

Large-Scale Secure Computation: Multi-party Computation for (Parallel) RAM Programs

Elette Boyle^{1*}, Kai-Min Chung², and Rafael Pass^{3**}

¹ Technion Israel, eboyle@alum.mit.edu

² Academia Sinica, kmchung@iis.sinica.edu.tw

³ Cornell University, rafael@cs.cornell.edu

Abstract. We present the first efficient (i.e., polylogarithmic overhead) method for securely and privately processing large data sets over multiple parties with *parallel, distributed algorithms*. More specifically, we demonstrate load-balanced, statistically secure computation protocols for computing Parallel RAM (PRAM) programs, handling $(1/3 - \epsilon)$ fraction malicious players, while preserving up to polylogarithmic factors the computation, parallel time, and memory complexities of the PRAM program, aside from a one-time execution of a broadcast protocol per party. Additionally, our protocol has **polylog** communication locality—that is, each of the n parties speaks only with **polylog**(n) other parties.

1 Introduction

Large data sets, such as medical data, genetic data, transaction data, the web and web access logs, and network traffic data, are now in abundance. Much of the data is stored or made accessible in a distributed fashion, having necessitated the development of efficient distributed protocols that compute over such data. In particular, novel programming models for processing large data sets with *parallel, distributed algorithms*, such as MapReduce (and its implementation Hadoop) are emerging as crucial tools for leveraging this data in important ways.

But these methods require that the data itself is revealed to the participating servers performing the computation—and thus blatantly violate the privacy of potentially sensitive data. As a consequence, such methods

* The research of the first author has received funding from the European Union’s Tenth Framework Programme (FP10/ 2010-2016) under grant agreement no. 259426 ERC-CaC, and ISF grant 1709/14.

** Pass is supported in part by a Alfred P. Sloan Fellowship, Microsoft New Faculty Fellowship, NSF Award CNS-1217821, NSF CAREER Award CCF-0746990, NSF Award CCF-1214844, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2- 0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

cannot be used in many critical applications (e.g., discovery of causes or treatments of diseases using genetic or medical data).

In contrast, methods such as secure multi-party computation (MPC), introduced in the seminal works of Yao [Yao86] and Goldreich, Micali and Wigderson [GMW87], enable securely and privately performing any computation on individuals private inputs (assuming some fraction of the parties are honest). However, despite great progress in developing these techniques, there are no MPC protocols whose efficiency and communication requirements scale to the modern regime of large-scale distributed, parallel data processing.

We are concerned with merging these two approaches. In particular,

We seek MPC protocols that efficiently (technically, with polylogarithmic overhead) enable secure and private processing of large data sets with parallel, distributed algorithms.

Explicitly, in this large-scale regime, the following properties are paramount:

1. *Exploiting Random Access.* Computations on large data sets are frequently “lightweight”: accessing a small number of dynamically chosen data items, relying on conditional branching, and/or maintaining small memory. This means that converting a program first into a circuit to enable its secure computation, which immediately obliterates these gains, will not be a feasible option.
2. *Exploiting Parallelism.* In fact, as mentioned, to effectively solve large-scale problems, modern programming models heavily leverage parallelism. The notion of a Parallel RAM (PRAM) better captures such computing models. In the PRAM model of computation, several (polynomially many) CPUs run simultaneously, potentially communicating with one another, while accessing the same shared external memory. We consider a PRAM model with a variable number of CPUs but with a fixed activation structure (i.e., what processors are activated at which time steps is fixed). Note that such a model simultaneously captures RAMs (a single CPU) and circuits (the circuit topology dictates the CPU activation structure).
3. *Exploiting Plurality of Users.* In the setting of MPC we would like to leverage not only parallelism within a single party (i.e., if a party has multiple CPUs that may run in parallel), but also that we have a large number of parties that can run in parallel. So, if we have n parties, each with k processors, we ideally would like to securely compute PRAMs that use nk CPUs (as opposed to just k CPUs).

Additionally, the following desiderata are often of importance:

3. *Load balancing.* When the data set contains tens or hundreds of thousands of users’ data, it is often unreasonable to assume that any single user can provide memory, computation, or communication resources on the order of the data of *all users*. Rather, we would like to *balance* the load across nodes.
4. *Communication Locality.* In many cases, establishing a secure communication channel with a large number of distinct parties may be costly, and thus we would like to minimize the *locality of communication* [BGT13]: that is, the number of total parties that each party must send and receive message to during the course of the protocol.

To date, no existing work addresses secure computation of Parallel RAM programs. Indeed, nearly all results in MPC require a *circuit* model for the function being evaluated (including the line of work on scalable MPC [DI06, DIK⁺08, DKMS12, ZMS14]), and thus inherit resource requirements that are linear in the circuit size. Even for (sequential) RAM, the only known protocols either only handle two parties [OS97, GKK⁺11, LO13, GGH⁺13], or in the context of multi-party computation require all parties to store *all inputs* [DMN11], rendering the protocol useless in a large-scale setting (even forgetting about computation load balancing and locality).

1.1 Our Results

We present a statistically secure MPC for (any sequence of) PRAMs handling $(1/3 - \epsilon)$ fraction static corruptions in a synchronous communication network, with secure point-to-point channels. In addition, our protocol is strongly *load balanced* and *communication local* (i.e., $\text{polylog}(n)$ locality). We state our theorem assuming each party itself is a k -processor PRAM, for parameter k .

Theorem 1 (Informal – Main Theorem). *For any constant $\epsilon > 0$ and polynomial parallelism parameter $k = k(n)$, there exists an n -party statistically secure (with error negligible in n) protocol for computing any adaptively chosen sequence of PRAM programs Π_j with fixed CPU activation structures (and that may have bounded shared state), handling $(1/3 - \epsilon)$ fraction static corruptions with the following complexities, where each party is a k -processor PRAM (and where $|x|, |y|$ denote per-party input and output size,⁴ $\text{space}(\Pi)$, $\text{comp}(\Pi)$, and $\text{time}(\Pi)$ denote the worst-case space, computation, and (parallel) runtime of Π , and $\text{CPUs}(\Pi)$ denotes the number of CPUs of Π):*

- Computation per party, per Π_j : $\tilde{O}(\text{comp}(\Pi_j)/n + |y|)$.
- Time steps, per Π_j : $\tilde{O}\left(\text{time}(\Pi_j) \cdot \max\left\{1, \frac{\text{CPUs}(\Pi)}{nk}\right\}\right)$.
- Memory per party: $\tilde{O}\left(|x| + |y| + \max_{j=1}^N \text{space}(\Pi_j)/n\right)$.
- Communication Locality: $\tilde{O}(1)$.

given a one-time preprocessing phase with complexity:

- Computation per party: $\tilde{O}(|x|)$, plus single broadcast of $\tilde{O}(1)$ bits.
- Time steps: $\tilde{O}\left(\max\left\{1, \frac{|x|}{k}\right\}\right)$.

