

Towards Generic MPC Compilers via Variable Instruction Set Architectures (VISAs)

Yibin Yang* Stanislav Peceny† David Heath‡ Vladimir Kolesnikov§

August 20, 2023

Abstract

In MPC, we usually represent programs as circuits. This is a poor fit for programs that use complex control flow, as it is costly to compile control flow to circuits. This motivated prior work to emulate CPUs inside MPC. Emulated CPUs can run complex programs, but they introduce high overhead due to the need to evaluate not just the program, but also the machinery of the CPU, including fetching, decoding, and executing instructions, accessing RAM, etc.

Thus, both circuits and CPU emulation seem a poor fit for general MPC. The former cannot scale to arbitrary programs; the latter incurs high per-operation overhead.

We propose *variable instruction set architectures* (VISAs), an approach that inherits the best features of both circuits and CPU emulation. Unlike a CPU, a VISA machine repeatedly executes entire program *fragments*, not individual instructions. By considering larger building blocks, we avoid most of the machinery associated with CPU emulation: we directly handle each fragment as a circuit.

We instantiated a VISA machine via garbled circuits (GC), yielding constant-round 2PC for arbitrary assembly programs. We use improved branching (Stacked Garbling, Heath and Kolesnikov, Crypto 2020) and recent Garbled RAM (GRAM) (Heath et al., Eurocrypt 2022). Composing these securely and efficiently is intricate, and is one of our main contributions.

We implemented our approach and ran it on common programs, including Dijkstra’s and Knuth-Morris-Pratt. Our 2PC VISA machine executes assembly instructions at 300Hz to 4000Hz, depending on the target program. We significantly outperform the state-of-the-art CPU-based approach (Wang et al., ESORICS 2016, whose tool we re-benchmarked on our setup). We run in constant rounds, use $6\times$ less bandwidth, and run more than $40\times$ faster on a low-latency network. With 50ms (resp. 100ms) latency, we are $898\times$ (resp. $1585\times$) faster on the same setup.

While our focus is MPC, the VISA model also benefits CPU-emulation-based Zero-Knowledge proof compilers, such as ZEE and EZEE (Heath et al., Oakland’21 and Yang et al., EuroS&P’22).

Keywords: MPC; General Purpose Programs; Garbled Circuits.

*Georgia Institute of Technology, yyang811@gatech.edu.

†Georgia Institute of Technology, stan.peceny@gatech.edu.

‡University of Illinois Urbana-Champaign, daheath@illinois.edu.

§Georgia Institute of Technology, kolesnikov@gatech.edu

Contents

1	Introduction	1
1.1	Case study: Dijkstra’s algorithm	2
1.2	Contribution	4
2	Overview	5
2.1	Notation	6
2.2	VISA Advantages	6
2.3	VISA Technical Challenges and Solutions	7
3	Related Work	8
4	Preliminaries	10
4.1	Garbled RAM	10
4.2	Stacked Garbling (SGC)	10
4.3	Cryptographic Assumptions	10
5	Our VISA	11
5.1	Base ISA	11
5.2	Fragments	13
5.3	Memory Hierarchy	13
5.4	Fragment Generation	14
6	GAR: Our VISA Machine	15
6.1	SGC with GRAM	15
6.2	Our Scheme: Formalization and Theorems	16
7	Evaluation	18
7.1	Implementation and Testing Environment	18
7.2	Benchmarks and Metrics	19
7.3	Overall Improvement	20
7.4	Performance Breakdown and Discussion	20
7.4.1	Breakdown of Wall-Clock Time	20
7.4.2	SGC Savings	21
7.4.3	Communication Rounds and Latency Impact	21
7.4.4	Active Set Sizes	22
7.4.5	Comparison with Straight-Line Circuit Evaluation	22
A	Assembly Benchmarks	27
A.1	PSI	27
A.2	Dijkstra’s	27
A.3	KMP	29
B	Full Garbling Scheme and Proofs	30
B.1	Scheme Definition	30
B.2	Correctness/Security Definitions and Proofs	30

1 Introduction

Secure multi-party computation (MPC) allows mutually untrusting parties to execute programs on their private inputs while revealing only the output. MPC has become relevant in academia and industry. It has been commercially deployed in online auctions, electronic voting, financial technology, and has found many use cases in medicine, privacy-preserving machine learning, and distributed databases.

Typically in MPC, we encode programs as circuits. While any bounded program can be compiled to a circuit, the compiled circuit is often *much* larger than the source program. Real world programs (1) access large arrays of data and (2) use complex control flow. Compiling these two program features often results in huge circuits, and MPC cost scales with the size of the circuit. If we wish to enable secure computation of real-world programs, we must circumvent the cost imposed by compiling these features to circuits.

While the issue of array access can be resolved via oblivious RAM (ORAM) [GO96] or garbled RAM (GRAM) [LO13], complex control flow has gone largely unaddressed.

Straight-line execution Indeed, most existing MPC tools “solve” the control flow problem by *disallowing* complex control flow. Most existing MPC toolchains require that the programmer hand-annotate each loop with a hard-coded upper bound on the number of loop iterations [HHNZ19]. With these annotations, the program becomes a simple *straight-line* program, compatible with the circuit model. A compiler can now unroll each loop precisely the specified number of times, then compile each iteration into gates.

This approach is problematic. At best, annotating programs is an annoyance. At worst, hard-coded loop bounds *ruin performance*, since we must execute each loop iteration, even if the loop should terminate early. Many programs are awkward to write and will have wildly wrong asymptotic complexity. In other words, the programmer is left with an incomplete programming environment where she cannot write every program she might need.

CPU Emulation There is another approach that addresses the control flow problem. Instead of evaluating the program directly, use MPC to *emulate a CPU*, then run the program on that CPU. To achieve this, we must fully emulate a CPU, including its program counter, register file, ALU, and main memory. At each step, the CPU will look up and decode the next instruction, load/store registers, run arithmetic operations, and read/write main memory. In this way, the parties can securely evaluate one instruction at a time.

CPU emulation can securely evaluate arbitrary programs, but at a cost. When we emulate a CPU, we pay to evaluate not just the program, but also the machinery inside the CPU. In comparison to straight-line execution, CPU emulation incurs *very* high per-operation cost. In straight-line execution, the arguments to each program operation are decided statically; in a CPU, the arguments are *dynamic* and must be moved into and out of the register file. Similarly, in straight-line execution, the operation to be performed at each step is decided statically; in a CPU, we must look up the next instruction from a large memory, then conditionally dispatch over each operation in the ISA. In short, while CPU emulation solves the control flow problem, it introduces high overhead.

```

1  #define MAX 100
2  #define MAX_INT 1000000
3  int dijkstra(int *a, int *b) {
4  |   int n = a[0];
5  |   int e = a[1];
6  |   int * node = a + 2;
7  |   int * edge = a + 2 + 101;
8  |   int * weight = a + 2 + 401;
9  |   int vis[MAX];
10 |   int dis[MAX];
11 |   dis[b[0]] = 0;
12 |   int i = 0;
13 |   while (i < n) {
14 |       int bestj = -1, bestdis = MAX_INT;
15 |       int j = 0;
16 |       while (j < n) {
17 |           if( vis[j] == 0 && dis[j] < bestdis ) {
18 |               bestj = j;
19 |               bestdis = dis[j];
20 |           }
21 |           j++;
22 |       }
23 |       vis[bestj] = 1;
24 |       j = node[bestj];
25 |       int bound = node[bestj+1];
26 |       while (j < bound) {
27 |           int newDis = bestdis + weight[j];
28 |           if(newDis < dis[edge[j]])
29 |               dis[edge[j]] = newDis;
30 |           j++;
31 |       }
32 |       i++;
33 |   }
34 |   return dis[b[1]];
35 | }

```

Figure 1: Dijkstra’s algorithm written in C. Each vertical line on the left denotes a contiguous string of instructions that are grouped into a fragment. I.e., this program has seven fragments.

1.1 Case study: Dijkstra’s algorithm

We illustrate the challenge of handling general programs in MPC. Consider Dijkstra’s algorithm¹ (Figure 1).

¹For performance, Dijkstra’s algorithm may be implemented with a *priority queue* containing partial solutions sorted by distance from the start node. Standard Dijkstra is based on a simple array, as is also done in [WGMK16]. We use standard Dijkstra for illustration and direct performance comparison with [WGMK16].

Dijkstra’s is a graph algorithm that computes the shortest path between a source node ($b[0]$) and a target node ($b[1]$). Assume that the graph, the source, and the target are private inputs (e.g., secret-shared between the parties). Both straight-line execution and CPU emulation struggle with this small program.

Straight-line execution Straight-line execution-based tools will generally achieve the wrong cost for the algorithm. Consider a graph (V, E) . Even if we suppose that $|V|$ and $|E|$ are public, this simple program presents a challenge to straight-line execution. The problem is that the loop on lines 26–31 dynamically accesses each edge from a graph node. The number of edges from this node is private, so the loop must be unrolled $|E|$ times to accommodate the maximum possible number of edges. In a cleartext execution of Dijkstra’s, this inner loop will in total iterate only $O(|E|)$ times; in this unrolled version, the inner loop will iterate $O(|E| \cdot |V|)$ times.

Even worse, suppose that Dijkstra’s is not the full program, but rather is a subprocedure of a larger program. Here, $|V|$ and $|E|$ are likely dynamic and should be kept secret. In this case, straight-line execution-based tools must pessimistically assume that $|V|$ and $|E|$ are maximal, ruining performance.

ObliVM [LWN⁺15] showed that for Dijkstra’s algorithm and if $|V|$ and $|E|$ are public, the straight-line approach can reclaim the loop asymptotics via *loop coalescing*. Using loop coalescing, we can flatten the nested loop on lines 13–33 into a single loop with an internal conditional. Then, the number of iterations of this top level loop is a function of $|V|$ and $|E|$, so it is possible to properly bound the loop. See further discussion in Section 3.

While loop coalescing can solve this particular problem, it places a significant burden on the programmer: the programmer must now reason about and properly specify upper bounds on *coalesced* loops. This may be expensive if $|V|$ and $|E|$ are secret, such as if Dijkstra’s is nested inside another data-dependent loop, requiring costly further coalescing or excessive padding. This syntactic transformation produces expensive code that is difficult to further optimize.

CPU emulation CPU emulation correctly implements Dijkstra’s asymptotics², but incurs significant concrete cost.

The state-of-the-art CPU emulator implements a sufficient subset of the MIPS instruction set [WGMK16] to handle Dijkstra’s. This CPU stores the compiled assembly program, the register file, and the main memory in three separate RAMs. [WGMK16] implements RAM using either Circuit ORAM [WCS15] or trivial linear scans, depending on the size of the needed array. Their CPU proceeds by continually fetching and executing instructions.

Storing the program in RAM and applying the fetch-and-execute paradigm discards all useful *static information*, some of which [WGMK16] *manually* reclaims by implementing various heuristics, such as periodic (rather than per-instruction) RAM access. Even applying this heuristic, their number of main memory accesses is suboptimal. Further, they must always access smaller memories to fetch instructions and to read/write registers. Their ALU decodes the instruction and conditionally executes the operation for each instruction type that is statically possible at a given step. As a result, each CPU step is a large circuit that often improves on the circuit-based computation only for problem instances where MPC is impractical.

²To be pedantic, the CPU emulation approach achieves the correct asymptotics modulo polylog factors imposed by ORAM/GRAM. Neither CPU emulation nor straight-line execution, nor indeed our approach, can avoid polylog overhead from ORAM/GRAM [LN18].

Our approach, discussed next, systematically optimizes away many of the principal inefficiencies of [WGMK16] and results in significantly improved performance. For instance, for Dijkstra’s with 100 nodes and 300 edges and when run on the *same* setup, our VISA machine uses $5.8\times$ fewer RAM accesses, consumes $7.3\times$ less bandwidth, and runs $44.9\times$ faster. We are $1585\times$ faster on a 100ms-latency network.

Our Solution: VISA Machines The state of the art presents a dichotomy: CPU emulation or straight-line programs.

In this work, we suggest and explore a hybrid approach to handling arbitrary programs inside MPC. Our *variable instruction set architecture machine*, or VISA machine, handles programs with arbitrary control flow, but avoids most of the overhead of the CPU emulation approach. It uses the statically available context to optimize the scope (and hence the cost) of each execution step.

In short, a VISA machine is distinct from a CPU in that it does not repeatedly execute instructions, but rather repeatedly executes entire *fragments* of the source program. Each fragment is an arbitrarily long straight-line portion of the source program text. The basic advantage of this is that we can cheaply handle each fragment as a circuit. While we still need CPU-like machinery to coordinate the execution of the fragments and ensure privacy, the amount of needed machinery is substantially reduced.

1.2 Contribution

We propose *variable instruction set architectures*, a basic approach to evaluating arbitrary programs inside MPC. We believe that VISAs are *the* sensible approach to executing arbitrary programs in MPC. VISAs do not limit the programmer to straight-line programs, and they do not incur the high overhead of a basic CPU. A VISA adapts to the target program of interest, an appropriate choice for MPC where we generally assume that the parties agree on a program.

In more detail, we:

- Introduce and motivate the VISA model.
- Construct a complete VISA-based 2PC toolchain for assembly programs. Our toolchain is implemented via garbled circuits (GC).
- Resolve technical issues needed to combine core components of a GC-based VISA machine: GC conditional branching [HK20a, HK21b] and Garbled RAM [HKO21].
- Formalize our instantiation as a *garbling scheme* [BHR12] and prove the resulting formalism secure. Our garbling scheme securely evaluates arbitrary assembly programs written in our ISA. Using garbling schemes as the underlying mechanism has two key benefits.
 - First, we dramatically decrease the number of communication rounds, resulting in orders of magnitude improvement (see Section 7.4.3). Prior work [WGMK16, Kel17] used tens of rounds *per CPU step*, while we require one message plus an OT for the entire 2PC.
 - Second, our technique can be elevated to the covert, PVC, and malicious models using standard techniques.
- We implemented VISA machine including, significantly, the first implementation of Garbled RAM [HKO21].

