

Secret-Shared Shuffle

Melissa Chase^{*}, Esha Ghosh^{**}, and Oxana Poburinnaya^{***}

Abstract. Generating additive secret shares of a *shuffled* dataset - such that neither party knows the order in which it is permuted - is a fundamental building block in many protocols, such as secure collaborative filtering, oblivious sorting, and secure function evaluation on set intersection. Traditional approaches to this problem either involve expensive public-key based crypto or using symmetric crypto on permutation networks. While public-key-based solutions are bandwidth efficient, they are computation-heavy. On the other hand, constructions based on permutation networks are communication-bound, especially when the dataset contains large elements, for e.g., feature vectors in an ML context.

We design a new 2-party protocol for this task of computing secret shares of shuffled data, which we refer to as secret-shared shuffle. Our protocol is secure against a static semi-honest adversary. At the heart of our approach is a new primitive we define (which we call “Share Translation”) that generates two sets of pseudorandom values “correlated via the permutation”. This allows us to reduce the problem of shuffling the dataset to the problem of shuffling pseudorandom values, which enables optimizations both in computation and communication. We then design a Share Translation protocol based on oblivious transfer and puncturable PRFs.

Our final protocol for secret-shared shuffle uses lightweight operations like XOR and PRGs, and in particular doesn’t use public-key operations besides the base OTs. As a result, our protocol is concretely more efficient than the existing solutions. In particular, we are two-three orders of magnitude faster than public-key-based approach and one order of magnitude faster compared to the best known symmetric-key approach when the elements are moderately large.

Keywords: secure shuffle, secure function evaluation, puncturable PRF

1 Introduction

Machine Learning algorithms are data-hungry: more data leads to more accurate models. On the other hand, privacy of data is becoming exceedingly important, for social, business and policy compliance reasons (e.g. GDPR). There has been decades of groundbreaking work in the academic literature in developing cryptographic technology for collaborative computation, but it still has some significant bottlenecks in terms of wide-scale adoption. Although theoretical results

^{*} Microsoft Research. Email: melissac@microsoft.com.

^{**} Microsoft Research. Email: Esha.Ghosh@microsoft.com.

^{***} University of Rochester/Ligero Inc. The work was partially done while doing internship at Microsoft Research. Email: oxanapob@bu.edu.

demonstrate the possibility of generic secure computation, they are not efficient enough to be adopted, both in terms of computation and communication size. For instance, Google cited network cost as a major hindrance in adopting cryptographic secure computation solution [13].

Secret-shared shuffle. In this work, we focus on computation and communication efficiency of a building block used in many important secure computation protocols, which we call “secret-shared shuffle”. Secret-shared shuffle is a protocol which allows two parties to jointly shuffle data and obtain additive secret shares of the result - without any party learning the permutation corresponding to the shuffle. (In the remainder of this paper, by secret sharing we will always mean *additive* secret sharing.)

Motivation. To see the importance of secret-shared shuffle, consider the task of securely evaluating some function on the intersection of two sets belonging to two parties - in particular, the intersection itself should also remain secret. As a concrete example, consider a merchant who wants to analyze efficiency of its online ads by running some ML algorithm on the data which contains the information about users who both (a) saw the ad and (b) made a purchase. Such data is split between the ad supplier (who knows which person clicked which add) and the merchant (who knows which person made a purchase). Thus, ML should be run on set intersection of the two databases - and both ML and set intersection have to be computed using secure multi-party computation protocols (MPC).

To do this securely, ideally we would use a private set intersection protocol which outputs an intersection in some “encrypted” form - e.g. by encrypting or secret sharing elements in the intersection - and then evaluate the ML function securely under MPC. However, currently known efficient protocols for private set intersection do not output an encrypted intersection: instead they output an encrypted indicator vector - i.e. a vector of bits indicating if each element is in the intersection or not [5]. This difference is very important, since in the former case one could run the ML function (under MPC) directly on the encrypted intersection, whereas in the latter case such MPC has to be run on the whole database, and the elements not in the intersection have to be filtered out under the MPC. Needless to say, this incurs unnecessary overhead, especially in cases where the intersection is relatively small compared to the input sets.

In other words, ideally we would want to get rid of non-intersection elements before running the rest of the MPC. A natural way to do this without compromising security is to shuffle the encrypted elements together with the encrypted indicator vector. Then parties can reveal the indicator vector and discard elements which are not in the intersection. Note that it is crucial that neither party learns how exactly the elements were permuted; otherwise this party could learn whether some of its elements are in the intersection or not. Also note that the requirement on the secrecy of the permutation implies that the result of the shuffle has to be in some encrypted or secret-shared form (in order to prevent linking original and shuffled elements), hence naturally leading to the notion of secret-shared shuffle.

Known techniques and their limitations. For convenience, let us look at “a half” of a secret-shared shuffle, which we call **Permute+Share**: in this protocol P_0 holds a permutation π and P_1 holds the database \mathbf{x} , and they would like to learn secret shares of permuted database¹. While this problem can be solved by any generic MPC, to the best of our knowledge, there are two specialized solutions for this problem, which differ in how exactly the permuting happens. One approach is to give P_0 ’s shares of \mathbf{x} to P_1 in some encrypted form, let P_1 permute them according to π under the encryption, rerandomize them, and return them to P_0 . This is a folklore solution that uses rerandomizable additively homomorphic public-key encryption, and because of that it is compute-intensive. We elaborately describe this solution in Section C. The other approach is to start with secret-shared \mathbf{x} and jointly compute atomic swaps, until all elements arrive to their target location. To prevent linking, each atomic swap should also rerandomize the shares. This approach is taken by [23, 16], who let parties jointly apply a permutation network to the shares, where each atomic swap is implemented using oblivious transfer (OT) in [23] and garbled circuit in [16]. The downside of this approach is its communication complexity which is proportional to $\ell \cdot N \log N$, where N is the number of elements in the database and ℓ is the bitlength of each element. This overhead seems to be inherent in approaches based on joint computation of atomic swaps, since each element has to be fully fed into at least $\log N$ swaps.

We also note that there exist efficient protocols for secure shuffle in the 3 party setting (e.g. see [4] and references within). We note that our 2 party setting is very different from 3 party setting, which allows for honest majority and thus for simpler and more efficient constructions.

Our Contribution We propose a novel approach to design a protocol for secret-shared shuffle, secure in the semi-honest model. Our protocol is parameterized by a value T , which can be chosen to optimize performance for a given tradeoff between network bandwidth and computation cost. Our protocol runs in 3 rounds (6 messages) with communication only proportional to $\lambda N \log N + N \ell \log N / \log T$, where λ is security parameter, N is the number of elements in the database and ℓ is the size of each element. In our experiments on databases of size 2^{20} - 2^{32} the optimal value for T is between 16 and 256, so we can think of $\log T$ as a number between 4 and 8. Note that the size ℓ of the element could be very large (e.g. each element could be a feature vector in ML algorithm), in which case the term $N \ell \log N / \log T$ dominates, and thus it could be a significant improvement compared to communication in permutation-network-based approach, which is proportional to $\ell N \log N$. While the computation cost of our protocol, dominated by $(NT \log N / \log T)(\ell/\lambda)$, is asymptotically worse than that of a PKE-based or permutation network-based approach, our protocol uses lightweight crypto primitives (XORs and PRGs) and does not require any public-key operations besides a set of base OTs, thus resulting in a concretely efficient protocol. We compute the concrete cost of our protocol and estimate its performance over different net-

¹ Note that one can get secret-shared shuffle by combining two instance of **Permute+Share**.

works (bandwidth 1Gbps, 100Mbps and 72Mbps). For large values of ℓ , we see a two to three orders of magnitude improvement over the best known public key based approach and an order of magnitude improvement over the best known symmetric key approach. The details of our experiment are in Section 7.

At the heart of our construction is a new primitive which we call Share Translation functionality. This functionality outputs two sets of pseudorandom values - one per party - with a special permutation-related dependency between them, and we show that this is enough to implement secret-shared shuffle. Conceptually, this functionality allows us to push the problem of permuting the *data* down to the problem of permuting *pseudorandom values*². This can be seen as the analogue of beaver triples or tiny tables for permutations rather than arithmetic or boolean computations.

Our Share Translation has quadratic running time (in N), and thus implementing secret-shared shuffle directly using Share Translation protocol becomes too prohibitive, even with lightweight operations like XOR and PRG. This brings us to the second crucial part of our construction: we devise a way to represent any permutation as a combination of several permutations π_i , where each π_i itself consists of several *disjoint* permutations, each acting on few elements. We find such decomposition using the special structure of Benes permutation network. This decomposition allows us to apply our Share Translation protocol to *small* individual disjoint permutations rather than *big* final permutation, allowing our protocol for secret-shared shuffle to achieve the claimed running time. We leverage the particular structure of our Share Translation protocol to make sure that this transformation doesn't increase the number of rounds.

1.1 Applications

Collaborative Filtering One immediate application of our shuffle protocol is to allow two parties who hold shares of a set of elements to filter out elements that satisfy a certain criterion. This could include removing poorly formed or outlier elements. Or it could be used after a PSI protocol [25, 26, 5] or in database join [22] to remove elements that were not matched. If we are willing to reveal the number of elements meeting this criterion, we can use a shuffle to securely remove these elements so that subsequent operations can be evaluated only on the resulting smaller set, which is particularly valuable if the subsequent computation is expensive (e.g. a machine learning task [21]). To do this, we first shuffle the set, then apply a 2PC to each element to evaluate the criterion, revealing the result bit in the clear, and finally remove those items whose result is 1.

Sorting under 2PC Our secret shared shuffle protocol can also be used to build efficient protocols for other fundamental operations. For example, in order to sort a list of secret shared elements and output the resulting secret shares, we can

² This in particular allows us to avoid the $\ell N \log N$ communication price of the permutation network-based approach (which stems from the fact that one has to feed the whole ℓ -bit element into each atomic swap of a permutation network, to retain security).

use the shuffle-and-reveal approach proposed by [15] together with our secret-shared shuffle. The idea in [15] is that if the data is shuffled first, then sorting algorithms can reveal the result of each comparison operation in the clear without compromising security. Thus their approach is to first shuffle the data, and then run a sorting algorithm where each comparison is done under 2PC, with the result revealed in the clear. This yields more efficient protocols than the standard oblivious sorting protocol based on sorting networks; those protocols either have huge constants [1] or require $O(N \log^2 N)$ running time (using Bitonic Sorting network), where N is the number of elements in the database. Note that in many cases we want to sort not just a set of elements, but also some associated data for each element.

Sort, in addition to being a fundamental operation, can be used to find the top k results in a list, to evaluate the median or quantiles, to find outliers, and so on.

Secure Computation for RAM programs There has been a line of work starting with [12, 10, 19, 18, 20, 30, 28, 9] that looks at secure computation for RAM programs (as opposed to circuits). The primary building block in these constructions is oblivious RAM (ORAM), which allows to hide memory accesses made by the computation. A naive way to initialize ORAM is to perform an ORAM write operation for each input item, but the concrete costs on this are very high. [18, 30] show that this can be made much more efficient using a shuffle: the parties simply permute their entries using a random secret shared permutation and then store them as the ORAM memory. [30] achieve significant improvements by using garbled circuits to implement a permutation network; as we will see in section 7 our solution far outperforms this approach, so we should get significant performance improvements for this application. Note that in ORAM it is often beneficial to have somewhat large block size, and our protocol for secret-shared shuffle is especially advantageous in the setting where elements are large.

1.2 Technical overview.

Notation. By bold letters $\mathbf{x}, \mathbf{a}, \mathbf{b}, \mathbf{r}, \mathbf{\Delta}$ we denote vectors of N elements, and by $\mathbf{x}[j]$ we denote the j -th element of \mathbf{x} . By $\pi(\mathbf{x})$, where π is a permutation, we denote the permuted vector $(\mathbf{x}[\pi(1)], \dots, \mathbf{x}[\pi(N)])$.

