference implementation of our initial prototype has already been publicly released [21], and wider licensing/availability of a more complete code base is under consideration.

1.4 Outline of the Paper

We presented the motivation and several use cases above in Sect. 1.1, and previewed our contribution and results in Sect. 1.2 and Sect. 1.3. We present preliminaries in Sect. 1.5 and discuss related work in Sect. 1.6. We provide a high-level overview of our approach in Sect. 2, including defining the logical players, the trust model, and security intuition. A generic security statement and theorem are outlined in Sect. 3, and specialized for our main GC-based protocol in Sect. 4. We discuss our prototype and demo implementations and present results and cost analysis in Sect. 5 and conclude in Sect. 6.

The full version of the paper [18] also includes several appendices that present some technical details of the GFSA approach used in our early implementations, the proof of the main security theorem, and additional details of our prototype implementations and test results.

A much more detailed report on the system design, implementation and demos, including reference source code for the initial prototype, is available as a Sandia National Laboratories technical report [21].

1.5 Preliminaries

Blockchain Technology. Bitcoin [28] is a stunningly successful technology, which generalizes public timestamping (ledger) and smart contracts work and uses proof-of-work, concepts considered before [19, 23, 29]. Bitcoin has limited support for programming digital smart contracts, i.e. it is not Turing-complete. This is in part due to the possibility of intentional or accidental resource overuse or even exhaustion as a result of programming issues. The need for a rich programming language for contracts was nevertheless recognized, and in 2013, the Ethereum blockchain was proposed [14]. Ethereum addressed the issues stemming from language Turing-completeness by putting the onus on the programmer/contract creator, and requiring payment per storage unit and execution step. We implement our system for Ethereum; our higher-level design is general and will fit most natural blockchain architectures.

MPC (including Two-Party Computation, 2PC, and the general p -Party Computation), is an area of cryptography that studies protocols which compute functions over players' private inputs while revealing nothing except the function output. MPC has improved dramatically over the past 15 years. The first proof-of-concept 2PC implementation, Fairplay [27], evaluated only 200 Boolean gates per second. Today, 2PC implementations can process up to 2–5 million gates/s [36]. Improvements in the malicious model and in 3PC and MPC are even more impressive. Recent work reports 3PC techniques that can evaluate as many as *seven billion* Boolean gates/second [6]. Research on algorithms and implementations has firmly transitioned MPC from a theoretical curiosity to a subject of practical engineering.

We note that 2PC protocols, and specifically Yao's garbled circuit (GC) techniques [34], are most suitable for use in our setting, because the blockchain itself can naturally serve as the GC evaluator.

FSA (finite-state automata) comprise a standard but simple model of computation that differs from Boolean circuits. The FSA model is weaker (some functions require exponentially larger representation in FSA as compared to circuits). The primary benefit of GFSA, in the case of sufficiently low-complexity applications, is simply that obscuring the structure of the computation becomes relatively trivial, since all computations reduce to a linear sequence of state transition table lookups, and relatively simple techniques suffice to obscure the topology of the state graph. As a result, the overhead to achieve functional privacy in GFSA becomes less than that of garbled (universal) circuits for computations operating on sufficiently small numbers of bits.

Garbling Schemes and Garbled Functions (GF). We build our approach around GC and GFSA. They are special cases of garbling schemes, as defined by Bellare et al. [10]. Informally, we will refer to garbling schemes as *garbled functions* (GF), similarly to how “GC” refers to both the GC garbling scheme and the GC approach. We will use the terms “garbled” and “encrypted” function interchangeably in this work. The BHR framework defines a garbling scheme as a tuple of algorithms $\mathcal{G} = (\text{ev}, \text{Gb}, \text{En}, \text{Ev}, \text{De})^2$ and requires that they satisfy the following properties:

In addition to the correctness property, BHR define relevant notions of security for garbled functions: privacy, obliviousness and authenticity. We refer the reader to [10] for precise definitions of these standard notions. Here, we informally summarize these notions.

GF *correctness* guarantees correct evaluation if all players behave honestly.

GF *privacy* guarantees that an adversary Adv who sees the garbled function (e.g. the GC), the encoded inputs and the output decoding information, does not learn anything beyond the result of the computation.

GF *obliviousness* guarantees that an Adv who sees the garbled function and the encoded inputs, does not learn anything. This notion is different from privacy, which gives Adv the decoding information and allows it to obtain the output of the computation (and nothing else).

GF *authenticity* captures Adv 's inability to create from a GF F and its garbled input X a garbled output $Y \neq F(X)$ that will be deemed authentic.

Note, we only need correctness and privacy. The authenticity requirement can be avoided if we choose to rely on the blockchain to honestly evaluate GF. Obliviousness becomes unnecessary if the output decoding information d is always published (e.g., d is provided to output receivers, who may be corrupted by Adv), in which case we are never in the setting without d available to Adv .