- Experimentally evaluate performance of our toolchain. We ran our VISA machine on a number of assembly benchmarks, including Dijkstra’s, Knuth-Morris-Pratt, and a private set intersection benchmark from [WGMK16]. Our results indicate significant improvement over the prior best approach to arbitrary assembly programs [WGMK16]: we run in constant rounds, use 4–7× less bandwidth, use 5–10× fewer RAM accesses, and run 40–70× faster (up to 1585× with 100ms latency), yielding a machine that executes assembly instructions at 300–4000Hz. We also experimentally show our work, as expected, overtakes circuit-based 2PC (EMP [WMK16]) even for small programs with non-trivial control flow.
- We plan to open source and maintain a cleaned version of our prototype toolchain.
- While our focus is on MPC, the VISA model also directly applies to CPU-emulation-based Zero-Knowledge Proof (ZKP) compilers, such as ZEE and EZEE [HYDK21, YHKD22]. Indeed, they face similar problems of more efficient CPU design (e.g., fragmentation and stacking), ZK ORAM integration with branching, etc., and the VISA approach is similarly beneficial to ZKP compiler work. We leave specific instantiations of ZKP VISA as exciting future work.

Recent breakthrough GC and MPC improvements on free branching [HK20a, HK21b, HKP20, HKP21] and efficient GRAM [HKO21] removed fundamental technical roadblocks needed to move away from straight-line circuit execution. We believe that our hybrid approach – contextual fragment-based execution engines – will underlie the next generation of 2PC and MPC toolchains. This paper initiates this direction and sets the stage for future cryptographic and interdisciplinary work that will likely involve programming language, static analysis, and compiler techniques, and that will interface with high-level programming languages.

2 Overview

In this section we at a high level introduce our model and explain the fundamental benefits of our approach. We then introduce lower-level technical challenges and briefly outline our approach to solving them.

Our basic observation is that CPU emulation is a blunt generic mechanism: CPUs in cleartext machines are static devices that can execute *each* step of *any* program. But in MPC, the program is public, and there is no need to use a fixed generic set of instructions. Instead, we can derive our machine’s ‘instruction types’, which we call *fragments*, from the target program itself.

Each fragment can be arbitrarily large and complex, so long as it does not contain data-dependent loops. We can generate custom circuitry tailored to each fragment, avoiding the need to mechanistically execute the fragment one instruction at a time. Thus, once our machine enters a fragment, we pay essentially no overhead to execute that fragment. In this sense, we obtain the benefit of straight-line execution.

At the same time, our machine dynamically dispatches over the fragments, so we can handle all possible execution paths. In this sense, we obtain the benefit of CPU emulation.

Our execution engine does not necessarily need to dynamically dispatch over each program fragment at each step. At each step it is sufficient to only guarantee execution of fragments that may occur at *this* step. In many useful programs, this *active set* is *much* smaller and consists of cheaper fragments than the full set.

Program fragments are generated by a compiler. There are many choices for how to fragment a program, and good fragmentation is crucial to performance. We discuss related trade-offs (see Section 5.4).

2.1 Notation

Our execution engine repeatedly conditionally dispatches over varying sets of fragments chosen from the target program. We call the specification of a machine that operates this way a *variable instruction set architecture* (VISA). A *VISA machine* instantiates a VISA specification. Our VISA machine, which we call GAR, is implemented via GC; of course, one could implement a VISA machine from different primitives, such as a secret-sharing-based protocol and off-the-shelf ORAM.

At each step i , a VISA machine can execute any fragment in the *active set* of step i . We compose each fragment from many *base instructions* in the program text. Note we thus consider two kinds of instructions: *base instructions* are typical low-level assembly instructions, whereas *fragments* are the instructions of a VISA and are composed from multiple base instructions. Fragments are automatically chosen by a type of compiler that we call a *fragmentation strategy*; our GAR construction includes a built-in fragmentation strategy.

In the remainder of this section, we explain and motivate VISA machines in more detail. We explain our advantages by referring to Dijkstra’s algorithm (Figure 1).

2.2 VISA Advantages

VISA machines do not repeatedly execute instructions, but rather repeatedly execute entire *fragments* of the source program. This leads to several important advantages:

Free register file As each fragment is a straight-line piece of code, we do not need to dynamically store and access local variables from a register file. Instead, like the straight-line approach, a VISA machine routes arguments to operations directly and without cryptographic cost.

We still pay to route the content of the register file *between* fragments, but *within* a single fragment, the register file is free.

Example 1. Consider line 18 of Dijkstra’s (Figure 1). Under CPU emulation, this simple assignment requires reading j from and writing `bestj` to the register file. In practice, these would be implemented by linear scans of a modest array. Linear scans are expensive. As a reference point, suppose the register file holds 16 32-bit registers. Using state-of-the-art GC, each linear scan of this register file costs $\approx 16\text{KB}$ of communication. In the CPU emulation approach, this cost is paid multiple times per CPU cycle. In our VISA machine, this overhead is erased: to handle line 18 the parties may simply agree to name certain wires in the fragment circuit `bestj`.

No instruction memory Programs execute fewer fragments than they do base instructions. Thus, when the VISA machine dynamically decides which fragment to execute next, the space of choices is smaller. This means that the VISA machine does not need to store fragments in an instruction memory. Instead, we conditionally dispatch over an integer that indicates which of the small number of statically known fragments should be executed next. This eliminates many usages of ORAM/GRAM.

Example 2. *In our ISA, Dijkstra’s has 56 instructions³ but only 7 fragments. (Our actual fragmentation is more nuanced; see Section 5.4.) At each step, we conditionally execute only those fragments that are possible. As a simple example, on the first cycle of Dijkstra’s, our VISA machine unconditionally executes the fragment on lines 4–12, since this is statically the only fragment possible. We track the fragments that are possible at each step by tracing the target program’s control flow graph.*

Fewer conditional choices Each fragment implements a larger portion of the overall execution than does each instruction. This is significant because there is overhead associated with conditionally executing code inside MPC, whether classically or by stacking [HK20a]. Since we execute fewer fragments than CPU emulation executes instructions, we make fewer conditional decisions, and hence pay the overhead of conditional branching fewer times. With SGC, this advantage manifests in the fact that we need fewer SGC multiplexer gadgets [HK20a, HK21b]. Importantly, for small branches, these gadgets dominate the cost of SGC.

Example 3. *Running Dijkstra’s with $|V| = 100$ and $|E| = 300$ involves executing 198,814 instructions, and hence making 198,814 conditional decisions. In contrast, we need only execute 21,800 fragments, and hence make only 21,800 conditional decisions.*

Fewer data RAM accesses Since each fragment is static, we know precisely how many times each fragment must move data to/from main memory. This allows a VISA machine to access memory less often than a CPU, since in a CPU it is possible that each instruction is a memory access.

Example 4. *Consider again line 18 of Dijkstra’s. Under CPU emulation, the CPU cannot statically deduce that the current instruction is not a RAM access, so when emulating line 18, it must perform a RAM access. Our VISA machine eliminates this access.*

The sum advantage of our approach as compared to CPU emulation is well illustrated by again considering line 18 of Dijkstra’s. Under CPU emulation, this instruction will involve fetching and decoding the instruction, linearly scanning the register file multiple times, conditionally executing the various instruction types, and accessing main memory. Each of these actions are expensive. In our VISA-based approach, line 18 is free of cryptographic cost.

2.3 VISA Technical Challenges and Solutions

Our core contribution is the introduction of VISA-based MPC. Efficiently implementing an MPC VISA machine presents crypto- and system-technical challenges; we discuss the main challenges here.

Managing the active set Inside a fragment, we have full static knowledge of the straight-line code, so we can directly and efficiently compile the code to a circuit. However, a VISA machine must conditionally execute fragments in the active set at each step.

The cost of this conditional dispatch is greatly improved thanks to the recent line of work on MPC conditional branching, in particular Stacked Garbling (SGC) [HK20a, HK21b]. By integrating

³For readability, Figure 1 is written in C; our machine manipulates low level assembly, and each line of C code can correspond to multiple assembly instructions. See Appendix A for the corresponding assembly.

SGC, we can conditionally dispatch over active set fragments with communication proportional to a single (largest) fragment. Although SGC improves communication, it still requires computation: for b fragments, the computational cost scales with $O(b \log b)$ [HK21b]. Thus, we must not allow the active set to grow too large.

Further, SGC-based conditional branching incurs communication cost that scales with the size of the conditional’s *interface*, i.e., the number of input/output wires, with additional factor dependent on the number of branches b . This cost imposes constraints on the efficiency of using small fragments, and impacts the utility of breaking down fragments, e.g., in alignment with RAM accesses.

In this work, we do not significantly optimize fragments, leaving it as crucial and significant future work. Our fragments are syntactically derived from the control flow structure of the assembly program. This choice is sufficient for modest programs. We envision that future work can use compiler techniques and static analysis to more intelligently select fragments. For example, a fragment can be split into pieces, or multiple fragments can be combined into one. We emphasize the complexity of this problem space: a good solution should simultaneously consider the size of each active set, the size of fragments, the number of RAM accesses, the per-fragment overhead, such as the size of the interface to SGC, etc.

Stacked Garbling with RAM Access Using SGC to conditionally evaluate fragments introduces a subtle technical challenge in handling RAM accesses within fragments. For multiple technical reasons, it is not possible or desired to access RAM directly from inside an SGC conditional branch. This is primarily because GRAM and ORAM reveal random-looking access patterns to the parties. If an access comes from an inactive conditional SGC branch, then SGC’s optimization will reveal information incompatible with the normal access pattern of the GRAM/ORAM. Thus, this use is insecure, as it allows the GC evaluator to identify the active branch in a conditional. See detailed discussion in Section 6. Other issues include the increased computational cost of processing GRAM’s expensive access procedure in each branch. Similar concerns may apply to accessing other types of resources, such as stacks, queues, expensive procedure calls (e.g. non-black-box crypto primitives), or recent improved and unstackable GC techniques [HK21a].

In Section 6, we design a novel mechanism for efficiently and securely handling RAM accesses from within SGC branches. In short, our mechanism allows us to cheaply escape the conditional branch, access the resource, and then re-enter *that same branch*. Each branch can access a resource multiple times. Our mechanism allows fragments to access RAM without paying high cost for SGC gadgets.

We also note the following lower-level contributions:

Entire System and Security Proof We package our approach as a garbling scheme and prove it secure.

Implementation Our system is a non-trivial systems-technical undertaking.

3 Related Work

In our review of related work, we focus on prior general purpose MPC tools.

Straight-line execution tools The vast majority of MPC tools use straight-line execution, e.g. [RHH14, ZE15, DSZ15, WMK16, ACC+22, LHS+14]. These tools require that each program loop has a hard-coded upper bound. CBMC-GC goes one step further by trying to infer loop bounds automatically, but still ultimately models the program as a straight-line circuit [FHK+14].

Straight-line execution cannot suitably support arbitrary programs where the number of loop iterations depends on the data. We note that [HHNZ19] is an excellent systematization of knowledge that explores the pros and cons of such tools.

CPU emulation tools We consider two works that operated in the CPU emulation paradigm [Kel17, WGMK16]. [Kel17] used SPDZ to implement a CPU-emulation-based protocol for malicious adversaries. While their online efficiency is competitive with the total cost of [WGMK16], their offline efficiency is $\approx 100\times$ slower. In our evaluation (Section 7), we accordingly focus our comparison on [WGMK16]. We described [WGMK16]’s approach in Section 1, and we compare to their performance in Section 7.

[WGMK16]’s uses Circuit ORAM [WCS15], which could be modularly swapped for a different ORAM, such as [Ds17], correspondingly affecting (improving) performance. We only compare to the existing system [WGMK16]. Constant-round complexity (and hence using EPIGRAM) is essential for CPU-emulation and VISA MPC due to the sequential nature of RAM accesses in these models. Interactive ORAMs incur latency cost proportional to the (large) number of steps of a typical program (cf. discussion in Section 7.4.3). Further, GRAM can be easily and cheaply upgraded to stronger security models, e.g. covert or malicious, using existing techniques. Such an upgrade for ORAM constructions, including [Ds17], is a challenge.

We note that TinyGarble implemented a MIPS ALU, but did not build on this to implement a working CPU emulation tool [SHS+15]. For example, they do not integrate RAM support to their prototype. Their main contributions are (1) better management of plaintext function by avoiding unrolling it into a plaintext circuit, and (2) applying hardware synthesis tools to reduce the size of the MIPS CPU, improving over naïve by up to 14.95%.

Loop Coalescing. Loop coalescing is a compiler technique explored in the MPC context by [LWN+15] (and in the proof system context by [WSR+15]). The basic idea is to combine the bodies of loops into a single loop with an internal conditional. [LWN+15, WSR+15] show that this can improve MPC (resp. proof system) performance by reducing the number of hard-coded loop bounds in the program (cf. Section 1.1). The technique does not suggest (nor do [LWN+15, WSR+15] explore) further optimization, such as fragment design.

There are common characteristics of loop coalescing and VISA. Both techniques conditionally dispatch over program fragments.

Crucially, VISA approaches MPC optimization holistically, providing a clean abstraction and vocabulary for general optimization of oblivious programs (e.g. include stacking, GRAM, our new gadgets, etc.) and for expressing optimization constraints. Indeed, VISA emphasizes fragment design as a crucial optimization problem. VISA also provides a convenient vocabulary for discussing low level details, such as the size of a register file and managing the active set. See further discussion in Sections 5.3 and 5.4. In contrast, coalescing is a source code transformation, and is at the wrong level of abstraction for fragmentation and low-level optimization.