Secret-Shared Shuffle. Recall that the goal of the secret-shared shuffle is to let parties learn secret shares of a shuffled dataset. More concretely, consider parties P_0, P_1 , where P_1 owns database \mathbf{x} . Our goal is to build a protocol which allows P_0 to learn \mathbf{r} and P_1 to learn $\mathbf{r} \oplus \pi(\mathbf{x})$, but nothing more; here \mathbf{r} is a random vector of the same size as the database, and π is a random permutation of appropriate size. Our protocol also works for the case when \mathbf{x} was secret shared between P_0 and P_1 to begin with (instead of being an input of one party).

Secret-shared shuffle can be easily built given its variant, which we call *Permute+Share*, where one of the parties *chooses* the permutation. That is, in this protocol P_0 holds π and P_1 holds \mathbf{x} , and as before, they would like to learn \mathbf{r} and $\mathbf{r} \oplus \pi(\mathbf{x})$, respectively. Indeed, secret-shared shuffle can be obtained by executing

Permute+Share twice, where first P_0 and then P_1 chooses the permutation (note that in the second execution the database is itself already secret-shared). Thus, in the rest of the introduction we describe how to build Permute+Share.

Our construction proceeds in several steps: first we explain how to build Permute+Share using another protocol called Share Translation. Then we build the latter from oblivious punctured vector primitive, which can be in turn implemented using a GGM-based PRF and oblivious transfer with low communication. Note that we are going to describe our protocols using \oplus (XOR) operation for simplicity, however, in the main body we instead use a more general syntax with addition and subtraction, to allow our protocols to work in different groups.

Building simplified Permute+Share from Share Translation. We first describe a simplified and inefficient version of Permute+Share; the running time of this protocol is proportional to the square of the size of the database. Later in the introduction we explain how we exploit the structure of Benes permutation network [2] to achieve our final protocol.

As a starting point, consider the following idea: P_1 chooses random masks $\mathbf{a} = (\mathbf{a}[1], \dots, \mathbf{a}[N])$ and sends its masked data $\mathbf{x} \oplus \mathbf{a}$ to P_0 . Now P_0 and P_1 together hold a secret-shared \mathbf{x} , albeit not permuted. Note that P_0 knows the permutation π and could easily locally rearrange its shares in order of $\pi(\mathbf{x} \oplus \mathbf{a})$. However, P_1 doesn't know π and thus cannot rearrange \mathbf{a} into $\pi(\mathbf{a})$. Further, any protocol which allows P_1 to learn $\pi(\mathbf{a})$ would immediately reveal π to P_1 , since P_1 also knows \mathbf{a} .

Therefore, instead of choosing a single set of masks, P_1 should choose two different and independent sets of masks, \mathbf{a} and \mathbf{b} , where \mathbf{a} , as before, is used to hide \mathbf{x} from P_0 , and \mathbf{b} will become the final P_1 's share of $\pi(\mathbf{x})$. However, now P_0 has a problem: since P_1 's share is \mathbf{b} , P_0 's share should be $\pi(\mathbf{x}) \oplus \mathbf{b}$; however, P_0 only receives $\mathbf{x} \oplus \mathbf{a}$ from P_1 , and has no way of “translating” it into $\pi(\mathbf{x}) \oplus \mathbf{b}$. Thus we additionally let parties execute a Share Translation protocol to allow P_0 obtain a “translation function” $\Delta = \pi(\mathbf{a}) \oplus \mathbf{b}$, as we explain next in more detail:

Share Translation protocol takes as input permutation π from P_0 and outputs vectors Δ to P_0 and \mathbf{a}, \mathbf{b} to P_1 , such that $\Delta = \pi(\mathbf{a}) \oplus \mathbf{b}$, and, roughly speaking, \mathbf{a}, \mathbf{b} look random³. A simple version of Permute+Share can be obtained from Share Translation as follows:

1. P_0 and P_1 execute a Share Translation protocol, where P_0 holds input π , receives output Δ , and P_1 receives output \mathbf{a}, \mathbf{b} .
2. P_1 sends $\mathbf{x} \oplus \mathbf{a}$ to P_0 and sets its final share to \mathbf{b} .
3. P_0 sets its share to $\pi(\mathbf{x} \oplus \mathbf{a}) \oplus \Delta$. Note that this is equal to $\pi(\mathbf{x}) \oplus \pi(\mathbf{a}) \oplus \pi(\mathbf{a}) \oplus \mathbf{b} = \pi(\mathbf{x}) \oplus \mathbf{b}$, and therefore the parties indeed obtain secret-shared $\pi(\mathbf{x})$.

In other words, the share translation vector Δ allows P_0 to translate “shares of x under \mathbf{a} ” into “shares of permuted x under \mathbf{b} ”.

³ More precisely, P_1 shouldn't learn anything about π , and P_0 shouldn't learn \mathbf{a}, \mathbf{b} , except for what is revealed by π and Δ (note that it still learns, e.g., $\mathbf{a}_{\pi(1)} \oplus \mathbf{b}_1$).

Note that the Share Translation protocol can be viewed as a variant of Permute+Share protocol, with a difference that the “data” which is being permuted and shared is pseudorandom and out of parties’ control (i.e. it is chosen by the protocol): indeed, in Share Translation protocol P_1 receives the “pseudorandom data” \mathbf{a} , and in addition P_0 and P_1 receive $\mathbf{\Delta} = \pi(\mathbf{a}) \oplus \mathbf{b}$ and \mathbf{b} , respectively, which can be thought of as shares of $\pi(\mathbf{a})$. In other words, we reduced the problem of permuting the fixed data \mathbf{x} to the problem of permuting some pseudorandom, out-of-control data \mathbf{a} . In the following paragraphs we explain how we can exploit pseudorandomness of \mathbf{a} and \mathbf{b} to build Share Translation protocol with reduced communication complexity.

Building Share Translation from Oblivious Punctured Vector. We start with defining an Oblivious Punctured Vector protocol (OPV), which is essentially an $(n-1)$ -out-of- n random oblivious transfer ⁴: this protocol, on input $j \in [N]$ from P_0 , allows parties to jointly generate vector \mathbf{v} with random-looking elements such that:

- P_0 learns all vector elements except for its j -th element $\mathbf{v}[j]$;
- P_1 learns the whole vector \mathbf{v} (but doesn’t learn index j)⁵.

We use OPV to build Share Translation as follows: the parties are going to run N executions of OPV protocol to generate N vectors $\mathbf{v}_1, \dots, \mathbf{v}_N$, where P_0 ’s input in execution i is $\pi(i)$. Consider an $N \times N$ matrix $\{\mathbf{v}_i[j]\}_{i,j \in [N]}$. By the properties of OPV protocol, P_1 learns the whole matrix, and P_0 learns the matrix except for elements corresponding to the permutation, i.e. it learns nothing about $\mathbf{v}_1[\pi(1)], \dots, \mathbf{v}_N[\pi(N)]$ (see fig. 1).

Then P_1 sets elements of \mathbf{a}, \mathbf{b} to be column- and row-wise sums of the matrix elements, i.e. for all $i \in N$ it sets $\mathbf{a}[i] \leftarrow \bigoplus_j \mathbf{v}_j[i]$, and for all $j \in N$ it sets $\mathbf{b}[j] \leftarrow \bigoplus_i \mathbf{v}_j[i]$. P_0 computes $\mathbf{\Delta}[i]$ by taking the the sum of column $\pi(i)$ (except the element $\mathbf{v}_i[\pi(i)]$ which it doesn’t know) and adding the sum of row i (again, except the element $\mathbf{v}_i[\pi(i)]$ which it doesn’t know), i.e. it sets
$$\mathbf{\Delta}[i] \leftarrow \left(\bigoplus_{j \neq i} \mathbf{v}_j[\pi(i)] \right) \oplus \left(\bigoplus_{j \neq \pi(i)} \mathbf{v}_j[i] \right).$$

Correctness of this protocol can be immediately verified: indeed, each $\mathbf{\Delta}[i] = \mathbf{a}[\pi(i)] \oplus \mathbf{b}[i]$, since the missing value $\mathbf{v}_i[\pi(i)]$ participates in the sum $\mathbf{a}[\pi(i)] \oplus \mathbf{b}[i]$ twice and therefore doesn’t influence the result. For security, note that P_0 doesn’t learn anything about \mathbf{a}, \mathbf{b} (except for $\mathbf{\Delta}$), since it is missing exactly one element from each row and column of the matrix; the missing element acts as a one-time pad and hides each $\mathbf{a}[i], \mathbf{b}[j]$ from P_0 . P_1 doesn’t learn anything about the permutation π due to index hiding property of the OPV protocol.

⁴ We note that similar definitions were developed independently in [27, 3].

⁵ Note that this is very similar to 1-out-of- N OT- except that j specifies which element P_0 doesn’t learn - and in fact is almost the same as $N-1$ -out-of- N OT. The difference is that in our primitive vector \mathbf{v} is pseudorandom and given by the protocol to the parties (rather than chosen by the sender as in standard OT). We use this fact to save on communication.

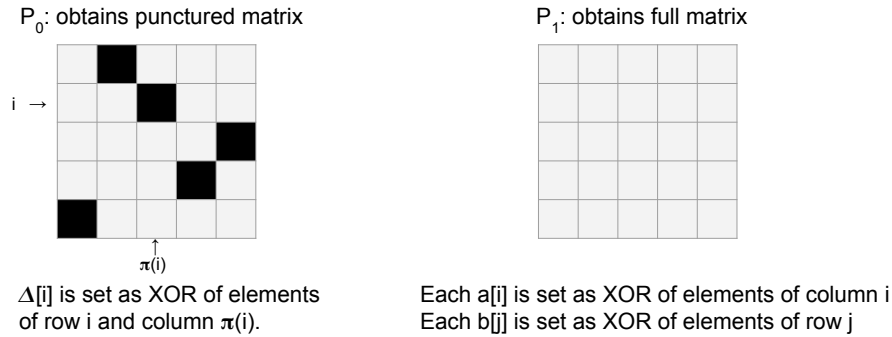


Fig. 1. (left) P_0 receives a “punctured” matrix, which is missing elements at positions $(i, \pi(i))$. Note that the missing elements are not needed to compute Δ . (right) P_1 receives the full matrix and uses it to compute masks \mathbf{a}, \mathbf{b} .

Note that this protocol has running time proportional to N^2 - we will show how to reduce this below.

Building Oblivious Punctured Vector from OT and PRFs. Oblivious Punctured Vector can be implemented using any $(n - 1)$ -out-of- n OT, but in order to make it communication-efficient, we devise a new technique which was inspired by the protocol for distributed point function by Doerner and Shelat [6]. The same technique appears in concurrent and independent works⁶ of Schoppmann et al and Boyle et al [27, 3] in the context of silent OT extension and vector-OLE.

In the beginning of the protocol P_1 computes \mathbf{v} by choosing key for GGM PRF at random, denoted seed_ϵ , and setting each $\mathbf{v}[i] \leftarrow \text{PRF}(\text{seed}_\epsilon; i)$, $i \in [N]$. Recall that in GGM construction the key is treated as a prg seed, which implicitly defines a binary tree with leaves containing PRF evaluations $F(1), F(2), \dots, F(N)$. In other words, we set vector \mathbf{v} to contain values at the leaves of the tree.

Let P_0 's input in the OPV protocol be j . This means that P_0 should learn leaves $F(i), i \neq j$, as a result of the protocol. This can be done as follows. Let us denote internal seeds in the tree by $\{\text{seed}_\gamma\}$, where γ is a string describing the position of the node in the tree (in particular, at the root $\gamma = \epsilon$, an empty string). Let's assume for concreteness that the first bit of j is 1. The parties are going to run 1-out of-2 OT protocol, where P_0 's input is the complement of the first bit of j , i.e. 0, and P_1 's inputs are $\text{seed}_0, \text{seed}_1$. This allows P_0 to recover seed_0 and therefore to locally compute the left half of the tree, i.e. all values $F(1), \dots, F(N/2)$, and corresponding intermediate seeds.

