Tri-State Circuits
A Circuit Model that Captures RAM

David Heath[1(✉)], Vladimir Kolesnikov[2], and Rafail Ostrovsky[3]

[1] UIUC, Champaign, USA
daheath@illinois.edu
[2] Georgia Tech, Atlanta, Georgia
kolesnikov@gatech.edu
[3] UCLA, Los Angeles, USA
rafail@cs.ucla.edu

Abstract. We introduce *tri-state circuits* (TSCs). TSCs form a natural model of computation that, to our knowledge, has not been considered by theorists. The model captures a surprising combination of simplicity and power. TSCs are simple in that they allow only three wire values ($0, 1$, and undefined – $\mathcal{Z}$) and three types of fan-in two gates; they are powerful in that their statically placed gates fire (execute) eagerly as their inputs become defined, implying orders of execution that depend on input. This behavior is sufficient to efficiently evaluate RAM programs.

We construct a TSC that emulates T steps of any RAM program and that has only $O(T \cdot \log^3 T \cdot \log \log T)$ gates. Contrast this with the reduction from RAM to Boolean circuits, where the best approach scans all of memory on each access, incurring quadratic cost.

We connect TSCs with Garbled Circuits (GC). TSCs capture the power of garbling far better than Boolean Circuits, offering a more expressive model of computation that leaves per-gate cost essentially unchanged.

As an important application, we construct *authenticated Garbled* RAM (GRAM), enabling constant-round maliciously-secure 2PC of RAM programs. Let λ denote the security parameter. We extend authenticated garbling to TSCs; by simply plugging in our TSC-based RAM, we obtain authenticated GRAM running at cost $O(T \cdot \log^3 T \cdot \log \log T \cdot \lambda)$, outperforming all prior work, including prior semi-honest GRAM.

We also give semi-honest garbling of TSCs from a one-way function (OWF). This yields OWF-based GRAM at cost $O(T \cdot \log^3 T \cdot \log \log T \cdot \lambda)$, outperforming the best prior OWF-based GRAM by more than factor λ.

Keywords: Garbled RAM · MPC · Models of Computation · Malicious Security

1 Introduction

Boolean circuits form perhaps our simplest complete model of computation. The model allows only a small set of gate types, each of which computes a basic

H. Handschuh and A. Lysyanskaya (Eds.): CRYPTO 2023, LNCS 14084, pp. 128–160, 2023.
https://doi.org/10.1007/978-3-031-38551-3_5

function. Moreover, a circuit's structure is static and explicit. This simplicity is ideal for both theory and practice, making them popular in complexity theory and, in particular, in cryptography.

On the other hand, random access machine programs (RAM programs)[1] form our most ubiquitous practical model of computation. The random access capability approximates the power of real-world devices, so theoretical advances in RAM can more readily translate to real-world impact.

Unfortunately, it is difficult to connect Boolean circuits and RAM. Indeed, the two models seem inherently at odds. RAMs are inherently *dynamic*, allowing the program to quickly and arbitrarily access one element in an immense array; circuits are inherently *static*, requiring that the program fix the order in which it manipulates data before input is known.

It is therefore unsurprising that reductions from RAMs to Boolean gates are expensive. The straightforward reduction emulates each memory access by *linearly scanning the entire RAM memory.* This simple approach is also the *best known.* Since RAMs access memory at each step, this reduction yields a circuit that grows quadratically in the RAM runtime T.

Motivation for Introducing Tri-State Circuits: Constant Round 2PC. It is unfortunate that reductions from RAMs to circuits are so expensive. Many technologies are more compatible with circuits than with RAMs, and a concretely efficient reduction would automatically connect real-world RAM programs with circuit-based technologies.

As our crucial example, we consider Yao's Garbled Circuit (GC) [Yao86], a multiparty computation (MPC) technology that achieves symmetric-key-based constant-round protocols.

The GC literature is extensive, see [NPS99, ZRE15, HJO+16, GLNP18, RR21] and many more. Most GC works, including all listed above, garble Boolean gates only, suggesting a natural connection between garbling and circuits. On the other hand, the goal of GC is to enable secure computation of arbitrary programs, and many programs are best handled by RAMs, not by circuits.

It is possible – though challenging – to garble RAM programs. *Garbled RAM* (GRAM) [LO13] does so *without* reducing to circuits. GRAM also has a rich literature [GHL+14, GLO15, GLOS15, CH16, CCHR16, LO17, HKO22, PLS22]. The basic observation of GRAM is that it is possible to garble interconnected *collections* of circuits that execute in an order decided *at runtime.* This dynamic ordering breaks from the circuit model, where the order of execution is static.

Advancing GRAM was challenging. The problem was that reasoning about GRAM required reasoning simultaneously about multiple complex topics, including gate garbling techniques, the above dynamic circuit execution, and constructions of *Oblivious RAM* [GO96]. Worse still, the community lacked an effective vocabulary for discussing the dynamic mechanisms of GRAM; prior work

[1] The RAM model we consider is called the *word RAM model* [Hag98]; it is a RAM with a fixed word size that incurs unit cost per random access. We state the definition of the in Sect. 3. Throughout this work, by 'RAM' we mean 'word RAM'.

explained these mechanisms via a concept they called *dynamic language translation*, see e.g. [GLO15,GLOS15,HKO22,PLS22]. Language translation is *deeply intertwined* with the specifics of circuit garbling, and thus understanding even the *basic* ideas of GRAM required intimate GC knowledge.

Matching the Power of Garbling to a Model of Computation. The mere existence of non-trivial GRAM, which leverages dynamic behavior, demonstrates that the circuit model poorly approximates the power of garbling. Clearly some additional expressive power is available, sufficient to efficiently execute RAM programs.

Thus, it is interesting to search for another model of computation – cheap to garble, simpler than RAM, and more expressive than Boolean circuits – that captures the dynamic power of garbling. Such a model would be useful, since it would decompose the GRAM problem into pieces, allowing us to think modularly about RAM constructions, untethered from garbling-specific concerns.

1.1 Our Contribution

We demonstrate that there exists a simple, circuit-like model of computation that closely approximates (within a polylogarithmic factor) the expressive power of RAM. Our *tri-state circuit* (TSC) model is strongly compatible with our target use-case of garbling, in the sense that it admits efficient and natural protocols.

Like a Boolean circuit, a TSC is composed from statically connected gates, each of which has one of only a small number (three) of possible gate types. Each gate computes a basic function of its two tri-value input wires. Despite their simplicity, TSCs are distinctly more powerful than Boolean circuits; they admit a primitive form of control flow where the order in which gates fire *depends on the input*. This basic control flow can *efficiently emulate RAM computation.* Thus, TSCs capture a surprising combination of conceptual simplicity and expressive power which, to our knowledge, has not been explored by theorists.[2,3]

We emphasize that while we feel the tri-state circuit model has intrinsic value, we are motivated by the concrete objective of improving symmetric-key-based constant-round secure computation (i.e., garbling).

Our contributions include:

[2] While tri-state circuits have not been theoretically explored, tri-state gates are used in practice. We chose our naming based on these real-world gates. A key gate in our model, which we call a *buffer*, exists as a digital logic element called a *tri-state buffer*. We show that RAM reduces to a relatively small number of such gates.

[3] Tri-state circuits are distinct from *ternary logic*. Ternary logic *has* been explored by theorists, even in the context of garbling [LY18,NPS99]. In ternary logic, wires can take three distinct values; however, the circuit executes in a standard topological order. In *tri-state circuits*, gates execute in a data-dependent order.

We formalize the tri-state circuit model.

We reduce RAM programs to (deterministic) tri-state circuits. Let T denote a runtime. We construct a tri-state circuit of size $O(T \cdot \log^4 T)$ that can emulate T steps of any RAM program.

We formalize randomized and *oblivious* tri-state circuits. In cryptography, data-independent orders of execution are useful for protecting privacy. *Basic* tri-state circuits discard input independence, losing cryptographic utility. *Oblivious* tri-state circuits reclaim this utility. A tri-state circuit is oblivious if its order of execution *appears* (to a distinguisher) independent of the input.

We reduce RAM programs to oblivious tri-state circuits. We construct an oblivious tri-state circuit of size $O(T \cdot \log^3 T \cdot \log \log T)$ that can simulate T steps of any RAM program. This oblivious reduction improves over our deterministic reduction by leveraging randomness.

We apply tri-state circuits to secure 2PC. Let λ denote the computational security parameter. We achieve two results:

- Our most exciting application is *authenticated GRAM*, a maliciously-secure constant-round 2PC RAM protocol. Our authenticated GRAM executes T RAM steps at cost $O(T \cdot \log^3 T \cdot \log \log T \cdot \lambda)$. Prior malicious GRAM relied on the expensive cut and choose technique, and was more than factor σ slower, for statistical security parameter σ.
- Our second application is improved semi-honest Garbled RAM from only one-way functions. Prior to our work, the best GRAMs were based on random-oracle-like assumptions [HKO22,PLS22]. The best GRAM avoiding such an assumption had quadratic scaling in λ [PLS22]. Our construction outperforms all prior RO-based GRAMs, and it relies only on one-way functions. It runs at cost $O(T \cdot \log^3 T \cdot \log \log T \cdot \lambda)$.

TSC garbling is lean. For example, Boolean circuits can be compiled to tri-state gates, and the communication cost of our resulting authenticated TSC protocol is less than $2\times$ that of state-of-the-art authenticated garbling of Boolean gates [DILO22], and with effort this overhead can likely be removed.

Impact on Garbled RAM. While [HKO22] and follow-on work [PLS22] substantially improved GRAM, these works left pressing and challenging open questions, including efficient malicious GRAM and standard-assumption-based GRAM.

We abstract garbled computation as tri-state circuits, not Boolean circuits, and the payoff is a modular approach to GRAM. This modularity allows us to make significant advances that would have been highly technically involved if expressed in the prior GRAM framework of language translation. We demonstrate that the above more challenging versions of GRAM can be constructed with overhead similar to basic GRAM. Going further, we discovered compatibility between a state-of-the-art Oblivious RAM construction [WCS15] and tri-state circuits, further improving GRAM's asymptotic cost.

Perhaps best of all, tri-state circuits markedly simplify GRAM fundamentals. Indeed, our new garbling procedures are extremely similar in complexity to their

Boolean-circuit-based counterparts. This reduced complexity will allow a broader cryptographic audience to understand, improve, and apply GRAM.

2 Background and Related Work

2.1 Garbled Circuits and Garbled RAM

Garbled Circuit (GC) [Yao86] is a fundamental MPC primitive that allows two parties – a garbler G and an evaluator E – to securely execute a program of their choice on their private inputs. GC is distinct from other secure computation primitives in that it allows for protocols that (1) run in a constant number of rounds and (2) rely almost entirely on fast symmetric key primitives.

Roughly speaking, GC splits program execution into two steps: garbling and evaluation. Garbling is independent of the input, and evaluation of the garbled program *appears* independent of the input. When these two steps are carried out by two different parties, we can arrange that each party's execution hides the input of the other, allowing privacy-preserving protocols.

While GC traditionally works in the Boolean circuit model, a number of works starting with [LO13] developed Garbled RAM (GRAM), an extension to the more expressive RAM model. [LO13] demonstrates that for a word-RAM program running in time T and for computational security parameter λ, we can garble the program at the following cost:

$$O(T \cdot \log^3 T \cdot \log^c \log T \cdot |\mathcal{C}_{prf}| \cdot \lambda) \qquad \text{[LO13]}$$

Here, c is an unspecified constant, and $|\mathcal{C}_{prf}|$ is the circuit size of a PRF with λ bits of output (asymptotic analysis by [PLS22]).