4 Preliminaries

We implement our VISA machine using garbled circuits (GC). GC allows for powerful protocols that achieve secure computation in only a constant number of protocol rounds. We build on the half-gates GC technique [ZRE15], which requires that the parties communicate two ciphertexts per AND gate and zero ciphertexts per XOR gate [KS08].

We combine the basic [ZRE15] scheme with recent improvements in Garbled RAM [HKO21] and with Stacked Garbling [HK20a, HK21b]. Garbled RAM is needed when accessing data from the VISA machine’s main memory, and Stacked Garbling improves the communication consumption incurred when conditionally handling fragments.

We use these GC improvements heavily, and we overcome technical problems needed to compose them.

4.1 Garbled RAM

Compiling large arrays to Boolean circuits is infeasible. The problem is that on each array access, the circuit must touch each element of the array. Hence, on each access we pay cost proportional to the size of the array. Garbled RAM (GRAM) [LO13] equips GC with random-access arrays that incur only *sublinear cost*. GRAM preserves GC’s important constant-round property.

A recent GRAM, called EPIGRAM, dramatically improved the concrete cost of the technique [HKO21]. We implemented EPIGRAM, and we use it to instantiate our VISA machine’s main memory.

Our formalism manipulates GRAM directly by using two gates provided by EPIGRAM:

- An ARRAY gate takes as input public natural numbers n and w . It outputs a zero-initialized size- n array of width- w elements. We initialize all of our arrays width $w = 32$.
- An ACCESS gate performs an array access. The gate accepts as input (1) an array A , (2) $\log_2 n$ bits that together encode an array index α , (3) a bit rw that indicates if this is a read or a write, and (4) a w -bit value y that indicates what to store in the array if this is a write. As output, the gate yields (1) $A[\alpha]$ and (2) the updated array where the content of index α has been replaced by y iff $rw = 1$.

4.2 Stacked Garbling (SGC)

Until recent breakthrough work [HK20a, HK21b], GC techniques required communication proportional to the computed program, including inactive branches. SGC [HK20a, HK21b] achieves communication proportional to only the single longest *execution path* of the program.

This improvement is a boon to our approach, because we repeatedly conditionally evaluate the target program’s fragments. SGC greatly improves the communication cost of fragments (see Section 7).

4.3 Cryptographic Assumptions

Our garbling scheme (Section 6.2) is secure under a typical GC assumption: We assume that the function H is a circular correlation robust hash function [CKKZ12, ZRE15].

As is standard in MPC (e.g., [GKK⁺12, WGMK16]), total runtime, i.e., the number of CPU emulation steps, is public. If desired, the steps can be padded.

We consider security in the presence of a semi-honest adversary. Since our construction is a garbling scheme, its security can be extended into covert, public verifiable covert (PVC), malicious models using standard techniques.

5 Our VISA

The general idea of a VISA is agnostic of low-level details. Of course, it is interesting to instantiate and experiment with a specific architecture. We formalize our specific VISA here.

Our VISA is built on top of a *base ISA*. Our base ISA is indeed basic, providing primitive instructions that (1) perform algebraic operations, (2) achieve dynamic control flow, and (3) read/write main memory. We first formalize this base ISA. We choose a custom base ISA for simplicity of presentation and implementation; it may be desirable in future work to replace the base ISA with an off-the-shelf ISA, such as MIPS.

Once we establish the base ISA, we formalize our VISA, which essentially aggregates base instructions into fragments.

5.1 Base ISA

The base ISA specifies the instructions that can appear in our supported assembly programs. We emphasize that we do not execute these instructions one by one; rather, our VISA groups base instructions into fragments, and our VISA machine treats fragments as its atomic units of computation.

The base ISA formalizes both the syntax and the semantics of instructions. Our instructions each provide a simple mechanism for performing algebra, achieving control flow, or accessing memory. To define instruction semantics, we define an abstract machine that executes instructions. Our ISA simultaneously defines our instruction set and the abstract machine that runs them.

Definition 1 (Base ISA). *Our instruction set is formally defined in Figure 2. The semantics of instructions are defined by reference to an abstract machine with a program counter pc , a register file \mathcal{R} , a main memory \mathcal{M} , and a program \mathcal{P} . pc is a 32-bit index that indicates which base instruction to execute next. \mathcal{R} is a length- m array of 32-bit integers. \mathcal{M} is a length- n array of 32-bit integers. \mathcal{P} is an array of instructions. Both n and m are configurable parameters of the abstract machine. A machine is initialized with an arbitrary program. At initialization, pc , \mathcal{R} , and \mathcal{M} are zero initialized. At each step, the machine updates itself based on the semantics of instruction $\mathcal{P}[\text{pc}]$.*

In our implementation, we instantiate a machine with a size-13 register file; we vary the size of RAM depending on the requirements of the executed program.

We emphasize that while both the register file and the memory are key-value data structures, our VISA machine handles them very differently. Our memory supports dynamic access and is implemented using Garbled RAM. On the other hand, our register file does not need to implement dynamic access: each usage of the register file is statically specified by an instruction, so each register is essentially just a named collection of 32 circuit wires. Inside a fragment, accessing the register file is free.

	Syntax	Semantics
Algebra	COPY $tar\ src$	$\mathcal{R}[tar] \leftarrow \mathcal{R}[src]$
	CCOPY $tar\ src_0\ src_1$	$\mathcal{R}[tar] \leftarrow \begin{cases} \mathcal{R}[src_1], & \text{if } \mathcal{R}[src_0] = 1 \\ \mathcal{R}[tar], & \text{otherwise} \end{cases}$
	ADD $tar\ src_0\ \{src_1\}$	$\mathcal{R}[tar] \leftarrow \mathcal{R}[src_0] + val(src_1)$
	SUB $tar\ src_0\ \{src_1\}$	$\mathcal{R}[tar] \leftarrow \mathcal{R}[src_0] - val(src_1)$
	MUL $tar\ src_0\ \{src_1\}$	$\mathcal{R}[tar] \leftarrow \mathcal{R}[src_0] \cdot val(src_1)$
	XOR $tar\ src_0\ \{src_1\}$	$\mathcal{R}[tar] \leftarrow \mathcal{R}[src_0] \oplus val(src_1)$
	AND $tar\ src_0\ \{src_1\}$	$\mathcal{R}[tar] \leftarrow \mathcal{R}[src_0] \wedge val(src_1)$
	EQ $tar\ src_0\ \{src_1\}$	$\mathcal{R}[tar] \leftarrow \begin{cases} 1, & \text{if } \mathcal{R}[src_0] = val(src_1) \\ 0, & \text{otherwise} \end{cases}$
	CMP $tar\ src_0\ src_1$	$\mathcal{R}[tar] \leftarrow 2 \cdot (\mathcal{R}[src_0] \stackrel{u}{<} \mathcal{R}[src_1]) + (\mathcal{R}[src_0] \stackrel{s}{<} \mathcal{R}[src_1])$
	SWAP $src_0\ src_1$	$\mathcal{R}[src_0], \mathcal{R}[src_1] \leftarrow \mathcal{R}[src_1], \mathcal{R}[src_0]$
	RS1 dst	$\mathcal{R}[dst] \leftarrow \lfloor \mathcal{R}[dst] / 2 \rfloor$
	IMM $dst\ imm$	$\mathcal{R}[dst] \leftarrow imm$
Control Flow	J imm	$pc \leftarrow imm$
	JE $src\ imm$	$pc \leftarrow \begin{cases} imm, & \text{if } \mathcal{R}[src] \neq 0 \\ pc + 1, & \text{otherwise} \end{cases}$
	JNE $src\ imm$	$pc \leftarrow \begin{cases} imm, & \text{if } \mathcal{R}[src] = 0 \\ pc + 1, & \text{otherwise} \end{cases}$
	JL $src\ imm$	$pc \leftarrow \begin{cases} imm, & \text{if } \mathcal{R}[src] \& 1 \neq 0 \\ pc + 1, & \text{otherwise} \end{cases}$
	JB $src\ imm$	$pc \leftarrow \begin{cases} imm, & \text{if } \mathcal{R}[src] \& 2 \neq 0 \\ pc + 1, & \text{otherwise} \end{cases}$
	HALT	– no effect, pc unchanged –
Memory	LOAD $tar\ src$	$\mathcal{R}[tar] \leftarrow \mathcal{M}[\mathcal{R}[src]]$
	STORE $tar\ src$	$\mathcal{M}[\mathcal{R}[tar]] \leftarrow \mathcal{R}[src]$

$$val(x) \triangleq \begin{cases} x, & \text{if } x \text{ is an immediate} \\ \mathcal{R}[x], & \text{if } x \text{ is a register id} \end{cases}$$

Figure 2: Our base ISA. Each instruction type handles between zero and three arguments. In general, arguments refer to *registers*, but some arguments, denoted $\{\cdot\}$, can also optionally be *immediates* (i.e., compile-time constants). val is a helper function that resolves an argument that can be either a register or an immediate. Unless the semantics otherwise mention an effect on the pc , each instruction also increments the pc . The symbol $<$ with an overset u (resp. s) denotes a comparison where the arguments are treated as an unsigned (resp. signed) integers.

5.2 Fragments

As discussed and motivated in Section 2, batching multiple instructions by creating fragments resolves the bulk of the cost of the CPU emulation approach.

Definition 2 (Fragment). *A fragment is a straight-line sequence of base ISA instructions where only the final instruction may be a Control Flow instruction (c.f. Figure 2).*

Definition 2 coincides with the notion of a program basic block. We still elect to use new terminology because the notion of a fragment can be (and, we expect, will be) generalized, for example by allowing intra-fragment control flow. The only limitation in extending the above definition is that a fragment should never contain a data-dependent loop, since this would break the straight-line nature of the fragment. For simplicity, we do not explore this direction here, but we believe that this can be exploited heavily in future work.

We now define the syntax/semantics of our VISA.

Definition 3 (Our VISA). *Like our base ISA, a VISA is a set of instructions together with the abstract machine that executes them. A VISA instruction is a fragment (Definition 2). The VISA abstract machine is identical to the base ISA machine, except that the program \mathcal{P} consists of fragments, and at each step the machine executes the semantics of the current fragment $\mathcal{P}[\text{pc}]$.*

Remark 1. *Note, a VISA program is thus viewed as including the corresponding variable instruction set. A VISA then specifies the interpretation of the program. A VISA machine instantiates the (secure) execution of the program. While a full toolchain starts from programs written in a base ISA, the VISA definition is about programs that have been fragmented. In practice, the VISA machine toolchain will generate the fragmentation and hence the program’s instruction set.*

While the above *specification* indicates an array lookup $\mathcal{P}[\text{pc}]$, our instantiation dispatches fragments via conditional branching. Note that to achieve the prescribed semantics, we do not need to conditionally dispatch over *each* fragment at each step. In general, not all fragments will be possible at a given step. We reduce the number of conditionally dispatched fragments by considering a control flow graph (CFG) representation of the target program. We maintain a set of pointers into the CFG that indicates the set of possible pc values. At each step, our VISA machine only dispatches over those fragments that are currently pointed to.

5.3 Memory Hierarchy

A VISA introduces the opportunity to distinguish three types of memory:

- **Main memory.** Most program state is stored in a large main memory that is accessed dynamically at high cost.
- **Persistent Registers.** The local state of a VISA machine is held in *persistent registers*. Inside the fragment, these registers are free. However, to conditionally dispatch over fragments, this local state must be passed to each branch. SGC imposes cost for each bit that crosses the interface to/from the conditional. It is sensible to store frequently used data in persistent registers, but the number of these registers should be kept in check.

- **Local Registers.** Since register access is free inside a fragment, a VISA program can introduce arbitrary numbers of local registers, allowing the fragment to store a large state without paying for it. At the exit of the fragment, the content of local registers is lost.

Allocating data to these levels of memory is a large and interesting optimization space. We use 13 persistent registers and a RAM of size up to 2^{13} 32-bit words in our experiments.

5.4 Fragment Generation

As discussed in Section 2, the choice of strategy for breaking a program into fragments can dramatically affect performance. In this work, we align fragments with program basic blocks (i.e., each control flow instruction maps to a fragment), with one exception: we introduce extra fragments for RAM accesses such that each fragment has at most one RAM access. We found that this simple strategy reduces the overall number of RAM accesses⁴, which we found is the performance bottleneck.

Note that for simplicity of presentation, Figure 1 does not show the extra fragments resulting from RAM accesses. Our actual fragmentation has 14 fragments.

While we leave further in-depth exploration of intelligently selecting fragments as significant future work, we outline several guidelines for such strategies. We note that these guidelines sometimes contradict one another, as fragment optimization is a challenging problem.

Generate fragments such that each conditional dispatch is over fragments of similar size and with a similar number of RAM accesses SGC, and other approaches to MPC free branching [HKP20, HKP21], achieves communication cost proportional to the single most expensive branch. To best take advantage of free branching, ensure that branches have similar cost. This can be achieved, e.g., by splitting large program basic blocks into more than one fragment and/or by merging multiple basic blocks into a single fragment.

RAM access is an expensive resource; an unbalanced allocation across dispatched fragments misses an opportunity to amortize accesses.

Prefer larger fragments This reduces the number of VISA machine steps. Hence, larger fragments further reduce the amount of CPU-emulation-style machinery.

Compress the interface to each fragment As explained in Section 5.3, we pay to transport the content of persistent registers into and out of branches. Using compiler techniques to reduce the number of needed persistent registers will reduce cost.