Next, assume the second bit of j is 0. Note that the parties could run 1-out of-4 OT to let P_0 learn seed_{11} and therefore locally compute the right quarter of the tree $F(3N/4), \dots, F(N)$, then run 1-out of-8 OT and so on. However, this

⁶ Our work was submitted to Eurocrypt 2020 on September 26, 2019, and [27, 3] appeared in the public domain (ePrint) roughly at the same time.

approach would require eventually sending 1-out of N OT, which defeats the initial purpose of having $\log N$ 1-out-of-2 OTs only.

Instead, we let P_0 learn seed_{11} in a different way: we let P_1 send only *two* values, via 1-out-of-2-OT: the first value is the sum of seeds which are left children, i.e. $\text{seed}_{00} \oplus \text{seed}_{10}$, and the second value is the sum of seeds which are right children, i.e. $\text{seed}_{01} \oplus \text{seed}_{11}$. Since P_0 already knows the whole left subtree and in particular seed_{00} and seed_{01} , it can receive $\text{seed}_{01} \oplus \text{seed}_{11}$ from the OT protocol and add seed_{01} to it to obtain seed_{11} . (We note that this idea of sending the sums of left and right children is coming from the work of Doerner and Shelat [6]).

More generally, the parties execute $\log N$ 1-out-of-2 OTs - one for each level of the tree - where at each level k the first input to OT is the sum of all odd seeds at that level, and the second input to OT is the sum of all even seeds at that level. It can be seen that each sum contains exactly one term which P_0 doesn't know yet, and therefore it can receive the appropriate sum (depending on the k -th bit of j) and subtract other seeds from it to learn the next seed of the subtree. Note that these OT's can be executed in parallel.

Note that the running time of the parties is proportional to the vector size, but their communication size only depends on its logarithm.

Applying Share Translation to the decomposed permutation. Recall that, while communication complexity in our protocol is low, computation complexity is proportional to the size of the database squared, and thus is only efficient for a small database. To deal with this issue, we change the way how **Permute+Share** is built from **Share Translation**: instead of applying **Share Translation** to the whole permutation π directly, we first split the permutation π into smaller permutations in a special way, then apply **Share Translation** to each separate permutation to get multiple shares, and then recombine these shares to obtain shares with respect to π .

More concretely, the idea is to split the permutation π into a composition of multiple permutations $\pi_1 \circ \dots \circ \pi_d$, such that each π_i is itself a composition of several disjoint permutations, each acting on T elements, for some parameter T . We refer to this as (T, d) -subpermutation representation of π . Such a representation can be found using a special structure of Benes permutation network. For instance, as shown on Fig. 2, the first two layers of a network on 8 elements can be split into two permutations, acting on $T = 4$ elements each, where the first permutation acts on odd elements and the second permutation acts on even elements. We present the full description of our decomposition in Section 6.2.

With such a decomposition in place, parties can run parallel executions of **Share Translation**, each acting on domain of size T . Note that, since the running time of a single **Share Translation** is proportional to the domain size squared, it is better to choose relatively small T . In our experiments, the typical optimal values of T were 16, 128, 256, depending on other parameters.

Note that setting $T = N$ corresponds to our simplified **Permute+Share** protocol described before, and setting $T = 2$ results in essentially computing the permutation network, where each swap is implemented in a somewhat-complicated




Fig. 2. (left) The first two layers of the Benes permutation network for 8 elements. A link indicates that the corresponding elements are potentially swapped, depending on the underlying permutation. (right) A grouping of these layers into two disjoint permutations acting on 4 elements each: one acting on white elements and the other acting on black elements.

way (using Share Translation protocol). Thus, this scheme can be thought of as a golden middle between the two approaches.

It remains to note that parties can run all executions of Share Translation in parallel (as opposed to taking multiple rounds, following the layered structure of the permutation network). To achieve this, in all execution except for the first ones, P_1 instead of sending initial masked data $\mathbf{x} \oplus \mathbf{a}$ should send correction vector $\mathbf{a}^{new} \oplus \mathbf{b}^{old}$, which can be added to the shares of P_0 in order to obtain $\mathbf{x} \oplus \mathbf{a}^{new}$. We refer the reader to Section 6.2 for more details.

Achieving simulation-based definition. We note that the protocols we described so far only achieve indistinguishability-based definition, but not simulation-based definition. The problem is that the output values are only *pseudo*-random, and parties in the protocols know their succinct “preimages” (like the GGM PRF root). Thus, the simulator, given a random string as an output of the protocol, cannot simulate internal state of that party since it would amount to compressing a random string.

To achieve simulation-based definition, we slightly modify the original Permute+Share protocol as follows: we additionally instruct P_1 to sample random string \mathbf{w} of the size of the database and send it to P_0 , together with $\mathbf{x} \oplus \mathbf{a}$. Then P_0 should set its share to be $\pi(\mathbf{x} \oplus \mathbf{a}) \oplus \mathbf{\Delta} \oplus \mathbf{w}$, and P_1 should set its share to be $\mathbf{b} \oplus \mathbf{w}$. In other words, P_1 should additionally secret-share its vector \mathbf{b} using random \mathbf{w} . Such a protocol can be simulated by a simulator who executes Share Translation protocol honestly (obtaining some \mathbf{a}' , \mathbf{b}' , $\mathbf{\Delta}'$) and then sets simulated \mathbf{w} to be $\mathbf{z} \oplus \mathbf{b}'$ (where \mathbf{z} is the output of Permute+Share protocol simulated by an external simulator)

2 Notations

We denote the security parameter as λ . The bit length of each element in the input set is ℓ , $\ell = \text{poly}(\lambda)$. We denote an upper bound on the size of the database as N . Ideal functionality is denoted as \mathcal{F} . We will denote vectors with bold fonts and individual elements with indices. For example, \mathbf{v} is a vector of N elements where each individual element is denoted as v_i . $\leftarrow^{\$}$ denotes selected uniformly at random from a domain. By S_N we denote the group of all permutations on N elements.

We also make use of the following notation:

Exec: Let Π be a two-party protocol. By $(\text{output}_0, \text{output}_1) \leftarrow \text{exec}^{\Pi}(\lambda; x_0, x_1; r_0, r_1)$ we denote the concatenated outputs of all parties after the execution of the protocol Π with security parameter λ on inputs x_0, x_1 using randomness r_0, r_1 .

View: Let Π be a two-party protocol. By $\text{view}_b^{\Pi}(\lambda; x_0, x_1; r_0, r_1)$ we denote the view of party b when parties P_0 and P_1 run the protocol Π with security parameter λ on inputs x_0, x_1 using randomness r_0, r_1 . The view of each party includes its inputs, random coins, all messages it receives, and its outputs. When the context is clear, we also write view_b for short.

Honest-but-curious security for a 2PC: Honest-but-curious security for a 2PC protocol Π evaluating function \mathcal{F} is defined in terms of the following two experiments:

$\text{IDEAL}_{\text{sim},b}^{\mathcal{F}}(\lambda, x_0, x_1)$ evaluates $\mathcal{F}(x_0, x_1)$ to obtain output (y_0, y_1) runs the stateful simulator $\text{sim}(1^\lambda, b, x_b, y_b)$ which produces a simulated view view_b for party P_b . The output of the experiment is (view_b, y_{1-b}) .

$\text{REAL}_b^{\Pi}(\lambda, x_0, x_1)$ runs the protocol with security parameter λ between honest parties P_0 with input x_0 and P_1 with input x_1 who obtain outputs y_0, y_1 respectively. It outputs (view_b, y_{1-b}) .

Definition 1. *Protocol Π realizes \mathcal{F} in the honest-but-curious setting if there exists a PPT simulator sim such that for all inputs x_0, x_1 , and corrupt parties $b \in \{0, 1\}$ the two experiments are indistinguishable.*

Pseudo Random Generator Let $\{\mathbf{G}\}_\lambda$ be a family of polynomial size circuits where each $\mathbf{G}_\lambda : \{0, 1\}^{m(\lambda)} \rightarrow \{0, 1\}^{l(\lambda)}$, $l(\lambda) \geq m(\lambda)$. $\{\mathbf{G}\}_\lambda$ is a PRG if the following distributions are computationally indistinguishable:

$$\{\mathcal{D}_1\}_\lambda = \{\mathbf{G}(\mathbf{s}) : \mathbf{s} \leftarrow \{0, 1\}^{m(\lambda)}\}, \{\mathcal{D}_2\}_\lambda = \{x : x \leftarrow \{0, 1\}^{l(\lambda)}\}$$

We will omit the dependence of m and l on λ for simplicity. When $l = 2m$, we call this a length doubling PRG.

Oblivious Transfer (OT) OT is a secure 2-party protocol that realizes the functionality $\mathcal{F}_{\text{OT}} : ((\text{str}_0, \text{str}_1), b) = (\perp, \text{str}_b)$ where $\text{str}_0, \text{str}_1 \in \{0, 1\}^k$, $b \in \{0, 1\}$.

3 Oblivious Punctured Vector (OPV)

3.1 Definition and Security Properties

An Oblivious Punctured Vector (OPV) for domain \mathbb{D} is an interactive protocol between two parties, P_0 and P_1 , where parties' inputs are $((1^\lambda, n), (1^\lambda, n, i))$ and their outputs are $(\mathbf{v}_0, \mathbf{v}_1)$, respectively. Here λ is the security parameter that determines the running time of the protocol, $\mathbf{v}_b, b \in \{0, 1\}$ are vectors of length n , $i \in [n]$ and $\mathbf{v}_b \in [\mathbb{D}]^n$.

This protocol lets the two parties jointly generate vector \mathbf{v} with random-looking elements such that: 1) P_0 learns the whole vector \mathbf{v} but doesn't learn index i . 2) P_1 learns all vector elements except for its i -th element $\mathbf{v}[i]$. So we define the protocol to be *correct* if $\mathbf{v}_1[j] = \mathbf{v}_0[j] \forall j \neq i$.

To capture the first property, we want to say that an adversarial P_0 , who is given two distinct indices $i, i' \in [n]$, $i \neq i'$ and participates in two executions of the protocol, one where party P_1 holds i , and the other, where P_1 holds i' , cannot tell the two executions apart. We call this property *Position hiding*. To capture the second property, we want to say that an adversarial P_1 , who, in addition to its view in the protocol execution, receives the vector \mathbf{v}_0 , cannot differentiate between the two cases: when \mathbf{v}_0 is generated according to `exec` and when \mathbf{v}_0 is generated according to `exec`, then $\mathbf{v}_0[i]$ is replaced a random string from the domain. We call this security property *Value hiding*. We define the properties formally below.

Correctness For any sufficiently large security parameter $\lambda \in \mathbb{N}$, for any $n \in \mathbb{N}$, $i \in [n]$, if $(\mathbf{v}_0, \mathbf{v}_1) \leftarrow \text{exec}^{\text{OPV}}((\lambda n), (\lambda, n, i))$ and $\mathbf{v}_b \in [\mathbb{D}]^n, b \in \{0, 1\}$, then $\mathbf{v}_1[j] = \mathbf{v}_0[j] \forall j \neq i$.