A sequence of works subsequently improved the Garbled RAM primitive, e.g. [GHL+14, GLOS15, GLO15, HKO22, PLS22]. The most recent garbled RAM constructions achieve the following costs:

$$O(T \cdot \log^4 T \cdot \lambda) \qquad \text{[HKO22]}$$
$$O(T \cdot \log^3 T \cdot (\log \log T)^2 \cdot \lambda) \qquad \text{[PLS22]}$$

[HKO22] and follow-on work [PLS22] far surpass prior GRAMs, bringing the technique's overhead in line with what is expected from more traditional Boolean-circuit-based GC.

Improving GRAM remains a crucial direction. In particular, it is interesting to (1) improve asymptotic cost, (2) extend GRAM to interesting and challenging settings, and (3) simplify the GRAM formalism, easing further exploration and application. Our work simultaneously achieves each of these goals.

Malicious GRAM. GC provides natural protection against *malicious* evaluator E, but protecting against *malicious* garbler G is more challenging. G can incorrectly garble the program, causing the program to, for instance, erroneously output bits of E's input. It is difficult to arrange that E can detect incorrectly

garbled programs, because E's inability to reason about garbled programs is exactly the property that protects G's input.

Despite this challenge, prior work developed powerful techniques for efficient handling of malicious garbled *circuits* (see later discussion of authenticated garbling). Until this work, malicious garbled *RAM* was far less effective.

Prior work, e.g. [GGMP16, HY16, Mia20], demonstrated feasibility of malicious GRAM, but performance was poor, especially as compared to semi-honest GRAM. The best prior GRAM could be constructed by combining semi-honest GRAM [HKO22] (or the asymptotically more efficient [PLS22], framed as a garbling scheme [BHR12]) with the classic *cut and choose* technique, see e.g. [Lin13].

Cut and choose upgrades semi-honest garbling to the malicious setting. The idea is to have G garble many copies of the same program, then allow E to challenge a randomly selected subset of those programs. While this works, G must garble a number of copies that grows with the statistical security parameter σ, leading to highly undesirable factor σ slowdown as compared to the semi-honest execution. The best malicious GRAM had the following asymptotic cost:

$$O(T \cdot \log^3 T \cdot (\log \log T)^2 \cdot \lambda \cdot \sigma) \qquad \text{[PLS22] with Cut and Choose}$$

We avoid this factor σ slowdown by implementing tri-state circuits via the techniques of *authenticated garbling* (see next). Our maliciously secure GRAM dramatically improves over prior state of the art, achieving the following cost:

$$O(T \cdot \log^3 T \cdot \log \log T \cdot \lambda) \qquad \text{Our maliciously secure GRAM}$$

Authenticated Garbling. The breakthrough work [WRK17] introduced a far superior approach to malicious GC. Their *authenticated garbling* technique achieves performance that asymptotically matches semi-honest garbling, incurring only $O(n \cdot \lambda)$ cost for an n-gate circuit. The approach is also practically performant.

In classic GC, each wire value is represented by a length-λ *label*. These labels are used as keys to encrypt/decrypt subsequent labels in a way that achieves the program semantics. *Authenticated* GC extends each GC label with an additional σ bits, forming a MAC on the wire value. These MACs allow G to reveal particular wire values to E such that E is confident the value is indeed correct.

To securely evaluate each AND gate, the parties require an *authenticated multiplication triple*. Multiplication triples can be computed offline in a function-independent preprocessing phase. Improving the efficiency of authenticated garbling is the subject of a growing body of works [KRRW18, YWZ20, DILO22].

We demonstrate natural compatibility between authenticated garbling and tri-state circuits. Our construction achieves cost $O(n \cdot \lambda)$ for an n-gate *tri-state* circuit. While formal treatment of *any* non-trivial malicious technique is complex, authenticated garbling of tri-state circuits is – at least at an intuitive level – a straightforward extension of the core ideas given by the above prior works.

Standard-Assumption-Based GRAM. In the semi-honest setting, the fastest garbling techniques rely on a non-standard random-oracle-like assumption called a

circular correlation robust hash (CCRH) function [CKKZ12]. This assumption stems from the classic "Free XOR" extension [KS08] whereby each GC wire has two labels related by a global correlation. The CCRH assumption is needed to achieve security in the presence of this correlation.

It is interesting to remove this assumption and to garble assuming only one-way functions (OWFs) [GLNP18].[4] Prior to our work, Garbled *RAM* from one-way functions was *far* inferior to GRAM based on Free XOR. Indeed, the best construction had the following cost (note the problematic scaling in λ):

$$O(T \cdot \log^3 T \cdot (\log \log T)^2 \cdot \lambda^2) \qquad \text{[PLS22]}$$

We demonstrate that classic OWF-based techniques from the literature can be almost directly applied to tri-state circuits. Indeed, our standard-assumption-based garbling scheme is relatively obvious, once the tri-state circuit model is understood. Applying this scheme in conjunction with our RAM constructions, we *as a corollary* achieve the best standard-assumption-based GRAM:

$$O(T \cdot \log^3 T \cdot \log \log T \cdot \lambda) \qquad \text{Our OWF-based semi-honest GRAM}$$

2.2 Oblivious RAM

Oblivious RAM (or ORAM, [GO96]) is a powerful technology that allows a weak client to outsource its database to a powerful untrusted server. The client can repeatedly query its sensitive database without the server learning what data is accessed, or even the pattern in which data elements are accessed. In ORAM, for each *logical* access, the client issues a sequence of queries to *physical* locations. These physical locations reveal nothing about the logical accesses, which can be formalized by showing that the server's view can be simulated.

ORAM is highly relevant to our notion of oblivious tri-state circuits. In particular, our reduction from RAM programs to oblivious tri-state circuits directly leverages the Circuit Oblivious RAM construction of [WCS15], implementing their ORAM algorithms via tri-state gates. Our reduction leverages this construction to hide memory access patterns, allowing for a circuit that executes RAM programs and whose gates execute in an order that can be simulated.

2.3 Other Models of Computation

The tri-state circuit model shows that, surprisingly, there exists a concrete set of gates that can efficiently (with polylog overhead) implement RAM. Said another

[4] Of course, full semi-honest GC protocols also use OT, which is not implied by OWFs. It is now traditional to view semi-honest GC as a *primitive*, independent of any particular protocol [BHR12]. This primitive, called a garbling scheme, can be meaningfully instantiated from OWFs alone.

way, tri-state circuits admit a small statically defined structure whose collection of use-once components jointly implement RAM. This capability distinguishes the model from other widely considered models.

Other models either *inefficiently* support RAM (e.g., Turing Machines, decision trees, Boolean circuits, arithmetic circuits, etc.), or have implicitly specified "static structure" that is either large or involves components that can be used repeatedly. For instance, while RAM can, of course, efficiently implement itself, it in some sense involves a very large static structure, where each RAM step is implicitly connected to each memory cell. Similarly *pointer machines* form a model of computation that proceeds by editing a directed graph, see e.g. [Sch80]; because of the large number of possible graphs that can emerge at runtime, pointer machines similarly have large implicit static structure.

Said yet another way, tri-state circuits require that we statically define a fixed "stage" that establishes explicit connections between computational elements and explicitly named memory cells. Runtime execution may only proceed within the connection constraints of this stage, and we measure cost in terms of the size of the stage, namely, the number of connections. Indeed, in GC, the garbler must account for each possible action and state of the evaluator. This accounting corresponds to generation of garbled tables – *garbling the stage*.

Despite these constraints, the model is expressive. It has sufficient freedom to (obliviously) implement RAM. This expressiveness in the presence of GC-compatible constraints is what makes tri-state circuits so useful in garbling.

While our focus is on secure computation and garbling, we envision that the tri-state circuit model may be interesting in other settings as well. For instance, it is intriguing that our tri-state circuit constructions can – at least in principle – be implemented via digital circuits, and the model may also have interesting connections to complexity theory.

3 Notation

Word RAM Model. In the word RAM model [Hag98], an abstract machine operates on length-w *words*. Basic operations, such as addition, comparisons, and, in particular, memory reads/writes are assumed to take constant time.

Let T denote RAM program runtime. We assume w is large enough to point into the program input (i.e., $w \geq \log_2 n$) and, for simplicity of analysis, we assume $w = \Theta(\log T)$. We assume that each non-memory-accessing instruction can be implemented by a Boolean circuit of size $O(w^2) = O(\log^2 T)$, sufficient to capture powerful operations such as multiplication. Throughout this work, we refer to word RAMs as RAMs.

Common Notation. $x \mid\mid y$ denotes the concatenation of strings x and y. We denote by $\langle x \rangle$ a Boolean encoding of the value x. E.g., if P is a RAM program, then $\langle P \rangle$ denotes a Boolean encoding of that RAM program. We leave the details of such encodings unspecified, as they are not interesting. σ denotes a statistical security parameter (e.g. 40 or 60). λ denotes a computational security parameter (e.g. 128). $X \stackrel{s}{=} Y$ denotes that distributions X and Y are *statistically close*.

$\oplus$	0	1	$\mathcal{Z}$
0	0	1	$\mathcal{Z}$
1	1	0	$\mathcal{Z}$
$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$

/	0	1	$\mathcal{Z}$
0	$\mathcal{Z}$	0	$\mathcal{Z}$
1	$\mathcal{Z}$	1	$\mathcal{Z}$
$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$

$\bowtie$	0	1	$\mathcal{Z}$
0	0	$\perp$	0
1	$\perp$	1	1
$\mathcal{Z}$	0	1	$\mathcal{Z}$

```
1  procedure notify(gate):
2    (f, in_0, in_1, out) ← gate
3    x ← wires[in_0]
4    y ← wires[in_1]
5    z ← f(x, y)
6    if z ≠ Z and wires[out] = Z:
7      wires[out] ← z
8      for gate' ∈ subscribers(out):
9        notify(gate')
```

Fig. 1. The semantics of tri-state circuits. Tri-state circuits have three types of fan-in two gates: XORs ($\oplus$), buffers (/), and joins ($\bowtie$). We define the function of each gate type (left), and we define a recursive procedure *notify* (right) which defines semantics. The array *wires* is a global object that stores the value of each wire. Each wire can take on three different values: 0, 1, or $\mathcal{Z}$. $\mathcal{Z}$ indicates that a wire has not yet been assigned. At initialization, all non-input wires are set to $\mathcal{Z}$. The function *subscribers* maps from wire ID w to the set of gates that take wire w as input. Circuit execution begins by calling *notify* on each gate subscribed to an input wire. The symbol $\perp$ denotes an illegal state; if any join evaluates to $\perp$, we set all wires to $\perp$, execution terminates, and the circuit outputs $\perp$.

Tri-State Notation. Section 4 introduces the following; we catalog for reference.

Based on notation from digital circuits, $\mathcal{Z}$ denotes the distinguished tri-state *high impedance* value. $\mathcal{Z}$ can be pronounced "nil", and can be informally understood as the value of a wire that is not yet defined. A wire carrying 0 or 1 is **set**; a wire carrying $\mathcal{Z}$ is **not set**. We denote buffers by division[5] (written x/y or $\frac{x}{y}$), and joins by $\bowtie$. The second argument to each buffer is called its *control*. The symbol $\oplus$ denotes XOR. XOR is extended to tri-state values in a natural manner. Namely, if either XOR input is $\mathcal{Z}$, then the output is $\mathcal{Z}$ (see Fig. 1).