Prefer fragmentation that leads to smaller active sets SGC computational and interface costs scale with the number of branches, so we should seek to reduce the number of branches per step (i.e., to shrink each active set). One way this guideline might be achieved is by artificially introducing periodicity into a program’s execution. For instance, we can split each loop into a number of fragments that is a power of two. Without periodicity in consecutive loops the active set will tend to grow with each step until it includes each program fragment. Artificially introducing periodicity groups fragments into “congruence classes” and ensures that most fragments never

⁴I.e., all active set fragments will have a same number of accesses.

coincide in the same active set. [WGMK16] considered a similar technique in their MIPS processor. Introducing periodicity for fragments introduces further opportunities to align code and amortize cost.

6 GAR: Our VISA Machine

This section introduces GAR (Garbled Assembly with RAM) our implementation of the VISA machine. GAR is formalized as a garbling scheme [BHR12]. As already mentioned, GAR conditionally dispatches fragments using SGC and implements main memory via GRAM.

We first discuss technical issues and our solution in combining our two main building blocks, SGC and GRAM. Then, in Section 6.2 we present the GAR scheme and state the main security theorem (proofs are presented in the Appendix).

6.1 SGC with GRAM

The incompatibility of SGC and GRAM SGC is compatible with many, but not all GC techniques. SGC requires that the string of material encoding each branch be indistinguishable from a uniform string. This restriction is needed to mask from the GC evaluator the identity of the conditional’s active branch: if a branch is inactive, SGC arranges that the evaluator obtains uniform *garbage material*.

Unfortunately, GRAM’s material *is distinguishable* from a uniform string. In short, GRAM will one-by-one reveal to the evaluator RAM indices that are randomly generated *without replacement* [HKO21]. These revealed indices are indistinguishable from a uniform permutation, but not from a uniform string. Thus it is not secure to use GRAM’s ACCESS gate inside an SGC conditional.

SGC’s uniform string requirement and GRAM’s revealed uniform permutations seem somewhat inherent to the techniques, and it is not clear that we can revise these techniques to make them compatible with one another. Even if it were possible to make the two techniques compatible, it would not be desirable. SGC requires that each party garble each branch multiple times, introducing added computational cost. Since the GRAM access procedure is large, we would like to avoid repeatedly garbling it. It is more pragmatic to simply garble each access once, as we end up doing.

Our Approach One way we could handle RAM access in a VISA machine would be to place each RAM access instruction in its own single-instruction fragment. While correct and secure, the approach violates several of our guidelines for program fragmentation (Section 5.4), and is undesirable for a number of performance reasons. In particular, the resulting fragments are smaller, more numerous, and each RAM access will service a smaller fragment. Ultimately, this discards many of the VISA’s benefits.

A much better way would be to temporarily escape a fragment just to perform the RAM access, then re-enter that fragment. This is the approach we take. We design a new scheme that allows us to temporarily escape an SGC conditional branch (i.e., a fragment), perform the access, then re-enter that same branch. Because we escape the SGC branch before accessing RAM, we avoid SGC’s uniform string requirement. Thus, RAM access is simulatable. Crucially for performance, our gadgets escape, and not fully exit SGC, and transfer across the SGC interface only those specific

bits that are directly related to the RAM access. Thus, we do not, for example, pay to transfer the full register file on each RAM access.

Instrumenting GRAM access in SGC SGC uses two garbled gadgets, the *demux* and the *mux*, to enter and exit a conditional, respectively. Each of these gadgets handles branch input/output wire-by-wire, where each wire is (indirectly) connected from the outside of the conditional to the internal circuit of each branch. We refer to each of these wire connections as a *port* of the demux/mux. There is one port per external wire.

Our observation is that, in contrast with standard SGC, the demux/mux need not be evaluated in one shot at the very beginning/end of the conditional. Instead, the GC evaluator can process ports of the gadgets in an *arbitrary* order, so long as data dependencies in the circuit are satisfied.

This in particular means that the evaluator can (1) process input to a branch by handling only some ports of the demux, (2) evaluate some gates in that branch, generating input to a RAM query, (3) feed the RAM query through ports in the mux to temporarily escape the branch, (4) execute the RAM access outside of SGC, in plain GC, (5) feed the RAM result through ports of the demux back into the branch, and (6) continue evaluation of the branch.

Interestingly, the GC generator’s order of building the corresponding GC material is different. Because each branch must be generated from a seed (this is a key trick behind SGC’s improvement), the generator garbles each branch all at once, before any RAM accesses are handled. As part of doing so, he assigns uniformly random GC labels to the branch side of each demux port. Only once each branch is fully generated, does he generate GC for RAM access(es). Labels of these GCs match the labels of the ports of the SGC conditional. Finally, he generates the GC material for the demux and mux.

We include a diagram illustrating this process in additional detail in Appendix C (Figure 14).

Our modification to SGC still uses the main ideas of Stacked Garbling [HK20a]: our GC generator garbles each branch starting from a distinct PRG seed and then stacks the material together using XOR. Our GC evaluator can decrypt the seed for each inactive branch and hence can reconstruct their garblings, unstack the material for the active branch, and evaluate. I.e., our scheme retains the important communication advantage of SGC.

Next, we formalize our full garbling scheme GAR, which includes the above trick.

6.2 Our Scheme: Formalization and Theorems

We formalize our VISA machine as a *garbling scheme* [BHR12]. SGC [HK21b] and GRAM [HKO21] are also formalized as garbling schemes; our scheme reorganizes and adjusts their procedures, making them compatible with each other and with our VISA (Section 5).

At a high level, our scheme should be understood as a new SGC scheme equipped with black-box GRAM. As an aside, it is possible to replace black-box GRAM with other garbled resources, for example a stack or queue [ZE13].

Program Description A garbling scheme securely handles any program from some specified language. Our goal is to support programs expressed in our base ISA (Figure 2). At the lowest level, we have primitive support for AND gates [ZRE15], XOR gates [KS08], SWITCH statements [HK21b], and ARRAY and ACCESS gates [HKO21]. The semantics of XOR and AND gates are natural; ARRAY and ACCESS gate semantics are specified in Section 4.1. A SWITCH executes only the

indicated branch and outputs the result. We group instructions from our base ISA, then compile these to our low level primitives. Thus, our formal garbling scheme consists of three major steps:

- **Compile base ISA program to VISA program.** Our scheme first groups base ISA instructions into fragments using the strategy described in Section 5.4.
- **Compile VISA program to primitives.** We compile each fragment primitive operation using standard techniques. Each basic instruction has a corresponding straight-line circuit, and our scheme stitches together each of the circuits in the fragment. To conditionally dispatch, the scheme wraps the fragment circuits in a SWITCH.
- **Evaluate primitives via GC.** The most interesting step is the evaluation of primitives, which is explained below.

Note that the first two steps of our handling are quite modular. It is easy to replace the ISA to VISA compiler with one that, for example, more intelligently selects fragments. Similarly, we could replace the compiler from fragments to circuits with more sophisticated techniques. From here, our scheme focuses on the handling of primitives, which is its crypto-technical component.

Definition 4 (Primitive Circuit Program). *A primitive circuit program is a circuit consisting of AND gates, XOR gates, ARRAY gates, ACCESS gates, and SWITCH statements. ARRAY and ACCESS gates are defined in Section 4.1. A SWITCH statement is recursively parameterized over b primitive circuit programs and $\lceil \log_2 b \rceil$ wires that indicate which branch to execute. Note that an ACCESS gate is allowed inside a SWITCH statement.*

Our GAR scheme handles arbitrary assembly programs by appropriately implementing the above circuit primitives. We note that our assembly compiler generates restricted classes of circuit programs, and we need not handle them in full generality of Definition 4. For example, the resulting primitive program will not feature nested conditionals, and the ARRAY gate will be used exactly once to initialize main memory *MEM* at the start of the program. Furthermore, each ACCESS gate will be parameterized by the specific array *MEM*. Looking ahead, our formalism will handle only the relevant special forms of primitive circuit programs.

We are now ready to present our main construction, the GAR⁵ (Garbled Assembly with RAM) garbling scheme [BHR12].

Construction 1 (GAR). *GAR consists of three components:*

- *A fragmentation strategy that specifies how to convert a base ISA program into a VISA program. GAR uses the strategy discussed in Section 5.4; we do not formally specify further.*
- *A compiler that transforms a VISA program (Definition 3) into a primitive circuit program (Definition 4); because each fragment has no data-dependent control flow, compiling each fragment to a primitive circuit program is straight-forward, and we do not specify further.*
- *A garbling scheme that securely executes primitive circuit programs (Definition 4).*

The GAR garbling scheme is the tuple of procedures:

$$(GAR.ev, GAR.Ev, GAR.Gb, GAR.En, GAR.De)$$

⁵A gar is a predatory and particularly menacingly looking fish.

Note, GAR’s functionality goes beyond simply instantiating a VISA machine; in particular GAR fragments programs written in base ISA and generates VISA programs. We could have treated this functionality separately as part of our toolchain.

We formally present procedures of the GAR garbling scheme in Appendix B, see Figures 10 to 13. Here, we review them at a high level.

GAR.ev This procedure defines the semantics of primitives (Definition 4). The semantics of AND gates and XOR gates are natural. ARRAY and ACCESS gate semantics are specified in Section 4.1. A SWITCH executes only the indicated branch and outputs the result.

GAR.Ev This procedure specifies the GC evaluator’s handling. In short, the handling of primitives is inherited from prior work [ZRE15, HKO21]. The exception is our new SWITCH primitive, which supports RAM ACCESS inside its branches. We discussed our method for handling ACCESS gates from within a SWITCH in Section 6.1.

GAR.Gb This procedure specifies the GC generator’s handling. Again, the handling of primitives is inherited from prior work [ZRE15, HKO21]. See Section 6.1 for the handling of our SWITCH primitive.

GAR.En This procedure specifies how cleartext GC inputs are mapped to GC labels. The procedure is standard: on each input wire, a zero maps to one label and a one maps to a different label.

GAR.De This procedure specifies how output GC labels are mapped to cleartext outputs. The procedure is standard.

GAR meets the standard garbling scheme definitions of *correctness*, *authenticity*, *obliviousness* and *privacy* [BHR12]. Meeting these is sufficient to instantiate 2PC/MPC protocols. We state the definitions and prove that GAR meets them in Appendix B.2. Theorems and lemmas in Appendix B.2 imply the following:

Theorem 1 (Main). *Assuming a circular correlation robust hash function, GAR’s garbling scheme is correct, authentic, oblivious, and private.*

7 Evaluation

7.1 Implementation and Testing Environment

We implemented GAR and used it to instantiate a semi-honest 2PC protocol in ≈ 5200 lines of C/C++. We instantiated Oblivious Transfer and Network I/O using the EMP Toolkit [WМК16]. We ran our experiments on two m6i.16xlarge⁶ machines in the same region of an Amazon EC2 cluster. One machine ran the GC generator and the other ran the GC evaluator. We also ran [WМК16]’s implementation on the exact same setup to establish our baseline.

We configured both systems with the same inputs and with RAM of the same size. GAR handles the program written in our assembly language. [WМК16] takes a MIPS binary compiled by an off-the shelf compiler, thus placing them somewhat at a disadvantage.

⁶Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz, 256GiB Memory, 25Gbps Network

Program	# Mem Ent		Time (s)			Comm. (GB)			# RAM Accesses		
			[WGMK16]	Ours	Impr.	[WGMK16]	Ours	Impr.	[WGMK16]	Ours	Impr.
PSI	64	128	29.4	0.6	49.0 ×	0.7	0.1	6.0 ×	2455	256	9.6 ×
	256	512	311.7	7.1	43.9 ×	7.1	1.8	3.9 ×	9755	1024	9.5 ×
	1024	2048	4378.6	61.9	70.7 ×	102.0	17.1	6.0 ×	38855	4096	9.5 ×
Dijkstra's	40	1024	2601.3	62.6	41.6 ×	58.5	9.5	6.2 ×	21579	3902	5.5 ×
	60		5462.7	127.3	42.9 ×	127.0	18.5	6.9 ×	46679	8302	5.6 ×
	80		9723.0	216.5	44.9 ×	221.0	30.5	7.2 ×	81379	14202	5.7 ×
	100		14910.0	332.2	44.9 ×	341.0	47.0	7.3 ×	125779	21802	5.8 ×

Figure 3: Comparison of our GAR system with [WGMK16]. We run PSI and Dijkstra’s for a range of input sizes. We ran both GAR and [WGMK16] on the same hardware setup. Our approach substantially improves wall-clock time, communication consumption, and RAM usage. Our count of [WGMK16]’s RAM accesses *does not* include instruction fetching; We list them separately in Figure 5.

G’s Pattern Length	E’s String Length	GRAM Size (32-bits)	# Executed Fragments	Gb. Time(s)	Trans. Time(s)	Ev. Time(s)	Total Time(s)	#Inst.	Speed #Inst/s	#Mem Access	Comm. (GB)
50	400	512	1400	9.7	5.3	4.4	19.4	12535	646Hz	1401	2.53
150	700	1024	2700	21.2	15.4	9.8	46.4	23635	509Hz	2701	7.44
250	1500	2048	5500	49.6	47.1	23.1	119.8	48735	407Hz	5501	22.70
500	7000	8192	23000	275.3	358.7	131.3	765.3	209485	274Hz	23001	173.00

Figure 4: GAR’s evaluation on KMP with different inputs and GRAM sizes.

7.2 Benchmarks and Metrics

As explained in Section 1 and further demonstrated in Section 7.4.5, straight-line execution is not feasible for programs with complex control flow. Accordingly, we focus comparison on [WGMK16]’s CPU emulation approach. We demonstrate significant improvement on three programs.