Additionally, our protocol achieves a strong “online” load-balancing guarantee: at all times during the protocol, all parties’ communication and computation loads vary by at most a constant multiplicative factor (up to a $\text{polylog}(n)$ additive term).

Remark 1 (Round complexity). As is the case with all general MPC protocols in the information-theoretic setting to date, the round complexity of our protocol corresponds directly with the time complexity (as when restricted to circuits, parallel complexity corresponds to circuit depth). That is, for each evaluated PRAM program Π_j , the protocol runs in $\tilde{O}(\text{time}(\Pi_j))$ sequential communication rounds to securely evaluate Π_j .

⁴ For simplicity of exposition, we assume all parties have the same input size and receive the same output.

Remark 2 (On the achieved parameters). Note that in terms of memory, each party only stores her input, output, and her “fair” share of the required space complexity, up to polylogarithmic factors. In terms of computation (up to polylogarithmic factors), each party does her “fair” share of the computation, receives her outputs, and in addition is required to read her entire input at an initial preprocessing stage (even though the computations may only involve a subset of the input bits; this additional overhead of “touching” the whole input once is necessary to achieve security).⁵ Finally, the time complexity corresponds to the parallel complexity of the PRAM being computed, as long as the combined number of available processors nk from all parties matches or exceeds the number of required parallel processes of the program (and degrades with the corresponding deficit).

Remark 3 (Instantiating the single-use broadcast). The broadcast channel can be instantiated either by the $O(\sqrt{n})$ -locality broadcast protocol of King *et al.* [KSSV06], or the $\text{polylog}(n)$ -average locality protocol of [BSGH13] at the expense of a cost of a one-time per-party computational cost of $O(\sqrt{n})$, or average cost of $\text{polylog}(n)$, respectively. We separate the broadcast cost from our protocol complexity measures to emphasize that any (existing or future) broadcast protocol can be directly plugged in, yielding associated desirable properties.⁶

1.2 Construction Overview

Our starting point is an *Oblivious PRAM (OPRAM)* compiler [BCP14b,GO96], a tool that compiles any PRAM program into one whose memory access patterns are independent of the data (i.e., “oblivious”). Such a compiler (with polylogarithmic overhead) was recently attained by [BCP14b]. Indeed, it is no surprise that such a tool will be useful toward our goal. It has been demonstrated in the sequential setting that Oblivious (sequential) RAM (ORAM) compilers can be used to build secure 2-party protocols for RAM programs [OS97,GKK⁺11,LO13,GGH⁺13]. Taking a similar approach, building upon the OPRAM compiler of [BCP14b] directly yields 2-party protocols for PRAMs.

However, OPRAM on its own does not directly provide a solution for *multi*-party computation (when there are many parties). While this approach gives protocols whose complexities scale well with the RAM (or

⁵ For general secure computation, and even if we restrict to functionalities that only access a few parties’ inputs, and only a few bits of their data, essentially all parties must perform computation at least $\Omega(|x|)$. To see this, consider secure computation of a “multi-party Private Information Retrieval (PIR)” functionality: each party $i > 1$ has as input some “big data” x_i , and party 1 has as input a party index i and an index j into their data x_i . The functionality returns $x_i[j]$ (i.e., the j ’th bit of party i ’s data) to party 1 and nothing to everyone else. We claim that each party $i > 1$ must access every bit of x_i ; if not, it learns that that particular bit of its data was not requested, which it cannot learn in an ideal execution of the functionality.

⁶ For instance, it remains open to achieve statistically secure broadcast with worst-case $\text{polylog}(n)$ locality.

PRAM) complexity of the programs, the complexities grow poorly with the number of parties. Indeed, the only current technique for securely evaluating a RAM program on *multiple* parties’ inputs [DMN11] is for all parties to hold secret shares of all parties’ inputs, and then jointly execute (using standard MPC for circuits) the trusted CPU instructions of the ORAM-compiled version of the program. This means each party must communicate and maintain information of size equivalent to *all parties’* inputs, and everyone must talk to everyone else for every time step of the RAM program evaluation.

One may attempt to improve the situation by first electing a small $\text{polylog}(n)$ -size representative committee of parties, and then only performing the above steps within this committee. This approach drops the total communication and computation of the protocol to reasonable levels. However, this approach does not save the subset of elected parties from carrying the burden of the entire computation. In particular, each elected party must memory storage equal to the size of *all parties’ inputs combined*, making the protocol unusable for “large-scale” computation. In this paper, we provide a new approach for dealing with this issue. We show how to use an OPRAM in a way that achieves balancing of memory, computation, and communication across all parties.

Our MPC construction proceeds in the following steps:

1. **From OPRAM to MPC.** Given an OPRAM, we begin by considering MPC in a “benign” adversarial setting, which we refer to as *oblivious multi-party computation*, where all parties are assumed to be honest, and we only require that an external attacker that views communication and activation (including memory and computation usages) patterns does not learn anything about the inputs. We show:
 - (a) OPRAM yields efficient *memory-balanced* oblivious MPC for PRAM.
 - (b) Using committee election techniques (à la [KLST11,DKMS12,BGT13]), any oblivious multi-party computation can be compiled into a standard secure MPC with only polylog overhead (and a one-time use of a broadcast channel per party).
2. **Load Balancing & Communication Locality.** We next show semi-generic compilers for “nice” (formally defined) *oblivious multi-party* protocols, each introducing only $\text{polylog}(n)$ overhead:
 - (a) From any “nice” protocol to one whose computation and communication are *load-balanced*.
 - (b) From any “nice” protocol to one that is both *load-balanced* and *communication local* (i.e., $\text{polylog}(n)$ locality).

Our final result is obtained by combining the above steps and observing that Step 1(b) preserves load-balancing and communication locality (and thus can be applied *after* Step 2). Let us mention that *just* Step 1 (together with existing construction of ORAMs) already yields the first MPC protocol for (sequential) RAM programs in which no party must store all parties’ inputs. Additionally, *just* Step 1 (together with the OPRAM construction of [BCP14b]) yields the first MPC for PRAMs.

We now expand upon each of these steps.

MPC from OPRAM Recall that our construction proceeds via an intermediate notion of *oblivious* security, in which we do not require security against corrupted parties, but rather against an external adversary who sees the activation patterns (i.e., accessed memory addresses and computation times) and communication patterns (i.e., sender/receiver ids and message lengths) of parties throughout the protocol.

Oblivious MPC from OPRAM. At a high level, our protocol will emulate a *distributed* OPRAM⁷ structure, where the CPUs and memory cells in the OPRAM are *each associated with parties*. (Recall that we need only achieve “oblivious” security, and thus can trust individual parties with these tasks). The “CPU” parties will control the evaluation flow of the (OPRAM-compiled) program, communicating with the parties emulating the role of the appropriate memory cells for each address to be accessed in the (OPRAM-compiled) database.

The distributed OPRAM structure will enable us to evenly spread the memory burden across parties, incurring only $\text{polylog}(n)$ overhead in total memory and computation, and while guaranteeing that the communication patterns between committees (corresponding to data access patterns) do not reveal information on the underlying secret values.