4 Tri-State Circuits

This section describes and formalizes the tri-state circuit model. Sections 5 and 6 later shows that tri-state circuits can efficiently implement RAM programs.

Tri-state circuits center on a non-Boolean gate that we call a *buffer*. A buffer takes two inputs, a *control wire* and a *data wire*:

[5] We chose division to denote buffers because buffer semantics produce the 'undefined' value $\mathcal{Z}$ when dividing by 0.

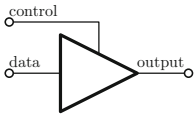

If the control is set to 1, then the output wire acquires the value of the data wire; if the control is set to 0, then the output wire *remains unassigned*, which we denote by stating the output wire has value $\mathcal{Z}$.[6] If the control is set to 1, we say that the buffer is **active** and that the output is **set**; else the buffer is **inactive** and the output is not set.

Because a buffer might not set its output, it is possible to implement interesting circuit arrangements, such as the following:

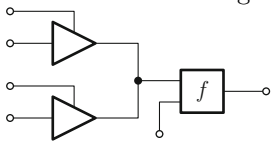

Here, we connect the outputs of two buffers, denoted by the black circle which we formalize as a gate that we refer to as a *join*. The join polls its two inputs, forwarding an input to its output as soon as some input is set. We connect the join's output to a subcircuit labelled f.

The crucial point is that the two buffers might be far apart in the circuit topology. Subcircuit f *eagerly fires* as soon as its inputs are set. Since the buffers fire at different times, the time at which f fires *depends on wire values*, not just the topology. This input-dependent order of execution is the key ingredient of tri-state circuits and is what distinguishes them from Boolean circuits.

Definition 1 (Tri-state Circuit). *A **tri-state circuit** is a circuit allowing cycles (i.e., its graph need not be acyclic) with three gate types: XORs, buffers, and joins. Each tri-state wire carries one of three values:* 0, 1, *or* $\mathcal{Z}$. *The semantics of each gate type and of circuit execution are formally specified in Fig. 1. Tri-state circuits may use two distinguished wires, named* 0 *and* 1, *which respectively carry the corresponding constants* 0 *and* 1.

Looking forward, we will consider constrained classes of tri-state circuits satisfying (combinations of) additional properties (see Definitions 4, 5 and 8).

The dynamic nature of tri-state circuits is formalized by *notify* (Fig. 1). When a wire is set to 0 or to 1 – i.e., when it is *not* $\mathcal{Z}$ – each gate **subscribed** to that wire (each gate taking the wire as input) is notified and fires.

At initialization, each non-input wire holds $\mathcal{Z}$. As gates fire, wire values change from $\mathcal{Z}$ to 0 or 1. Once set, a wire value cannot change again. Thus, the state of the wires converges to a final configuration, the **halt-time state**.

Definition 2 (Halt-time state). *The **halt-time state** of a tri-state circuit* $\mathcal{C}$ *is a wiring* w *(i.e., a map from circuit wires to wire values) such that there is no gate* $g \in \mathcal{C}$ *where notify*(g) *changes* w.

[6] We use $\mathcal{Z}$, pronounced 'nil,' to denote 'no signal'. In digital circuits, the 'no signal' value is called *high impedance*, and is denoted 'hi-Z'. In GC, $\mathcal{Z}$ on a wire corresponds to E holding *no* key on that wire; see Sect. 7.

A gate only notifies its subscribers if it sets its output. This, combined with the fact that each gate has only two inputs, means that each gate is notified at most twice, tightly bounding the total runtime. I.e., it is a straightforward fact that a random access machine (e.g., a computer evaluating the TSC) can emulate a size-n tri-state circuit in time $O(n)$ by simply running *notify*.

The halt-time state of a tri-state circuit $\mathcal{C}$ is *unique*, even when we allow calls to *notify* to occur in an arbitrary order. Indeed, in Appendix A of the full version of this paper[7] we prove the following:

Lemma 1 (Halt-Time State Unique). *Let $\mathcal{C}$ be a TSC that, on input x and for some sequence of calls to notify, reaches a halt-time state w.* Any *sequence of calls to notify reaching a halt-time state will reach the same state w.*

Circuits with Cycles. Definition 1 explicitly allows circuit graphs with cycles. Indeed, cycles seem to be *essential* for implementing efficient RAM with TSCs.

Consider two executions of a RAM program. In the first execution, suppose we first access some index i, then we access some index j; in the second execution, suppose we first access index j, then index i. Ideally, we would save indexes i and j on particular collections of wires such that the two executions read the same two collections of wires, just in different orders. To achieve this, we *must* admit cycles in our circuits: there is a possible data path from the i wires to the j wires, and from the j wires to the i wires.

There is no inherent inconsistency in allowing circuits with cycles, so long as we are careful in our circuit designs. Namely, tri-state circuits are allowed to have cycles, but their runtime data paths *are not*. Consider the following example:

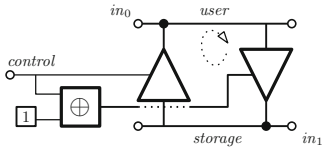

Here, we connect a wire named *user* to a wire named *storage*, allowing *user* to read from/write to *storage*. At first glance, the circuit appears to allow *user* to *write to itself*, a potentially problematic arrangement (especially when proving GC security). On closer inspection, it becomes clear that the wire *control* statically rules out this possibility: at most one buffer can activate, so there is no way for *user* to write to itself. This circuit has a cycle, but there is no possible cycle in the runtime data paths through the circuit.

We rule out runtime cycles by considering circuits that are *runtime acyclic*:

Definition 3 (Runtime Dependency). *A tri-state gate g is* ***runtime dependent*** *on another gate g' with respect to a circuit input x if:*

- *g is an XOR, join, or a buffer with control* 1*, and g is subscribed to the output wire of g'.*
- *g is a buffer with control* 0 *or $\mathcal{Z}$ (at halt-time), and g is subscribed to g' w.r.t. g's control wire (i.e., g' outputs the control of g).*

[7] https://eprint.iacr.org/2023/455.

We explicitly emphasize that a buffer with control 0 *or* $\mathcal{Z}$ *(at halt-time)* is not *runtime dependent on the gate that outputs its data wire.*

Definition 4 (Runtime Acyclic). *A tri-state circuit* $\mathcal{C}$ *is* ***runtime acyclic*** *if for all inputs* x*, there exists a winning strategy to a graph pebbling game with the following rules:*

- *The player is allowed to place a pebble on each circuit input.*
- *The player is allowed to place a pebble on a gate* g *iff there is a pebble on each of* g*'s runtime dependencies with respect to* x *(Definition 3).*
- *The player wins if it successfully places a pebble on each gate.*

Roughly speaking, Definition 4 states that for any input, there is no data cycle; if there were, then it would be impossible to win the pebbling game, since pebbling a gate requires first pebbling each of that gate's runtime dependencies.[8] Note, the above example circuit *is* runtime acyclic. Indeed, if $control = 0$, then the left buffer is inactive, and we can pebble the cycle by first pebbling this left buffer; if instead $control = 1$, then we can first pebble the right buffer.

For the rest of this work, we only consider tri-state circuits that are runtime acyclic, and our formal security theorems (see Appendices C and D of the full version) require runtime acyclicity.

Subcircuit Sharing. Because tri-state gates run dynamically, we can arrange a trick that we call **subcircuit sharing**. Consider the following circuit:

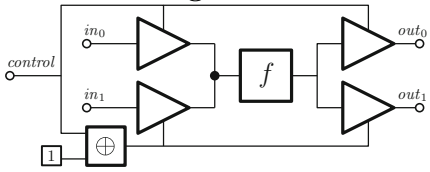

For sake of argument, suppose subcircuit f is composed from a large number of gates. Our example allows f to be called in two different ways: we can either set $control = 1$, running circuit f on **input port** in_0 and setting **output port** out_0, or we can symmetrically set $control = 0$, running f on input port in_1 and setting output port out_1. Since we can only set *control* to either 0 or 1, we can only activate *one* of the pairs of buffers, and so f is used only once. We again emphasize that the time at which f fires depends on *control*: in_0 and in_1 might each be set by the calling circuit at an arbitrary time.

Thus, f can be used in a *conditional manner*, solving a subproblem at one of two very different points in time. Crucially, our example is *efficient* in the sense that it contains only enough gates to implement f *once*; the gates in f are *shared* across the two call sites.

[8] One might wish to consider simpler definitions of runtime acyclicity, such as removing inactive buffers from the circuit, then requiring that the remaining graph is acyclic. Unfortunately, our attempts at such a definition admitted circuit designs for which we cannot prove GC security. Such designs feature cycles which set their own control wires. Our pebbling-game-based definition leads to a natural proof of GC security; it is inspired by pebbling-based techniques from adaptively secure GC [HJO+16].

RAM from Cyclic Circuits with Subcircuit Sharing. Subcircuit sharing is the **key idea** of our RAM reductions. In short – and as we later explain in detail – we arrange our RAM memory as a collection of small subcircuits, each of which stores RAM elements and is shared across many accesses. By sharing each such subcircuit, we allow each data-dependent access to consume only the subcircuit storing its desired element. Thus, the number of required gates is amortized across accesses; in total, we only need a number of gates that grows quasilinearly in the number of accesses. Each gate in our RAM can be used to satisfy a variety of different accesses because our RAM circuits feature *cycles*, allowing accessed memory elements to "flow backwards through the topology" to the particular RAM step where it is needed.

The complexity of our RAM constructions arises from arranging subcircuit sharing at a large scale. We ultimately share each of a large number of subcircuits across a large number of memory accesses. This is achieved by arranging subcircuits in a binary tree where each node is itself a shared subcircuit providing shared access to further subcircuits. Section 5 explains in detail.

Preventing Short Circuits. Definition 1 includes the possibility of *illegal states*, denoted $\perp$. One can erroneously join two wires where one wire holds 0 and the other holds 1. This causes a 'short circuit', and is ill defined. We must restrict ourselves to circuit designs that cannot enter an illegal state. For this reason, we focus on tri-state circuits that compute Boolean functions:

Definition 5 (Computing a Boolean function). *Let $f : \{0,1\}^n \to \{0,1\}^m$ denote a Boolean function and $\mathcal{C}$ denote a tri-state circuit. We say that $\mathcal{C}$* ***computes*** *f if for all $x \in \{0,1\}^n$, $\mathcal{C}(x) = f(x)$.*

This definition rules out illegal states, because entering an illegal state causes the circuit to output $\perp$, which is not a possible output of a Boolean function.

Note, the property of computing a Boolean function (Definition 5) neither implies nor is implied by runtime acyclicity (Definition 4).

Completeness. Definition 1 does not include AND gates. Even without AND, tri-state circuits are as expressive as Boolean circuits. Indeed, for every Boolean circuit, there is a similarly-sized tri-state circuit computing the same function:

Theorem 1 (Emulating Boolean circuits; tri-state AND gates). *For any* Boolean *circuit $\mathcal{C}$, there exists a* tri-state *circuit $\mathcal{C}'$ such that:*

$$\mathcal{C}' \text{ computes } \mathcal{C} \quad \text{and} \quad |\mathcal{C}'| = O(|\mathcal{C}|)$$

Proof. By constructing Boolean gates from tri-state gates.