Position hiding For any sufficiently large security parameter $\lambda \in \mathbb{N}$, $n \in \mathbb{N}$, $i, i' \in [n]$, the following distributions are computationally indistinguishable:

$$\begin{aligned} \mathcal{D}_1 &= \{(\mathbf{v}_0, \mathbf{v}_1) \leftarrow \text{exec}^{\text{OPV}}((1^\lambda, n), (1^\lambda, n, i)) : (1^\lambda, n, i, i', \text{view}_0)\} \\ \mathcal{D}_2 &= \{(\mathbf{v}_0, \mathbf{v}_1) \leftarrow \text{exec}^{\text{OPV}}((1^\lambda, n), (1^\lambda, n, i')) : (1^\lambda, n, i, i', \text{view}_0)\} \end{aligned}$$

Value hiding For any sufficiently large security parameter $\lambda \in \mathbb{N}$, for any $n \in \mathbb{N}$, $i \in [n]$, the following distributions are computationally indistinguishable:

$$\begin{aligned} \mathcal{D}_1 &= \{(\mathbf{v}_0, \mathbf{v}_1) \leftarrow \text{exec}^{\text{OPV}}((1^\lambda, n), (1^\lambda, n, i)) : (1^\lambda, n, i, \mathbf{v}_0, \text{view}_1)\} \\ \mathcal{D}_2 &= \{((\mathbf{v}_0, \mathbf{v}_1) \leftarrow \text{exec}^{\text{OPV}}((1^\lambda, n), (1^\lambda, n, i)), \mathbf{v}_0[i] := r \text{ where } r \leftarrow^{\$} \mathbb{D} : \\ &\quad (1^\lambda, n, i, \mathbf{v}_0, \text{view}_1)\}\} \end{aligned}$$

Construction: We defer the formal construction and security proof of Theorem 1 to the Sec A). For an informal description of the construction, please refer to Section 1.2.

Please note that we only count the cryptographic operations while analyzing the computation complexity of our protocols.

Theorem 1. *The OPV construction satisfies position and value hiding as defined in Definition 3.1. The protocol runs n (1-out-of-2) OT on messages of length λ bits in parallel. The communication cost is that of the OTs and the computation cost is the cost of these OTs + n length-doubling PRG computations for each party⁷, where λ is a security parameter and n is the number of elements in the vector.*

3.2 OPV construction for longer strings

Let $\text{OPV}_{\mathbb{D}}$ denote the interactive protocol between two parties, P_0 and P_1 , where parties' inputs are $((1^\lambda, n), (1^\lambda, n, i))$ and their outputs are $(\mathbf{v}_0, \mathbf{v}_1)$, where $\mathbf{v}_b \in [\mathbb{D}]^n$ and \mathbb{D} is strings of length λ . We construct $\text{OPV}_{\mathbb{D}'}$ where \mathbb{D}' is strings of length $\ell \geq \lambda$ using $\text{OPV}_{\mathbb{D}}$ and a PRG $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\ell$ as follows.

- Run $(\mathbf{v}_0, \mathbf{v}_1) \leftarrow \text{exec}^{\text{OPV}_{\mathbb{D}}}((1^\lambda, n), (1^\lambda, n, i))$
- Party P_b , $b \in \{0, 1\}$ does the following: for each $\mathbf{v}_b[j], j \in [1, n]$, expand it to a ℓ -bit string using $G(\mathbf{v}_b[j])$, i.e., $\mathbf{v}'_b[j] \leftarrow G(\mathbf{v}_b[j])$. P_b 's output is \mathbf{v}'_b .

Theorem 2. *If $\text{OPV}_{\mathbb{D}}$ satisfies correctness, position and value hiding as defined in Definition 3.1, and G is a secure PRG, then our construction for $\text{OPV}_{\mathbb{D}'}$ satisfies correctness, position and value hiding as well. The round complexity and communication cost is the same as the cost of $\text{OPV}_{\mathbb{D}}$. The computation cost includes the computation cost of $\text{OPV}_{\mathbb{D}}$ + n λ -bit-to- ℓ -bit PRGs.*

Proof. Correctness: By the correctness of $\text{OPV}_{\mathbb{D}}$, $\mathbf{v}_0[j] = \mathbf{v}_1[j], \forall j \neq i$. Therefore, by our construction, $\mathbf{v}'_0[j] = \mathbf{v}'_1[j], \forall j \neq i$.

Position hiding: For the sake of contradiction, suppose not. Then, there exists a distinguisher D that breaks the position hiding property of $\text{OPV}_{\mathbb{D}'}$. We use D to build a distinguisher \mathcal{A} that breaks the position hiding property of $\text{OPV}_{\mathbb{D}}$ as follows. \mathcal{A} receives $(1^\lambda, n, i, i', \text{view}_0^{\text{OPV}_{\mathbb{D}}})$ as input, where $\text{view}_0^{\text{OPV}_{\mathbb{D}}}$ contains \mathbf{v}_0 . For every $\mathbf{v}_0[j], j \in [1, n]$, \mathcal{A} computes $\mathbf{v}'_0[j] = G(\mathbf{v}_0[j])$. Then it constructs $\text{view}_0^{\text{OPV}_{\mathbb{D}'}}$, which is $\text{view}_0^{\text{OPV}_{\mathbb{D}}}$, augmented with $\mathbf{v}'_0[j]$. \mathcal{A} forwards $(1^\lambda, n, i, i', \text{view}_0^{\text{OPV}_{\mathbb{D}'}})$ to D . Thus, \mathcal{A} directly inherits the success probability D .

Value hiding: Recall that we are trying to prove the following two distributions are computationally indistinguishable.

$$\begin{aligned} \mathcal{D}_1 &= \{(\mathbf{v}'_0, \mathbf{v}'_1) \leftarrow \text{exec}^{\text{OPV}_{\mathbb{D}'}}((1^\lambda, n), (1^\lambda, n, i)) : (1^\lambda, n, i, \mathbf{v}'_0, \text{view}_1^{\text{OPV}_{\mathbb{D}'}})\} \\ \mathcal{D}_2 &= \{((\mathbf{v}'_0, \mathbf{v}'_1) \leftarrow \text{exec}^{\text{OPV}_{\mathbb{D}'}}((1^\lambda, n), (1^\lambda, n, i)), \mathbf{v}'_0[i] := r \text{ where } r \leftarrow^{\$} \mathbb{D}' : \\ &\quad (1^\lambda, n, i, \mathbf{v}'_0, \text{view}_1^{\text{OPV}_{\mathbb{D}'}})\} \end{aligned}$$

The proof will proceed through a series of hybrid steps. We define a series of distributions as follows.

$$H_0: \mathcal{D}_1 = \{(\mathbf{v}'_0, \mathbf{v}'_1) \leftarrow \text{exec}^{\text{OPV}_{\mathbb{D}'}}((1^\lambda, n), (1^\lambda, n, i)) : (1^\lambda, n, i, \mathbf{v}'_0, \text{view}_1^{\text{OPV}_{\mathbb{D}'}})\}$$

⁷ We give the concrete cost of OT and OPV in Section 7.

- H_1 : Identical to the previous distribution except the following: generate $(\mathbf{v}_0, \mathbf{v}_1) \leftarrow \text{exec}^{\text{OPV}_{\mathbb{D}}}((1^\lambda, \mathbf{n}), (1^\lambda, \mathbf{n}, i))$, then set $v_0[i] := r$ where $r \leftarrow^{\$} \mathbb{D}$ and set $\mathbf{v}'_0[i] \leftarrow \mathbf{G}(\mathbf{v}'_0[i])$. By the value-hiding property of $\text{OPV}_{\mathbb{D}}$, H_0, H_1 are identical.
- H_2 : Identical to the previous distribution except the following: instead of computing $\mathbf{v}'_0[i] \leftarrow \mathbf{G}(\mathbf{v}'_0[i])$, set $\mathbf{v}'_0[i] := r'$ where $r' \leftarrow^{\$} \mathbb{D}'$. By the security property of PRG, H_1, H_2 are identical. Note that distribution H_2 is identical to \mathcal{D}_2 . So this concludes the proof of value hiding. \square

4 Share Translation Protocol

4.1 Definition

Share Translation (ST) protocol with parameters (N, ℓ) is an interactive protocol between two parties, P_0 and P_1 , where parties' inputs are (π, \perp) and their outputs are $(\mathbf{\Delta}, (\mathbf{a}, \mathbf{b}))$, respectively. Here π is a permutation on N elements, and $\mathbf{\Delta}, \mathbf{a}, \mathbf{b}$ are all vectors of N elements in group \mathbb{G} , where each element can be represented with ℓ bits. The protocol should satisfy the following correctness and security guarantees:

Correctness: For each sufficiently large security parameter λ , for each $\pi \in S_N$, and for each r_0, r_1 of appropriate length, let $(\mathbf{\Delta}, (\mathbf{a}, \mathbf{b})) \leftarrow \text{exec}^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)$. Then it should hold that $\mathbf{\Delta} = \mathbf{b} - \pi(\mathbf{a})$.

This definition can be modified in a straightforward way for statistical or computational correctness.

Permutation hiding: For all sufficiently large λ it should hold that for all $\pi, \pi' \in S_N$,

$$\text{view}_1^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1) \approx \text{view}_1^{\text{ST}}(\lambda; \pi', \perp; r_0, r_1),$$

where indistinguishability holds over uniformly chosen r_0, r_1 .

Share hiding: For all sufficiently large λ it should hold that for any $\pi \in S_N$,

$$(\mathbf{a}, \mathbf{b}, \text{view}_0^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)) \approx (\mathbf{a}', \mathbf{b}', \text{view}_0^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)),$$

where $(\mathbf{\Delta}, \mathbf{a}, \mathbf{b}) = \text{exec}_{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)$, $\mathbf{a}' \leftarrow^{\$} \mathbb{G}^N$, $\mathbf{b}' = \mathbf{\Delta} + \pi(\mathbf{a}')$, and indistinguishability holds over uniformly chosen r_0, r_1 .

4.2 Construction

We build Share Translation protocol out of an Oblivious Punctured Vector (OPV) protocol for domain $\mathbb{D} = \mathbb{G}$. Let π be P_0 's input in Share Translation protocol. The protocol proceeds as follows:

1. P_0 and P_1 run N executions of the OPV protocol in parallel, where P_0 uses $\pi(i)$ as its input in execution i , for $i \in [N]$. Denote $\mathbf{v}'_i, \mathbf{v}_i$ to be the outputs of the OPV protocol in execution i , for parties P_0 and P_1 , respectively, and denote $\mathbf{v}'_i[j], \mathbf{v}_i[j]$ to be j -th elements of these vectors.

2. For each $i \in [N]$ P_0 sets $\Delta[i] \leftarrow \sum_{j \neq \pi(i)} \mathbf{v}'_i[j] - \sum_{j \neq i} \mathbf{v}'_j[\pi(i)]$. It sets its output to be $\Delta = (\Delta[1], \dots, \Delta[N])$.
3. For each $i \in [N]$ P_1 sets $\mathbf{b}_i \leftarrow \sum_j \mathbf{v}_i[j]$, $\mathbf{a}_i \leftarrow \sum_j \mathbf{v}_j[i]$. It sets (\mathbf{a}, \mathbf{b}) as its output, where $\mathbf{a} = (\mathbf{a}[1], \dots, \mathbf{a}[N])$, $\mathbf{b} = (\mathbf{b}[1], \dots, \mathbf{b}[N])$.

Theorem 3. *The construction described above satisfies correctness, permutation hiding and share hiding, assuming underlying OPV protocol satisfies correctness, value hiding and position hiding. The round complexity, communication and computation cost of this protocol are equal to those of N instances of OPV run in parallel.*

Correctness. For any $i \in [N]$ we have

$$\begin{aligned} \Delta_i &= \sum_{j \neq \pi(i)} \mathbf{v}'_i[j] - \sum_{j \neq i} \mathbf{v}'_j[\pi(i)] \stackrel{(1)}{=} \sum_{j \neq \pi(i)} \mathbf{v}_i[j] - \sum_{j \neq i} \mathbf{v}_j[\pi(i)] \stackrel{(2)}{=} \\ &\stackrel{(2)}{=} \sum_{j \in [N]} \mathbf{v}_i[j] - \sum_{j \in [N]} \mathbf{v}_j[\pi(i)] = \mathbf{b}_i - \mathbf{a}_{\pi(i)}. \end{aligned}$$

Here (1) follows from correctness of the OPV protocol, and (2) holds since we add and subtract the same value $\mathbf{v}_i[\pi(i)]$. Note that a computationally (resp., statistically, perfectly) correct OPV protocol results in a computationally (resp., statistically, perfectly) correct ST protocol.