MPC from GF. Bellare et al. [10] do not systematize the ways one can obtain an MPC protocol from GF. However, the following (informally presented) natural 2PC construction works, and is proven secure against semi-honest adversaries in [10], assuming GF is private:

² In the BHR notation, ev is a reference evaluator for plaintext functions, Gb is the Garbler, En is the input encoder, Ev is the garbled function evaluator, and De is the output decoder.

Construction 1 (2PC from GF, informal). *Gen generates GF F , encoding information e , and decoding information d by running G_b . Gen sends F, d to Eval. Gen and Eval securely (e.g. via Oblivious Transfer (OT)), deliver to Eval the labels of e corresponding to players' inputs. Eval evaluates GF by running E_v , and obtains the plaintext output from garbled output labels and d by running D_e .*

We note that the above construction is secure against malicious Eval, as long as label delivery remains secure against a malicious Eval. We will use this property in our security argument.

1.6 Related Work

We briefly discussed relevant MPC and FSA preliminaries in Sect. 1.5. In this section, we review several systems addressing privacy on the blockchain and compare them to our approach.

MPC+Blockchain. As we discuss next, many works explore the interplay of blockchain and MPC. To our knowledge, only YOSO (You Only Speak Once) MPC [11, 22] formally models a *public* blockchain executing MPC. YOSO deviates from the typical blockchain architecture (e.g., of Ethereum) of all nodes sharing the same view. Instead, YOSO nodes have private data and are selected to perform MPC subtasks. If sufficiently large fraction of selected players are honest, MPC is secure. To protect against adversarial corruption, these players are hidden: they are unpredictably self-selected (e.g., via mining-like process), and each MPC subtask consists of computing and sending a *single* message, after which they erase their relevant private state. The main technical challenge of YOSO MPC is sending encrypted messages (e.g. containing internal state and subtask computation output) to *unidentified* players who are self-selected in the future. While YOSO MPC has attractive asymptotic complexity, unfortunately, it is concretely prohibitively expensive due to the cost of its building blocks. Our solution, at the cost of much stronger corruption and trust model (e.g., we only handle non-adaptive corruptions, while YOSO supports adaptive), is far more efficient and aligns with Ethereum architecture.

On the other hand, permissioned networks, such as Hyperledger, may be run by a small number of semi-trusted servers, and MPC can naturally be executed among the servers to achieve full privacy of transactions and contracts. This direction is explored in [12]. Their approach does not extend to public blockchains, since an arbitrary number of adversarial nodes may participate in the public network. Our work can be seen as general MPC on a public ledger for a restricted use case, where the encrypted function is generated by an organization trusted by the participants (and whose honesty can be later audited).

Hawk [26] is an architecture for a blockchain that can support private data. It handles private data using a trusted *manager*, realized using trusted hardware, such as Intel SGX. The trusted enclave *may* be implemented via MPC. It is not clear who would be the MPC principals to achieve a reasonable trust model; further (and [26] acknowledges this) this would cause an impractical overhead.

The Enigma system [37] uses MPC protocols to implement support for private data on a blockchain. They use MPC off-chain to perform computation on shares of data. We aim to run MPC on-chain for resilience, availability and auditability; Enigma’s techniques will not achieve these properties.

A line of work explores the interplay of blockchain and (separately executed) MPC to achieve fairness in MPC or connect MPC to financial mechanisms directly [5, 13, 17]. Works such as [33] use blockchain to manage encrypted inputs to MPC performed by a separate trusted network. Reference [16] considers a blockchain-hybrid MPC model (plain model with available ledger), and addresses foundational issues of MPC, such as concurrent composability, in this model. In contrast, in our work, the blockchain itself executes MPC.

Zero-knowledge proofs (ZKP) are widely used both in MPC and in blockchain. We note that public ledger nodes never prove anything (indeed, the underlying secret would then be known to everyone). Instead, ZKPs are used by off-chain entities, such as wallets, to prove correctness of their actions. Several ledgers, such as ZCash, provide transaction privacy based on ZKPs. This line of work is orthogonal to the privacy protection work we consider.

Solidus [15] uses a publicly verifiable ORAM machine to generalize and scale up the ZKPs for the use case where financial institutions representing many accounts interact with a ledger.

Blockstream CA [31] use simple ZKPs in conjunction with additive homomorphic commitments to manipulate secret data on the ledger. Partial privacy can be achieved for very simple functionalities (for efficiency, we are constrained by additively-homomorphic encryption).

In contrast to the above approaches, our solution is general MPC.

Trusted Enclaves. As in the Hawk example above, privacy can be achieved if one is willing to entrust hardware enclaves, such as SGX. Nodes of the blockchain network may be equipped with enclaves, which would execute encrypted contracts on encrypted data. Several other systems, such as Secret Network [4], also implement this approach. We note that enclave security is a cat-and-mouse game; in this work, we do not rely on secure enclaves.