Indeed, it suffices to construct AND gates; XORs and constants are part of the tri-state circuit definition, and $\{\wedge, \oplus, 1\}$ is a complete Boolean basis. We use division to denote the buffer operation (Sect. 3). An AND gate can be constructed as follows:

$$AND(x, y) \triangleq \left(\frac{x}{y}\right) \bowtie \left(\frac{0}{y \oplus 1}\right)$$

The above definition can be read as follows: When the value of y is 1, the result is x; when the value of $y \oplus 1$ is 1, the result is 0. □

4.1 Randomized and Oblivious Tri-State Circuits

In cryptographic settings, one of the principal advantages of non-tri-state circuits is their input independent order of execution. However, the entire point of the tri-state circuit model is its dependence on the input. This leads to a natural question: can we construct tri-state circuits where orders of execution – which depend on inputs – *appear* to be independent of the input?

Indeed, we can meaningfully define *oblivious* tri-state circuits. This definition is sufficient for cryptographic applications, as we demonstrate in Sect. 7. The definition of oblivious tri-state circuits is analogous to that of oblivious Turing Machines and of oblivious RAMs (ORAM, [GO96]).

In short, a tri-state circuit is *oblivious* if we can *simulate* all of its buffer control wires. I.e., there exists a poly-time simulator that outputs a distribution of control bits which – on every input – is close to the distribution of the *real* controls. This idea is reasonable because the order in which gates are executed can be deduced from the controls alone. Indeed, buffer control wires are the *only* mechanism in a tri-state circuit that can set a wire conditionally, and hence they determine the order of execution. Thus, if the controls can be simulated, then the order of gate execution hides the input.

Given our definitions so far, we cannot construct non-trivial oblivious circuits. So far, there is no mechanism for deviating from an order of execution that is deterministically prescribed by the input. Thus, the value on each control wire is a determined by the input, and we cannot simulate. Somehow we must *mask* each sensitive wire value before using it to control a buffer, e.g. by applying a one-time pad. Thus, we consider tri-state circuits with *randomized* inputs:

Definition 6 (Randomized Tri-State Circuit). *A **randomized tri-state circuit** is a pair consisting of a tri-state circuit $\mathcal{C}$ and a distribution of bit-strings $\mathcal{D}$. The execution of a randomized tri-state circuit on input x is defined by randomly sampling a string r from $\mathcal{D}$, then running $\mathcal{C}$ on x and r:*

$$(\mathcal{C}, \mathcal{D})(x) \triangleq \mathcal{C}(x; r) \qquad \textbf{where } r \in_{\$} \mathcal{D}$$

A randomized tri-state circuit *obliviously* computes f if its *controls* (Definition 7) can be simulated:

Definition 7 (Controls). *Let $\mathcal{C}$ be a tri-state circuit with input $x \in \{0,1\}^n$. The **controls of $\mathcal{C}$ on** x, denoted controls$(\mathcal{C}, x) \in \{0, 1, \mathcal{Z}\}^*$, is the set of all buffer control wire values (each labeled by its gate ID) at halt-time.*

Definition 8 (Obliviously computing a function). *Let $f : \{0,1\}^n \to \{0,1\}^m$ be a Boolean function. Let $\sigma \in \mathbb{N}$ be the statistical security parameter. Let $(\mathcal{C}, \mathcal{D})_{i \in [\mathbb{N}]}$ denote a* family *of randomized tri-state circuits. The family **obliviously computes** f if:*

1. *For all* $x \in \{0,1\}^n$, $(\mathcal{C}, \mathcal{D})_\sigma$ *outputs* $f(x)$ *with overwhelming probability:*

$$\Pr_{r \in_\$ \mathcal{D}}[\ \mathcal{C}(x;r) = f(x)\ ] > 1 - \mathrm{negl}(\sigma)$$

2. *The distribution of controls of* $(\mathcal{C}, \mathcal{D})_\sigma$ *can be simulated. I.e., there exists a simulator* $\mathcal{S}$ *such that for all inputs* $x \in \{0,1\}^n$ *the following holds:*

$$\mathcal{S}(1^\sigma) \stackrel{s}{=} \{\ \mathit{controls}(\mathcal{C}, (x;r))\ \mid\ r \in_\$ \mathcal{D}\ \}$$

While only tri-state circuit *families* obliviously compute functions, we will sometimes slightly abuse notation and omit the explicit mention of families.

As a warm up, we show that for every Boolean circuit, there exists a similarly-sized randomized tri-state circuit obliviously computing the same function:

Theorem 2 (Obliviously Emulating Boolean Circuits). *For any* Boolean *circuit* $\mathcal{C}$*, there exists a randomized tri-state circuit* $(\mathcal{C}', \mathcal{D})$ *s.t.:*

$$(\mathcal{C}', \mathcal{D})\ \textit{obliviously computes}\ \mathcal{C} \quad \textit{and} \quad |\mathcal{C}'| = O(|\mathcal{C}|)$$

Proof. By reducing Boolean gates to tri-state gates with randomized input.

As in Theorem 1, we need only demonstrate how to build an oblivious AND gate, since XOR gates are part of the tri-state circuit definition.

To construct each AND gate, we use the classic idea of Beaver multiplication triples [Bea92]. For each AND gate, we define a distribution $\mathcal{D}$ as follows:

$$\mathcal{D} \triangleq \{\ \alpha, \beta, \alpha \cdot \beta\ \mid\ \alpha, \beta \in_\$ \{0,1\}\ \}$$

Our oblivious AND gate uses the multiplication triple to mask its input bits before using them as buffer controls:

$$\mathit{AND}_{obv}(x, y; \alpha, \beta, \gamma = \alpha \cdot \beta) \triangleq \left(\left(\frac{y}{x \oplus \alpha} \bowtie \frac{0}{(x \oplus \alpha) \oplus 1}\right) \oplus \left(\frac{\alpha}{y \oplus \beta} \bowtie \frac{0}{(y \oplus \beta) \oplus 1}\right)\right) \oplus \gamma$$

Both $x \oplus \alpha$ and $y \oplus \beta$ (and their complements) are controls, so to prove this gate is oblivious, we must simulate these values. This is straightforward: α and β act as one-time pads, masking x and y:

$$\mathcal{S}(1^\sigma) \triangleq \{\ r_0, r_0 \oplus 1, r_1, r_1 \oplus 1\ \mid\ r_0, r_1 \in_\$ \{0,1\}\ \}$$

Formally, the full circuit $(\mathcal{C}, \mathcal{D})$ consists of *many* such AND gates, each with its own triple, and we must jointly simulate *all* controls. This is trivial: multiplication triples are mutually independent and each is used only once. □

Simple Distributions. The formal definition of randomized tri-state circuits allows *arbitrary* distributions $\mathcal{D}$. In practice, we cannot handle *any* distribution. For instance, some distributions are not computable. Moreover, in some settings – and in particular in the authenticated garbling setting – we wish to

consider distributions that are as simple as possible, such that they are easy to sample.

Constructions presented in this work use simple distributions. In particular, our distributions can be described as the concatenation of independent copies of the following two sub-distributions: (1) a uniformly sampled bit $r \in_\$ \{0,1\}$ and (2) a uniform multiplication triple $\{ \alpha, \beta, \alpha \cdot \beta \mid \alpha, \beta \in_\$ \{0,1\} \}$. Uniform bits and multiplication triples suffice for our oblivious tri-state RAM.

Simple distributions are important because, as we will see, our approach to authenticated garbled tri-state circuits samples $\mathcal{D}$ via a (malicious) preprocessing functionality. Efficient protocols exist for our considered class of distributions [WRK17, KRRW18, YWZ20, DILO22].

5 Deterministic Tri-State RAM

In this section, we reduce RAM execution to deterministic tri-state circuits. We emphasize that this section constructs only RAM, not *oblivious* RAM.

Our focus is our later oblivious reduction (Sect. 6), which has utility in 2PC. We give a deterministic reduction here for two reasons. First, it explores the theoretical capabilities of TSCs. Second – and more importantly – our deterministic reduction is simpler than our oblivious reduction. Our deterministic RAM sets the stage for our more complex oblivious reduction, which mixes the same high-level ideas with the Oblivious RAM construction of [WCS15].

We set the stage by defining what it means for a TSC to emulate a RAM:

Definition 9 (T-Emulation). *Let $T \in \mathbb{N}$ denote a runtime. A tri-state circuit $\mathcal{C}$ **T-emulates a RAM** if $\mathcal{C}$ computes (Definition 5) the following function: Let P denote a word RAM program and $x \in \{0,1\}^n$ denote a string. $\langle x \rangle$ denotes a Boolean encoding of value x.*

$$\mathcal{C}(\langle P \rangle, x) = \begin{cases} P(x) & \text{if } P \text{ halts on input } x \text{ within } T \text{ steps} \\ \langle \bot \rangle & \text{otherwise} \end{cases}$$

Theorem 3 (Deterministic Tri-State RAM). *For any runtime $T \in \mathbb{N}$, there is a tri-state circuit $\mathcal{C}$ s.t. $\mathcal{C}$ T-emulates a RAM and $|\mathcal{C}| = O(T \cdot \log^4 T)$.*

We describe our deterministic RAM. Our oblivious construction (Sect. 6) is more sophisticated, but builds on the ideas developed in this section.

The challenge of emulating a RAM is in accessing a large main memory. Other details – including operating on machine words and managing internal state – are straightforward, even without tri-state-specific capabilities. Thus we focus on repeatedly and arbitrarily accessing main memory.

Our approach is strongly inspired by the GRAM construction of [HKO22]; we show that their high level ideas are compatible with tri-state circuits, replacing their complex language translation mechanism by simple tri-state gates. We later asymptotically improve over [HKO22]'s construction.

Throughout the following discussion, we refer the reader to Fig. 2, which depicts our deterministic RAM construction.

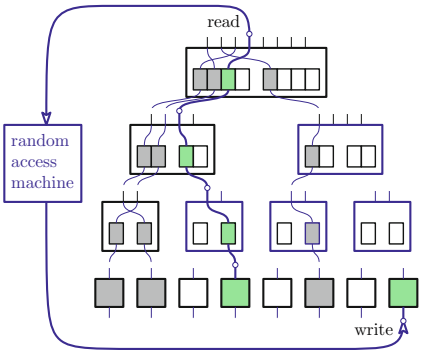

Fig. 2. Our deterministic tri-state RAM is arranged as a binary tree where memory elements are stored in the leaves. Each inner node has two stacks (concatenated rectangles) which allow the node to dynamically communicate with its two children. On each access, the RAM sets the desired leaf address on a top port of the RAM. This causes the circuit to dynamically traverse a path to the addressed leaf (example depicted in green). At each node, the RAM pops one stack and not the other, allowing the RAM to proceed either left or right. This establishes a path through which the requested element will flow back up to the root. Traversing the tree uses up parts of the circuit (previously used up components are in grey). Since the accessed memory element might be needed again later, the RAM writes an element back to a statically chosen and unused leaf. (Color figure online)

5.1 Deterministic Tri-State RAM Overview

Our main idea is to construct inside a tri-state circuit a binary tree of *nodes*, where each leaf holds one memory element, and where each internal node contains machinery needed to access its descendants.

On an access, the emulated RAM uses this machinery to dynamically traverse a path towards the particular leaf holding the target memory element. By leveraging **subcircuit sharing**, we ensure that while each traversal uses up some gates, it crucially does *not* use any gates off of its path. Thus, those gates can be used later. This basic idea leads to single circuit structure that is amortized across all RAM accesses.

Note that *Boolean* circuits cannot realize the above amortization since there is no mechanism by which to set aside a portion of the circuit for later use. In contrast, TSCs can, based on their dynamic order of execution and support for cyclic graphs. **This is the expressive advantage of TSCs**.