Permutation hiding. Recall that we need to show that for all $\pi, \pi' \in S_N$,

$$\text{view}_1^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1) \approx \text{view}_1^{\text{ST}}(\lambda; \pi', \perp; r_0, r_1).$$

We show this indistinguishability in a sequence of hybrids H_0, H_1, \dots, H_N , where:

- $H_0 = \text{view}_1^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)$, for uniformly chosen r_0, r_1 ,
- $H_N = \text{view}_1^{\text{ST}}(\lambda; \pi', \perp; r_0, r_1)$, for uniformly chosen r_0, r_1 ,
- For $1 \leq i < N$, $H_i = \text{view}_1^{(i)}(\lambda; (\pi, \pi'), \perp; r_0, r_1)$, where $\text{view}_1^{(i)}(\lambda; (\pi, \pi'), \perp; r_0, r_1)$ is a view of P_1 in the modified Share Translation protocol where party P_0 uses $\pi'(j)$ as its input in OPV executions $1 \leq j \leq i$ and $\pi(j)$ as its input in OPV executions $i < j \leq N$. r_0, r_1 are uniformly chosen.

We argue that for each $1 \leq i \leq N$ $H_i \approx H_{i-1}$ due to position-hiding property of the OPV protocol, and therefore $H_0 \approx H_N$.

Indeed, note that the only difference between H_i and H_{i-1} is that in i -th execution of OPV party P_0 uses input $\pi'(i)$ instead of $\pi(i)$. Therefore if some PPT adversary distinguishes between H_i and H_{i-1} , then we break position hiding of OPV as follows. Given the challenge in the OPV position hiding game

($\pi(i), \pi'(i), \text{view}_1^{\text{OPV}}(\lambda; x, \perp; r_0^{\text{OPV}}, r_1^{\text{OPV}})$), where $r_0^{\text{OPV}}, r_1^{\text{OPV}}$ are uniformly chosen randomness of P_0 and P_1 in the OPV protocol, and $\text{view}_1^{\text{OPV}}$ is a view of P_1 in OPV protocol (which uses randomness $r_0^{\text{OPV}}, r_1^{\text{OPV}}$ and P_0 's input x which is either $\pi(i)$ or $\pi'(i)$), we execute the rest $N - 1$ OPV protocols honestly using uniform randomness for each party and setting P_0 's input to $\pi'(j)$ (for executions $j < i$) and $\pi(j)$ (for executions $j > i$). Let $\mathbf{v}_j, j = 1, \dots, N$, be the output of P_1 in j -th execution of OPV.

We give the adversary P_1 's view in all N OPV executions (including $\text{view}_1^{\text{OPV}}(\lambda; x, \perp; r_0^{\text{OPV}}, r_1^{\text{OPV}})$ of i -th execution which we received as a challenge). Depending on whether challenge input x was $\pi(i)$ or $\pi'(i)$, the distribution the adversary sees is either H_{i-1} or H_i . Therefore, if the adversary distinguishes between the two distributions, we can break position hiding of OPV protocol with the same success probability.

Share hiding. Recall that we need to show that for any $\pi \in S_N$,

$$(\mathbf{a}, \mathbf{b}, \text{view}_0^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)) \approx (\mathbf{a}', \mathbf{b}', \text{view}_0^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)),$$

where \mathbf{a}, \mathbf{b} are true shares produced by the protocol, and \mathbf{a}', \mathbf{b}' are uniformly random, subject to $\mathbf{\Delta} = \mathbf{b} - \pi(\mathbf{a})$.

We show this indistinguishability in a sequence of hybrids H_0, H_1, \dots, H_N , where:

- $H_0 = (\mathbf{a}, \mathbf{b}, \text{view}_0^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1))$, for uniformly chosen r_0, r_1 ,
- $H_N = (\mathbf{a}', \mathbf{b}', \text{view}_0^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1))$, for uniformly chosen r_0, r_1, \mathbf{a}' , and $\mathbf{b}' = \mathbf{\Delta} + \pi(\mathbf{a})$, where $(\mathbf{\Delta}, \mathbf{a}, \mathbf{b}) = \text{exec}_{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)$,
- $H_i = (\mathbf{a}^{(i)}, \mathbf{b}^{(i)}, \text{view}_0^{\text{ST}}(\lambda; \pi, \perp; r_0, r_1))$, where $(\mathbf{\Delta}, \mathbf{a}, \mathbf{b}) = \text{exec}_{\text{ST}}(\lambda; \pi, \perp; r_0, r_1)$ is the output of the Share Translation protocol for random r_1, r_2 , $\mathbf{a}^{(i)} = (\mathbf{a}_1^{(i)}, \dots, \mathbf{a}_N^{(i)})$ is such that $\mathbf{a}_j^{(i)}$ is uniformly chosen for $1 \leq j \leq i$, $\mathbf{a}_j^{(i)} = \mathbf{a}_j$ for $i < j \leq N$, and $\mathbf{b}^{(i)} = \mathbf{\Delta} + \pi(\mathbf{a}^{(i)})$.

We argue that for each $1 \leq i \leq N$ $H_i \approx H_{i-1}$, by reducing it to value hiding of OPV protocol. Indeed, note that the only difference between H_i and H_{i-1} is that $\mathbf{a}_i^{(i)}$ is generated uniformly at random, rather than set to the true output of the protocol. Therefore if some PPT adversary distinguishes between H_i and H_{i-1} , then we break security of OPV as follows. Assume we are given the challenge $(\mathbf{v}_i, \text{view}_0^{\text{OPV}}(\lambda; \pi(i), \perp; r_0^{\text{OPV}}, r_1^{\text{OPV}}))$, where $r_0^{\text{OPV}}, r_1^{\text{OPV}}$ are uniformly chosen randomness of P_0 and P_1 in the OPV protocol, and $\text{view}_0^{\text{OPV}}$ is a view of P_0 in OPV protocol (which uses randomness $r_0^{\text{OPV}}, r_1^{\text{OPV}}$ and P_0 's input $\pi(i)$), and challenge \mathbf{v}_i is either the true output of P_1 , or the output of P_1 except that $\mathbf{v}_i[\pi(i)]$ is set to a uniform value. We execute the rest $N - 1$ OPV protocols honestly using uniform randomness for each party and setting P_0 's input to $\pi(j)$, for $j \neq i$. Let's denote the outputs of each OPV execution $j \neq i$ as $(\mathbf{v}_j, \mathbf{v}'_j)$.

Then we compute $\mathbf{a}^{(i)}, \mathbf{b}^{(i)}$ as follows:

- $\mathbf{b}^{(i)}[k] \leftarrow \sum_j \mathbf{v}_k[j]$, for each $k \in [N]$,
- $\mathbf{a}^{(i)}[k] \leftarrow \sum_j \mathbf{v}_j[k]$, for each $k \in [N]$,

Then we give the adversary $\mathbf{a}^{(i)}, \mathbf{b}^{(i)}$, and the views of party P_0 in all N OPV executions (including the challenge view $\text{view}_0^{\text{OPV}}(\lambda; \pi(i), \perp; r_0^{\text{OPV}}, r_1^{\text{OPV}})$ of i -th execution). Depending on whether challenge $\mathbf{v}_i[\pi(i)]$ was uniform or not, the distribution the adversary sees is either H_{i-1} or H_i .

Thus, we showed that H_0 and H_N are indistinguishable, as required.

5 (T, d) -subpermutation representation based on Benes permutation network

In this section we describe how to obtain (T, d) -subpermutation representation, which is used in our final construction of Share Translation and secret-shared shuffle in section 6. That is, we show how to represent any permutation $\pi \in S_N$, where $N = 2^n$ for some integer n , as a composition of permutations $\pi_1 \circ \dots \circ \pi_d$, such that each π_i is itself a composition of several disjoint permutations, each acting on T elements, for some parameter T . In our construction $d = 2^{\lceil \frac{\log N}{\log T} \rceil} - 1$.

Our decomposition is based on the special structure of the Benes permutation network. This network has $2 \log N - 1$ layers, each containing $N/2$ 2-element permutations (that is, each is either an identity permutation or a swap). Specifically, if inputs are numbered with index $1, \dots, N$, where each index is expressed in binary as $\sigma_1, \dots, \sigma_n$, then the j -th layer and the $2 \log N - j$ -th layer contain 2-element permutations, each acting on elements number $\sigma_1 \dots \sigma_{j-1} 0 \sigma_{j+1}, \dots, \sigma_n$ and $\sigma_1 \dots \sigma_{j-1} 1 \sigma_{j+1}, \dots, \sigma_n$, for all $\sigma_1, \dots, \sigma_{j-1}, \sigma_{j+1}, \dots, \sigma_n \in \{0, 1\}^{n-1}$.

Now we describe our decomposition of π into $\pi_1 \circ \dots \circ \pi_d$. For any parameter $T = 2^t$, $t \in \mathbb{N}$, set $d = 2^{\lceil \frac{n}{t} \rceil} - 1$, and consider Benes network for π . We set π_1 to consist of first t layers $1, \dots, t$ of this network, π_2 to consist of next t layers $t+1, \dots, 2t$, and so on, except for the middle permutation $\pi_{\lfloor \frac{d}{2} \rfloor + 1}$ which consists of $2t - 1$ layers in the middle⁸. That is, we set each π_i , for $i = 1, \dots, \lfloor \frac{d}{2} \rfloor$, to consist of t consecutive layers number $i \cdot t - (t - 1), \dots, i \cdot t - 1, i \cdot t$, and π_i for $i = \lfloor \frac{d}{2} \rfloor + 2, \dots, d$ are defined symmetrically. From the description of Benes layers above, it follows that these t consecutive layers do not permute all N elements together, but instead only permute elements within each group of the form $\sigma_1, \dots, \sigma_{i(t-1)} x \sigma_{i \cdot t + 1}, \dots, \sigma_n$, where x includes all t -bit strings, and the remaining $n - t$ bits $\sigma_1, \dots, \sigma_{i(t-1)}, \sigma_{i \cdot t + 1}, \dots, \sigma_n$ are fixed. Therefore it follows that each π_i , $i \neq \lfloor \frac{d}{2} \rfloor + 1$, consists of $2^{n-t} = N/T$ disjoint permutations, each acting on $2^t = T$ elements. Similarly, the middle permutation $\pi_{\lfloor \frac{d}{2} \rfloor + 1}$, consisting of $2t - 1$ layers in the middle of the network, only permutes elements within each group of the form $\sigma_1, \dots, \sigma_{n-t} x$, and thus can also be represented as a combination of N/T disjoint permutations each acting on T elements.

Finally, note that the total number of permutations is $\lceil \frac{(2n-1)-(2t-1)}{t} \rceil + 1 = 2^{\lceil \frac{n}{t} \rceil} - 1$. Therefore, $\pi = \pi_1 \circ \dots \circ \pi_d$ is indeed a (T, d) -subpermutation representation of π , for $d = 2^{\lceil \frac{n}{t} \rceil} - 1$.

⁸ For a more general case, when $\log T$ doesn't divide $\log N$, there are $2n - 1 - t(d - 1)$ layers in the middle.

6 Permute and Share and Secret-Shared Shuffle

Recall that we use $\pi(\mathbf{x})$ for a permutation π and vector \mathbf{x} to mean the permutation which produces $x_{\pi(1)}, \dots, x_{\pi(N)}$.

We will use the **Share Translation** scheme we presented in the previous scheme to construct first a secure computation for permuting and secret sharing elements where one party chooses the permutation and the other the elements, and then a construction for a full secret-shared shuffle.