- Private Set Intersection (PSI): Two parties each hold a sorted integer array and wish to compute the number of common elements. While fast tailored PSI protocols exist, we use this benchmark for direct comparison with prior work [WGMK16].
- Dijkstra’s Shortest Path (Dijkstra): One party holds a directed graph while the other holds a pair of source and destination nodes (This is the setting of [WGMK16]; other input configurations, e.g., all inputs secret-shared, incur no extra cost). Parties wish to compute the shortest path between the two nodes. This benchmark was used by [WGMK16] and [LWN+15].
- Knuth–Morris–Pratt String Search (KMP): One party inputs a pattern string and the other inputs a search string. They wish to compute the number of occurrences of the pattern in the search string. This benchmark was suggested by [LHS+14].

We note that our reported runtimes for [WGMK16] are in some cases slower than what was reported in [WGMK16] itself. We believe this is due to the fact that program runtime is variable and depends on the program input. Crucially, we ran our GAR system on the same input as [WGMK16], and thus our reported numbers are directly comparable.

Assembly code for each benchmark is included in Appendix A.

We report the following metrics:

- **Wall-clock Time:** Wall-clock time includes the time needed for the GC generator to garble, for network transmission, and for the GC evaluator to evaluate. Figure 6 provides a breakdown of these three metrics.

Input Size	PSI			Dijkstra			
	64	256	1024	40	60	80	100
Thousands of Fetches	2.5	9.8	38.9	21.6	46.7	81.4	125.8

Figure 5: [WGMK16]’s number of instruction fetch RAM accesses.

- **Communication:** Both [WGMK16] and GAR communicate through one TCP/IP connection. We directly measure communication from the TCP port and report the amount of transmitted data.
- **# RAM accesses:** RAM accesses are the most expensive operation in our approach. Recall that [WGMK16] uses ORAM while GAR uses GRAM. We report the number of times we and [WGMK16] access RAM. Recall that [WGMK16] uses RAM to fetch instructions; we do not. Figure 3 *does not* include [WGMK16]’s RAM accesses to fetch instructions. We list them separately in Figure 5.

7.3 Overall Improvement

Figure 3 tabulates GAR’s improvement over [WGMK16] for PSI and Dijkstra’s.

PSI We ran the PSI benchmark on three different pairs of input arrays: two 64-element arrays, two 256-element arrays, and two 1024-element arrays. The PSI program primarily consists of a loop that compares a single element from each array. Our VISA approach captures PSI’s loop in a *single fragment*. This results in simple control flow and high performance. In total, we use only three fragments: one that initializes state before the loop, one that implements the body of the loop, and one that handles the end of the program. Because the loop is captured by one fragment, our approach uses precisely the number of RAM accesses that are prescribed by the program’s execution path.

GAR is 44–70× faster than [WGMK16] on each input, uses 4–6× less bandwidth, and uses ≈ 10× fewer RAM accesses.

Dijkstra’s We ran Dijkstra’s (Figure 1) on a graph with 40, 60, 80, and 100 nodes. For a graph with n nodes, we set the number of edges $|E| = 3n$. The sparse graph is stored in the adjacency list. Our program is split into 14 fragments.

For each input, GAR is 42–45× faster than the baseline [WGMK16], uses 6–7× less bandwidth, and uses 5–6× fewer RAM accesses.

7.4 Performance Breakdown and Discussion

7.4.1 Breakdown of Wall-Clock Time

Figure 6 breaks down the wall-clock-time for each of our runs of PSI and Dijkstra. No one cost clearly stands out as the bottleneck. We note that we did not stream the GC from the generator to the evaluator; we expect that proper streaming would allow to overlap garbling and evaluation with transmission, essentially eliminating the separate cost of garbling and evaluation. We also include detailed GAR costs in Figure 4.

Bench.		Gb.	Trans.	Ev.	Total
		Time(s)	Time(s)	Time(s)	Time(s)
PSI	64	0.2	0.3	0.1	0.6
	256	2.3	3.7	1.1	7.1
	1024	18.3	35.9	7.7	61.9
Dijkstra	40	29.4	19.5	13.7	62.6
	60	61.0	38.3	28.0	127.3
	80	103.6	62.9	50.0	216.5
	100	159.1	97.0	76.1	332.2

Figure 6: Breakdown of our total wall-clock time into the time needed to (1) garble the circuit, (2) transmit the GC across the network, and (3) evaluate the GC.

Program		#Inst.	#Fragments	Impr.	Speed #Inst/s
PSI	64	2419	129	18.8 ×	4032Hz
	256	9715	513	18.9 ×	1368Hz
	1024	38899	2049	19.0 ×	628Hz
Dijkstra’s	40	33914	3900	8.7 ×	542Hz
	60	73714	8300	8.9 ×	579Hz
	80	128614	14200	9.1 ×	594Hz
	100	198814	21800	9.1 ×	598Hz

Figure 7: Comparison of the number of fragments with the number of base ISA instructions for PSI and Dijkstra’s with different input sizes. The improvement ($\frac{\#instructions}{\#fragments}$) illustrates that we reduce the number of execution steps by an order of magnitude. Our Hz rate is base instructions per second.

7.4.2 SGC Savings

Recall that we use SGC to stack GC material from fragments in the active set. In our Dijkstra experiments, we observed that SGC improved communication by roughly 3×, excluding the cost of GRAM access. We expect that this improvement will become more significant for larger and more complex programs where the total number of fragments and likely the size of the active set will be larger.

7.4.3 Communication Rounds and Latency Impact

GAR is implemented via a garbling scheme, and our instantiation in the semi-honest model only requires performing (parallel) OTs and sending a single message from the generator to the evaluator. In contrast, [WGMK16]’s ORAM-based CPU uses multiple rounds of communication *per RAM access*.

This distinction is not significant in the ultra-low latency setup we have explored so far, but even modest latency harshly penalizes multi-round approaches. We evaluate the impact by executing one program (PSI-256) with various latencies. We used the Linux traffic control tool `tc` to configure a network with 2Gbps bandwidth and 0/50/100ms latencies. Figure 8 tabulates wall-clock time

	Round-trip Delay		
	0ms	50ms	100ms
[WGMK16] (s)	314.5	11495.4 (+11180.9)	22984.8 (+22670.3)
GAR (s)	11.5	12.8 (+1.3)	14.5 (+3.0)
Improvement	27.3×	898.1×	1585.2×

Figure 8: Runtime comparison of GAR and [WGMK16] when solving PSI-256 for different latency settings. To tightly control latency, we ran these experiments on a single machine with a simulated (via the Linux `tc` command) 2Gbps network. For clarity, we note the added cost of latency in parentheses. Note, the sliding window in TCP implicitly forces latency-like delays on GAR.

<i>G</i> 's Pattern Length	<i>E</i> 's String Length	EMP #AND Gates ($\times 10^9$)	EMP Total Time(s)	GAR Total Time(s)	Impr.
50	400	0.142	7.4	19.4	0.4×
150	700	2.264	120.0	46.4	2.6×
250	1500	13.11	732.7	119.8	6.1×
500	7000	233.0	12843.3	765.3	16.8×

Figure 9: Comparison of GAR with EMP’s straight-line execution on KMP. Our improvement over EMP increases with larger inputs. Note, EMP implements array lookup with linear scans, not GRAM. For very small arrays, linear scans outperform EPIGRAM, which explains EMP’s performance in the smallest instance.

performance. (In this experiment we run both parties on a single machine, so measurements in Figure 8 are not identical to our other experiments.) With higher latency, GAR’s execution speed is almost unchanged, but [WGMK16] becomes significantly slower; GAR’s advantage grows from 27× on a 0ms latency network to 898× (resp. 1585×) with 50ms (resp. 100ms) latency.

7.4.4 Active Set Sizes

GAR (and VISA) performance declines with the increase of sizes of active sets (i.e., sets of fragments that can be possibly executed in the corresponding step). Let M be the total number of fragments. In our experiments we observe that the active set size starts with 1 and quickly grows to $M - 1$ as execution proceeds: $M = 4$ (resp. 15 and 11) for PSI (resp. Dijkstra and KMP). We view active set optimization as crucial future work.

7.4.5 Comparison with Straight-Line Circuit Evaluation

Finally, we illustrate the advantage of the VISA approach over straight-line circuit evaluation by comparing with the semi-honest 2PC of the widely adopted EMP Toolkit [WMK16]. We implemented the Knuth-Morris-Pratt (KMP) string-searching algorithm in both GAR and EMP. (We do not include [WGMK16]’s performance, because their repository did not include this benchmark.) Figure 9 tabulates the results, and Figure 4 presents a fine-grained analysis of GAR’s performance results.

KMP searches for occurrences of a length- k pattern held by G in a length- m string held by E and outputs the number of occurrences. An important feature of KMP is its $O(m + k)$ time complexity, rather than the naive $O(m \cdot k)$. Circuits are not a suitable representation as KMP contains an inner loop that must be pessimistically unrolled a total of $O(m \cdot k)$ times when in fact only $O(m + k)$ total iterations are needed.

Acknowledgment

This Work is supported in part by Cisco research award and NSF awards CNS-2246353, CNS-2246354, and CCF-2217070. This material is also based upon work supported in part by DARPA under Contract No. HR001120C0087. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA.

References

- [ACC⁺22] Abdelrahman Aly, Benjamin Coenen, Kelong Cong, Karl Koch, Marcel Keller, Dragos Rotaru, Oliver Scherer, Peter Scholl, Nigel P. Smart, Titouan Tanguy, and Tim Wood. SCALE-MAMBA software. <https://homes.esat.kuleuven.be/~nsmart/SCALE/>, 2022.
- [BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, October 2012.
- [CKKZ12] Seung Geol Choi, Jonathan Katz, Ranjit Kumaresan, and Hong-Sheng Zhou. On the security of the “free-XOR” technique. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 39–53. Springer, Heidelberg, March 2012.
- [Ds17] Jack Doerner and abhi shelat. Scaling ORAM for secure computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 523–535. ACM Press, October / November 2017.
- [DSZ15] Daniel Demmler, Thomas Schneider, and Michael Zohner. ABY - A framework for efficient mixed-protocol secure two-party computation. In *NDSS 2015*. The Internet Society, February 2015.
- [FHK⁺14] Martin Franz, Andreas Holzer, Stefan Katzenbeisser, Christian Schallhart, and Helmut Veith. Cbmc-gc: An ansi c compiler for secure two-party computations. In Albert Cohen, editor, *Compiler Construction*, pages 244–249, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- [GKK⁺12] S. Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Fernando Krell, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure two-party computation in sublinear (amortized) time. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 513–524. ACM Press, October 2012.

- [GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 43(3):431–473, 1996.
- [HHNZ19] Marcella Hastings, Brett Hemenway, Daniel Noble, and Steve Zdancewic. SoK: General purpose compilers for secure multi-party computation. In *2019 IEEE Symposium on Security and Privacy*, pages 1220–1237. IEEE Computer Society Press, May 2019.
- [HK20a] David Heath and Vladimir Kolesnikov. Stacked garbling - garbled circuit proportional to longest execution path. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 763–792. Springer, Heidelberg, August 2020.
- [HK20b] David Heath and Vladimir Kolesnikov. Stacked garbling for disjunctive zero-knowledge proofs. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 569–598. Springer, Heidelberg, May 2020.
- [HK21a] David Heath and Vladimir Kolesnikov. One hot garbling. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 574–593. ACM Press, November 2021.
- [HK21b] David Heath and Vladimir Kolesnikov. LogStack: Stacked garbling with $O(b \log b)$ computation. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 3–32. Springer, Heidelberg, October 2021.
- [HKO21] David Heath, Vladimir Kolesnikov, and Rafail Ostrovsky. Practical garbled RAM: GRAM with $O(\log^2 n)$ overhead. Cryptology ePrint Archive, Report 2021/1519, 2021. <https://eprint.iacr.org/2021/1519>.
- [HKP20] David Heath, Vladimir Kolesnikov, and Stanislav Peceny. MOTIF: (almost) free branching in GMW - via vector-scalar multiplication. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 3–30. Springer, Heidelberg, December 2020.
- [HKP21] David Heath, Vladimir Kolesnikov, and Stanislav Peceny. Masked triples - amortizing multiplication triples across conditionals. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 319–348. Springer, Heidelberg, May 2021.
- [HYDK21] David Heath, Yibin Yang, David Devecsery, and Vladimir Kolesnikov. Zero knowledge for everything and everyone: Fast ZK processor with cached ORAM for ANSI C programs. In *2021 IEEE Symposium on Security and Privacy*, pages 1538–1556. IEEE Computer Society Press, May 2021.
- [Kel17] Marcel Keller. The oblivious machine - or: How to put the C into MPC. In Tanja Lange and Orr Dunkelman, editors, *LATINCRYPT 2017*, volume 11368 of *LNCS*, pages 271–288. Springer, Heidelberg, September 2017.
- [KS08] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 486–498. Springer, Heidelberg, July 2008.