Communicating Nodes. Leveraging this ability to amortize gates, the next crucial insight of our deterministic RAM is to view each tree node as an *object* that can dynamically send messages to and receive messages from its two children, consuming only "on-the-path" TSC gates. (Boolean circuits do not have this ability, and each access to a child requires processing both children, ultimately resulting in a linear scan.) Messages are passed by *setting* particular collections

of wires that we refer to as **ports**. By setting a child's **input port**, a parent can send a message to its child; by setting its own **output port**, the child can respond to its parent. The challenge is in allowing each parent to communicate with its children *dynamically*. Namely, we must arrange that a parent sends a message to its child if and only if that child is on a dynamically traversed path.

Circuit-Based Stacks. Like [HKO22], we arrange this dynamic communication via **circuit-based stacks** [ZE13]. Namely, there exists a Boolean-circuit-based analog of a stack data structure. These stacks are created with n elements and can support up to m *conditional pop* (`cpop`) operations. On each `cpop`, a stack takes as argument a single control bit p. If $p = 1$, then the stack indeed pops, returning and removing its top element; if $p = 0$, then the stack instead returns the all zeros string and its contents remain unchanged.

We can port stacks to tri-state gates by simply substituting ANDs by buffers and XORs by joins.[9] A stack with n w-bit elements supporting m `cpop` operations requires $O(w \cdot m \cdot \log n)$ tri-state gates. Thus, each of the m `cpop` operations requires only an amortized log number of gates. This straightforward substitution unlocks *substantial* utility. Leveraging tri-state semantics, we can use stacks as dynamic *communication channels* between nodes; see next.

Using Stacks. A parent node and its children communicate via a stack. The parent manages the stack's control bits and outputs. The child manages the stack's content: it appropriately connects the stack's content wires to its input/output ports. To communicate with its child, the parent pops the stack (the call to `cpop` with $p = 1$ is non-oblivious). This establishes a chain of buffers whose control wires are each set to 1, but whose data wires are not yet set. As soon as we set the data wire of the first buffer in the chain, the buffers will one-by-one fire, sending the data through the chain, from parent to child.

To enable two-way communication, we place two kinds of buffers in the stack. Some buffers are oriented from parent to child, allowing messages to flow from the parent into the input port of the child; other buffers are oriented from child to parent, allowing the child to set messages on its output port which will flow through the stack to its parent. (In Appendix B.1 of the full version, we formalize this notion by giving two variants of a stack, one that sends messages from inputs to outputs and the other that sends messages from outputs to inputs.) Thus, by calling `cpop` with $p = 1$, the parent dynamically connects itself to one input port and one output port of its child.

Now that ports are connected, the parent can send a message to its child by setting data wires on its side of the stack. These values automatically flow through the stack into the child's input port, causing gates in the child to fire, compute the relevant response, and load the response onto an output port, where it, again, automatically flows through the stack back to the parent.

[9] In addition to wires that store data elements, stacks include control logic that tracks element positions. Here, we do *not* simply replace ANDs by buffers (XORs by joins), but rather translate Boolean control logic into tri-state gates via Theorem 1.

Crucially, if the parent instead calls `cpop` with $p = 0$, no communication occurs. The parent does not connect to its child's port, and hence no gates inside the child fire, so all gates in the unused child remain ready for later use.

Inner Nodes. Let tree level 0 denote the root; level i has 2^i nodes. Let level ℓ denote the tree's largest level. Each node on level i has two stacks, each supporting $2^{\ell-i}$ calls to `cpop`, of which half can be called with $p = 1$. I.e., each stack allows the node to communicate with its respective child up to $2^{\ell-i-1}$ times.

Each node also consists of $2^{\ell-i}$ subcircuits, each of which performs the following task: (1) receive the address of some leaf from an input port, (2) use the first bit of this address as a stack control bit such that we pop only the stack corresponding to the subtree that stores the requested address, (3) save the remaining bits of the address on the output wires of each stack (sending the bits to the active child), (4) read the response from each child, (5) join the responses together, and (6) save the joined response on the output port. We note that as we inspect nodes closer and closer to the leaves, the nodes become progressively smaller, until the leaves are subcircuits capable of handling exactly one request.

Read Traversals. The RAM can read elements from memory by traversing full root-to-leaf paths through the tree. To do so, the RAM loads into the root the address of the target leaf. The root strips the most significant bit from this address and uses it to conditionally communicate with its two children, indeed popping (i.e. calling `cpop` with $p = 1$) the stack for the child on the target path, and not popping (i.e. calling `cpop` with $p = 0$) the other child's stack. The root can now forward a message to its child, so it forwards the address's remaining bits. The child then recursively computes this same procedure, and so on, until we reach the target leaf.

This leaf stores a single element, and it sets its single output port to this element. Based on the semantics of tri-state circuits, this automatically triggers a *cascade of events*. The element flows through the stack of its immediate parent, causing the parent to fire and set its own output port to this newly received element. This causes the element to flow through a stack in the next level of the tree, and so on until the element reaches the root. At this point, the RAM can read the element from the root, completing the memory read.

Direct Writes. Now that the RAM has read its desired element, it *must* write something back. Note, we require this even if the goal of the memory access is simply to read. The problem is that we have now used up the gates associated with the accessed leaf, so we cannot reach that same leaf again. Thus, we need to write back to a fresh leaf, allowing later reads to access the same element again.

It is straightforward to arrange that each step of the RAM is statically and directly connected to one leaf, allowing it to directly write back without a dynamic traversal. Note, these connections induce cycles in the circuit graph, but not at runtime (Definition 4).

Recursive Position Map. As just discussed, each time we read a memory element, we write it back to a fresh location. This introduces a problem: how does the RAM *remember* where it last placed a particular element? This problem is typical in Oblivious RAM constructions, e.g. [SvS+13, WCS15], and can be solved via recursion. Namely, we explicitly store the current position of each memory element in a smaller, recursively-instantiated *position map.*

We can ensure that each recursively instantiated memory holds half the number of elements as the last, so only $O(\log T)$ levels of memory are needed. To achieve this, we arrange that each position map element holds the positions of (at least) two elements in memory. To terminate the recursion, we instantiate the smallest, constant-sized memory via Boolean-logic-based linear scans.

In sum, the RAM construction is binary tree where each node on level i is capable of handling $2^{\ell-i}$ RAM read requests. We dynamically traverse the tree via tri-state-circuit-based stacks; each node holds two stacks, and on each access we pop only the stack on the path to the desired element.

5.2 Sources of Logarithmic Overhead

Our deterministic tri-state RAM has $O(T \cdot \log^4 T)$ gates. We characterize four distinct sources of cost, each of which adds a logarithmic factor:

1. **Word size.** The first source of scaling is unavoidable, as it stems simply from the size of RAM words. Words are assumed to have size $\Theta(\log T)$, and tri-state gates operate on only one bit at a time. Hence, each action on a word requires $O(\log T)$ gates. This factor highlights that the comparison between the word RAM model and the circuit model is "unfair". Word RAMs can manipulate entire words at unit cost; tri-state circuits cannot.
2. **Binary Tree.** On each access, our RAM traverses a path through a binary tree of size $O(T)$. Each traversal touches $O(\log T)$ nodes.
3. **Stacks.** During each traversal and at each tree node, our RAM calls `cpop` on a constant number of stacks, each of size $O(T)$. Circuit-based stacks of size n have $O(\log n)$ overhead per `cpop`, yielding an additional $O(\log T)$ factor. **Our oblivious construction** leverages randomness to reduce the size of stacks from $O(T)$ to only $O(\text{poly}(\log T))$, and hence reduces stack overhead from $O(\log T)$ to only $O(\log \log T)$. This is how our oblivious construction is able to improve over our deterministic RAM.
4. **Recursion.** To track positions of elements, we use $O(\log T)$ position maps.

It is difficult to foresee methods for achieving a tri-state RAM with fewer than $O(T \cdot \log^3 T \cdot \log \log T)$ gates. Indeed, each above source of scaling seems relatively inherent to our constructions, so further asymptotic improvement will likely require fundamentally new techniques.

5.3 Formal Construction

We present our formal reduction from RAM to deterministic tri-state circuits in Appendix B.2 of the full version of this paper. We emphasize that the circuit described there is simply a formalism of the key ideas explained in Sect. 5.1.

6 Oblivious Tri-State RAM

In this section, we reduce RAM programs to *oblivious* tri-state circuits. At the highest level, we demonstrate that techniques in Sect. 5 can be combined with the *Circuit Oblivious RAM construction* of [WCS15].

We note that as a proof of concept, one can achieve oblivious tri-state RAM by simply employing off-the-shelf ORAM. Namely, use our deterministic RAM construction to emulate an ORAM server, and use oblivious tri-state Boolean gates (Theorem 2) to emulate an ORAM client. This works, but introduces high overhead which we would like to avoid. Here, we give a direct construction that is far more efficient than this proof of concept.

We begin by formalizing our claim:

Definition 10 (Oblivious T-Emulation). *Let $T \in \mathbb{N}$ denote a runtime. A randomized tri-state circuit family $(\mathcal{C}, \mathcal{D})_{i \in [\mathbb{N}]}$* ***obliviously T-emulates a RAM*** *if it obliviously computes the following function: Let P denote a word RAM program and $x \in \{0,1\}^n$ denote a string.*

$$\mathcal{C}(\langle P \rangle, x) = \begin{cases} P(x) & \text{if } P \text{ halts on input } x \text{ within } T \text{ steps} \\ \langle \bot \rangle & \text{otherwise} \end{cases}$$

Theorem 4 (RAM to Oblivious Tri-State Circuits). *For any runtime $T = \Theta(\mathrm{poly}(\sigma))$, there is a randomized tri-state circuit family $(\mathcal{C}, \mathcal{D})_{i \in \mathbb{N}}$ such that $(\mathcal{C}, \mathcal{D})_\sigma$ obliviously T-emulates a RAM and $|\mathcal{C}| = O(T \cdot \log^3 T \cdot \log \log T)$*

6.1 Circuit ORAM [WCS15] Review

Our key idea is to implement inside a tri-state circuit the Circuit Oblivious RAM construction of [WCS15]. We thus review the relevant ideas of Circuit ORAM.

Circuit ORAM is a *statistically-secure* ORAM: it hides memory access patterns *without* computational assumptions. This property is achieved because the Circuit ORAM client does not use cryptographic primitives to choose its queries. The simplicity of the ORAM client is compatible with the tri-state circuit setting where implementing cryptographic primitives via gates is expensive.

Circuit ORAM arranges memory elements in a binary tree with $O(T)$ leaves. Each node holds up to a constant number (e.g., 3) of memory elements. The root is the only exception: it stores a larger *stash* with capacity $\Theta(\log T \cdot \log \log T)$.

When the ORAM client accesses an element, that element is retrieved from its node and moved to the stash. To prevent the stash from overflowing, Circuit ORAM consistently moves elements away from the root in a process called *eviction*. The key invariant – originally proposed by Path ORAM [SvS+13] – is that even as an element is evicted, it remains on the path to a *fixed leaf*.

To access an element, we scan only those nodes along that element's path. By the invariant, this scan is guaranteed to find the target element, and because each non-root node holds only a few elements, the scan is relatively cheap: the entire path – including the stash – holds only $O(\log T \cdot \log \log T)$ total elements.

Once the element is accessed, the ORAM places that element in the stash and reassigns the element to a fresh, uniformly chosen (with replacement) path.

Because each path is chosen randomly, it is easy to simulate Circuit ORAM's access pattern: the simulator handles each access by choosing a uniform path.