2 Overview: Approach and Trust Model

As discussed in Sect. 1, we wish to add privacy of both computations and data to the process of contract execution on the Ethereum network. Data and function privacy is normally achieved using an appropriate secure computation protocol. However, in the public blockchain setting, the number of network nodes is unspecified, and MPC privacy guarantees cannot be achieved. Instead, we take the following approach:

2.1 Logical Players and Evaluation Pattern

We consider several logical players:

- The *Contract creator* or *Garbler* sets up encryptions of functions and inputs. It initializes the contract and sends encrypted labels to corresponding input providers. Garbler can be run by an MPC, e.g. over a private chain.
- *Input provider* or *writer*. This player is authorized to interact with a published contract (which implements a GF) and provide (garbled) input into the contract based on the plaintext input it has.
- (*Input*) *unlocker*. This player facilitates secure input provision by establishing an extra decryption step (performed by the unlocker) of the submitted garbled input. This prevents input providers, who possess *both* input labels on each input wire, from decrypting the internals of the encrypted computation. Effectively, use of the unlocker (who we assume does not collude with input providers) implements a secure OT of the input label based on the input.
- *Evaluator*. This player (implemented by the blockchain itself) evaluates the GF by executing the contract created by the contract creator on garbled inputs provided by the input providers. (By its nature, the blockchain also generates an indelible public archive of the contract’s execution, including garbled inputs and outputs.)
- *Output recipient* or *reader*. This player is authorized to receive the output of the computation. It is also possible to make the output available to all.

2.2 Approach

In our approach, the blockchain network itself plays the role of the *Evaluator* **Eval** of the GF (either a garbled circuit or a garbled FSA, in this work).

GF Generation and Contract Publishing. The computed function is first represented as a Boolean circuit or FSA. Then it is garbled within the BHR framework [10], resulting in a GF (e.g., GC or GFSA).

The GF is assumed to be *honestly* generated by an agent of a contract creator, **Gen**. We note that **Gen** possesses all secrets of the encrypted function and therefore is able to infer the internal state of the (plaintext) computation, should it ever gain access to the encrypted evaluation. Therefore we assume that all the secrets of the (small and self-contained) computation performed by **Gen** are *securely deleted*³. That is, we assume that **Gen** produces GF F , encoding information e and decoding information d . Upon delivery (as we discuss next) of e and d to the blockchain network players, and of F to the contract, **Gen** securely erases all its state (perhaps except F and d). We note that secure deletion of **Gen**’s state is not needed if audit may be desired or it is allowable for **Gen** to inspect the details of the evaluation, such as inputs, intermediate states, etc.

³ If we require auditability of MPC, this information must be securely stored instead of being deleted. Then, upon audit, the generated GC can be reconstructed, and its correctness and correctness of MPC execution verified. See Sect. 2.3 for details.

Input Provision. As plaintext input becomes available to input providers, they may enter the corresponding garbled input labels into the contract. To do this, in the GC case, they must have access to *both* garbled labels for each Boolean input. This would present a serious security problem if not addressed. A player who knows more than one label of a wire may infer unallowed plaintext information. In addition to passively learning private information, the attacker may adaptively substitute its input, thereby affecting the correctness of the computation as well.

We address this issue by introducing and using *unlockers*, logical players who help manage input labels. Thus, the process of input provision proceeds as follows (we specify it for the case of GC; the GFSA case is analogous):

1. **Gen** generates GF and corresponding input labels, w_i^0 and w_i^1 , representing two labels for each Boolean input wire W_i . **Gen** encrypts these labels with unlocker key k_u . For each input wire W_i , **Gen** gives the two encryptions $Enc_{k_u}(w_i^0)$, $Enc_{k_u}(w_i^1)$ to the input provider responsible for the wire W_i . **Gen** gives the unlocker key k_u to the unlocker associated with W_i .
2. When the input provider is ready to submit the (encrypted) input $b \in \{0, 1\}$ on wire W_i , it publishes to the contract $Enc_{k_u}(w_i^b)$, the encrypted label corresponding to its input b , received from **Gen**.
3. When notified (e.g., off-chain or by the blockchain, or in response to monitoring the blockchain), the unlocker retrieves $Enc_{k_u}(w_i^b)$ from the contract, decrypts it with the key k_u received from **Gen**, and publishes w_i^b to the contract.

Secure Evaluation and Output Delivery. Once all inputs are provided to the contract, the contract is evaluated by the blockchain and the (encrypted) output is produced. Anyone may inspect the encrypted output, and only authorized players (those who received d , or corresponding portions of d from **Gen**) may decrypt and obtain the plaintext output.