- [LHS⁺14] Chang Liu, Yan Huang, Elaine Shi, Jonathan Katz, and Michael W. Hicks. Automating efficient RAM-model secure computation. In *2014 IEEE Symposium on Security and Privacy*, pages 623–638. IEEE Computer Society Press, May 2014.
- [LN18] Kasper Green Larsen and Jesper Buus Nielsen. Yes, there is an oblivious RAM lower bound! In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 523–542. Springer, Heidelberg, August 2018.
- [LO13] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 719–734. Springer, Heidelberg, May 2013.
- [LWN⁺15] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. ObliVM: A programming framework for secure computation. In *2015 IEEE Symposium on Security and Privacy*, pages 359–376. IEEE Computer Society Press, May 2015.
- [RHH14] Aseem Rastogi, Matthew A. Hammer, and Michael Hicks. Wysteria: A programming language for generic, mixed-mode multiparty computations. In *2014 IEEE Symposium on Security and Privacy*, pages 655–670. IEEE Computer Society Press, May 2014.
- [SHS⁺15] Ebrahim M. Songhori, Siam U. Hussain, Ahmad-Reza Sadeghi, Thomas Schneider, and Farinaz Koushanfar. TinyGarble: Highly compressed and scalable sequential garbled circuits. In *2015 IEEE Symposium on Security and Privacy*, pages 411–428. IEEE Computer Society Press, May 2015.
- [WCS15] Xiao Wang, T.-H. Hubert Chan, and Elaine Shi. Circuit ORAM: On tightness of the Goldreich-Ostrovsky lower bound. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 850–861. ACM Press, October 2015.
- [WGMK16] Xiao Shaun Wang, S. Dov Gordon, Allen McIntosh, and Jonathan Katz. Secure computation of MIPS machine code. In Ioannis G. Askoxylakis, Sotiris Ioannidis, Sokratis K. Katsikas, and Catherine A. Meadows, editors, *ESORICS 2016, Part II*, volume 9879 of *LNCS*, pages 99–117. Springer, Heidelberg, September 2016.
- [WMK16] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. EMP-toolkit: Efficient Multi-Party computation toolkit. <https://github.com/emp-toolkit>, 2016.
- [WSR⁺15] Riad S. Wahby, Srinath T. V. Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *NDSS 2015*. The Internet Society, February 2015.
- [YHKD22] Yibin Yang, David Heath, Vladimir Kolesnikov, and David Devecsery. Ezee: Epoch parallel zero knowledge for ansi c. In *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*, pages 109–123. IEEE, 2022.
- [ZE13] Samee Zahur and David Evans. Circuit structures for improving efficiency of security and privacy tools. In *2013 IEEE Symposium on Security and Privacy*, pages 493–507. IEEE Computer Society Press, May 2013.

- [ZE15] Samee Zahur and David Evans. Obliv-C: A language for extensible data-oblivious computation. Cryptology ePrint Archive, Report 2015/1153, 2015. <https://eprint.iacr.org/2015/1153>.
- [ZRE15] Samee Zahur, Mike Rosulek, and David Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 220–250. Springer, Heidelberg, April 2015.

SUPPLEMENTARY MATERIAL

A Assembly Benchmarks

A.1 PSI

```
1 BB_BEGIN:
2     IMM    %GR0, 0
3     IMM    %GR1, 64 (resp. 256, 1024)
4     IMM    %GR2, 0
5 BB0:
6     LOAD   %GR3, %GR0
7     LOAD   %GR4, %GR1
8     CMP    %GR5, %GR3, %GR4
9     RS1    %GR5
10    ADD    %GR0, %GR0, %GR5
11    CMP    %GR5, %GR4, %GR3
12    RS1    %GR5
13    ADD    %GR1, %GR1, %GR5
14    EQ     %GR5, %GR3, %GR4
15    ADD    %GR0, %GR0, %GR5
16    ADD    %GR2, %GR2, %GR5
17    EQI    %GR5, %GR0, 64 (resp. 256, 1024)
18    EQI    %GR6, %GR1, 128 (resp. 512, 2048)
19    XOR    %GR5, %GR5, %GR6
20    JNE    %GR5, BB0
21 BB_END:
22    HALT
23
```

A.2 Dijkstra's

```
1 BB_BEGIN:
2     IMM    %GR5, 2
3     LOAD   %GR5, %GR5
4     IMM    %GR6, 3
5     LOAD   %GR6, %GR6
6     IMM    %GR4, 0
7     LOAD   %GR4, %GR4
8     ADDI   %GR4, %GR4, 900
9     IMM    %GR0, 0
10    STORE  %GR4, %GR0
11 BB_0:
12    IMM    %GR2, 4294967295
13    IMM    %GR3, 2147483647
14    IMM    %GR1, 0
15 BB_1:
16    ADDI   %GR4, %GR1, 800
17    LOAD   %GR8, %GR4
18    EQI    %GR8, %GR8, 0
19 BB_SP0:
20    ADDI   %GR4, %GR1, 900
```

```

21      LOAD   %GR4, %GR4
22      CMP    %GR7, %GR4, %GR3
23      RS1    %GR7
24      ADD    %GR7, %GR7, %GR8
25      RS1    %GR7
26      CCOPY %GR2, %GR7, %GR1
27      CCOPY %GR3, %GR7, %GR4
28      ADDI   %GR1, %GR1, 1
29      EQ     %GR4, %GR1, %GR5
30      JNE    %GR4, BB_1
31  BB_2:
32      ADDI   %GR4, %GR2, 800
33      IMM    %GR7, 1
34      STORE  %GR4, %GR7
35  BB_SP1:
36      ADDI   %GR4, %GR2, 4
37      LOAD   %GR1, %GR4
38  BB_SP2:
39      ADDI   %GR2, %GR2, 5
40      LOAD   %GR2, %GR2
41  BB_3:
42      ADDI   %GR4, %GR1, 405
43      LOAD   %GR4, %GR4
44  BB_SP3:
45      ADD    %GR4, %GR4, %GR3
46      ADDI   %GR7, %GR1, 105
47      LOAD   %GR7, %GR7
48  BB_SP4:
49      ADDI   %GR8, %GR7, 900
50      LOAD   %GR8, %GR8
51      CMP    %SPO, %GR4, %GR8
52      RS1    %SPO
53      CCOPY %GR8, %SPO, %GR4
54      ADDI   %GR7, %GR7, 900
55  BB_SP5:
56      STORE  %GR7, %GR8
57      ADDI   %GR1, %GR1, 1
58      EQ     %GR4, %GR1, %GR2
59      JNE    %GR4, BB_3
60  BB_4:
61      ADDI   %GRO, %GRO, 1
62      EQ     %GR4, %GRO, %GR5
63      JNE    %GR4, BB_0
64  BB_END:
65      IMM    %GRO, 1
66      LOAD   %GRO, %GRO
67  BB_SP6:
68      ADDI   %GRO, %GRO, 900
69      LOAD   %GRO, %GRO
70      HALT

```

A.3 KMP

```
1  BB_BEGIN:
2      IMM    %GR5, 0
3      LOAD   %GR5, %GR5
4      IMM    %GR6, 1
5      LOAD   %GR6, %GR6
6      IMM    %GR0, 1
7  BB0:
8      SUBI   %GR1, %GR0, 1
9      ADDI   %GR1, %GR1, 50
10     LOAD   %GR1, %GR1
11  BBSP0:
12     ADDI   %GR4, %GR0, 100
13     LOAD   %GR4, %GR4
14  BB1:
15     ADDI   %GR3, %GR1, 100
16     LOAD   %GR3, %GR3
17     EQI    %GR8, %GR1, 0
18     EQ     %GR2, %GR4, %GR3
19     ADD    %GR2, %GR2, %GR8
20     EQI    %GR2, %GR2, 0
21     JNE    %GR2, BB2
22     SUBI   %GR1, %GR1, 1
23     ADDI   %GR1, %GR1, 50
24     LOAD   %GR1, %GR1
25     J      BB1
26  BB2:
27     EQ     %GR2, %GR4, %GR3
28     ADD    %GR2, %GR2, %GR1
29     ADDI   %GR3, %GR0, 50
30     STORE  %GR3, %GR2
31     ADDI   %GR0, %GR0, 1
32     EQ     %GR2, %GR0, %GR5
33     JNE    %GR2, BB0
34  BB3:
35     IMM    %GR7, 0
36     IMM    %GR1, 0
37     IMM    %GR0, 0
38  BB4:
39     ADDI   %GR3, %GR0, 150
40     LOAD   %GR3, %GR3
41  BB5:
42     ADDI   %GR4, %GR7, 100
43     LOAD   %GR4, %GR4
44     EQI    %GR8, %GR7, 0
45     EQ     %GR2, %GR3, %GR4
46     ADD    %GR2, %GR2, %GR8
47     EQI    %GR2, %GR2, 0
48     JNE    %GR2, BB6
49     SUBI   %GR7, %GR7, 1
50     ADDI   %GR7, %GR7, 50
51     LOAD   %GR7, %GR7
52     J      BB5
53  BB6:
```



```

54      EQ      %GR2, %GR3, %GR4
55      ADD     %GR7, %GR7, %GR2
56      EQ      %GR2, %GR7, %GR5
57      ADD     %GR1, %GR1, %GR2
58      SUBI    %GR8, %GR7, 1
59      ADDI    %GR8, %GR8, 50
60      LOAD    %GR8, %GR8
61      CCOPY   %GR7, %GR2, %GR8
62      ADDI    %GR0, %GR0, 1
63      EQ      %GR2, %GR0, %GR6
64      JNE     %GR2, BB4
65 BB_END:
66      HALT

```

B Full Garbling Scheme and Proofs

B.1 Scheme Definition

We provide the formal procedures for GAR. Figure 10 lists the scheme procedures (i.e., Construction 1) of GAR. Figure 11 explains how we handle ACCESS gates internal to SGC branches. Figure 12 and Figure 13 are unrolled modifications of the COND gate procedures from LOGSTACK.

In Definition 4, we define primitive programs as having explicit AND gates, XOR gates, etc. For brevity and to closely match the procedures of [HK21b], which are conceptually quite similar, we use slightly different syntax our figures. I.e., a *netlist* is a sequence of AND and XOR gates. Netlists are handled via the [ZRE15] garbling scheme. ‘Cond’ statements correspond to the SWITCH keyword. ‘Seq’ statements denote two circuit components that are run in sequence. We emphasize that this language-level difference does not change the meaning of primitive circuit programs.

B.2 Correctness/Security Definitions and Proofs

In this appendix, we prove our garbling scheme correct and secure. Security definitions are adapted from [BHR12]. Our garbling scheme is proved secure assuming the existence of a *circular correlation robust hash function*. The following definition is adapted from [ZRE15]:

Definition 5 (Circular Correlation Robustness). *Let H be a function. We define two oracles:*

- $\text{circ}_\Delta(x, i, b) = H(x \oplus \Delta, i) \oplus b\Delta$ where $\Delta \in \{0, 1\}^{\kappa-1}$.
- $\mathcal{R}(x, i, b)$ is a random function with κ -bit output.

A sequence of oracle queries (x, i, b) is legal when the same value (x, i) is never queried with different values of b . H is a circular correlation robust hash function if for all poly-time adversaries A :

$$\left| \Pr_{\Delta} [\mathcal{A}^{\text{circ}_\Delta}(1^\kappa) = 1] - \Pr_{\mathcal{R}} [\mathcal{A}^{\mathcal{R}}(1^\kappa) = 1] \right| \text{ is negligible.}$$

Circular correlation robustness ensures that a valid sequence of queries is computationally indistinguishable from uniform randomness, even when we use the important Free XOR optimization [KS08].

```

GAR.ev( $\mathcal{C}, \vec{x}$ ) :
  ▷  $\vec{m}$  is a global array initialized to all 0.
  ▷ What are the circuit semantics?
  Switch  $\mathcal{C}$  :
    case Netlist( $\cdot$ ) : return HALFGATES.ev( $\mathcal{C}, \vec{x}$ )
    case Seq( $\mathcal{C}_0, \mathcal{C}_1$ ) : return GAR.ev( $\mathcal{C}_1, \text{GAR.ev}(\mathcal{C}_0, \vec{x})$ )
    case Cond( $\vec{\mathcal{C}}$ ) :
      ▷ split branch index from input
       $\alpha \mid \vec{x}' \leftarrow \vec{x}$ 
      ▷ Run the active branch.
      return GAR.ev( $\vec{\mathcal{C}}[\alpha], \vec{x}'$ )
    case Access( $i\vec{d}x, \vec{val}, rw, \vec{dst}$ ) :
      if  $\vec{x}[rw] == 0$  :
        ▷ Load
         $\vec{x}[\vec{dst}] \leftarrow \vec{m}[\vec{x}[i\vec{d}x]]$ 
      else :
        ▷ Store
         $\vec{m}[\vec{x}[i\vec{d}x]] \leftarrow \vec{x}[\vec{val}]$ 
      return  $\vec{x}$ 

GAR.Ev( $\mathcal{C}, M, \vec{X}$ ) :
  ▷ MEM is a unique global Array object by EPIGRAM.init( $\cdot$ )
  ▷ How does E evaluate the GC?
  Switch( $\mathcal{C}$ ) :
    case Netlist( $\cdot$ ) : return HALFGATES.Ev( $\mathcal{C}, M, \vec{X}$ )
    case Seq( $\mathcal{C}_0, \mathcal{C}_1$ ) :
       $M_0 \mid M_{tr} \mid M_1 \leftarrow M$ 
      return GAR.Ev( $\mathcal{C}_1, M_1, \text{trans.Ev}(\text{GAR.Ev}(\mathcal{C}_0, M_0, \vec{X}), M_{tr})$ )
    case Cond( $\vec{\mathcal{C}}$ ) : return EvCond( $\vec{\mathcal{C}}, M, \vec{X}$ )
    case Access( $i\vec{d}x, \vec{val}, rw, \vec{dst}$ ) :
      ▷ Call GRAM access procedure
       $MEM, \vec{Y} \leftarrow \text{EPIGRAM.access}(MEM, \vec{X}[i\vec{d}x], \vec{X}[\vec{val}], \vec{X}[rw], M)$ 
       $\vec{X}[\vec{dst}] \leftarrow \vec{Y}$ 
      return  $\vec{X}$ 

GAR.Gb( $1^\kappa, \mathcal{C}, S$ )
  ▷ MEM is a unique global Array object by EPIGRAM.init( $\cdot$ )
  ▷ How does G garble the GC?
  ▷ S is an explicit seed.
  Switch  $\mathcal{C}$  :
    case Netlist( $\cdot$ ) :
      return HALFGATES.Gb( $1^\kappa, \mathcal{C}, S$ )
    case Seq( $\mathcal{C}_0, \mathcal{C}_1$ ) :
      ▷ Derive seeds for two circuits.
       $S_0 \leftarrow F_S(0)$ 
       $S_1 \leftarrow F_S(1)$ 
       $(M_0, e_0, d_0) \leftarrow \text{GAR.Gb}(1^\kappa, \mathcal{C}_0, S_0)$ 
       $(M_1, e_1, d_1) \leftarrow \text{GAR.Gb}(1^\kappa, \mathcal{C}_1, S_1)$ 
      ▷ Labels out of  $\mathcal{C}_0$  must be translated
      ▷ to labels into  $\mathcal{C}_1$ .
       $M_{tr} \leftarrow \text{trans.Gb}(d_0, e_1)$ 
       $M \leftarrow M_0 \mid M_{tr} \mid M_1$ 
      return  $(M, e_0, d_1)$ 
    case Cond( $\vec{\mathcal{C}}$ ) : return GbCond( $\vec{\mathcal{C}}, S$ )
    case Access( $i\vec{d}x, \vec{val}, rw, \vec{dst}$ ) :
      ▷ Generate encodings for access arguments
       $e_{i\vec{d}x} \leftarrow \text{repeatedly-sample } F_S(\cdot)$ 
       $e_{\vec{val}} \leftarrow \text{repeatedly-sample } F_S(\cdot)$ 
       $e_{rw} \leftarrow \text{repeatedly-sample } F_S(\cdot)$ 
      ▷ Call GRAM access procedure
       $d, MEM, M_{acc} \leftarrow \text{EPIGRAM.access}(MEM, e_{i\vec{d}x}, e_{\vec{val}}, e_{rw})$ 
      return  $(M_{acc}, e_{i\vec{d}x} \mid e_{\vec{val}} \mid e_{rw}, d)$ 

```

Figure 10: GAR’s garbling scheme. The included algorithms are typical except for the handling of conditionals. *Ev* and *Gb* delegate the core of our approach: *EvCond* (Figure 13) and *GbCond* (Figure 12). *En* and *De* are not listed as they are standard.