Remembering Paths. To access an element, the ORAM client must somehow *remember* that element's path. Recall that each element's path was chosen when it was last accessed, and there might be long gaps between accesses of a particular element. There are too many data elements for the client to remember paths locally, so the client remembers paths by recursively instantiating a smaller ORAM called the *position map*. See also our discussion of recursion in Sect. 5.

Eviction. After each access, the RAM deterministically chooses two paths and evicts elements along those paths. Each node on a chosen path evicts up to one RAM element to its child. To ensure that the RAM does not get 'stuck' with too many elements in the stash, the identity of evicted elements must be chosen carefully. The goal is to move elements towards the leaves – where there is more space – as quickly as possible. [WCS15]'s key contribution is an efficient procedure for deciding *which* element each node should evict.

Some details of this eviction strategy are *highly relevant* here, because we must implement the procedure with tri-state gates within our asymptotic budget.

[WCS15]*'s Eviction Strategy.* During eviction, [WCS15] first computes metadata, deciding for each path node which element to evict. This metadata computation scans the path twice, starting at the root, performing a (cheap) step of computation at each node towards the leaf, then performing a second scan starting from the leaf and returning to the root. Crucially, this metadata computation has *high locality*: each step only considers local information stored in the currently considered node, plus $O(\log T)$ bits from the previous step.

Jumping ahead to our construction, this locality is *absolutely essential*, because it bounds the amount of information that needs to be passed from a tree node to its parent/child, and hence bounds the amount of information that needs to pass through circuit-based stacks (see discussion in Sect. 5). Thus, our reduction can use stacks of small items, each of size $O(\log T)$ bits.

The remaining details of metadata computation are not crucial for understanding our construction, except that they ensure eviction prevents the stash from overflowing (except with negligible probability). For further detail, we refer the reader to [WCS15] (see their Algorithms 2 and 3 as well as their Fig. 2).

After metadata is computed, Circuit ORAM again performs a scan from root to leaf where each node evicts (up to) one element to its child. The identity of this child is chosen according to the metadata.

By evicting elements this way, Circuit ORAM maintains its crucial path invariant while ensuring that the root will never overflow.

6.2 Overview of Our Oblivious Tri-State RAM

In short, our oblivious tri-state RAM reuses almost every idea explained in Sect. 5. It similarly maintains a binary tree of nodes, each of which conditionally communicates with its two children via circuit-based stacks, and our RAM reads elements by traversing paths through the tree. Our oblivious construction improves over our deterministic RAM in two ways: it is *oblivious*, making it suitable for cryptographic use, and it is asymptotically smaller.

These properties are achieved by using tri-state circuits to directly implement the Circuit ORAM construction [WCS15]. We also leverage an insight described by [PLS22] that allows us to use smaller circuit-based stacks, reducing asymptotic cost. We describe our oblivious tri-state RAM by highlighting the differences as compared to our deterministic reduction (Sect. 5).

Storage in Every Node. In our deterministic RAM, only the leaves store memory elements. Our oblivious construction follows Circuit ORAM, where *each* node can hold $O(1)$ elements and where the root stores $O(\log T \cdot \log \log T)$ elements. When the RAM accesses an element, that element is written back to the root (and not written directly to a leaf). These elements subsequently move down the tree via Circuit ORAM's eviction strategy, implemented via tri-state gates.

Multi-purpose Node Subcircuits. In our deterministic construction, each node holds $O(T)$ subcircuits, each of which completes a basic task: conditionally pop both stacks, then join the resulting values and send them back to the parent. Our oblivious construction's subcircuits are more complex. They each conditionally perform various tasks, depending on the current need of the ORAM construction.

Each subcircuit conditionally performs one of three tasks: (1) **read**, including scanning the node's local content, (2) **evict**, including computing appropriate metadata and sending an element to a child, or (3) **do nothing** (the need for this option is explained when we discuss "smaller stacks"). While each subcircuit must include enough circuitry to complete *any* of these tasks, the subcircuit is small. This is achieved by reusing parts of the circuit across the different possible tasks. In particular, we need only two total calls to `cpop` per subcircuit.

Multiple Scans. In our deterministic RAM, each access scans a path twice, from root to leaf and then back to the root. Our *oblivious* tri-state RAM performs *three* scans. While only two scans are needed to read, three are needed to evict. When evicting, the RAM uses two scans to compute Circuit ORAM's relevant metadata, and it uses the third scan to evict elements from parent to child.

Our tri-state stacks are thus used multiple times per access: the parent sends a message to its child, receives a message back, and then sends a second message. We emphasize that there is no technical challenge in using a circuit-based stack to communicate more than once: just increase the size of stack elements and leverage tri-state semantics to send bits at the right time.

Even though our nodes use three scans, the total information flowing through stacks remains small. In total, each node sends/receives $O(\log T)$ bits of information. Keeping this amount of information small is crucial, because transmitted

bits pass through stacks, and hence we must pay in additional gates for every bit of information transmitted between parent and child.

Smaller Stacks. In our deterministic RAM, each node communicates with each of its children via a circuit-based stack of size $O(T)$. Our oblivious construction improves on this by leveraging an elegant idea demonstrated by [PLS22], allowing much smaller stacks that hold only $O(\text{poly}(\log T))$ elements. The smaller stacks account for our oblivious construction's improved asymptotic size.

We explain [PLS22]'s observation – which is derived from an observation of [FNR+15] – in the context of Circuit ORAM. Recall that in Circuit ORAM, each memory access scans a uniformly chosen path.

Let $B = \Theta(\log^{1+\epsilon} \sigma)$ denote a parameter super-logarithmic in the security parameter for constant $\epsilon > 0$. We call B the *batch parameter*. Let level 0 denote the root of the RAM tree; each level i has 2^i nodes. Consider: how often will a particular node on level i be scanned over the course of $2^i \cdot B$ accesses?

[FNR+15]'s insight is that because elements are randomly assigned to leaves, accesses should be roughly evenly distributed amongst nodes on level i. Indeed, it is incredibly unlikely that a particular node will be scanned significantly more often than its peers. [FNR+15] proved that it is only negligibly likely that over $2^i \cdot B$ accesses any node on level i will be used more than $2 \cdot B$ times.

The upshot is that we need never instantiate a stack with more than $O(B)$ entries (e.g., 256 entries in practice), since it is unlikely that we will exhaust its entries over the course of $2^i \cdot B$ accesses. Instead, every $2^i \cdot B$ accesses, we insert a *reset* step, forcibly clearing all stacks on level i and instantiating fresh stacks. Since each `cpop` operation is made to a smaller stack, this strategy reduces stack overhead from factor $\log T$ to factor $\log \log T$.

Smaller stacks introduce nuance in implementing node subcircuits. Consider a particular node, consisting of many sequentially composed subcircuits. [PLS22]'s strategy partitions these subcircuits into *generations* of size $2 \cdot B$. After each generation, we insert a statically scheduled reset, preparing for the next generation. **For this to work**, we must ensure that over the course of $2^i \cdot B$ accesses, *every* subcircuit in the current generation is consumed. If not, the circuit is not well defined, since our reset will manipulate wires coming out of the generation's last subcircuit, and the wires of this last subcircuit are defined only if it and all of its predecessors have been used. Thus we must *ensure* that each subcircuit is ultimately used. Since subcircuits are used only if they are on a randomly chosen path, it is highly unlikely that *every* subcircuit will be used up naturally.

To account for this problem, we insert additional logic allowing a parent to burn through subcircuits in its children's current generations. This is the role of the **do nothing** subcircuit task. When a parent calls its child with a particular flag set, the child's subcircuit simply calls `cpop` on each of its respective stacks with $p = 0$, and no further action is taken. This burns the subcircuit.

Because we reset level i every $2^i \cdot B$ accesses, we reset level $i+1$ in synchrony with one out of every *two* resets of level i. On each second reset of level i, we add circuitry that causes each node on level i to call `cpop` on each of its stacks $2 \cdot B$ additional times, sending a message that instructs the corresponding child to

burn a subcircuit (once all subcircuits are burned, the parent stops forwarding this message by instead calling `cpop` with $p = 0$). A statically known state, and we can correctly wire gates.

Oblivious Circuitry. To allow a simulator, our oblivious reduction uses oblivious Boolean gates (see Theorem 2). There is nuance here: circuit-based stacks continue to elide obliviousness, and the role each subcircuit ends up executing (read, evict, or do nothing) is leaked by the circuit. This leakage is fine, however, since this information is implied by the RAM's physical access pattern, and Circuit ORAM ensures that the physical access pattern hides the logical access pattern.

On the other hand, some oblivious gates *are* required. In particular, any circuitry that actually scans the content of a RAM node is oblivious. It is cheap to instantiate these components with oblivious ANDs (Theorem 2).

In sum, our oblivious reduction builds on the basic ideas of Sect. 5, and then layers in the key ideas of Circuit ORAM [WCS15] and of [PLS22]. Circuit ORAM's eviction procedure can be implemented by tri-state gates, allowing for a lean memory structure whose access pattern can be simulated. The resulting memory features an access pattern that touches nodes on each tree level uniformly, allowing us to use smaller stacks, reducing the size of circuitry required to support intra-node communication. Together, these ideas yield a circuit that obliviously simulates RAM and that has low poly-logarithmic overhead.

6.3 Formal Construction

We present our formal reduction from RAM to oblivious tri-state circuits in Appendix B.3 of the full version of this paper. We emphasize that the circuit described there is simply a formalism of the key ideas explained in Sect. 6.2.

7 Garbling Tri-State Circuits

In this section, we demonstrate how to garble tri-state circuits. We give two constructions. Our first construction garbles tri-state gates based only on one-way functions, achieving a *garbling scheme* [BHR12] suited to semi-honest protocols. Our second construction builds on *authenticated garbling* [WRK17] to achieve malicious security. Both constructions leverage similar high level ideas.

Intuition. In short, garbling of tri-state circuits is similar to classic garbling of Boolean circuits. Just as in classic garbling, the garbler G chooses two keys per wire. One key encodes logical zero, the other encodes one. To garble the circuit, G proceeds gate by gate. At each gate, G uses appropriate combinations of input keys to encrypt output keys according to the gate's function.

At runtime, the evaluator E obtains *at most* one key per wire. E walks the circuit, using keys to decrypt subsequent keys until obtaining output keys.

The crucial point is this: G *only* chooses keys that encode 0 and 1; the distinguished value $\mathcal{Z}$ is encoded by *the lack of a key.* If a particular wire holds

$\mathcal{Z}$ at halt-time, then E will *never* learn a key for that wire. The inability to decrypt certain wires differentiates tri-state garbling from classic garbling.

Throughout evaluation, E will keep track of which wires are set and which are not set. To arrange this, we *reveal to E the cleartext value of every buffer control wire.* (It is easy to arrange that E learns the cleartext values of particular wires.) Because E knows which wires hold $\mathcal{Z}$ and which do not, E can execute the circuit in a dynamic order, at each step handling those gates for which input keys are available. It is safe to reveal controls because the *obliviousness* (Definition 8) of the circuit ensures that these bits give E no information about the input.

Note the fit between garbling and the out-of-order nature of tri-state circuits: E can, of course, decrypt each GC gate *as soon as matching keys are obtained*, making it easy to execute gates in an order prescribed by *notify* (Fig. 1).

The following sections show how we garble tri-state gates. Our approaches build on known techniques for garbling from one-way functions (e.g., see [LP09]) and for authenticated garbling [WRK17].