Reactive Functionalities. We stress that the computation need not be one-shot. It is natural to consider multi-staged evaluation, where intermediate outputs may be provided to output recipients, and function state propagated across the stages. This is easy to achieve with obvious variations of GF evaluation. One approach to this is illustrated in Fig. 1. We prove security only for one-shot functionalities. Proofs can be naturally extended to the reactive case.

2.3 Trust Model

After having described the players and their actions, we are now ready to specify the trust model. There are two main assumptions:

- We assume that the contract generator acts honestly and *securely erases* its state after completion of its task. Note, this is immediately achieved if garbler is implemented as MPC e.g., run on a private chain.
- Input providers *do not collude* with corresponding unlockers. That is, we allow arbitrary collusions of players, but a set of colluding parties may not include an input provider and an unlocker for the same wire/GFSA step.

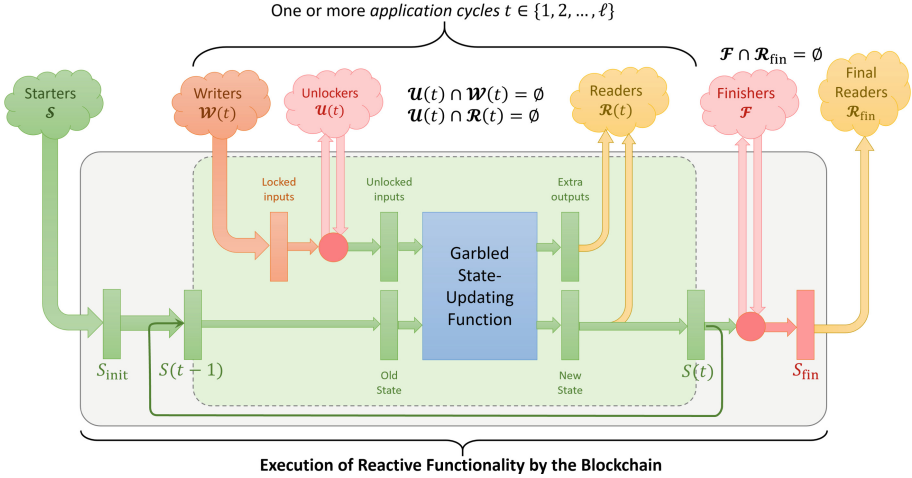


Fig. 1. Reactive execution of garbled machines by blockchain contracts. The gray region represents the full machine execution, which may require one or several contracts. The green sub-region represents operation within one *application cycle*, each of which may accept new inputs and produce new intermediate outputs. The blue rectangle represents a garbled state-update function that maps (old state, input) to (new state, output); this block can be implemented either using a (monolithic) garbled FSA transition table, or as a traditional GC (the latter requiring a *projective* or bit-vectorized state encoding). It must be garbled separately for each cycle. In this conception, \mathcal{S} denotes a set of players called *Starters* authorized to configure the initial state, and each cycle t may have its own sets of authorized input providers $\mathcal{W}(t)$, unlockers $\mathcal{U}(t)$, and output recipients $\mathcal{R}(t)$. The *Finishers* \mathcal{F} and *Final Readers* \mathcal{R}_{fin} are only required if there is a final output that is supposed to be visible to a broad audience including the other players, but where the other players may have a disincentive to reveal the specific output value to that audience. (Color figure online)

Security Against Cheating Gen: Audits, Covert [8] and PVC [7,24] Security. We assume that **Gen** behaves honestly and, further, erases its state. In some scenarios, it may be desired to open the computation at a later stage, e.g., for audit purposes. This can increase trust in **Gen** and the transparency of the process. Of course, the auditor (or the public, if the computation is opened to the public) will learn the inputs of all players. Release of this information may be acceptable, e.g., in situations where inputs are sensitive only for a certain duration of time.

Auditing of **Gen** is easily achieved by requiring **Gen** to generate everything from a PRG seed and to securely store the seed. During audit, the seed is revealed and the auditor verifies that all actions of **Gen** are consistent with the seed (this may require participation of unlockers and input providers). Because GC is secure against a malicious evaluator, honest generation of GC implies correctness and security against malicious players in the collusion model described above.

In the case when the function is public, we can also easily achieve covert [8] or even publicly verifiable covert (PVC) [7, 24] security. Following the ideas on [8], covert security can be achieved by requiring **Gen** to produce two GFs and sets of inputs; the blockchain network, e.g. via a randomness beacon, challenges to open one of them, verifies its correctness, and evaluates the unopened GF. PVC, a strengthening of the covert model introduced by [7], requires the ability to prove cheating, in case a cheater was caught. Because the GF and all inputs are published on the chain, it is easy to collect evidence of cheating. Firstly, we can require **Gen** to publish a seed (failure to do so will automatically imply guilt). Further, it is easy to verify that **Gen**'s actions are consistent with the seed and punish it (e.g. via funds slashing) if a violation is detected.