<p>$GAR.Ev^*(\mathcal{C}, M, \vec{X}) :$</p> <ul style="list-style-type: none"> ▷ Similar as Ev until access case case $Access(\vec{id}x, \vec{val}, rw, \vec{dst}) :$ <li style="padding-left: 20px;">$M_{acc} \mid M_{tr} \leftarrow M$ <li style="padding-left: 20px;">▷ The access happened outside conditional <li style="padding-left: 20px;">▷ M should be \perp if \vec{X} has \vec{dst} ready : return \vec{X} else : ▷ HALT and return \vec{X} to top layer 	<p>$GAR.Gb^*(1^\kappa, \mathcal{C}, S)$</p> <ul style="list-style-type: none"> ▷ Similar as Gb until access case case $Access(\vec{id}x, \vec{val}, rw, \vec{dst}) :$ <li style="padding-left: 20px;">▷ Generate encodings <li style="padding-left: 20px;">$e, d \leftarrow repeatedly\text{-sample } F_S(\cdot)$ <li style="padding-left: 20px;">▷ Treat it as inputs and outputs return (\perp, e, d)
--	--

Figure 11: Variants for Gb and Ev . These variants are called inside $GbCond$ and $EvCond$. We implicitly use $function^*$ to denote $function$ from where the underlying calling to Gb (resp. Ev) is replaced by Gb^* (resp. Ev^*).

Definition 6 (Correctness). *A garbling scheme is **correct** if for all valid circuits \mathcal{C} , all input strings \vec{x} of length $\text{inpSize}(\mathcal{C})$, and all pseudorandom seeds S :*

$$De(d, Ev(\mathcal{C}, M, En(e, \vec{x}))) = ev(\mathcal{C}, \vec{x})$$

where $(M, e, d) = Gb(1^\kappa, \mathcal{C}, S)$

Theorem 2. *GAR's garbling scheme is correct.*

Proof. By induction on the structure of \mathcal{C} .

In short, each of the cases of induction, except the handling of conditionals, is trivially inherited from prior work.

- When initializing the main memory MEM , correctness is inherited from EPIGRAM's array initialize procedure.
- Suppose \mathcal{C} is a Netlist. Correctness is inherited from [ZRE15].
- Suppose \mathcal{C} is a sequence $Seq(\mathcal{C}_0, \mathcal{C}_1)$. Correctness holds trivially by induction.
- Suppose \mathcal{C} is an access $Access$. Correctness is inherited from EPIGRAM's array access procedure.
- Suppose \mathcal{C} is a conditional $Cond(\vec{\mathcal{C}})$.

We proceed by further induction on the maximum number of accesses across each branch, i.e., $c \triangleq \text{ComputeMaxAcc}(\vec{\mathcal{C}})$:

- **Base Case:** Suppose c is 0. I.e., there is no access in any branch. In this case, we delegate to a normal conditional gate as specified by LOGSTACK. Specifically, $GbCond$ and $EvCond$ will delegate to the functions defined in LOGSTACK. Correctness is inherited.

```

GbCond( $\vec{C}, S$ ) :
   $b \leftarrow |\vec{C}|$ 
  ▷ Find maximum RAM access count among all branches from the circuits.
   $c \leftarrow \text{ComputeMaxAcc}(\vec{C})$ 
  ▷ Recursively derive all 'good' seeds for the entire tree.
   $s \leftarrow \text{DeriveSeedTree}(S, b)$ 
  ▷ Sample input/output encodings for the conditional.
   $e \leftarrow \text{GenProjection}(S, \text{inpSize}(\text{Cond}(\vec{C})))$ 
   $d \leftarrow \text{GenProjection}(S, \text{outSize}(\text{Cond}(\vec{C})) + c * (n + 33))$ 
  ▷ Parse encoding into encoding of  $\alpha$  and encoding of rest of input.
   $e_\alpha \mid e' \leftarrow e$ 
  ▷ Parse encoding into encoding of  $c$  access arguments and encoding of rest of outputs.
   $d_0 \mid \dots \mid d_{c-1} \mid d' \leftarrow d$ 
  ▷ Garble SortingHat based on the encoding of  $\alpha$ .
  ▷ This outputs material as well as the tree of all 'bad' seeds  $s'$ .
   $M_{\text{SortingHat}}, s' \leftarrow \text{SortingHat.Gb}(e_\alpha, s)$ 
  ▷ Construct the stacked material and input encodings for each branch.
   $M_{\text{cond}}, e_{\text{cond}}, d_{\text{cond}} \leftarrow \text{GbSubtreeFromSeed}^*(\vec{C}, 0, b - 1, s_{0,b-1})$ 
  ▷ Parse encodings for each branch into encodings on  $c$  access results and rest of inputs.
   $e_{\text{cond},0} \mid \dots \mid e_{\text{cond},c-1} \mid e'_{\text{cond}} \leftarrow e_{\text{cond}}$ 
  ▷ Parse encodings for each branch into encodings on  $c$  access arguments and rest of outputs.
   $d_{\text{cond},0} \mid \dots \mid d_{\text{cond},c-1} \mid d'_{\text{cond}} \leftarrow d_{\text{cond}}$ 
  ▷ The demux conditionally translates the input encoding  $e'$ 
  ▷ to one of the branch encodings in  $e'_{\text{cond}}$  based on  $e_\alpha$ .
   $M_{\text{dem}}, \Lambda_{\text{in}} \leftarrow \text{demux.Gb}(e_\alpha, e', e'_{\text{cond}})$ 
  for  $i \in 0..c-1$  :
     $\Lambda_{\text{out}} \leftarrow \text{ComputeGarbage}^*(\vec{C}, M_{\text{cond}}, \Lambda_{\text{in}}, s, s')$ 
    ▷ Extract garbage on  $i$  access argument wires.
     $\Lambda_{\text{arg}} \leftarrow \text{ExtractArg}(\vec{C}, i, \Lambda_{\text{out}})$ 
     $M_{\text{accum},i} \leftarrow \text{mux.Gb}(e_\alpha, d_i, d_{\text{cond},i}, \Lambda_{\text{arg}})$ 
    ▷ Get encoding of access outputs from GRAM procedure.
     $e_i, \text{MEM}, M_{\text{acc},i} \leftarrow \text{EPIGRAM.access}(\text{MEM}, d_i)$ 
     $M_{\text{accdemux},i}, \Lambda'_{\text{in}} \leftarrow \text{demux.Gb}(e_\alpha, e_i, e_{\text{cond},i})$ 
    ▷ Merge existing garbage with new garbage encodings.
     $\Lambda_{\text{in}} \leftarrow \Lambda_{\text{in}} \cup \Lambda'_{\text{in}}$ 
  ▷ Compute all possible garbage outputs.
   $\Lambda_{\text{out}} \leftarrow \text{ComputeGarbage}^*(\vec{C}, M_{\text{cond}}, \Lambda_{\text{in}}, s, s')$ 
  ▷ The demultiplexer collects garbage outputs.
   $M_{\text{mux}} \leftarrow \text{mux.Gb}(e_\alpha, d', d'_{\text{cond}}, \Lambda_{\text{out}})$ 
   $M_{\text{ram}} \leftarrow M_{\text{accum},0} \mid M_{\text{acc},0} \mid M_{\text{accdemux},0} \mid \dots \mid M_{\text{accum},c-1} \mid M_{\text{acc},c-1} \mid M_{\text{accdemux},c-1}$ 
  return  $(M_{\text{SortingHat}} \mid M_{\text{dem}} \mid M_{\text{cond}} \mid M_{\text{mux}} \mid M_{\text{ram}}, e, d')$ 

```

Figure 12: The algorithm for garbling a conditional with b branches where each branch has at most c RAM accesses. Main memory is a length- 2^n GRAM with 32-bit entries. Our procedure follows the structure of LOGSTACK's procedure of the same name. Our colored boxes highlight differences as compared to LOGSTACK, and the green box highlights the most important modification.

```

EvCond( $\vec{C}, M, X$ ) :
   $b \leftarrow |\vec{C}|$ 
   $\triangleright$  Find maximum RAM access count among all branches from the circuits.
   $c \leftarrow \text{ComputeMaxAcc}(\vec{C})$ 
   $\triangleright$  Parse the active branch index from the rest of the input.
   $\alpha \mid X' \leftarrow X$ 
   $\triangleright$  Parse material for gadgets and body of conditional.
   $M_{\text{SortingHat}} \mid M_{\text{dem}} \mid M_{\text{cond}} \mid M_{\text{mux}} \mid M_{\text{ram}} \leftarrow M$ 
   $M_{\text{accmux},0} \mid M_{\text{acc},0} \mid M_{\text{accdemux},0} \mid \dots \mid M_{\text{accmux},c-1} \mid M_{\text{acc},c-1} \mid M_{\text{accdemux},c-1} \leftarrow M_{\text{ram}}$ 
   $\triangleright$  Run SortingHat to compute all of  $E$ 's seeds.
   $\text{es} \leftarrow \text{SortingHat.Ev}(\alpha, M_{\text{SortingHat}})$ 
   $\triangleright$  Run the demultiplexer to compute input (except access results) for each branch  $C_i$ .
   $\vec{X}_{\text{cond}} \leftarrow \text{demux.Ev}(\alpha, X, M_{\text{dem}})$ 

   $\triangleright$  We define a recursive subprocedure that evaluates  $C_i - C_j$  using material  $M$ .
  EvCond'( $i, j, M_{i,j}$ ) :
    if  $i = j$  :
       $\triangleright$  Base case: compute output by evaluating the branch normally.
       $\triangleright$  This base case corresponds to  $\text{guess} = i$ .
       $\triangleright$  Accumulate output labels into the vector  $\vec{Y}_{\text{cond}}$  (for later garbage collection).
       $\vec{Y}_{\text{cond}}[i] \leftarrow \text{Ev}^*(C_i, M, \vec{X}_{\text{cond}}[i])$ 
    else :
       $k \leftarrow \text{halfway}(i, j)$ 
       $\triangleright$  Garble the right subtree using the available seed,
       $\triangleright$  unstack, and recursively evaluate the left subtree.
       $M_{k+1,j}, \cdot, \cdot \leftarrow \text{GbSubtreeFromSeed}^*(\vec{C}, k+1, j, \text{es}_{k+1,j})$ 
      EvCond'( $i, k, M_{i,j} \oplus M_{k+1,j}$ )
       $\triangleright$  Symmetrically evaluate the right subtree.
       $M_{i,k}, \cdot, \cdot \leftarrow \text{GbSubtreeFromSeed}^*(\vec{C}, i, k, \text{es}_{i,k})$ 
      EvCond'( $k+1, j, M_{i,j} \oplus M_{i,k}$ )

  for  $i \in 0..c-1$  :
     $\triangleright$  Start recursive process from the top of the tree.
    EvCond'(0,  $b-1, M_{\text{cond}}$ )
     $\triangleright$  Extract garbage on  $i$  access argument wires.
     $\vec{A}_{\text{cond}} \leftarrow \text{ExtractArg}(\vec{C}, i, \vec{Y}_{\text{cond}})$ 
     $\triangleright$  Mux to get encodings of the access arguments.
     $A \leftarrow \text{mux.Ev}(\alpha, \vec{A}_{\text{cond}}, M_{\text{accmux},i})$ 
     $\triangleright$  Get encodings of the access results.
     $\text{MEM}, R \leftarrow \text{EPIGRAM.access}(\text{MEM}, A, M_{\text{acc},i})$ 
     $\triangleright$  Demux to get encodings of the access arguments for each branch.
     $\vec{X}'_{\text{cond}} \leftarrow \text{demux.Ev}(\alpha, R, M_{\text{accdemux},i})$ 
     $\triangleright$  Merge to get updated encodings of inputs.
     $\vec{X}_{\text{cond}} \leftarrow \vec{X}_{\text{cond}} \cup \vec{X}'_{\text{cond}}$ 
   $\triangleright$  Start recursive process from the top of the tree.
  EvCond'(0,  $b-1, M_{\text{cond}}$ )
   $\triangleright$  Eliminate garbage and propagate  $\vec{Y}_\alpha$  via the multiplexer.
  return  $\text{mux.Ev}(\alpha, \vec{Y}_{\text{cond}}, M_{\text{mux}})$ 

```

Figure 13: E 's procedure, `EvCond`, evaluates a conditional with b branches and at most c RAM accesses. Our procedure follows the structure of `LOGSTACK`'s procedure of the same name. Colored boxes highlight the major differences as compared to `LOGSTACK`, and the green box highlights the most important modification.