7.1 Tri-State Garbling from One-Way Functions

Recall that tri-state circuits include XORs, buffers, and joins. We present our semi-honest garbling of each gate type from one-way functions.

Wire Keys; Point and Permute. For each wire w, G uniformly samples two length-λ keys K_w^0 and K_w^1. The first key encodes logical zero; the second encodes one.

We use the classic *point and permute* trick [NPS99]. In GC, each gate uses several ciphertexts, of which E should decrypt one. Point and permute allows E to decrypt the *correct* ciphertext without using awkward tricks like trying to decrypt a ciphertext and then checking if it decrypted correctly or not.

The trick requires that for each wire w, the least significant bits of K_w^0 and K_w^1 differ. G conditionally flips the least significant bit of K_w^1 to ensure it differs from that of K_w^0. G then *permutes* gate ciphertexts according to these keys.

In the following, we elide details of point and permute, opting for a simpler presentation. Appendix C of the full version presents a formal construction.

Gate Handling. With keys chosen, G garbles each gate one at a time. Consider a gate with input wires x and y and with output wire z. Let $Enc(\cdot,\cdot)$ denote CPA-secure encryption (which can be instantiated from one-way functions).

– **XOR.** For each XOR gate, G classically garbles the gate by encrypting each output key according to the appropriate combination of inputs keys:

$$\begin{array}{ll} Enc(K_x^0, Enc(K_y^0, K_z^0)) & Enc(K_x^0, Enc(K_y^1, K_z^1)) \\ Enc(K_x^1, Enc(K_y^0, K_z^1)) & Enc(K_x^1, Enc(K_y^1, K_z^0)) \end{array}$$

These four ciphertexts are shuffled according to point and permute and sent to E. At runtime and when two input keys become available, E decrypts one ciphertext, obtaining the corresponding output key.

– **Buffer.** For each buffer, G ensures that E can obtain an output key iff E holds the one key for the control. Recall, E must *learn* the value of each control. To achieve this, G send to E the least significant bit of K_y^0; E can compare this to its own lsb and learn y. Altogether, G sends:

$$lsb(K_y^0) \qquad Enc(K_x^0, Enc(K_y^1, K_z^0)) \qquad Enc(K_x^1, Enc(K_y^1, K_z^1))$$

The two encryptions are shuffled according to point and permute.
– **Join.** For each join, G ensures E obtains an output key if E holds *any* input key. G sends the following rows (shuffled wire-wise with point-and-permute):

$$Enc(K_x^0, K_z^0) \qquad Enc(K_x^1, K_z^1) \qquad Enc(K_y^0, K_z^0) \qquad Enc(K_y^1, K_z^1)$$

As soon as E obtains *any* input key, E decrypts the appropriate ciphertext, obtaining a corresponding output key. Here it is *essential* that $(\mathcal{C}, \mathcal{D})$ computes a Boolean function, preventing the possibility of a short circuit (see discussion near Definition 5) which would allow E to learn *both* output keys.

Sampling $\mathcal{D}$. Recall that an oblivious tri-state circuit includes a distribution on bits $\mathcal{D}$. In the semi-honest setting, it is trivial to handle input randomness: G locally samples $r \in_\$ \mathcal{D}$, then sends to E wire keys corresponding to r.

In sum, our OWF-based garbling scheme is simple: we gate-by-gate garble the circuit, and each garbled gate has size $O(\lambda)$ bits. E evaluates the circuit as keys become available, implementing the dynamic behavior of the tri-state model.

Formal Construction. We present our OWF-based tri-state *garbling scheme* in Appendix C of the full version of this paper. We emphasize that the presentation there is just a formalization of the ideas presented above.

In terms of security, our proof gives a simulator and a hybrid argument demonstrating that the garbled circuit hides the input. This proof is similar to classic proofs of GC security, e.g. [LP09], with *two key exceptions*: we use our oblivious tri-state circuit's simulator (Definition 8) to argue that the E's observed order of execution can be simulated, and we use runtime acyclicity (Definition 4) to guide our hybrid argument.

By combining facts of Appendix C with Theorem 4, we obtain:

Corollary 1 (Garbled RAM from one-way functions). *Assuming one-way functions and in the OT-hybrid model, there exists a constant-round, semi-honest secure 2PC protocol for word-RAM programs such that for any program halting in T steps, communication cost is $O(T \cdot \log^3 T \cdot \log\log T \cdot \lambda)$ bits.*

7.2 Authenticated Garbling of Tri-State Circuits

In this section, we extend authenticated garbling [WRK17] to tri-state circuits. This extension implies an efficient constant-round maliciously-secure 2PC protocol for RAM programs.

Our authenticated handling is similar to that of our standard-assumption -based garbling (Sect. 7.1). We highlight the key differences:

- **Doubly-authenticated labels.** In standard GC, we encode each wire by a pair of keys chosen by G. In authenticated GC, each key contains components chosen by *each* party. In particular, each key includes a MAC, allowing E to authenticate that certain values are well formed.
- **Preprocessing.** In Sect. 7.1, we allowed G to sample randomness $r \in_\$ \mathcal{D}$. In the authenticated setting, we instead require that G and E *jointly* sample r via a preprocessing functionality. This achieves two goals. First, it ensures that r is indeed sampled from $\mathcal{D}$, and not arbitrarily chosen by malicious G. Second it ensures that neither G nor E knows r. This (1) prevents E from learning wire values and (2) prevents G from performing *selective abort attacks*. Note, prior work on authenticated garbling also leverages preprocessed randomness in the form of *doubly authenticated multiplication triples*.
- **Correlations; Random Oracle Assumption.** In Sect. 7.1, we garbled tri-state gates using only one-way functions. Our authenticated approach uses a function H modeled as a random oracle. The use of RO stems from *correlations* in wire labels. It is typical to use RO for authenticated GC.

We next describe authenticated garbling in more detail.

Garblings. The crucial authenticated GC invariant [WRK17] is that on each wire, G and E hold XOR secret shares of two MACs, one that authenticates the cleartext value to G and one that authenticates the cleartext value to E.

These *garblings* are defined over two global secrets:

- $\Delta \in_\$ \{0,1\}^\lambda$ is a global key drawn by G and hidden from E. G uses Δ as a key with which to encrypt gates.
- $\mu \in_\$ \{0,1\}^\sigma$ is a global MAC drawn by E and hidden from G. E uses μ to check that values opened by G are honestly constructed.

For convenience of notation, we define a value $\Gamma \triangleq \Delta \,||\, \mu \,||\, 1$.

As the circuit executes, for each wire holding value x, G and E will hold XOR secret shares of the value $x \cdot \Gamma$. We refer to these shares as *garblings*:

Notation 1 (Distributed Pair). *We denote by $\langle\langle x, y \rangle\rangle$ a* distributed pair *of values, where G holds value x and E holds value y.*

Definition 11 (Garbling). *Let $x \in \{0,1,\mathcal{Z}\}$ be a tri-state value. The **garbling** of x is a secret share held between G and E. G's share is a string $X \in \{0,1\}^{\lambda+\sigma+1}$. E's share is either (1) the symbol $\mathcal{Z}$ if $x = \mathcal{Z}$ or (2) the following:*

$$X \oplus (x \cdot \Delta \,||\, x \cdot \mu \,||\, x) = X \oplus x \cdot \Gamma$$

We denote a garbling of x by $[\![x]\!]$:

$$[\![x]\!] \triangleq \left\langle\!\!\left\langle X, \begin{cases} \mathcal{Z} & \text{if } x = \mathcal{Z} \\ X \oplus x \cdot \Gamma & \text{otherwise} \end{cases} \right\rangle\!\!\right\rangle$$

*For values $x \neq \mathcal{Z}$, we refer to the Δ component of a garbling as the **key part** of the garbling, to the μ component as the **MAC part**, and to the third component*

as the ***value part****. We use key, mac, val to denote appropriate projections. I.e., when $x \neq \mathcal{Z}$, we define the following projections:*

$$\begin{aligned} key(\llbracket x \rrbracket) &= \langle\!\langle X_0, X_0 \oplus x \cdot \Delta \rangle\!\rangle && \textbf{\textit{where}}\ X_0 \in \{0,1\}^\lambda \\ mac(\llbracket x \rrbracket) &= \langle\!\langle X_1, X_1 \oplus x \cdot \mu \rangle\!\rangle && \textbf{\textit{where}}\ X_1 \in \{0,1\}^\sigma \\ val(\llbracket x \rrbracket) &= \langle\!\langle X_2, X_2 \oplus x \rangle\!\rangle && \textbf{\textit{where}}\ X_2 \in \{0,1\} \\ & && \textbf{\textit{and}}\ X_0 \mid\mid X_1 \mid\mid X_2 = X \end{aligned}$$

Garbling XORs. Garblings are linearly homomorphic. Namely, if each party locally XORs the shares of two garbled bits, the result is itself a garbled bit. XOR gates are 'free' [KS08]:

Lemma 2 (Free XOR). $\llbracket x \rrbracket \oplus \llbracket y \rrbracket = \llbracket x \oplus y \rrbracket$

Revealing Values to E. Recall that in tri-state execution we reveal to E the cleartext value of each control bit. In the authenticated setting, we must be careful when revealing values. In particular, we must preserve two properties:

- **Privacy.** Revealed values should not leak G's input to E. Data privacy is preserved via tri-state *obliviousness* (Definition 8).
- **Authenticity.** Even a malicious G should not be able to reveal the wrong value. We prevent G from cheating via the MAC μ on the revealed wire.

It is relatively straightforward for G to reveal a circuit value $\llbracket x \rrbracket$ to E. The parties first compute:

$$\langle\!\langle X_1, X_1 \oplus x \cdot \mu \rangle\!\rangle \leftarrow mac(\llbracket x \rrbracket) \qquad \langle\!\langle X_2, X_2 \oplus x \rangle\!\rangle \leftarrow val(\llbracket x \rrbracket)$$

If G is honest, G can reveal x by sending to E the strings X_1 and X_2. Of course, G might attempt to cheat, so it may be the case that G sends $X_1' \neq X_1$ and $X_2' \neq X_2$. Thus, E must check that the strings are well formed. Recall that E knows the MAC μ. E checks the following:

$$X_1' \oplus x' \cdot \mu \stackrel{?}{=} X_1 \oplus x \cdot \mu \qquad \text{where } x' = (X_2 \oplus x) \oplus X_2' \tag{1}$$

If this passes, E is convinced that the wire indeed holds x'; otherwise, E aborts.

The above check passes whenever G indeed sends X_1 and X_2. Moreover, if G attempts to cheat, then the above check will only pass with probability negligible in σ. Indeed, to successfully reveal $x \oplus 1$, G must send $X_2' = X_2 \oplus 1$ and $X_1' = X_1 \oplus \mu$. However, G does not know μ, and so G's attempt to send $X_1 \oplus \mu$ requires guessing μ, which succeeds with probability at most $1/2^\sigma$.

Note that obliviousness ensures that each revealed value x can be simulated. This is important not only for protecting G's privacy, but also for protecting E's. In particular, G cannot employ a selective abort attack. Without obliviousness, G could cause E's check to fail iff x has a particular value; E's choice to abort/not abort reveals to G information about x. Indeed, G can still attempt such an "attack". However, the attempt is useless, since it only reveals information about a control wire which, by obliviousness, can be simulated anyway. Note that this argument crucially relies on the fact that G does not know the circuit randomness $r \in_\$ \mathcal{D}$, which is one reason r must be jointly computed in preprocessing.