6.1 Definitions

We consider the following functionality, which we call **Permute+Share**, in which one party provides as input a permutation π , and the other party provides as input a set of elements \mathbf{x} in group \mathbb{G} , and the output is secret shares of the permuted elements:

$$\mathcal{F}_{\text{Permute+Share}[N,\ell]}(\pi, \mathbf{x}) = (\mathbf{r}, \pi(\mathbf{x}) - \mathbf{r}), \text{ where } \mathbf{r} \xleftarrow{\$} \mathbb{G}^N.$$

We can also consider the equivalent functionality when the permutation or the initial database is secret shared as input. (Here we consider a secret sharing of permutation π which consists of two permutations π_0, π_1 such that $\pi = \pi_0 \circ \pi_1$.)

Finally, we define the secret shared shuffle functionality:

$$\mathcal{F}_{\text{SecretSharedShuffle}[N,\ell]}(\mathbf{x}_0, \mathbf{x}_1) = (\mathbf{r}, \pi(\mathbf{x}_0 + \mathbf{x}_1) - \mathbf{r}),$$

where $\mathbf{r} \xleftarrow{\$} \mathbb{G}^N$ and π is a random permutation over N elements.

6.2 Permute + Share from Share Translation

Let ShareTrans_T be a protocol satisfying the definition in Section 4 for permutations on T elements in group \mathbb{G} , where each element can be represented in ℓ bits. Let T, d be some parameters such that any permutation in S_N has (T, d) -subpermutation representation (e.g. $d = 2^{\lceil \frac{\log N}{\log T} \rceil} - 1$ for any $T = 2^t$, as described in section 5). We construct our permute and share protocol **Permute + Share** using (T, d) -subpermutation representation as follows.

1. P_0 computes the (T, d) -subpermutation representation π_1, \dots, π_d of its input π .
2. For each layer i , the parties run N/T instances of ShareTrans_T , with P_0 providing as input the N/T permutations making up π_i . (Note that all of these instances and layers can be run in parallel.) For each i , P_1 obtains $\mathbf{a}^{(i,1)}, \dots, \mathbf{a}^{(i,N/T)}$ and $\mathbf{b}^{(i,1)}, \dots, \mathbf{b}^{(i,N/T)}$. Call the combined vectors $\mathbf{a}^{(i)}$ and $\mathbf{b}^{(i)}$. Similarly, P_0 obtains $\mathbf{\Delta}^{(i,1)}, \dots, \mathbf{\Delta}^{(i,N/T)}$, which we will call $\mathbf{\Delta}^{(i)}$.
3. For each $i \in 1, \dots, d-1$, P_1 computes $\mathbf{\delta}^{(i)} = \mathbf{a}^{(i+1)} - \mathbf{b}^{(i)}$ and sends it to P_0 . P_1 also sends $\mathbf{m} = \mathbf{x} + \mathbf{a}^{(1)}$, and samples and sends random \mathbf{w} . P_1 outputs $\mathbf{b} = \mathbf{w} - \mathbf{b}^{(d)}$.
4. P_0 computes $\mathbf{\Delta} = \mathbf{\Delta}^{(d)} + \pi_d(\mathbf{\delta}^{(d-1)} + \mathbf{\Delta}^{(d-1)}) + \pi_{d-1}(\mathbf{\delta}^{(d-2)} + \mathbf{\Delta}^{(d-2)}) + \dots + \pi_2(\mathbf{\delta}^{(1)} + \mathbf{\Delta}^{(1)})$ and outputs $\pi(\mathbf{m}) + \mathbf{\Delta} - \mathbf{w}$.

Theorem 4. *Let N and ℓ be the number of elements in the database and the size of each element, respectively, and let T, d be arbitrary parameters such that any permutation in S_N has (T, d) -subpermutation representation. Then the construction described above is a **Permute+Share** protocol secure against static semi-honest corruptions with the following efficiency:*

- The communication cost is $(d + 1)N\ell$ bits together with the cost of dN/T **Share Translation** protocols on T elements each, run in parallel,
- The computation cost is equal to the cost of dN/T **Share Translation** protocols on T elements each, run in parallel.⁹

Correctness By correctness of ShareTrans_T , for all i $\Delta^{(i)} = \mathbf{b}^{(i)} - \pi_i(\mathbf{a}^{(i)})$. This means that for all i , $\delta^{(i)} + \Delta^{(i)} = \mathbf{a}^{(i+1)} - \mathbf{b}^{(i)} + \mathbf{b}^{(i)} - \pi_i(\mathbf{a}^{(i)}) = \mathbf{a}^{(i+1)} - \pi_i(\mathbf{a}^{(i)})$.

Thus, the final Δ produced by P_0 is

$$\begin{aligned}
& \Delta^{(d)} + \pi_d(\delta^{(d-1)} + \Delta^{(d-1)} + \pi_{d-1}(\delta^{(d-2)} + \Delta^{(d-2)} + \dots + \pi_2(\delta^{(1)} + \Delta^{(1)}))) \\
&= \Delta^{(d)} + \pi_d(\mathbf{a}^{(d)} - \pi_{d-1}(\mathbf{a}^{(d-1)})) + \pi_{d-1}(\mathbf{a}^{(d-1)} - \pi_{d-2}(\mathbf{a}^{(d-2)})) + \dots + \pi_2(\mathbf{a}^{(2)} - \pi_1(\mathbf{a}^{(1)})) \\
&= \Delta^{(d)} + \pi_d(\mathbf{a}^{(d)} - \pi_{d-1}(\dots \pi_2(\pi_1(\mathbf{a}^{(1)})))) \\
&= \mathbf{b}^{(d)} - \pi_d(\mathbf{a}^{(d)}) + \pi_d(\mathbf{a}^{(d)} - \pi_{d-1}(\dots \pi_2(\pi_1(\mathbf{a}^{(1)})))) \\
&= \mathbf{b}^{(d)} - \pi_d(\pi_{d-1}(\dots \pi_2(\pi_1(\mathbf{a}^{(1)})))) \\
&= \mathbf{b}^{(d)} - \pi(\mathbf{a}^{(1)})
\end{aligned}$$

The output for P_0, P_1 is:

$$\begin{aligned}
& \pi(\mathbf{m}) + \Delta - \mathbf{w}, & \mathbf{w} - \mathbf{b}^{(d)} \\
& = \pi(\mathbf{x} + \mathbf{a}^{(1)}) + \Delta - \mathbf{w}, & \mathbf{w} - (\Delta + \pi(\mathbf{a}^{(1)})) \\
& = \pi(\mathbf{x}) + \pi(\mathbf{a}^{(1)}) + \Delta - \mathbf{w}, & -\Delta - \pi(\mathbf{a}^{(1)}) + \mathbf{w}
\end{aligned}$$

If we let $\mathbf{r} = \pi(\mathbf{x}) + \pi(\mathbf{a}^{(1)}) + \Delta - \mathbf{w}$, we see that this has the correct distribution.

Security. Our simulator behaves as follows: If $b = 0$ (i.e. P_0 is corrupt): $\text{sim}(1^\lambda, 0, \pi, \mathbf{y}_0)$ will first generate the subpermutations for π as described above, and then internally run all of the ShareTrans_T protocols to obtain simulated view for P_0 and $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(d)}, \mathbf{b}^{(1)}, \dots, \mathbf{b}^{(d)}$. Let $\Delta^{(1)}, \dots, \Delta^{(d)}$ be the corresponding values computed by P_0 in these protocols. Choose random $\delta^{(1)}, \dots, \delta^{(d-1)}$. It then computes Δ as in step 4 of the protocol and sets $\mathbf{w} = -\mathbf{y}_0 + \pi(\mathbf{m}) + \Delta$. It outputs the views from the ShareTrans_T protocols and the messages $\mathbf{m}, \mathbf{w}, \delta^{(1)}, \dots, \delta^{(d)}$.

If $b = 1$ (i.e. P_1 is corrupt): $\text{sim}(1^\lambda, 1, \mathbf{x}, \mathbf{y}_1)$ will pick random π' , compute the subpermutations, internally run the ShareTrans_T protocols with these permutations to obtain the views for P_1 , and compute $\mathbf{b}^{(d)}$ from these runs as in the real protocol. It will set the random tape $\mathbf{w} = \mathbf{y}_1 + \mathbf{b}^{(d)}$. It outputs the view from the ShareTrans_T protocols and the random tape \mathbf{w} .

⁹ We give a concrete cost analysis on Section 7.

We show that this simulator produces an ideal experiment that is indistinguishable from the real experiment. We start with the case where $b = 0$ and show this through a series of games:

Real Game : Runs the real experiment. The output is P_0 's view (its input, the view_0 s from the `Share Translation` protocols and the messages \mathbf{m} , \mathbf{w} , and $\delta^{(1)}, \dots, \delta^{(d-1)}$ it receives), and the honest P_1 's input \mathbf{x} and output $\mathbf{w} - \mathbf{b}$.

Game 1: As in the previous game except in step 2, compute $\Delta^{(i)}$ as $\mathbf{b}^{(i)} - \pi_i(\mathbf{a}^{(i)})$ instead of through the `ShareTransT` protocols. This is identical by correctness of `Share Translation` .

Game 2: As in the previous game except after step 2 for each i we sample random $\mathbf{a}'^{(i)}$ and compute $\mathbf{b}'^{(i)} = \pi_i(\mathbf{a}'^{(i)}) + \Delta^{(i)}$, and then use these values in place of $\mathbf{a}^{(i)}, \mathbf{b}^{(i)}$ in steps 3 and 4.

We can show that this is indistinguishable via a series of hybrids, where in hybrid H_i , we use $\mathbf{a}'^{(j)}, \mathbf{b}'^{(j)}$ for the output of the first i `ShareTransT` protocols and $\mathbf{a}^{(j)}, \mathbf{b}^{(j)}$ for the rest. Then H_i, H_{i+1} are indistinguishable by the share hiding property of `ShareTransT`.

Game 3: As above, but choose random $\mathbf{m}, \delta^{(1)}, \dots, \delta^{(d-1)}$. Set $\mathbf{a}'^{(1)} = \mathbf{m} - \mathbf{x}$. For $i = 1 \dots d$, compute $\mathbf{b}'^{(i)} = \pi_i(\mathbf{a}'^{(i)}) + \Delta^{(i)}$ as above, and then set $\mathbf{a}'^{(i+1)} = \delta^{(i)} - \mathbf{b}^{(i)}$. Note that this is distributed identically to Game 2.

Game Simulated: The only difference between the simulated game and Game 3 is that in Game 3, \mathbf{w} is chosen at random, and P_1 's output is computed as $\mathbf{w} - \mathbf{b}'^{(d)}$, while in Game Simulated, P_1 's output is random \mathbf{r} and \mathbf{w} is set to $-\mathbf{y}_0 + \pi(\mathbf{m}) + \Delta = -(\pi(\mathbf{x}) - \mathbf{r}) + \pi(\mathbf{m}) + \Delta = \pi(\mathbf{a}'^{(1)}) + \mathbf{r} + \Delta = \mathbf{b}'^{(d)} + \mathbf{r}$ by construction of Δ . Thus, the two games are identical.

[nolistsep,noitemsep]

We argue the case when $b = 1$ as follows:

Real Game : Runs the real experiment. The output is P_1 's view (it's input \mathbf{x} , view_1 from the `Share Translation` protocol and the random string \mathbf{w} it chooses) and the honest P_0 's input π and output $\pi(\mathbf{m}) + \Delta - \mathbf{w}$ where Δ is as computed in step 4 of the protocol.

Game 1: As in the previous game, but P_0 's output is $\pi(\mathbf{x}) + \mathbf{b}^{(d)} - \mathbf{w}$. Note that $\pi(\mathbf{x}) + \mathbf{b}^{(d)} - \mathbf{w} = \pi(\mathbf{x} + \mathbf{a}^{(1)}) + \mathbf{b}^{(d)} - \pi(\mathbf{a}^{(1)}) - \mathbf{w} = \pi(\mathbf{m}) + \Delta - \mathbf{w}$ where $\mathbf{a}^{(1)}, \mathbf{b}^{(d)}$ are the values P_1 obtains from the first and last layer `ShareTranst` protocols.