- **Induction Step:** Suppose the scheme is correct for $c = k$ accesses. That is, after k iterations of the green boxes in GbCond and EvCond, there will be correct encodings on each wire of the active branch \mathcal{C}_α and correct garbage on each wire of each inactive branch $\mathcal{C}_{i \neq \alpha}$. Now, consider $c = k + 1$. I.e., there is one additional access. W.l.o.g, suppose *each* branch has $k + 1$ access gates. By induction, after k iterations of each green box, there will be correct encodings and garbage on each wire except those directly connected to each branch’s last access gate response. Notice that for each such last access gate, the input wires do not depend on the output wires by construction. Thus for the last ACCESS gate of the active branch \mathcal{C}_α , the encodings on these wires in are correct. Focusing on the last iteration of the green boxes, ComputeGarbage* will produce correct garbage on the inputs to the last ACCESS gate on each inactive branch $\mathcal{C}_{i \neq \alpha}$. With these values as inputs, outputs of *mux* (for this ACCESS query) will correctly encode the arguments of this single extra access. Together with the correctness of EPIGRAM, the outputs of the access procedure correctly encode the accessed value. Crucially, the MEM object will also be updated correctly. With these results as input, the *demux* (for this ACCESS response) will generate correct encodings of the accessed element for the active branch and correct garbage for each inactive branch.

Thus, the overall conditional is correct.

GAR is correct. □

Definition 7 (Obliviousness). *A garbling scheme is **oblivious** if there exists a simulator \mathcal{S}_{obv} such that for any circuit \mathcal{C} and all inputs \vec{x} of length $\text{inpSize}(\mathcal{C})$, the following are indistinguishable:*

$$(\mathcal{C}, M, \vec{X}) \stackrel{c}{\equiv} \mathcal{S}_{obv}(1^\kappa, \mathcal{C})$$

where S is uniform, $(M, e, \cdot) = Gb(1^\kappa, \mathcal{C}, S)$ and $\vec{X} = En(e, \vec{x})$.

Obliviousness ensures that the garbled circuit with input labels can be simulated, and hence reveals no extra information to E .

To prove obliviousness, we use the following lemma, which follows from [ZRE15]’s, [HK20b]’s, and [HKO21]’s security proofs.

Lemma 1. *If H is circular correlation robust (Definition 5), then there exist simulators \mathcal{S}_{AND} , $\mathcal{S}_{\text{SortingHat}}$, $\mathcal{S}_{\text{trans}}$, $\mathcal{S}_{\text{demux}}$, \mathcal{S}_{mux} , $\mathcal{S}_{\text{ARRAY}}$, $\mathcal{S}_{\text{ACCESS}}$ that securely simulate the material generated by the respective GC procedures.*

For the most part, these simulators are straightforward, since most of these GC procedures are implemented as straightforward and standard encrypted truth tables. We note that EPIGRAM’s ARRAY and ACCESS simulators are significantly more involved, but are given in [HKO21].

Theorem 3. *If H is circular correlation robust (Definition 5), GAR’s garbling scheme is oblivious.*

Proof. By construction of a simulator $\mathcal{S}_{obv}(1^\kappa, \mathcal{C})$.

For the most part, the simulator precisely follows the generator’s GC procedure *Gb*, except that the simulator simulates each GC subprocedure by invoking the appropriate simulator.

Technically, our proof proceeds by a hybrid argument where we one-by-one substitute the real handling of each program component by its simulator. The details of this hybrid argument are standard, so we focus on simulating each component.

In more detail, \mathcal{S}_{obv} starts by simulating each input label with a uniform string. This is indistinguishable from real since the real garbling procedure assigns input labels from a PRG. \mathcal{S}_{obv} then steps through the circuit gate by gate. Each gate is handled by calling the appropriate simulator and updating the gate's output wire:

- For each ARRAY gate, simulate by calling \mathcal{S}_{ARRAY} from EPIGRAM.
- For the Netlist, simulate each internal AND gate by calling \mathcal{S}_{AND} from [ZRE15] and simulate each XOR gate by xoring the two simulated inputs.
- For the sequence $\text{Seq}(\mathcal{C}_0, \mathcal{C}_1)$ gate, recursively simulate left part and right one. Use \mathcal{S}_{trans} to simulate the translation gate that forwards the output of \mathcal{C}_0 as input to \mathcal{C}_1 .
- For the access Access gate, simulate by calling \mathcal{S}_{access} from EPIGRAM and update MEM accordingly.
- For the conditional $\text{Cond}(\vec{\mathcal{C}})$ gate, simulate **GbCond**:
 - Compute static information about $\vec{\mathcal{C}}$ by calling appropriate procedures. For instance, compute c by calling **ComputeMaxAcc**.
 - Simulate input/output encodings for the conditional with uniform strings. They are indistinguishable from real since **GbCond** chooses encodings by sampling from a PRG.
 - Simulate each seed s with a uniform string. These are indistinguishable from real since **GbCond** samples them from a PRG.
 - Simulate $M_{\text{SortingHat}}$ by calling $\mathcal{S}_{\text{SortingHat}}$. Notice that seeds s' can be computed by evaluating on the simulated material for the gadget and the simulated labels for α .
 - Simulate $M_{cond}, e_{cond}, d_{cond}$ with uniform strings. Indistinguishability holds M_{cond} due to the *strong stackability* of circuits [HK21b].
 - Simulate M_{dem} by calling \mathcal{S}_{dem} . Notice that the garbage values on inputs of the conditional can be calculated from simulated M_{dem} and simulated labels of conditional inputs.
 - For each iteration, do following steps in order: simulate garbage outputs of this iteration by calling the procedure **ComputeGarbage***; simulate M_{accmux} by calling \mathcal{S}_{mux} ; simulate M_{acc} by calling \mathcal{S}_{access} from EPIGRAM; simulate $M_{accdemux}$ by calling \mathcal{S}_{demux} . Each step is updates simulated labels, garbage, and the simulated MEM object.
 - Simulate M_{mux} by calling \mathcal{S}_{mux} .

GAR is oblivious. □

Definition 8 (Privacy). *A garbling scheme is **private** if there exists a simulator \mathcal{S}_{prv} such that for any circuit \mathcal{C} and all inputs \vec{x} of length $\text{inpSize}(\mathcal{C})$, the following are computationally indistinguishable:*

$$(M, \vec{X}, d) \stackrel{c}{\approx} \mathcal{S}_{prv}(1^\kappa, \mathcal{C}, \vec{y}),$$

where S is uniform, $(M, e, d) = \text{Gb}(1^\kappa, \mathcal{C}, S)$, $\vec{X} = \text{En}(e, \vec{x})$, and $\vec{y} = \text{ev}(\mathcal{C}, \vec{x})$.

Privacy ensures that E , who is given access to (M, \vec{X}, d) , learns nothing except what can be learned from the output \vec{y} . I.e., G 's input is protected.

We note that for most garbling schemes, privacy holds by a straightforward and unsurprising argument. For instance, our proof is very similar to that of [HKO21].

Theorem 4. *If H is circular correlation robust (Definition 5), GAR’s garbling scheme is private.*

Proof. By construction of a simulator \mathcal{S}_{prv} .

In short, privacy follows from obliviousness (Theorem 3) and from the definition of De . \mathcal{S}_{prv} performs the following steps:

- Simulate a GC by invoking the obliviousness simulator (Theorem 3):

$$(\mathcal{C}, M', \vec{X}') \leftarrow \mathcal{S}_{obv}(1^\kappa, \mathcal{C})$$

- Evaluate \mathcal{C} using M' and \vec{X}' :

$$\vec{Y}' \leftarrow \text{Ev}(\mathcal{C}, M', \vec{X}')$$

- For an output decoding string d' that causes Y' to correctly decode to y . Specifically, for the i -th output, let $nonce_i$ denote the nonce used in the decoding truth table. Let $r_i \xleftarrow{\$} \{0, 1\}^\kappa$ be a uniform string. Set:

$$d'_i \leftarrow \begin{cases} (H(Y_i, nonce_j), r_i) & \text{if } y_i = 0 \\ (r_i, H(Y_i, nonce_j)) & \text{if } y_i = 1 \end{cases}$$

- Output (M', \vec{X}', d') .

The simulated output is indistinguishable from (M, \vec{X}, d) because (1) (M', \vec{X}') and (M, \vec{X}) are indistinguishable due to obliviousness, (2) the decrypted rows of d and d' each yield output y , and (3) the undecrypted rows of d and d' are indistinguishable thanks to the circular correlation robust hash function H .

GAR is private. □

Definition 9 (Authenticity). *A garbling scheme is **authentic** if for all circuits \mathcal{C} , all inputs \vec{x} of length $\text{inpSize}(\mathcal{C})$, and all poly-time adversaries \mathcal{A} the following probability is negligible in κ :*

$$\Pr\left(\vec{Y}' \neq \text{Ev}(\mathcal{C}, M, \vec{X}) \wedge De(d, \vec{Y}') \neq \perp\right)$$

where S is uniform, $(M, e, d) = Gb(1^\kappa, \mathcal{C}, S)$, $\vec{X} = En(e, \vec{x})$, and $\vec{Y}' = \mathcal{A}(\mathcal{C}, M, \vec{X})$

Authenticity ensures that an adversary cannot compute GC output labels except by running the scheme as intended. For most garbling schemes, authenticity holds by a standard and unsurprising argument. Indeed, the following proof is very similar to that of [HKO21].

Theorem 5. *If H is circular correlation robust, GAR’s garbling scheme is authentic.*

Proof. By the definition of the privacy simulator (Theorem 4) and the definition of De .

Authenticity allows \mathcal{A} access to material M and to encoded input \vec{X} . To derive a contradiction, let (M', \vec{X}', d') be a GC constructed by \mathcal{S}_{prv} . Now, suppose \mathcal{A} is given (M', \vec{X}') as input. By the definition of \mathcal{S}_{prv} , it should be infeasible for \mathcal{A} to forge an encoded output \vec{Y}' that successfully decodes. Indeed, to produce a value \vec{Y}' that successfully decodes, \mathcal{A} must *guess* a security parameter-length uniform string r_i that was sampled by the simulator and that is independent of \mathcal{A} ’s view. This is clearly infeasible. That is, \mathcal{A} cannot forge an output when given a simulated GC.




Figure 14: Sketch of our procedure for handling RAM accesses from within SGC branches.

If \mathcal{A} can forge output when given a *real* GC, then we can clearly construct a poly-time *privacy distinguisher*. I.e., there is a poly-time procedure that on input \mathcal{C} , M , \vec{X} distinguishes real from simulation with non-negligible probability. To do so, this distinguisher simply does the following: (1) Evaluate \mathcal{C} normally with material M and input \vec{X} to obtain encoded output \vec{Y} , (2) invoke the adversary \mathcal{A} to produce forged output \vec{Y}' , and (3) output 1 if and only if \vec{Y}' successfully decodes to a value that is different from decoded \vec{Y} .

We already concluded that \mathcal{A} cannot successfully forge an output (except with negligible probability) when given simulated input, so the above succeeds with non-negligible probability on real input, then it is indeed a privacy distinguisher. But GAR is private, so the distinguisher should not exist. We have reached a contradiction. It must be that \mathcal{A} cannot forge an output given a real-world input (except with negligible probability).

GAR is authentic. □

C SGC and GRAM – Illustration

Figure 14 illustrates the order of events in which we handle RAM accesses from within SGC branches. The dashed RAM box indicates that each access will happen outside conditional. The leftmost (resp. rightmost) circuit component represents the *demux* (resp. *mux*). Each \bullet (resp. \times) denotes a garbled label (resp. a SGC-style *garbage* label).

We proceed as follows. We describe G 's handling which, crucially, also includes E 's handling, since G emulates E as a part of SGC.

- (a) G holds labels that encode the input to the overall conditional. G assigns to each branch (fragment) a random seed. From this seed, he draws labels for the input wires of each branch. For now, G does not handle the RAM access. Instead, he treats the wires out of the ACCESS

gate as input to each branch, so these are also assigned GC labels. G can now garble the $demux$, which produces garbage labels for each branch. Notice that G cannot yet garble the $demux$ for the access response because of not knowing garbling encodings on the response outside SGC. That is, there is no \bullet on solid RAM box outputs.

- (b) G garbles each branch. Notice that this is possible since all input wires already have input labels. G xors materials from garbling of each branch. G can now generate garbage material for each branch properly.
- (c) G plays E in his head. G evaluates each branch with appropriate garbage material and garbage inputs. Notice that G can only emulate E up to the wires of the access query in each branch. I.e., we cannot yet compute \times garbage values on the output wires of the dashed RAM boxes.
- (d) G garbles the mux to produce GC labels on the solid RAM box query wires. G garbles the GRAM access, generating garbled labels on the RAM's output wires.
- (e) G garbles $demux$ with GC labels on each branch's access response. The gadget also determines the garbage on each branch's dashed access response wires. Thus, G can now continue to play E in his head, allowing him to compute garbage labels on each branch's outputs.
- (f) G garbles the mux for the conditional and puts the GC labels on the output of the conditional.