3 Generic Security Statement and Proof

We state the general security theorem for functions implemented as garbled functions and present the proof. The security of our specific construction presented next in Sect. 4 is an immediate corollary of this general theorem.

Let $\mathcal{G} = (\text{ev}, \text{Gb}, \text{En}, \text{Ev}, \text{De})$ be a garbling scheme, satisfying correctness and privacy as defined by [10] (as noted in Sect. 1.5, obliviousness and authenticity are not needed). We additionally require that the decoding information d is projective⁴, and decoding each bit calls a hash function, modeled as a *Random Oracle* (RO). Note, standard GC constructions in fact do implement d this way: output wire's plaintext value, for example, can be obtained by computing $\text{low_bit}[H(w_i)]$, where w_i is the output label⁵. Similarly, other garbling schemes, such as GFSA, can have a decoding function d incorporate a call to RO. We will use the RO programmability [20] in our simulation.

Theorem 1. *Let $\mathcal{G} = (\text{ev}, \text{Gb}, \text{En}, \text{Ev}, \text{De})$ be a garbling scheme as above. Let $(y_0, \dots, y_p) = f(x_0, \dots, x_q)$ be the function desired to be computed, such that each bit of the function output depends on all inputs⁶. Let **Gen** be the contract generator, IP_1, \dots, IP_n be the input providers, U_1, \dots, U_m be the unlockers, and R_1, \dots, R_ℓ be the output receivers. Assume **Gen** is honest and generates $(F, e, d) = \mathcal{G}.\text{Gb}$ and distributes (F, e, d) to players as described above. Let $I \subset \{IP_i, U_j, R_k\}$ be the set of*

⁴ As defined in [10], in a projective garbling scheme, the encoding information is represented as a list of tokens, one denoting 0, and one denoting 1, for each bit of the input; an encoding of a player's input is a collection of the tokens corresponding to its plaintext input. Similarly, for the output decoding, we say it is projective if the plaintext output is decoded bitwise in a similar manner.

⁵ To use $\text{low_bit}[H(w_i)]$ as the decoding function, **Gen** needs to ensure that $\text{low_bit}[H(w_i^b)] = b$. This is easy to do by choosing the output labels from corresponding domains. We stress that this is but one way of implementing d with these properties.

⁶ While some functions of interest do not meet this requirement, the functions we consider in this work will: indeed, universal circuit and FSA function outputs depend on all their inputs.

colluding malicious players, such that for no input wire W_i both its input provider and unlocker are in I .

Then blockchain evaluation of f which computes $\mathcal{G}.ev$ as described above, is secure against a malicious adversary corrupting I .

Proof. For lack of space, we present the proof in the full version of this paper [18].

Remark 1. If an unlocker U_j colludes with a reader R_k , together they can learn the output of the computation and abort based on it. This is not considered a vulnerability in the standard notion of simulation-based security. Note, if we wish to avoid such adaptive abort, we can require that no unlocker colludes with any reader.

4 Instantiations and Security Proofs

Construction 2 (UC GC-based). *Our main construction is the instantiation of the generic GF-based construction described above in Sect. 2 based on the following choice of underlying primitives/schemes:*

Let f be a function to be computed on the blockchain. Let C be a Universal Circuit computing f . Let \mathcal{G} be the classic Yao GC garbling scheme with point-and-permute [9] and projective decoding function as specified in assumptions of Theorem 1.

Having proven a generic security theorem (Theorem 1) for computing functions represented by arbitrary garbled functions, the proof of security of our main protocol, which is GC-based, is an immediate corollary of Theorem 1.

Theorem 2. *Assume all assumptions of Theorem 1 hold, including the collusion assumptions. Then Construction 2 is secure in the malicious model against collusions specified in Theorem 1.*

Proof. Proof is an immediate corollary of Theorem 1 and the fact that the underlying GC scheme used in 2 satisfies the required assumptions of Theorem 1. \square

Other Instantiations and Proofs are Analogous. In particular, GC-based instantiation is the same as UC GC with the exception of garbling the circuit C , and not necessarily a UC. A garbled FSA offers a reasonable performance for simple functions compared to UC GC. In an appendix of the full version [18] we cast a one-shot evaluation of GFSA as a GF in the [10] notation. A GFSA-based GF satisfying privacy and correctness can be used as a basis of our general construction.

5 Prototype Implementations and Test Results

To illustrate our approach and assess its real-world cost, we implemented a simple prototype and several demonstration applications. Specifically, an implementation of the GFSA approach (see full version [18]) was developed, for simplicity, and applied to several demo applications represented as finite state automata.




Fig. 2. Base FSA used in Millionaire’s Problem demo. In this version, players AB supply bits of their input values simultaneously, least-significant bit first, on successive cycles. In the final cycle, after L time steps have passed, a Finisher (as in Fig. 1) supplies a special “Finish” symbol \ominus which makes the final result readable by both parties.