This framework shares a similar flavor to the protocols of [DKMS12,BGJK12], which assign committees to each of the gates of a circuit being evaluated, and to [BGT13], which uses CPU and input committees to direct program execution and distributedly store parties’ inputs. The distributed OPRAM idea improves and conceptually simplifies the input storage handling of Boyle *et al.* [BGT13], in which n committees holding the n parties’ inputs execute a distributed “oblivious input shuffling” procedure to break the link between which committees are communicating and which inputs are being accessed in the computation.

Compiling from “Oblivious” Security to Malicious Security. We next present a general compiler taking an oblivious protocol to one that is secure against $(1/3 - \epsilon)n$ statically corrupted malicious parties. (This step can be viewed as a refinement and generalization of ideas from [KLST11,DKMS12,BGT13].) We ensure the compiler tightly preserves the computation, memory, load-balancing, and communication locality of the original protocol, up to $\text{polylog}(n)$ factors (modulo a one-time broadcast per party). This enables us to apply the transformation to any of the oblivious protocols resulting from the intermediate steps in our progression.

At a high level, the compiler takes the following form: (1) First, the parties collectively elect a large number of “good” committees, each of size $\text{polylog}(n)$, where “good” means each committee is composed of at least $2/3$ honest parties, and that parties are spread roughly evenly across committees. (2) Each party will verifiably secret share his input among the corresponding committee C_i . (3) From this point on, the role of each party P_i in the original protocol will be emulated by the corresponding

⁷ We remark that the term “distributed OPRAM” was used with a different meaning in [LO13], in regard to an OPRAM that was split across two users.

committee C_i . That is, each local P_i computation will be executed via a small-scale MPC among C_i , and each communication from P_i to P_j will be performed via an MPC among committees C_i and C_j .

The primary challenge in this step is how to elect such committees while incurring only $\text{polylog}(n)$ locality and computation per party. To do so, we build atop the “almost-everywhere” scalable committee election protocol of King *et al.* [KSSV06] to elect a single good committee, and then show that one may use a $\text{polylog}(n)$ -wise independent function family $\{F_s\}_{s \in S}$ to elect the remaining committees with small description size (in the fashion of [KLST11,BGT13], for the case of combinatorial samplers and computational pseudorandom functions), with committee i defined as $C_i := F_s(i)$ for fixed random seed s .

We remark that, aside from the one-time broadcast, this compiler preserves load balancing and $\text{polylog}(n)$ locality. Indeed, load balancing is maintained since the committee setup procedure is computationally inexpensive, and each party appears in roughly the same number of “worker” committees. The locality of the resulting protocol increases by an additive $\text{polylog}(n)$ for the committee setup, and a multiplicative $\text{polylog}(n)$ term since all communications are now performed among $\text{polylog}(n)$ -size committees instead of individual parties.

Load Balancing Distributed Protocols

Load-balancing (Without Locality). We now show how to modify our protocol such that the total computational complexity and memory balancing are preserved, while additionally achieving a strong *computation* load balancing property—with high probability, at all times throughout the protocol execution, every party performs close to $1/n$ fraction of current total work, up to an additive $\text{polylog}(n)$ amount of work. This will hold simultaneously for both computation and communication.⁸

We present and analyze our load-balancing solution in the intermediate *oblivious* MPC security setting (recall that one can then apply the compiler from Step 2(b) above to obtain malicious MPC with analogous load-balancing). Let us mention that there is a huge literature on “load-balanced distributed computation” (e.g., [ACMR95,MPS02,MR98,AAK08]): As far as we can tell, our setting differs from the typical studied scenarios in that we must load balance an underlying distributed protocol, as opposed to a collection of independent “non-communicating jobs”. Indeed, the main challenge in our setting is to deal with the fact that “jobs” talk to one another, and this communication must remain efficient also be made load balanced. Furthermore, we seek a load-balanced solution with communication locality.

We consider a large class of arbitrary (potentially load-unbalanced and large-locality) distributed protocols Π , where we view each party in this underlying protocol as a “job”. Our goal is to load-balance Π by passing

⁸ Note that while our current protocol is memory balanced, it is currently rather imbalanced in computation: e.g., the parties emulating OPRAM CPUs are required to perform computation that is proportional to the whole PRAM computation.

“jobs” between “workers” (which will be the actual parties in the new protocols). More precisely, we start off with any protocol Π that satisfies the following (natural) “nice” properties:

- Each “job” has $\text{polylog}(n)$ size state;
- In each round, each “job” performs at most $\text{polylog}(n)$ computation and communication;
- In each round, each “job” communicates (either sending or receiving a message) to at most one other “job”.

It can be verified that these properties hold for our oblivious MPC for PRAM protocol.

Our load-balanced version of such a protocol first randomly⁹ efficiently assigns “workers” (i.e., parties) to “jobs”. Next, whenever a worker W has performed “enough” work for a particular job J , it randomly selects a replacement worker W' and passes the job over to it (that is, it passes over the state of the job J —which is “small” by assumption). The key obstacle in our setting is that the job J may later communicate with many other jobs, and all the workers responsible for those jobs need to be informed of the switch (and in particular, who the new worker responsible for the job J is). Since the number of jobs is $\Omega(n)$, workers cannot afford to store a complete directory of which worker is currently responsible for each job.

We overcome this obstacle by first modifying Π to ensure that it has small locality—this enables each job to only maintain a short list of the workers currently responsible for the “*neighboring*” jobs. We achieve this locality by requiring that parties (i.e., jobs) in the original protocol Π route their messages along the hypercube. Now, whenever a worker W for a job J is being replaced by some worker W' , W informs all J 's neighboring jobs (i.e., the workers responsible for them) of this change. We use the Valiant-Brebner [VB81] routing procedure to implement the hypercube routing because it ensures a desirable “low-congestion property,” which in our setting translates to ensuring that the overhead of routing is not too high for any individual worker.

The above description has not yet mentioned what it means for a worker to have done “enough” work for a job J . Each round a job is active (i.e., performing some computation), its “cost” increases by 1—we refer to this as an *emulation cost*. Additionally, each time a worker W is switched out from a job J , then J 's and each of J 's neighboring jobs' costs are increased by 1—we refer to this as a *switch cost*. Finally, once a job's (total) cost has reached a particular threshold τ , its cost is reset to 1 and the worker responsible for the job is switched out. The threshold τ is set to $2 \log M + 1$ where M is the number of jobs.

We show: (1) This switching does not introduce too much overhead. We, in fact, show that the total induced switching cost is bounded above by the emulation cost. (2) The resulting total work is load balanced across workers—we show this by first demonstrating that the protocol is load-balanced in expectation, and then using concentration to argue our stronger online load-balancing property.

⁹ In the actual analysis, we show that it also suffices to use $\text{polylog}(n)$ -wise independent randomness to pick this and subsequent assignments.