Game 2: As in the previous game except run the `ShareTransT` protocols with π'_1, \dots, π'_d derived from a random permutation π' .

We can show that this is indistinguishable via a series of hybrids, where in hybrid H_i , we use the subpermutations derived from π' for the first i protocols, and the subpermutations derived from π for the rest. Then H_i, H_{i+1} are indistinguishable by the permutation hiding property of `ShareTransT`.

Game Simulated: As in the previous game except choose random \mathbf{r} and set $\mathbf{w} = \pi(\mathbf{x}) - \mathbf{r} + \mathbf{b}^{(d)}$. This is identically distributed to Game 1 and identical to the ideal experiment.

6.3 Secret Shared Shuffle from Permute+Share

The Secret Shared Shuffle protocol proceeds as follows:

0. P_0 and P_1 each choose a random permutation $\pi_0, \pi_1 \leftarrow S_N$.
1. P_0 and P_1 run the Permute+Share protocol to apply π_0 to \mathbf{x}_1 , resulting in shares $\mathbf{x}_0^{(1)}$ for P_0 and $\mathbf{x}_1^{(1)}$ for P_1 .
2. P_0 computes $\mathbf{x}_0^{(2)} = \pi_0(\mathbf{x}_0) + \mathbf{x}_0^{(1)}$.
3. P_1 and P_0 run the Permute + Share protocol to apply π_1 to $\mathbf{x}_0^{(2)}$, resulting in shares $\mathbf{x}_1^{(3)}$ for P_1 and $\mathbf{x}_0^{(3)}$ for P_0 .
4. P_1 computes $\mathbf{x}_1^{(4)} = \pi_1(\mathbf{x}_1^{(1)}) + \mathbf{x}_1^{(3)}$.
5. P_0 outputs $\mathbf{x}_0^{(3)}$ and P_1 outputs $\mathbf{x}_1^{(4)}$.

Theorem 5. *The construction above is a Secret Shared Shuffle protocol secure against static semi-honest corruptions. Its communication and computation cost is that of invokes 2 sequential Permute+Share's.*¹⁰

Correctness. The output for P_0, P_1 is:

$$\begin{aligned}
& \mathbf{x}_0^{(3)}, & \mathbf{x}_1^{(4)} \\
= & \mathbf{x}_0^{(3)}, & \pi_1(\mathbf{x}_1^{(1)}) + \mathbf{x}_1^{(3)} \\
= & \pi_1(\mathbf{x}_0^{(2)}) - \mathbf{r}^{(3)}, & \pi_1(\mathbf{x}_1^{(1)}) + \mathbf{r}^{(3)} \\
= & \pi_1(\pi_0(\mathbf{x}_0) + \mathbf{x}_0^{(1)}) - \mathbf{r}^{(3)}, & \pi_1(\mathbf{x}_1^{(1)}) + \mathbf{r}^{(3)} \\
= & \pi_1(\pi_0(\mathbf{x}_0) + \mathbf{r}^{(1)}) - \mathbf{r}^{(3)}, & \pi_1(\pi_0(\mathbf{x}_1) - \mathbf{r}^{(1)}) + \mathbf{r}^{(3)} \\
= & \pi_1(\pi_0(\mathbf{x}_0)) + \pi_1(\mathbf{r}^{(1)}) - \mathbf{r}^{(3)}, & \pi_1(\pi_0(\mathbf{x}_1)) - (\pi_1(\mathbf{r}^{(1)}) - \mathbf{r}^{(3)})
\end{aligned}$$

Where $\mathbf{r}^{(1)}$ and $\mathbf{r}^{(3)}$ are the values generated by the first and second invocations of Permute+Share. If we let $\mathbf{r} = \pi_1(\pi_0(\mathbf{x}_0)) + \pi_1(\mathbf{r}^{(1)}) - \mathbf{r}^{(3)}$ and $\pi = \pi_1 \circ \pi_0$ we see that this has the correct distribution.

Security. Our simulator behaves as follows:

If $b = 0$ (i.e. P_0 is corrupt): $\text{sim}(1^\lambda, 0, \mathbf{x}_0, \mathbf{y}_0)$ will choose random $\pi_0, \mathbf{x}_0^{(1)}$, set $\mathbf{x}_0^{(2)} = \pi_0(\mathbf{x}_0) + \mathbf{x}_0^{(1)}$, simulate the view from the first Permute+Share with $\text{sim}^{\text{Permute+Share}}(1^\lambda, 0, \pi_0, \mathbf{x}_0^{(1)})$, and simulate the view from the second Permute+Share with $\text{sim}^{\text{Permute+Share}}(1^\lambda, 1, \mathbf{x}_0^{(2)}, \mathbf{y}_0)$.

If $b = 1$ (i.e. P_1 is corrupt): $\text{sim}(1^\lambda, 1, \mathbf{x}_1, \mathbf{y}_1)$ will choose random $\pi_1, \mathbf{x}_1^{(1)}$, set $\mathbf{x}_1^{(3)} = \mathbf{y}_1 - \pi_1(\mathbf{x}_1^{(1)})$, simulate the view from the first Permute+Share with $\text{sim}^{\text{Permute+Share}}(1^\lambda, 1, \mathbf{x}_1, \mathbf{x}_1^{(1)})$, and simulate the view from the second Permute+Share with $\text{sim}^{\text{Permute+Share}}(1^\lambda, 0, \pi_1, \mathbf{x}_1^{(3)})$.

We show that this simulator produces an ideal experiment that is indistinguishable from the real experiment. We start with the case where $b = 0$ and show this through a series of games:

¹⁰ We give a concrete cost analysis on Section 7

Real Game : Runs the real experiment.

The output is P_0 's view (its input \mathbf{x}_0 , $\text{view}_0^{(1)}$, $\text{view}_0^{(2)}$ from the two Permute+Share protocols including the outputs $\mathbf{x}_0^{(1)}$, $\mathbf{x}_0^{(3)}$, and the honest P_1 's input \mathbf{x}_1 and output $\mathbf{x}_1^{(4)} = \pi_1(\mathbf{x}_1^{(1)}) + \mathbf{x}_1^{(3)}$

Game 1 : In step 1, first compute $\mathcal{F}_{\text{Permute+Share}}(\pi_0, \mathbf{x}_1)$, i.e. choose random $\mathbf{r}^{(1)}$, and set $\mathbf{x}_0^{(1)} = \mathbf{r}^{(1)}$ and $\mathbf{x}_1^{(1)} = \pi_0(\mathbf{x}_1) - \mathbf{r}^{(1)}$. Then run the Permute+Share simulator to generate the view $\text{view}_0^{(1)'}$ for the first Permute+Share.

The output is P_0 's view (its input \mathbf{x}_0 , $\text{view}_0^{(1)'}$, $\text{view}_0^{(2)}$ from the two Permute+Share protocols including its outputs from those protocols $\mathbf{x}_0^{(1)} = \mathbf{r}^{(1)}$ and $\mathbf{x}_0^{(3)}$, and the honest P_1 's input \mathbf{x}_1 and output $\mathbf{x}_1^{(4)} = \pi_1(\mathbf{x}_1^{(1)}) + \mathbf{x}_1^{(3)} = \pi_1(\pi_0(\mathbf{x}_1) - \mathbf{r}^{(1)}) + \mathbf{x}_1^{(3)}$.

This is indistinguishable by security of the Permute+Share protocol.

Game 2 : In step 3, first compute $\mathcal{F}_{\text{Permute+Share}}(\pi_1, \mathbf{x}_0^{(2)})$, i.e. choose random $\mathbf{r}^{(3)}$ and set $\mathbf{x}_1^{(3)} = \mathbf{r}^{(3)}$ and $\mathbf{x}_0^{(3)} = \pi_1(\mathbf{x}_0^{(2)}) - \mathbf{r}^{(3)}$. Then run the Permute+Share simulator to generate the view $\text{view}_0^{(2)'}$ for the second Permute+Share.

The output is P_0 's view (its input \mathbf{x}_0 , $\text{view}_0^{(1)'}$, $\text{view}_0^{(2)'}$ from the two Permute+Share protocols including its outputs from those protocols $\mathbf{x}_0^{(1)} = \mathbf{r}^{(1)}$ and $\mathbf{x}_0^{(3)} = \pi_1(\mathbf{x}_0^{(2)}) - \mathbf{r}^{(3)}$, and the honest P_1 's input \mathbf{x}_1 and output $\mathbf{x}_1^{(4)} = \pi_1(\pi_0(\mathbf{x}_1) - \mathbf{r}^{(1)}) + \mathbf{x}_1^{(3)} = \pi_1(\pi_0(\mathbf{x}_1) - \mathbf{r}^{(1)}) + \mathbf{r}^{(3)}$.

This is again indistinguishable by security of the Permute+Share protocol.

Game 3 : Choose random $\pi, \mathbf{r}, \mathbf{x}_0^{(1)}$. Set $\pi_1 = \pi \circ \pi_0^{-1}$, $\mathbf{r}^{(1)} = \mathbf{x}_0^{(1)}$ and $\mathbf{r}^{(3)} = \pi_1(\pi_0(\mathbf{x}_0)) + \pi_1(\mathbf{r}^{(1)}) - \mathbf{r}$. Other than that, proceed as in Game 2.

The output is P_0 's view (its input \mathbf{x}_0 , $\text{view}_0^{(1)'}$, $\text{view}_0^{(2)'}$ from the two Permute+Share protocols including its outputs from those protocols $\mathbf{x}_0^{(1)} = \mathbf{r}^{(1)}$ and $\mathbf{x}_0^{(3)}$, and the honest P_1 's input \mathbf{x}_1 and output $\mathbf{x}_1^{(4)}$).

This is identically distributed to Game 2. P_1 's output in this game is

$$\begin{aligned} \mathbf{x}_1^{(4)} &= \pi_1(\mathbf{x}_1^{(1)}) + \mathbf{x}_1^{(3)} \\ &= \pi_1(\mathbf{x}_1^{(1)}) + \pi_1(\mathbf{x}_0^{(2)}) - \mathbf{x}_0^{(3)} \\ &= \pi_1(\mathbf{x}_1^{(1)}) + \pi_1(\pi_0(\mathbf{x}_0) + \mathbf{x}_0^{(1)}) - \mathbf{x}_0^{(3)} \\ &= \pi_1(\pi_0(\mathbf{x}_1) - \mathbf{x}_0^{(1)}) + \pi_1(\pi_0(\mathbf{x}_0) + \mathbf{x}_0^{(1)}) - \mathbf{x}_0^{(3)} \\ &= \pi_1(\pi_0(\mathbf{x}_1 + \mathbf{x}_0)) - \mathbf{x}_0^{(3)} \\ &= \pi(\mathbf{x}_1 + \mathbf{x}_0) - \mathbf{x}_0^{(3)} \end{aligned}$$

Thus, this is identical to the ideal experiment.

Next, we turn to the case where $b = 1$.