The prototype Garbler, implemented in Python, takes a simple JSON-format description of an FSA transition function and translates it to a sequence of garbled tables, one for each state update cycle (time step). After garbling, another Python script translates the garbled machine data to source code in the Solidity programming language for a smart contract for the Ethereum platform; this contract includes the garbled tables as static data, together with a generic *Executor* which accepts garbled input values from input providers and evaluates the garbled machine, producing garbled outputs which can then be interpreted by authorized output recipients.

We used the popular Truffle tool suite, which provides a framework for Ethereum development, to develop, test and deploy (on a private test network, and later on the Ethereum mainnet) several prototypes and demonstrations, which we now discuss.

We implemented a provenance tracking demo, presented in detail in the full version [18]. Next we discuss our implementation of the millionaires problem.

5.1 Millionaires’ Problem Demo

This demo executed a 5-state machine (Fig. 2) implementing MPC for Yao’s classic 2-party “Millionaires’ Problem” [35] with bit-serial inputs. Before we had added Unlockers, to achieve MPC fairness (the last player to move possesses an informational advantage due to his ability to look ahead at final outputs), we invoked a special extra player “Finisher” (separate from the 2 normal parties) that acts to reveal the result. (See also Fig. 1).

Extending this line of argument, we observe that every step (input provision and corresponding GFSA state update) of the GFSA execution may exhibit the following similar vulnerability: the input provider may see the immediate effect of its input, such as whether the next FSA state depends on its input. This issue can be resolved by state-graph transformations, which increase the size of the state machine. In one version that we tested, each additional bit of input length cost almost exactly 2 million gas units to store the additional garbled machine data (since each time step has to be garbled separately), which was about \$0.50 worth of Ethereum on the day of that test. With some overhead, the total cost to run an FSA for a 32-bit, two-party Millionaire’s Problem was 75 million gas, corresponding to roughly US\$75 or so given typical prices at that time. That demo required spreading out the GFSA data over multiple smart contracts, due to Ethereum contract size limits.⁷ Reimplementing this same demo using Unlockers for each input and an optimized 2-state FSA allowed us to reduce the gas cost to ~\$12.

5.2 Configurable Garbled Universal Circuit (GUC) Method

To let us handle applications of a complexity beyond the reach of the GFSA approach in future implementations of GABLE, a simple approach was designed to implement a garbled circuit (GC) for (configurable) universal circuits (UCs). Figure 3 illustrates our basic UC approach. Input values here are activated using Unlockers (not shown), as we described earlier in the paper.

Although implementation of this method is still in progress, careful analysis of the approach allowed us to already compare its costs to those of the existing GFSA technique for an example problem, a multi-party auction (generalized from the Millionaire’s Problem). Figure 4 shows comparison results. As we expected, cost scales up exponentially with the number of bidders B for GFSA, but only as $\Theta(B \log^2 B)$ for the GUC. (Circuit width scales as $\Theta(B)$, circuit depth scales as $\Theta(\log B)$, and the depth of the Thompson network for each application circuit layer also scales as $\Theta(\log B)$.) The break-even point with our implementation falls at $B = 7$ bidders, where the cost of both approaches is roughly \$20 per input bit.

⁷ Per EIP-170 (<https://github.com/ethereum/EIPs/blob/master/EIPS/eip-170.md>), a contract’s deployed bytecode size cannot exceed 24,576 bytes.

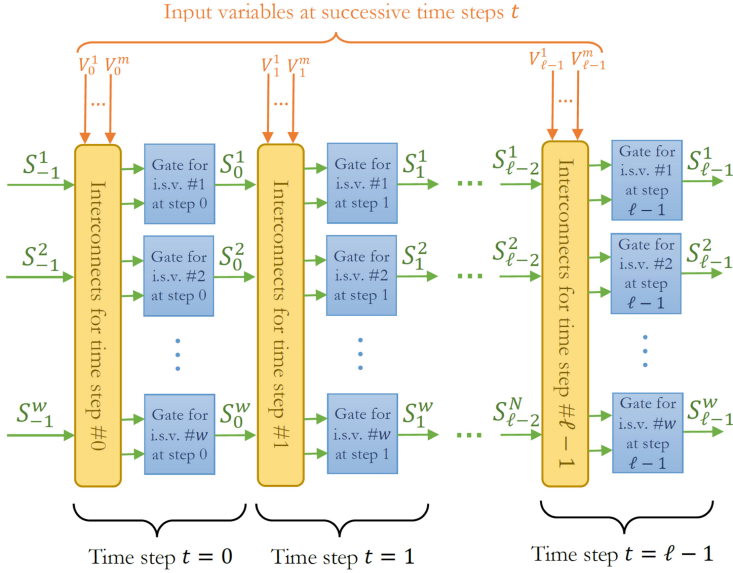


Fig. 3. Concept for configurable universal circuits in multi-step computations. In this approach, for each layer of application logic, a generalized connection network such as a Thompson network [30] obscures the interconnect topology. The elements of that network, together with generic application gates for computing new values of internal state variables (i.s.v.'s), are configured via truth tables during the garbling process. Thus, no separate programming input is required for this type of UC, yet the function of the network remains obscured.