Finally, note that although communication between *jobs* is being routed through the hypercube, and thus the job communication protocol has small locality, the final load-balanced protocol, being run by *workers*, does *not* have small locality. This is because workers are assigned the role of many different jobs over time, and may possibly speak to a new set of neighbors for each position. (Indeed, over time, each worker will eventually need to speak to every other worker). We next show how to modify this protocol to achieve *locality*, while preserving load-balancing.

Achieving Both Load-Balancing and Locality. In our final step, we show how to modify the above-mentioned protocol to also achieve locality. We modify the protocol to also let *workers* route messages through a low-degree network (on top of the routing in the previous step). This immediately ensures locality. But, we must be careful to ensure that the additional message passing does not break load-balancing.

A natural idea is to again simply pass messages between *workers* along a low-degree hypercube network via Valiant-Brebner (VB) routing [VB81]. Indeed, the low-congestion property will ensure (as before) that routing does not incur too large an overhead for each worker.

However, when analyzing the overall load balance (for workers), we see an inherent distinction between this case and the previous. Previously, the nodes of the hypercube corresponded to *jobs*, each emulated by workers who swap in and out over time. When the underlying jobs protocol required job s to send a message to job t , the resulting message routing induced a cost along a path of neighboring jobs (that is, the workers emulating them), *independent of which workers are currently emulating them*. This independence, together with the fact that a worker passes his job after performing “enough” work for it, enabled us to obtain concentration bounds on overall load balancing over the random assignment of workers to jobs.

Now, the nodes correspond directly to *workers*. When the underlying jobs protocol requires a message transferred from job s to job t , routing along the workers’ graph must traverse a path from the *worker currently emulating job s* to the *worker currently emulating job t* , removing the crucial independence property from above. Even worse, workers along the routing path can now incur costs *even if they are not assigned to any job*. In this case, it is not even clear that job passing in of itself will be sufficient to ensure balancing.

To get around these issues, we add an extra step in the VB routing procedure (itself inspired by [VB81]) to break potential bad correlations. The idea is as follows: To route from the worker W_s emulating job s to the worker W_t emulating job t , we first route (as usual) from W_s to a *random* worker W_u , and then from W_u to W_t ; i.e., travel from W_s to W_t by “walking into the woods” and back. We may now partition the cost of routing into these two sub-parts, each associated with a single active job (s or t). Now, although workers along the worker-routing path will still incur costs from this routing (even though their jobs may be completely unrelated), the *distribution* of these costs on workers depends only on the identity of the initiating worker (W_s or W_t). We may thus

generalize the previous analysis to argue that if the expectation of work is load-balanced, then it still has concentration in this case.

For a modular analysis, we formalize the required properties of the underlying communication network and routing algorithm (to be used for the s -to- u and u -to- t routing) as a *local load-balanced routing network*, and show that the hypercube network together with VB routing satisfies these conditions.

1.3 Discussion and Future Work

With the explosive growth of data made available in a distributed fashion, and the growth of efficient parallel, distributed algorithms (such as those enabled by MapReduce) to compute on this data, ensuring privacy and security in such large-scale parallel settings is of fundamental importance. We have taken the first steps in addressing this problem by presenting the first protocols for secure multi-party computation, that with only polylogarithmic overhead, enable evaluating PRAM programs on a (large) number of parties' inputs. Our work leaves open several interesting open problems:

Honest Majority. We have assumed that $2/3$ of the players are honest.

In the absence of a broadcast channel,¹⁰ it is known that this is optimal. But if we assume the existence of a broadcast channel, it may suffice to assume $1/2$ fraction honest players.

Asynchrony. Our protocol assumes a synchronous communication network. We leave open the handling of asynchronous communication.

Trading efficiency for security. An interesting avenue to pursue are various tradeoffs between boosted efficiency and partial sacrifices in security. For example, in some settings, it is not detrimental to leak which parties' inputs were used within the computation; in such scenarios, one could then hope to remove the one-time $\Theta(n|x|)$ input preprocessing cost. Similarly, it may be acceptable to reveal the input-specific resources (runtime, space) required by the program on parties inputs; in such cases, we may modify the protocol to take only input-specific runtime and use input-specific memory.

In this work we focus only on achieving standard "full" security. However, we remark that our protocol can serve as a solid basis for achieving such tradeoffs (e.g., a straightforward tweak to our protocol results in input-specific resource use).

Communication complexity. As with all existing generic multi-party computation protocols in the information-theoretic setting, the communication complexity of our protocol is equal to its computation complexity. In contrast, in the computational setting (based on cryptographic assumptions), protocols with communication complexity below the complexity of the evaluated function have been constructed by relying on *fully homomorphic encryption (FHE)* [Gen09] (e.g., [Gen09,AJLA⁺12,MSS13]).

We leave as an interesting open question whether FHE-style techniques can be applied also to our protocol to improve the communication complexity, based on computational assumptions.

¹⁰ While the statement of our result makes use of a broadcast channel, as we mention, this channel can also be instantiated with known protocols.

1.4 Overview of the Paper

Section 2 contains preliminaries. In Section 3 we provide our ultimate theorem, and the sequence of intermediate notions and theorems which combine to yield this final result. We refer the reader to the full version of this work [BCP14a] for a complete descriptions and proofs.

2 Preliminaries

2.1 Multi-party Computation (MPC)

Protocol Syntax. We model parties as (parallel) RAM machines. An n -party protocol Φ is described as a collection of n (parallel) RAM programs $(P_i)_{i \in [n]}$, to be executed by the respective parties, containing additional special communication instructions $\text{Comm}(i, \text{msg})$, indicating for the executing party to send message msg to party i .

The per-party space, computation, and time complexities of the protocol $\Phi = (P_i)_{i \in [n]}$ are defined directly with respect to the corresponding party's PRAM program P_i , where each Comm is charged as a single computation time step. (See Section 2.2 for a definition of $\text{CPUs}(P)$, $\text{space}(P)$, $\text{comp}(P)$, $\text{time}(P)$ for PRAM P). The analogous total protocol complexities are defined as expected: Namely, $\text{space}(\Phi)$ and $\text{comp}(\Phi)$ are the *sums*, $\text{space}(\Phi) = \sum_{i \in [n]} \text{space}(P_i)$, $\text{comp}(\Phi) = \sum_{i \in [n]} \text{comp}(P_i)$, and $\text{time}(\Phi)$ is the *maximum*, $\text{time}(\Phi) = \max_{i \in [n]} \text{time}(P_i)$.

MPC Security. We consider the standard notion of (statistical) MPC security. We refer the reader to e.g. [BGW88] for more a more complete description of MPC security within this setting.

2.2 Parallel RAM (PRAM) Programs

A Concurrent Read Concurrent Write (CRCW) m -processor *parallel random-access machine (PRAM)* with memory size n consists of numbered processors $\text{CPU}_1, \dots, \text{CPU}_m$, each with local memory registers of size $\log n$, which operate synchronously in parallel and can make access to shared “external” memory of size n .