Authenticated Buffers. Consider a buffer with data input $[\![x]\!]$ and control $[\![s]\!]$. Suppose $s \neq \mathcal{Z}$ and $x \neq \mathcal{Z}$. We show how the parties compute $[\![x \;/\; s]\!]$.

First, the parties reveal the control s to E, as described above. Now, let:

$$\langle\!\langle S, S \oplus s \cdot \Delta \rangle\!\rangle = key([\![s]\!]) \qquad \langle\!\langle X, X \oplus x \cdot \Gamma \rangle\!\rangle = [\![x]\!]$$

Honest G wishes to let E propagate $[\![x]\!]$ to the gate output iff $s = 1$. The parties publicly agree on gate-specific nonce ν, and G sets its output share as follows:

$$Y \triangleq H(S \oplus \Delta, \nu) \oplus X$$

At runtime, E checks if $s = 1$. Note that if $s = 1$, in an honest execution E holds $S \oplus \Delta$. In this case, E computes:

$$H(S \oplus \Delta, \nu) \oplus (X \oplus x \cdot \Gamma) = (Y \oplus X) \oplus (X \oplus x \cdot \Gamma) = Y \oplus x \cdot \Gamma$$

E places this value on the output wire, matching G's share Y. Thus, if $s = 1$, the output wire holds $[\![x]\!]$, as prescribed by buffer semantics. If instead $s = 0$ then the gate is indeed inactive: E cannot compute $H(S \oplus \Delta, \nu)$ without $S \oplus \Delta$.

Thus, a buffer is implemented by G revealing one value and having each party compute H at most once.

Authenticated Joins. Consider a join with inputs $[\![x]\!]$ and $[\![y]\!]$ and suppose that at runtime at least one input is set. We show how the parties compute $[\![x \bowtie y]\!]$.

Let the shares of $[\![x]\!]$, $[\![y]\!]$ be as follows:

$$[\![x]\!] = \langle\!\langle X, X_E \rangle\!\rangle = \langle\!\langle X, X \oplus x \cdot \Gamma \rangle\!\rangle \qquad [\![y]\!] = \langle\!\langle Y, Y_E \rangle\!\rangle = \langle\!\langle Y, Y \oplus y \cdot \Gamma \rangle\!\rangle$$

G sets its output share as X. Thus, if $x \neq \mathcal{Z}$ is set, then E simply copies its share $X \oplus x \cdot \Gamma$ onto the gate output wire. If instead $x = \mathcal{Z}$, then E must *translate* its share $Y \oplus y \cdot \Gamma$ to a format compatible with G's share. Hence, G includes in the GC the message $X \oplus Y$, which E can use to compute a matching share:

$$\left(\begin{cases} X_E & \text{if } y = \mathcal{Z} \\ Y_E \oplus X \oplus Y & \text{if } x = \mathcal{Z} \end{cases}\right) = \left(\begin{cases} X \oplus x \cdot \Gamma & \text{if } y = \mathcal{Z} \\ X \oplus y \cdot \Gamma & \text{if } x = \mathcal{Z} \end{cases}\right) = X \oplus (x \bowtie y) \cdot \Gamma$$

Thus, a join is implemented by having G send one correction value, which E conditionally XORs with its local value.

In sum, the authenticated garbling of tri-state gates is a relatively straightforward extension of existing techniques. Adding a MAC to each key prevents G from arbitrarily flipping wire values. Crucially, it remains possible for G to reveal values to E, allowing E to execute the dynamic behavior of the tri-state model.

Communication cost of our protocol is low. For example, plugging in Theorem 2, we evaluate an oblivious AND gate using an authenticated triple and $2\lambda + 4\sigma + 2$ additional communicated bits (two buffers, two joins). This is only slightly worse than the authenticated AND garbling of [KRRW18], which consumes an authenticated triple and $2\lambda + 1$ additional bits. For typical parameters $\lambda = 128$ and $\sigma = 40$, our approach is less than $2\times$ worse. Optimizations of [KRRW18] can likely be integrated with TSC handling, improving performance.

Formal Construction. We defer our full malicious protocol to Appendix D of the full version. Our protocol and its security proof are very similar to those of [WRK17], with the crucial difference that we use the above gate handling.

By combining facts proved in Appendix D with Theorem 4, we trivially obtain the following corollary:

Corollary 2 (Authenticated Garbled RAM). *In the random oracle/OT-hybrid model, there exists a constant-round, maliciously-secure 2PC protocol for word RAMs such that for any program halting in T steps, communication cost is $O(T \cdot \log^3 T \cdot \log\log T \cdot \lambda)$ bits.*

Acknowledgements. Distribution Statement "A": (Approved for Public Release, Distribution Unlimited). This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA), supported in part by DARPA under Cooperative Agreement HR0011-20-2-0025, and DARPA Contract No. HR001120C0087, Algorand Centers of Excellence programme managed by Algorand Foundation, NSF grants CNS-2246353, CNS-2246354, CNS-2246355, CNS-2001096 and CCF-2220450, US-Israel BSF grant 2015782, Amazon Faculty Award, Cisco Research Award and Sunday Group. Any views, opinions, findings, conclusions or recommendations contained herein are those of the author(s) and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, the Department of Defense, the Algorand Foundation, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes not withstanding any copyright annotation therein.

References

[Bea92] Beaver, D.: Efficient multiparty protocols using circuit randomization. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 420–432. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_34

[BHR12] 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

[CCHR16] Canetti, R., Chen, Y., Holmgren, J., Raykova, M.: Adaptive succinct garbled RAM or: how to delegate your database. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9986, pp. 61–90. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53644-5_3

[CH16] Canetti, R., Holmgren, J.: Fully succinct garbled RAM. In: Sudan, M. (ed.) ITCS 2016, pp. 169–178. ACM, January 2016

[CKKZ12] Choi, S.G., Katz, J., Kumaresan, R., Zhou, H.-S.: On the security of the "Free-XOR" technique. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 39–53. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28914-9_3

[DILO22] Dittmer, S., Ishai, Y., Steve, L., Ostrovsky, R.: Authenticated garbling from simple correlations. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022. LNCS, vol. 13510, pp. 57–87. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15985-5_3

[FNR+15] Fletcher, C., Naveed, M., Ren, L., Shi, E., Stefanov, E.: Bucket ORAM: single online roundtrip, constant bandwidth oblivious RAM. Cryptology ePrint Archive, Report 2015/1065 (2015). https://eprint.iacr.org/2015/1065

[GGMP16] Garg, S., Gupta, D., Miao, P., Pandey, O.: Secure multiparty RAM computation in constant rounds. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9985, pp. 491–520. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53641-4_19

[GHL+14] Gentry, C., Halevi, S., Lu, S., Ostrovsky, R., Raykova, M., Wichs, D.: Garbled RAM revisited. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 405–422. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_23

[GLNP18] Gueron, S., Lindell, Y., Nof, A., Pinkas, B.: Fast garbling of circuits under standard assumptions. J. Cryptol. **31**(3), 798–844 (2018)

[GLO15] Garg, S., Lu, S., Ostrovsky, R.: Black-box garbled RAM. In: Guruswami, V. (ed.) 56th FOCS, pp. 210–229. IEEE Computer Society Press, October 2015

[GLOS15] Garg, S., Lu, S., Ostrovsky, R., Scafuro, A.: Garbled RAM from one-way functions. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC, pp. 449–458. ACM Press, June 2015

[GO96] Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious RAMs. J. ACM **43**(3), 431–473 (1996)

[Hag98] Hagerup, T.: Sorting and searching on the word RAM. In: Morvan, M., Meinel, C., Krob, D. (eds.) STACS 1998. LNCS, vol. 1373, pp. 366–398. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0028575

[HJO+16] Hemenway, B., Jafargholi, Z., Ostrovsky, R., Scafuro, A., Wichs, D.: Adaptively secure garbled circuits from one-way functions. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 149–178. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53015-3_6

[HKO22] Heath, D., Kolesnikov, V., Ostrovsky, R.: EPIGRAM: practical garbled RAM. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022. LNCS, vol. 13275, pp. 3–33. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-06944-4_1

[HY16] Hazay, C., Yanai, A.: Constant-round maliciously secure two-party computation in the RAM model. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9985, pp. 521–553. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53641-4_20

[KRRW18] Katz, J., Ranellucci, S., Rosulek, M., Wang, X.: Optimizing authenticated garbling for faster secure two-party computation. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10993, pp. 365–391. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96878-0_13

[KS08] Kolesnikov, V., Schneider, T.: Improved garbled circuit: free XOR gates and applications. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008. LNCS, vol. 5126, pp. 486–498. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-70583-3_40

[Lin13] Lindell, Y.: Fast cut-and-choose based protocols for malicious and covert adversaries. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 1–17. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_1

[LO13] Lu, S., Ostrovsky, R.: How to garble RAM programs? In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 719–734. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_42

[LO17] Lu, S., Ostrovsky, R.: Black-box parallel garbled RAM. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10402, pp. 66–92. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_3

[LP09] Lindell, Y., Pinkas, B.: A proof of security of Yao's protocol for two-party computation. J. Cryptol. **22**(2), 161–188 (2009)

[LY18] Lindell, Y., Yanai, A.: Fast garbling of circuits over 3-valued logic. In: Abdalla, M., Dahab, R. (eds.) PKC 2018. LNCS, vol. 10769, pp. 620–643. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76578-5_21

[Mia20] Miao, P.: Cut-and-choose for garbled RAM. In: Jarecki, S. (ed.) CT-RSA 2020. LNCS, vol. 12006, pp. 610–637. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-40186-3_26

[NPS99] Naor, M., Pinkas, B., Sumner, R.: Privacy preserving auctions and mechanism design. In: Proceedings of the 1st ACM Conference on Electronic Commerce, pp. 129–139. ACM (1999)

[PLS22] Park, A., Lin, W.-K., Shi, E.: NanoGRAM: garbled RAM with $\widetilde{O}(\log N)$ overhead. Cryptology ePrint Archive, Report 2022/191 (2022). https://eprint.iacr.org/2022/191

[RR21] Rosulek, M., Roy, L.: Three halves make a whole? Beating the half-gates lower bound for garbled circuits. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12825, pp. 94–124. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84242-0_5

[Sch80] Schönhage, A.: Storage modification machines. SIAM J. Comput. **9**(3), 490–508 (1980)

[SvS+13] Stefanov, E., et al.: Path ORAM: an extremely simple oblivious RAM protocol. In: Sadeghi, A.-R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013, pp. 299–310. ACM Press, November 2013

[WCS15] Wang, X., Hubert Chan, T.-H., Shi, E.: Circuit ORAM: on tightness of the Goldreich-Ostrovsky lower bound. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015, pp. 850–861. ACM Press, October 2015

[WRK17] Wang, X., Ranellucci, S., Katz, J.: Authenticated garbling and efficient maliciously secure two-party computation. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 21–37. ACM Press, October/November 2017

[Yao86] Yao, A.C.-C.: How to generate and exchange secrets (extended abstract). In: 27th FOCS, pp. 162–167. IEEE Computer Society Press, October 1986

[YWZ20] Yang, K., Wang, X., Zhang, J.: More efficient MPC from improved triple generation and authenticated garbling. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020, pp. 1627–1646. ACM Press, November 2020

[ZE13] Zahur, S., Evans, D.: Circuit structures for improving efficiency of security and privacy tools. In: 2013 IEEE Symposium on Security and Privacy, pp. 493–507. IEEE Computer Society Press, May 2013

[ZRE15] Zahur, S., Rosulek, M., Evans, D.: Two halves make a whole - reducing data transfer in garbled circuits using half gates. 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