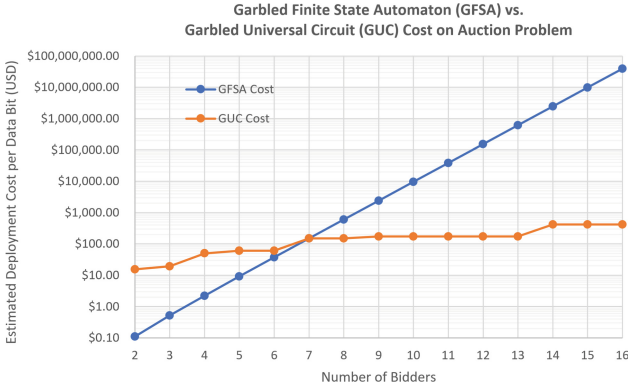


Fig. 4. Cost comparison for GFSA vs. configurable universal GCs for multi-party auctions. The break-even point occurs for $B = 7$ bidders, where the cost of both techniques is about 800 million gas per bit of input length. Prices here assumed optimistically that we are paying only 1 mETH (or in the ballpark of \$0.20) per million gas; however, after our tests, the average gas price grew substantially higher.

6 Conclusion

In this paper, we described a novel approach to performing secure computation (including functional privacy) on a blockchain. The general approach has two basic embodiments that we discuss, based on the garbling of finite-state automata (FSA) and Boolean circuits, respectively. We gave an overview of the basic structure of the approach, including its participant roles and high-level procedures, outlined a proof of its basic security properties, and discussed early implementations and test results.

We found that simple FSA-based applications can be executed privately at moderate dollar costs on the Ethereum blockchain. For more complex applications, a simple approach based on a construction we call configurable Garbled Universal Circuits (GUC) achieves complete functional privacy with costs that scale as $\Theta(wd \log w)$ in the width w and depth d of the application circuit, with reasonable constant factors. We carried out a detailed cost comparison for a multi-bidder auction application, for which GUC outperforms garbled FSA for $B > 7$ bidders, and remains arguably feasible to perform on-chain with full functional privacy for up to hundreds or even thousands of bidders.

We deployed and executed two GABLE demos on the Ethereum mainnet in late July and mid-September of 2020.⁸ The purpose of these tests was to 1) ensure that there were no unforeseen difficulties with real-world deployment, and 2) validate our cost estimation methodology. Both purposes were realized, with no surprises. The first deployment [2,3] (July) was a very simple four-state machine, similar to the supply-chain example discussed in the full version [18]. The second deployment ([1], Sep.) was for a GFSA implementation of a two-party auction as in Fig. 4.

References

1. Auction contract. Ethereum address `0x98ccd7e190ac28a36d4f065a4f14dc5e0b67f5c7`
2. Simple executor contract. Ethereum address `0xc8a54a72f187ec444ed08968901284bbd6d2ec06`
3. Simple storage contract. Ethereum address `0x57f1c190982d0a9ecdf7c4703e134d9eaf347de0`
4. The Secret Network. <https://scrt.network/>. Accessed 25 Jun 2020
5. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multi-party computations on bitcoin. In: 2014 IEEE Symposium on Security and Privacy, pp. 443–458. IEEE Computer Society Press (May 2014)
6. Araki, T., Furukawa, J., Lindell, Y., Nof, A., Ohara, K.: High-throughput semi-honest secure three-party computation with an honest majority. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016, pp. 805–817. ACM Press (October 2016)

⁸ Source code for these contracts is at <https://github.com/sandialabs/GABLE>.