A PRAM program Π (given m, n , and some input x stored in shared memory) provides CPU-specific execution instructions, which can access the shared data via commands $\text{Access}(r, v)$, where $r \in [n]$ is an index to a memory location, and v is a word (of size $\log n$) or \perp . Each $\text{Access}(r, v)$ instruction is executed as:

1. **Read** from shared memory cell address r ; denote value by v_{old} .
2. **Write** value $v \neq \perp$ to address r (if $v = \perp$, then take no action).
3. **Return** v_{old} .

In the case that two or more processors simultaneously initiate $\text{Access}(r, v_i)$ with the same address r , then all requesting processors receive the previously existing memory value v_{old} , and the memory is rewritten with the value v_i corresponding to the lowest-numbered CPU i for which $v_i \neq \perp$.

We more generally support PRAM programs with a dynamic number of processors (i.e., m_i processors required for each time step i of the computation), as long as this sequence of processor numbers m_1, m_2, \dots is fixed, public information. The complexity of our OPRAM solution will scale with the number of required processors in each round, instead of the maximum number of required processors.

We consider the following *worst-case* metrics of a PRAM (over all inputs):

- $\text{CPUs}(II)$: number of parallel processors required by II .
- $\text{space}(II)$: largest database address accessed by II .
- $\text{time}(II)$: maximum number of time steps taken by any processor to evaluate II (where each `Access` is charged as a single step).¹¹
- $\text{comp}(II)$: the total sum of all computation steps of active CPUs evaluating II (which, for programs with fixed activation schedules as we consider, is a fixed value).

3 Local, Load-Balanced MPC for PRAM

Ultimately, we construct a protocol that securely realizes the ideal functionality $\mathcal{F}_{\text{PRAM}_s}$ (Figure 1) for evaluating a sequence of PRAM programs (with bounded state maintained between program) on parties' fixed inputs. For simplicity of exposition, we assume each party has equal input size and receives the same output. We further assume the total remnant state from one program execution to the next is bounded in size by the combined input size of all parties.¹²

Theorem 2 (Main Theorem). *For any constant $\epsilon > 0$ and polynomial parallelism parameter $k = k(n)$, there exists an n -party statistically secure (with error negligible in n) protocol realizing the functionality $\mathcal{F}_{\text{PRAM}_s}$, handling $(1/3 - \epsilon)$ fraction static corruptions with the following complexities, where each party is a k -processor PRAM (and where $|x|, |y|$ denote per-party input and output size, $\text{space}(II)$, $\text{comp}(II)$, and $\text{time}(II)$ denote the worst-case space, computation, and (parallel) runtime of II , and $\text{CPUs}(II)$ denotes the number of CPUs of II):*

- *Computation per party, per II_j :* $\tilde{O}(\text{comp}(II_j)/n + |y|)$.
- *Time steps, per II_j :* $\tilde{O}\left(\text{time}(II_j) \cdot \max\left\{1, \frac{\text{CPUs}(II)}{nk}\right\}\right)$.
- *Memory per party:* $\tilde{O}\left(|x| + |y| + \max_{j=1}^N \text{space}(II_j)/n\right)$.
- *Communication Locality:* $\tilde{O}(1)$.

given a one-time preprocessing phase with complexity:

- *Computation per party:* $\tilde{O}(|x|)$, plus single broadcast of $\tilde{O}(1)$ bits.
- *Time steps:* $\tilde{O}\left(\max\left\{1, \frac{|x|}{k}\right\}\right)$.

¹¹ We remark that the PRAM time complexity of any function f is bounded above by its circuit depth complexity (where the PRAM complexity of f is defined as the minimal value of $\text{time}(II)$ of any PRAM II which evaluates f).

¹² To support larger shared state size $\text{space}^{\text{Remnant}}$, the memory requirements of the protocol must grow with an extra additive $\tilde{O}(\text{space}^{\text{Remnant}})$.

Ideal Functionality $\mathcal{F}_{\text{PRAMs}}$:
 $\mathcal{F}_{\text{PRAMs}}$ running with parties P_1, \dots, P_n and an adversary proceeds as follows. The functionality maintains longterm storage of parties' inputs $\{x_i\}_{i \in [n]}$ (each of equal size $|x|$), per-CPU state information state_i , and remnant memory $\text{data}^{\text{Remnant}}$ of total size $\text{space}^{\text{Remnant}} \in O(n \cdot |x|)$ transferred from computation to computation.

- Initialize $\text{data}^{\text{Remnant}} \leftarrow \emptyset$ and $\text{state}_i \leftarrow \emptyset$ for each processor $i \in [m]$.
- Input Submission: Upon receiving an input $(\text{commit}, \text{sid}, \text{input}, x_i)$ from party P_i , record the value x_i as the input of P_i .
- Computation: Upon receiving a tuple $(\text{compute}, \text{sid}, \Pi, \text{space}, \text{time})$ consisting of an m -processor PRAM program Π , a space bound space , and a time bound time , execute Π as $(\text{output}, \text{state}_1, \dots, \text{state}_m, \text{data}^{\text{Remnant}}) \leftarrow \Pi(x_1, \dots, x_n, \text{state}_1, \dots, \text{state}_m, \text{data}^{\text{Remnant}})$ with the current value of state_i for each CPU $i \in [m]$. Send output to all parties.

Fig. 1: The ideal functionality $\mathcal{F}_{\text{PRAMs}}$, corresponding to secure computation of a sequence of adaptively chosen PRAMs on parties' inputs.

Additionally, the protocol achieves $\text{polylog}(n)$ communication locality, and a strong “online” load-balancing guarantee:

Online Load Balancing: For every constant $\delta > 0$, with all but negligible probability in n , the following holds at all times during the protocol: Let cc and $\text{cc}(P_j)$ denote the total communication complexity and communication complexity of party P_j , comp and $\text{comp}(P_j)$ denote the total computation complexity and computation complexity of party P_j , we have

$$\frac{(1 - \delta)}{n} \text{cc} - \text{polylog}(n) \leq \text{cc}(P_j) \leq \frac{(1 + \delta)}{n} \text{cc} + \text{polylog}(n)$$

$$\frac{(1 - \delta)}{n} \text{comp} - \text{polylog}(n) \leq \text{comp}(P_j) \leq \frac{(1 + \delta)}{n} \text{comp} + \text{polylog}(n).$$

3.1 Proof of Main Theorem

At a very high level, the proof takes three steps: We first obtain MPC realizing $\mathcal{F}_{\text{PRAMs}}$ with a weaker notion of *oblivious* security. We then show how to attain communication locality and load balancing, while preserving oblivious security. (This combines two steps described within the introduction). Finally, we convert the obliviously secure protocol to one secure in the malicious setting. We now proceed to describe these steps in greater technical detail.

Step 1: Oblivious-Secure MPC for PRAM. Intuitively, an adversary in the oblivious model is not allowed to corrupt any parties, and instead is restricted to seeing the “externally measurable” properties of the protocol (e.g., party response times, communication patterns, etc).

Definition 1 (Oblivious secure MPC). *Secure realization of a functionality F by a protocol in the oblivious model is defined by the following real-ideal world scenario:*

Ideal World: Same as standard MPC without corrupted parties. That is, the adversary learns only public outputs of the functionality F evaluated on honest-party inputs.

Real World: Instead of corrupting parties, viewing their states, and controlling their actions (as in the standard malicious adversarial setting), the adversary is now limited as an external observer, and is given access only to the following information:

- *Activation Patterns:* Complete list of tuples of the form
 - (timestep, party-id, compute-time): Specifying all local computation times of parties.
 - (timestep, party-id, local-mem-addr): Specifying all memory access patterns of parties.
- *Communication Patterns:* Complete list of tuples of the form
 - (timestep, sndr-id, rcvr-id, msg-len): Specifying all sender-receiver pairs, in addition to the corresponding communicated message bit-length.

The output of the real-world experiment consists of the outputs of the (honest) parties, in addition to an arbitrary PPT function of the adversary’s view at the conclusion of the protocol.

(Statistical) Security: For every PPT adversary \mathcal{A} in the real-world execution, there exists a PPT ideal-world adversary \mathcal{S} for which for every environment \mathcal{Z} , we have $\text{output}_{\text{Real}}(1^k, \mathcal{A}, \mathcal{Z}) \stackrel{s}{\cong} \text{output}_{\text{Ideal}}(1^k, \mathcal{S}, \mathcal{Z})$.

Toward our result, it will be advantageous to think of computations as composed of several sub-parts, or “jobs,” that each maintain and compute on small polylogarithmic-size state (Note that this is natural in the PRAM setting, where each CPU has polylogarithmic-size local memory). Later, to achieve load balancing, jobs will be assigned to and passed around between “workers,” so that each worker roughly performs the same amount of work. (The small state requirement per job will guarantee that “job passing” is not too expensive). Then, to obtain malicious security, each worker will ultimately be emulated by a committee of parties via small-scale MPCs; because of the polynomial overhead in the underlying MPC protocol, it will be important that this is only done for computations of $\text{polylog}(n)$ size on $\text{polylog}(n)$ -size memory.

We now define the notion of a protocol in the *jobs model*.

Definition 2 (Jobs Model). Let n be a security parameter. A jobs protocol consists of a $\text{poly}(n)$ -size set Jobs of agents (called jobs), and a distributed protocol description $\Pi_{\mathcal{J}}$, instructing each job to perform local computations and to communicate over a synchronized network (via point-to-point communication), with the following properties:

- *Bounded memory:* each job’s space complexity is $w \in \text{polylog}(n)$.
- *Bounded per-round computation and communication:* the computation and communication complexity of each job at each round is upper bounded by $w \in \text{polylog}(n)$.

A job is active in a round if it performs computation within this round.

A jobs protocol is further said to have injective communication if the following property is satisfied:

- *Injective communication:* each round, a set of jobs are activated, and each sends a single $\text{polylog}(n)$ -sized message to a distinct job.

By convention, we assume the first m_{in} jobs of a jobs protocol are *input jobs*, the last m_{out} are *output jobs*, and the remaining jobs are *helper jobs*. Each input job J_i holds a single-word input $x_i \in \{0, 1\}^w$ (for $w \in \text{polylog}(n)$); output and helper jobs have no input. We then have a canonical correspondence between functionalities in the standard n -party setting and the equivalent functionalities in the Worker-Jobs Model:

- Functionality \mathcal{F} : In the n -party setting. Accepts inputs x_i from each party P_i , evaluates $y \leftarrow F(x_1 || \dots || x_n)$, outputs the resulting value y to all parties P_i .
- Functionality $\mathcal{F}^{\text{Jobs}}$: In the Jobs Model. Accepts (short) inputs x_u^i from each Input Job, evaluates $y \leftarrow F(x_1 || \dots || x_\ell)$, and distributes the resulting value y (in short pieces) to the Output Jobs.

We may analogously define *oblivious security* of a jobs protocol (where jobs are honest and the adversary sees only “externally measurable” properties of the protocol, as in Definition 1). Within the jobs model, we thus wish to securely realize the functionality $\mathcal{F}_{\text{PRAMs}}^{\text{Jobs}}$, equivalent to $\mathcal{F}_{\text{PRAMs}}$ with the above syntactic change. Note that in the regime of oblivious security, a jobs protocol yields a *memory*-balanced protocol in the standard n -party model, by simply assigning jobs to the n parties evenly.

Theorem 3. *There exists an oblivious-secure protocol in the Jobs Model realizing the functionality $\mathcal{F}_{\text{PRAMs}}^{\text{Jobs}}$ for securely computing a sequence of N adaptively chosen PRAM programs Π_j , with the following complexities (where $n \cdot |x|, |y|$ denote the total input and output size, and $\text{space}(\Pi)$, comp , and $\text{time}(\Pi)$ denote the worst-case space, computation, and (parallel) runtime of Π over all inputs):*

- Number of jobs: $\tilde{O}(n \cdot |x| + |y| + \max_{j \in [N]} \text{space}(\Pi_j))$.
 - Computation complexity, per Π_j : $\tilde{O}(\text{comp}(\Pi_j))$.
 - Time steps, per Π_j : $\tilde{O}(\text{time}(\Pi_j))$.
 - The number of active jobs in each round is $O(\max_{j \in [N]} \text{CPU}(\Pi_j))$.
- given a one-time preprocessing phase with complexity
- Computation complexity: $\tilde{O}(n \cdot |x|)$.
 - Time steps: $\tilde{O}(1)$.

Further, the protocol has injective communication: in each round, each activated job sends a single $\text{polylog}(n)$ -size message to a distinct job.

Recall within the Jobs Model each job is limited to maintaining state of size $\text{polylog}(n)$; thus the *memory requirement* of the above protocol is

$$\tilde{O}\left(n \cdot |x| + |y| + \max_{j \in [N]} \text{space}(\Pi_j)\right),$$

based on the number of required jobs.

Idea of proof. The result builds upon the existence of an Oblivious PRAM compiler with $\text{polylog}(n)$ time and space overhead that is *collision-free* (i.e., where no two CPUs must access the same memory address in the same timestep), which is guaranteed to exist unconditionally based on [BCP14b]. In addition to the standard Input and Output jobs, our protocol will have one Helper job for each of the CPUs and each memory cell in the database of the OPRAM-compiled program. The CPU jobs

store the local state and perform the computations of their corresponding CPU. In each round that the i th CPU's instructions dictate a memory access at location $\text{addr}^{(i)}$, the CPU job i will communicate with the Memory job $\text{addr}^{(i)}$ to perform the access. (Thus, in each round, at most $2 \cdot \text{CPUs}(\text{OPRAM}(\Pi))$ jobs are active, where $\text{OPRAM}(\Pi)$ denotes the OPRAM-compilation of Π). Activation and communication patterns in the resulting protocol are simulatable directly by the OPRAM security. The preprocessing phase of the protocol corresponds to inserting all inputs into the OPRAM-protected database in parallel (i.e., emulating the OPRAM-compiled input insertion program that simply inserts each input x_i into address i of the database).