7. Asharov, G., Orlandi, C.: Calling out cheaters: covert security with public verifiability. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 681–698. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34961-4_41
8. Aumann, Y., Lindell, Y.: Security against covert adversaries: efficient protocols for realistic adversaries. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 137–156. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70936-7_8
9. Beaver, D., Micali, S., Rogaway, P.: The round complexity of secure protocols (extended abstract). In: 22nd ACM STOC, pp. 503–513. ACM Press (May 1990)
10. Bellare, M., Hoang, V.T., Rogaway, P.: Foundations of garbled circuits. In: Yu, T., Danezis, G., Gligor, V.D. (eds.) ACM CCS 2012, pp. 784–796. ACM Press (October 2012)
11. Benhamouda, F., et al.: Can a public blockchain keep a secret? In: Pass, R., Pietrzak, K. (eds.) TCC 2020. LNCS, vol. 12550, pp. 260–290. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64375-1_10
12. Benhamouda, F., Halevi, S., Halevi, T.: Supporting private data on hyperledger fabric with secure multiparty computation. In: 2018 IEEE International Conference on Cloud Engineering (IC2E), pp. 357–363 (2018)
13. Bentov, I., Kumaresan, R.: How to use bitcoin to design fair protocols. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8617, pp. 421–439. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44381-1_24
14. Buterin, V.: Ethereum Whitepaper (2013). <http://ethereum.org/en/whitepaper>
15. Cecchetti, E., Zhang, F., Ji, Y., Kosba, A.E., Juels, A., Shi, E.: Solidus: confidential distributed ledger transactions via PVORM. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 701–717. ACM Press (October 2017)
16. Choudhuri, A.R., Goyal, V., Jain, A.: Founding secure computation on blockchains. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11477, pp. 351–380. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_13
17. Choudhuri, A.R., Green, M., Jain, A., Kaptchuk, G., Miers, I.: Fairness in an unfair world: fair multiparty computation from public bulletin boards. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 719–728. ACM Press (October 2017)
18. Cordi, C., et al.: Auditable, available and resilient private computation on the blockchain via MPC. Cryptology ePrint Archive, Report 2022/398 (2022). <https://eprint.iacr.org/2022/398>
19. Dwork, C., Naor, M.: Pricing via processing or combatting junk mail. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 139–147. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-48071-4_10
20. Fischlin, M., Lehmann, A., Ristenpart, T., Shrimpton, T., Stam, M., Tessaro, S.: Random Oracles with(out) programmability. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 303–320. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_18
21. Frank, M.P., et al.: The GABLE report: garbled autonomous bots leveraging Ethereum. Technical report SAND2020-5413, Sandia National Laboratories (2020). <https://www.osti.gov/biblio/1763537>
22. Gentry, C., et al.: YOSO: you only speak once. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12826, pp. 64–93. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84245-1_3
23. Haber, S., Stornetta, W.S.: How to time-stamp a digital document. J. Cryptol. 3(2), 99–111 (1991). <https://doi.org/10.1007/BF00196791>

24. Kolesnikov, V., Malozemoff, A.J.: Public verifiability in the covert model (almost) for free. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015. LNCS, vol. 9453, pp. 210–235. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48800-3_9
25. Kolesnikov, V., Schneider, T.: A practical universal circuit construction and secure evaluation of private functions. In: Tsudik, G. (ed.) FC 2008. LNCS, vol. 5143, pp. 83–97. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85230-8_7
26. Kosba, A.E., Miller, A., Shi, E., Wen, Z., Papamanthou, C.: Hawk: the blockchain model of cryptography and privacy-preserving smart contracts. In: 2016 IEEE Symposium on Security and Privacy, pp. 839–858. IEEE Computer Society Press (May 2016)
27. Malkhi, D., Nisan, N., Pinkas, B., Sella, Y.: Fairplay - secure two-party computation system. In: Blaze, M. (ed.) USENIX Security 2004, pp. 287–302. USENIX Association (August 2004)
28. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2008). <https://bitcoin.org/bitcoin.pdf>. Accessed 25 Jun 2020
29. Szabo, N.: Secure property titles with owner authority (1998). <https://nakamoinstitute.org/secure-property-titles/>. Accessed 25 Jun 2020
30. Thompson, C.: Generalized connection networks for parallel processor intercommunication. Technical report, Carnegie-Mellon University, Pittsburgh, PA (1977)
31. van Wirdum, A.: “Confidential assets” brings privacy to all blockchain assets: Blockstream. Bitcoin Magazine (April 2017). Accessed 25 Jun 2020
32. Wikipedia. Transactive Energy. https://en.wikipedia.org/wiki/Transactive_energy
33. Yang, A., Wei, L., Wu, J., Long, C.: Block-SMPC: a blockchain-based secure multi-party computation for privacy-protected data sharing. In: Proceedings of the 2020 the 2nd International Conference on Blockchain Technology, ICBCT 2020, pp. 46–51. Association for Computing Machinery, New York (2020)
34. Yao, A.: How to generate and exchange secrets. In: Proceedings of the 27th Annual Symposium on Foundations of Computer Science, pp. 162–167. IEEE (1986)
35. Yao, A.C.: Protocols for secure computations. In: 23rd Annual Symposium on Foundations of Computer Science, SFCS 1982, Chicago, IL, USA, pp. 160–164 (1982). <https://doi.org/10.1109/SFCS.1982.38>
36. Zahur, S., Rosulek, M., Evans, D.: Two halves make a whole. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 220–250. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_8
37. Zyskind, G., Nathan, O., Pentland, A.: Decentralizing privacy: using blockchain to protect personal data. In: IEEE Symposium on Security and Privacy Workshops, pp. 180–184 (2015)