Step 2: Locality and Load Balancing. This step attains $\text{polylog}(n)$ communication *locality*,¹³ and computation *load balancing* from any jobs protocol $\Pi_{\mathcal{J}}$ with injective communication. We do so by emulating $\Pi_{\mathcal{J}}$ by a fixed set of parties (which we sometimes refer to as “workers”), where each worker is assigned several jobs, and will pass jobs to other workers once he has performed a certain amount of work. This yields a standard N -party protocol with a special *decomposable state* structure: i.e., parties’ memory can be decomposed into separate $\text{polylog}(n)$ -size memory blocks, which are only ever computed on independently or in pairs, in steps of $\text{polylog}(n)$ computation per round. This is because parties’ computation is limited to individual jobs to which it was assigned.¹⁴

Definition 3 (Decomposable State). *An N -party protocol Π is said to have decomposable state if for every party P , the local memory mem of P can be decomposed into $\text{polylog}(n)$ -size blocks $\text{mem} = (\text{mem}_1, \text{mem}_2, \dots, \text{mem}_m)$ such that: In each round of Π , the (parallel) local computation performed by party P is described as a list $\{(i, j, f_{i,j})\}_{(i,j) \in I}$ for some $I \subseteq [m] \times [m]$, such that each $f_{i,j}$ has complexity $\text{polylog}(n)$. For each $(i, j) \in I$, party P executes $(\text{mem}_i, \text{mem}_j) \leftarrow f_{i,j}(\text{mem}_i, \text{mem}_j)$.¹⁵ By convention, received communication messages are stored in local memory.*

We achieve the following “fully load-balanced” properties. Note that the first two properties correspond directly to our final load-balancing goal. The final property will be used to ensure that no individual worker is ever assigned drastically more than the expected number of simultaneous parallel computation tasks; this is important since workers will eventually be emulated by (technically, committees of) parties, who themselves may have bounded parallelism capability (i.e., small number of CPUs).

Definition 4 (Fully Load Balanced). *An N -party protocol Π is said to be fully load balanced with respect to security parameter n if the following properties hold:*

¹³ Recall a protocol has (communication) locality $\ell(n)$ if during the course of the protocol every party communicates with at most $\ell(n)$ other parties.

¹⁴ Looking ahead, pairwise computation will be used when emulating job-to-job communication, and will be sufficient when the original jobs protocol has *injective communication*, so that each job communicates with at most one other job per round.

¹⁵ With some canonical resolution for write conflicts. (In our constructions, the sets (i, j) will be disjoint).

- *Memory load balancing:* Let $\text{space}(\Pi)$ denote the total space complexity of protocol Π . For every constant $\delta > 0$, with all but negligible probability in n , every party P_j has space complexity

$$\text{space}(P_j) \leq \frac{(1 + \delta)}{N} \text{space}(\Pi) + \text{polylog}(n).$$

- *Online computation/communication load balancing:* For every constant $\delta > 0$, with all but negligible probability in n , the following holds at all times during the protocol: Let cc and $\text{cc}(P_j)$ denote the total communication complexity and communication complexity of party P_j , comp and $\text{comp}(P_j)$ denote the total computation complexity and computation complexity of party P_j , we have

$$\begin{aligned} \frac{(1 - \delta)}{N} \text{cc} - \text{polylog}(n) &\leq \text{cc}(P_j) \leq \frac{(1 + \delta)}{N} \text{cc} + \text{polylog}(n) \\ \frac{(1 - \delta)}{N} \text{comp} - \text{polylog}(n) &\leq \text{comp}(P_j) \leq \frac{(1 + \delta)}{N} \text{comp} + \text{polylog}(n). \end{aligned}$$

- *Per-round per-party efficiency:*¹⁶ Let A be an upper bound on the number of active jobs at each round in $\Pi_{\mathcal{J}}$. With all but negligible probability in n , the per-round per-party computation complexity is upper bounded by $\tilde{O}(1 + (A/N))$.

Theorem 4. *Let $\Pi_{\mathcal{J}}$ be an M -job protocol with computation complexity comp and injective communication, realizing functionality $\mathcal{F}^{\text{Jobs}}$. Then there exists a fully load-balanced (Definition 4) $\tilde{O}(n)$ -party protocol $\Pi_{\mathcal{W}}$ with decomposable states (Definition 3) that realizes \mathcal{F} with total computation $\tilde{O}(\text{comp})$, space complexity $\tilde{O}(M)$, and $\text{polylog}(n)$ locality. If $\Pi_{\mathcal{J}}$ satisfies oblivious security, so does $\Pi_{\mathcal{W}}$.*

Idea of proof. Recall that in our construction of $\Pi_{\mathcal{W}}$ (in the introduction), at any point of the protocol execution, each job is assigned to a random worker¹⁷ and is stored in at most 2 workers. This is sufficient to imply memory load balancing by standard concentration and union bounds. Online computation/communication load balancing follows by observing that (i) the job-passing pattern is independent of the worker-job assignment, and (ii) jobs are passed frequently enough before accumulating large cost. This allows us to think of the execution as partitioned into “job chunks” each of which is assigned to a random worker, thus amenable to concentration bounds. The last load-balanced property follows again by the fact that each job is independently assigned to a random worker and that each job only performs $\text{polylog}(n)$ amount of work per round. To obtain locality, we consider a fixed low-degree communication network between workers, and pass messages using a load-balanced

¹⁶ We note that the last two properties are related but incomparable. The online load balancing property focuses on accumulated work, whereas the per-round per-party efficiency concerns upper bounds on per-round work, which is used to bound the required amount of parallelism to execute the protocol with efficient parallel time.

¹⁷ Technically, the initial job-worker assignment is only K -wise independent for $K = \log^3 n$. Nevertheless, this is sufficient for concentration bounds to go through.

routing algorithm. Load balancing of this modified scheme follows by similar, but more delicate analysis.

The resulting protocol has *decomposable state*, since parties’ memory and computation are completely local to individual jobs, or pairs of jobs in the case of emulating job-to-job communication (since the starting jobs protocol has injective communication).

Step 3: From Oblivious to Malicious Security. Finally, we present a general transformation that produces an n -party MPC protocol securely realizing a functionality \mathcal{F} against $(1/3-\epsilon)n$ static corruptions, given any $\tilde{\Theta}(n)$ -party protocol with decomposable states (see Definition 3) realizing the corresponding jobs-model functionality $\mathcal{F}^{\text{jobs}}$ with only *oblivious* security. This step can be viewed as a refinement and generalization of ideas from [KLST11,DKMS12,BGT13].

Theorem 5 (From Oblivious Security to Malicious Security).

Suppose there exists an $N \in \Theta(n \cdot \text{polylog}(n))$ -party oblivious protocol with decomposable state, realizing functionality $\mathcal{F}^{\text{jobs}}$ in space, computation, and (parallel) time complexity `space, comp, time`. Then for any constant $\epsilon > 0$ there exists an n -party MPC protocol (with error negligible in n) securely realizing the corresponding functionality \mathcal{F} against $(1/3 - \epsilon)n$ static corruptions, with the following complexities (where each party is a PRAM with possibly many processors), given a one-time preprocessing phase with a single broadcast of $\tilde{O}(1)$ bits per party:

- Per-party memory: $\tilde{O}(\text{space}/n)$.
- Total computation: $\tilde{O}(\text{comp})$.
- Time complexity: $\tilde{O}(\text{time})$.

In addition, if the original protocol has $\tilde{O}(1)$ locality and is fully load-balanced (i.e., satisfying all properties of Definition 4), then the resulting protocol additionally possesses the following properties:

- Communication locality $\tilde{O}(1)$.
- Online computation load balancing, as in Definition 4(c).
- Time complexity $\tilde{O}(\text{time} \cdot \max\{1, \frac{A}{nk}\})$ when each party is limited to being a k -processor PRAM, where A denotes the maximum per-round per-party computation complexity of any party in the original oblivious-secure protocol.¹⁸

Idea of Proof. The compiler takes the following form: First, parties collectively elect a large number of “good” committees, each of size $\text{polylog}(n)$, where “good” means each committee is composed of at least $2/3$ honest parties, and that parties are spread roughly evenly across committees. The one-time broadcast is used to reach full agreement on the first committee. These committees will then emulate each of the *decomposable sub-computations* of the original protocol Π (see Definition 3), via small-scale MPCs. That is, committees are initialized with inputs by having the parties in Π' split their inputs into $\text{polylog}(n)$ -size pieces and verifiably secret share them to the appropriate committee(s). Each local

¹⁸ In particular, for our MPC for PRAMs protocol formed by combining Steps 1 and 2, the parameter A will correspond to the number of CPUs required in the evaluated PRAM Π , with polylog overhead.

computation (and communication) in Π decomposes as a collection of $f_{i,j}$, each affecting only two committees (emulating mem_i and mem_j). Since committees are only size $\text{polylog}(n)$, and *each small-scale MPC has only $\text{polylog}(n)$ memory and computation* (because of decomposability), the memory, computation, and time complexity overhead is small. Since parties are spread across committees, the protocol remains load balanced. Finally, by using a perfectly secure underlying MPC protocol (such as [BGW88]), the only information revealed corresponds directly to the “observable” properties (communication patterns, etc.), thus reducing directly to oblivious security (as per Definition 1).

References

- [AAK08] Baruch Awerbuch, Yossi Azar, and Rohit Khandekar. Fast load balancing via bounded best response. In *SODA 2008*, pages 314–322, 2008.
- [ACMR95] Micah Adler, Soumen Chakrabarti, Michael Mitzenmacher, and Lars Eilstrup Rasmussen. Parallel randomized load balancing (preliminary version). In *STOC 1995*, pages 238–247, 1995.
- [AJLA⁺12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold fhe. In *EUROCRYPT*, pages 483–501, 2012.
- [BCP14a] Elette Boyle, Kai-Min Chung, and Rafael Pass. Large-scale secure computation. Cryptology ePrint Archive, Report 2014/404, 2014.
- [BCP14b] Elette Boyle, Kai-Min Chung, and Rafael Pass. Oblivious parallel ram. Cryptology ePrint Archive, Report 2014/594, 2014.
- [BGJK12] Elette Boyle, Shafi Goldwasser, Abhishek Jain, and Yael Tsa-man Kalai. Multiparty computation secure against continual memory leakage. In *STOC*, pages 1235–1254, 2012.
- [BGT13] Elette Boyle, Shafi Goldwasser, and Stefano Tessaro. Communication locality in secure multi-party computation - how to run sublinear algorithms in a distributed setting. In *TCC*, pages 356–376, 2013.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *STOC*, pages 1–10, 1988.
- [BSGH13] Nicolas Braud-Santoni, Rachid Guerraoui, and Florian Huc. Fast byzantine agreement. In *PODC*, pages 57–64, 2013.
- [DI06] Ivan Damgård and Yuval Ishai. Scalable secure multiparty computation. In *CRYPTO*, pages 501–520, 2006.
- [DIK⁺08] Ivan Damgård, Yuval Ishai, Mikkel Krøigaard, Jesper Buus Nielsen, and Adam Smith. Scalable multiparty computation with nearly optimal work and resilience. In *CRYPTO*, pages 241–261, 2008.

- [DKMS12] Varsha Dani, Valerie King, Mahnush Movahedi, and Jared Saia. Breaking the $o(nm)$ bit barrier: Secure multiparty computation with a static adversary. *CoRR*, abs/1203.0289, 2012.
- [DMN11] Ivan Damgård, Sigurd Meldgaard, and Jesper Buus Nielsen. Perfectly secure oblivious ram without random oracles. In *TCC*, pages 144–163, 2011.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [GGH⁺13] Craig Gentry, Kenny A. Goldman, Shai Halevi, Charanjit S. Jutla, Mariana Raykova, and Daniel Wichs. Optimizing oram and using it efficiently for secure computation. In *Privacy Enhancing Technologies*, pages 1–18, 2013.
- [GKK⁺11] S. Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure computation with sublinear amortized work. *IACR Cryptology ePrint Archive*, 2011:482, 2011.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *J. ACM*, 43(3):431–473, 1996.
- [KLST11] Valerie King, Steven Lonargan, Jared Saia, and Amitabh Trehan. Load balanced scalable byzantine agreement through quorum building, with full information. In *ICDCN*, pages 203–214, 2011.
- [KSSV06] Valerie King, Jared Saia, Vishal Sanwalani, and Erik Vee. Scalable leader election. In *SODA*, pages 990–999, 2006.
- [LO13] Steve Lu and Rafail Ostrovsky. Distributed oblivious ram for secure two-party computation. In *TCC*, pages 377–396, 2013.
- [MPS02] Michael Mitzenmacher, Balaji Prabhakar, and Devavrat Shah. Load balancing with memory. In *(FOCS), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 799–808, 2002.
- [MR98] S. Muthukrishnan and Rajmohan Rajaraman. An adversarial model for distributed dynamic load balancing. *SPAA '98*, pages 47–54, 1998.
- [MSS13] Steven Myers, Mona Sergi, and Abhi Shelat. Black-box proof of knowledge of plaintext and multiparty computation with low communication overhead. In *TCC*, pages 397–417, 2013.
- [OS97] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In *STOC*, pages 294–303, 1997.
- [VB81] Leslie G. Valiant and Gordon J. Brebner. Universal schemes for parallel communication. In *STOC*, pages 263–277, 1981.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.
- [ZMS14] Mahdi Zamani, Mahnush Movahedi, and Jared Saia. Millions of millionaires: Multiparty computation in large networks. *Cryptology ePrint Archive*, Report 2014/149, 2014.