Real Game : Runs the real experiment

- Game 1** : In step 1, first compute $\mathcal{F}_{\text{Permute+Share}}(\pi_0, \mathbf{x}_1)$, i.e. choose random $\mathbf{x}_0^{(1)}$, and then compute $\mathbf{x}_1^{(1)} = \pi_0(\mathbf{x}_1) - \mathbf{x}_0^{(1)}$. Then run the Permute+Share simulator to generate the view for the first Permute+Share. *This is indistinguishable by security of the Permute+Share protocol.*
- Game 2** : In step 3, first compute $\mathcal{F}_{\text{Permute+Share}}(\pi_1, \mathbf{x}_0^{(2)})$, i.e. choose random $\mathbf{x}_1^{(3)}$, and then compute $\mathbf{x}_0^{(3)} = \pi_1(\mathbf{x}_0^{(2)}) - \mathbf{x}_1^{(3)}$. Then run the Permute+Share simulator to generate the view for the second Permute+Share. *This is again indistinguishable by security of the Permute+Share protocol.*
- Game 3** : Choose random $\mathbf{x}_0^{(3)}$. Set $\mathbf{x}_1^{(3)} = \pi_1(\mathbf{x}_0^{(2)}) - \mathbf{x}_0^{(3)}$. Other than that, proceed as in Game 2. *This is identically distributed to Game 2.*
- Game 4**: Choose random π , set $\pi_0 = \pi_1^{-1} \circ \pi$ and set $\mathbf{x}_1^{(3)} = \pi(\mathbf{x}_0 + \mathbf{x}_1) - \pi_1(\mathbf{x}_1^{(1)}) - \mathbf{x}_0^{(3)}$. *Note that this means $\mathbf{x}_1^{(4)} = \pi(\mathbf{x}_0 + \mathbf{x}_1) - \mathbf{x}_0^{(3)}$ so this is distributed identically to the ideal experiment. Note also that this is distributed identically to Game 3, because:*

$$\begin{aligned}
& \pi_1(\mathbf{x}_0^{(2)}) - \mathbf{x}_0^{(3)} \\
&= \pi_1(\pi_0(\mathbf{x}_0) + \mathbf{x}_0^{(1)}) - \mathbf{x}_0^{(3)} \\
&= \pi_1(\pi_0(\mathbf{x}_0) + \pi_0(\mathbf{x}_1) - \mathbf{x}_1^{(1)}) - \mathbf{x}_0^{(3)} \\
&= \pi_1(\pi_0(\mathbf{x}_0 + \mathbf{x}_1)) - \pi_1(\mathbf{x}_1^{(1)}) - \mathbf{x}_0^{(3)} \\
&= \pi(\mathbf{x}_0 + \mathbf{x}_1) - \pi_1(\mathbf{x}_1^{(1)}) - \mathbf{x}_0^{(3)}
\end{aligned}$$

7 Experimental Evaluation

In this section, we compare the solution for our Permute + Share with public key based solution and with the best previous permutation network based solution [23]. We consider Permute + Share where party P_0 starts with a permutation π and party P_1 starts with a input vector \mathbf{x} of N strings in $\{0, 1\}^\ell$.

We take a microbenchmarking approach to estimating the cost of the two protocols, where we first empirically estimate the cost of the individual operations (AES computations or RSA group operations), and then use that number to estimate the cost of the full protocol, by plugging the time of individual operations into the formula for execution time of the protocol. For example, for our protocol and the protocol of [23], both of which are dominated by AES operations, we estimate the cost as follows:

1. We compute the computation cost in terms of number of AES calls.
2. We empirically estimate the cost for a computing fixed key AES (per 128-bit block).
3. We compute the communication cost in bits.
4. Then we compute the time to communicate the calculated number of bits using various networks (bandwidth 72Mbps, 100Mbps and 1Gbps).
5. The total time reported is number of AES calls \times the cost of a single AES + the size of communication / bandwidth.

In the following we will describe our cost estimates in more detail and then present a detailed comparison. First we discuss some specifics on how we implement AES and how we analyze the cost of OT, then we present formulas for the number of basic operations required for each solution, then we describe how we estimate the cost of these operations, and finally we present the detailed efficiency comparison.

7.1 More detail on cost of OT and AES

Fixed key Block Ciphers The symmetric key based protocols (ours and the one described in [23]) rely on two fundamental building blocks, namely, Oblivious Transfer extension (OTe) [17] and GGM PRG [11]. Typically, published OTe protocols are based on a hash function that is modeled as a random oracle. However, in most of the recent implementations, the hash function is instantiated, somewhat haphazardly, using fixed key block ciphers (AES). In a recent work [14], the authors provided a principled way of implementing [17] using fixed key AES and formally proved that it is secure. The authors also propose that the length doubling PRG used in GGM [11] can be implemented using fixed key AES for better efficiency, though they do not prove it. Here, we first prove that it is safe to use this optimized PRG construction [14], and then use it in our experiments. In our experiments, we will also use the fixed-key AES-based length extension technique for stretching short messages into longer ones (both for OTe and for OPV message length extension) described in Section 6.1 in [14].

The optimized PRG construction is based on correlation-robust hash (CRH) function [17, 14]. Roughly, H is said to be correlation-robust if the keyed function $f_R(x) = H(x \oplus R)$ is pseudorandom, as long as R is sufficiently random. Given a CRH H , the length doubling PRG is constructed as follows: $G(x) = H(1 \oplus x) \circ H(2 \oplus x)$. We give more details in the Sec B).

In our experiments, we will use the following concrete instantiation of CRH [14]: $H(x) = \pi(x) \oplus x$ where $\pi(\cdot)$ is a fixed key block cipher, e.g. AES.

OT extension costs The computation in OT extension consists of $O(m\ell)$ bitwise operations (ANDs and XORs), running λ public key OTs, and $O(m + m\ell/\lambda)$ AES operations as discussed above. This means that for sufficiently large m , like those we consider, the cost is dominated by the AES operations, as can be verified empirically using any standard OT library. For example, we benchmark Naor-Pinkas base OT (dubbed NPOT) using [29] and the average time to run 128 base OTs is 13ms. As a result, we can focus our analysis of computational costs on the AES operations.

In our experiments, we simulate the cost of OT-extension as follows. The cost is reported in number of fixed-key AES calls for sender and receiver and communication is reported in number of bits. For random OT's on strings of length $\ell > \lambda = 128$ bits, we use IKNP OT-extension protocol with fixed-key AES optimization [14]. The cost of m Random OTs on messages of length ℓ bits is shown in Table 7.1, where the terms $2m\ell/\lambda$ for sender and $m\ell/\lambda$ for receiver are for extending the random messages from λ to ℓ bits. We denote this functionality as ROT_ℓ^m . For $\ell = \lambda$, no message length extension is required (both for ROT and

SOT). Fixed message OT's or standard OTs (SOT) are obtained from ROT by using the ROT messages as one-time pads for the actual messages. So SOT_ℓ^m adds an additional $2m\ell$ bits of communication over ROT_ℓ^m , i.e., the communication cost of SOT_ℓ^m is $m(\lambda + 2\ell)$ bits. There is no additional computation overhead (except some additional XORs, which we ignore).

OT	Sender	Receiver	Communication (bits)
ROT_λ^m	$3m$	$3m$	$m\lambda$
ROT_ℓ^m	$3m + 2m\ell/\lambda$	$3m + m\ell/\lambda$	$m\lambda$
SOT_ℓ^m	$3m + 2m\ell/\lambda$	$3m + m\ell/\lambda$	$m(\lambda + 2\ell)$
SOT_λ^m	$3m$	$3m$	$3m\lambda$

Table 1. The computation and communication cost for variants of OT extension.

7.2 Analyzing the cost of each solution

As discussed above, in the following estimates we only count computation time of AES, not base OTs or XORs, since the latter are fairly small.

Let N be the number of elements in the database, ℓ be the length of each element, λ be the security parameter, and T be the size of subpermutations. Let $d = 2\lceil \log N / \log T \rceil - 1$.

Our protocol: The compute cost of our **Permute + Share** protocol is the compute cost of dN/T **ShareTrans_T**'s, where $d = 2\lceil \log N / \log T \rceil - 1$. The communication includes the cost of dN/T **ShareTrans_T**'s + $(d + 1)N\ell$ bits.

Each **ShareTrans_T** protocol requires $\text{SOT}_\lambda^{T \log T}$ and $T^2(2 + \ell/\lambda)$ local fixed key AES calls (for both parties) which includes PRG calls in the GGM tree and message length extension and for the underlying **OPV** protocol. There is no additional communication over the cost of $\text{SOT}_\lambda^{T \log T}$.

Computation Cost The number of AES calls (for each of sender and receiver) is the following:

$$3dN \log T + dNT(2 + \ell/\lambda)$$

Communication Cost Communication in number of bits is the following:

$$3dN\lambda \log T + (d + 1)N\ell$$

Protocol from [23]: This **Permute + Share** requires $\text{SOT}_{2\ell}^{N \log N - N/2}$ and has an additional $2N\ell$ bits communication overhead.

So, the total computation and computation costs are the following:

Computation Cost Number of AES calls for receiver (receiver is slightly more efficient than sender) in the protocol of [23] is the following:

$$3(N \log N - N/2) + 2(\ell/\lambda)(N \log N - N/2)$$

Communication Cost Communication in number of bits in the protocol of [23] is the following:

$$(N \log N - N/2)(\lambda + 4\ell) + 2N\ell$$

Paillier based solution. In appendix C we describe a solution based on additively homomorphic encryption in which P_1 encrypts his data and sends it to P_0 , who permutes the ciphertexts, randomizes them, and adds a random share to each before returning them to P_0 ; P_0 outputs the decryptions and P_1 outputs the random shares he added.

In this protocol, since every element of \mathbf{x} has to be encrypted and the encryption message space is defined to be \mathbb{Z}_n , each element has to be broken into blocks of size n . This means that P_0 computes $N * \lceil \ell/n \rceil$ encryptions and P_1 computes $N * \lceil \ell/n \rceil$ ciphertext randomizations and ciphertext-plaintext multiplications. The communication for this protocol is $N * \lceil \ell/n \rceil * 2n$ bits. To get a very rough estimate of the cost of Paillier encryption and randomization+multiplication, we measure the cost of an RSA signing operation with modulus n . Note that this is a significant underestimate since Paillier operations actually happen in \mathbb{Z}_n^2 , and since the RSA signer knows the factorization of n , while P_1 does not.

7.3 Microbenchmarking

To estimate the per block cost of AES, we use the *permute_block* function in prp of [29] to benchmark the cost of a fixed key AES-ECB 128 per 128-bit block (we use security parameter $\lambda = 128$ for our experiments). To get this cost, we run fixed key AES for different numbers of blocks (4096, 8192, 12288) to get the amortized cost of a single AES. We repeat each experiment 100 times and then report the average amortized cost per 128-bit block (no significant variance was noticeable).

For estimating the cost of a single encryption and a single ciphertext randomization for the Paillier based protocol, we use the RSA signing cost for modulus of size 4096. We get this cost using the OpenSSL benchmark [8] by running the command *openssl speed*.

The costs we get are the following: *AES-ECB 128*: 3.5 ns, *RSA 4096 signing* 0.17s. All the benchmarks are run on a Macbook Pro 2017 with a 3.1 GHz Intel core i-7 processor and 16GB of 2133MHz LPDDR3 RAM.

7.4 Performance Comparison

We estimate the performance of the different constructions described above. For this estimation, we experiment with three different database sizes, $N = 2^{20}, 2^{24}$ and 2^{32} elements and three different network bandwidths, *72Mbps*, *100Mbps* and *1Gbps*. We vary the length of each element in the database from 640 bits to 64000 bits. This range of values is roughly inspired from Machine Learning training applications which has 100s to 1000s of features (with each feature represented by a 64 bit integer).

In the following graphs on fig. 3, 4, 5 we report the estimated running time of our protocol and the protocol from [23]. We do not report the running time of

the PKE based protocol in the graphs since they are 2 - 3 orders of magnitude slower compared to our protocol. Instead we summarize their performance in Table 2.

N	T	Bandwidth	PKE time/Our time
2^{20}	16	1Gbps	3000 – 7000x
2^{20}	128	72bps	400 – 600x
2^{32}	16	1Gbps	1900 – 4000x
2^{32}	256	72Mbps	260 – 400x

Table 2. Comparative performance of our protocol vs PKE based protocol

Bandwidth	[23] time/Our time
1Gbps	3 – 5x
100bps	4 – 11x
72Mbps	5 – 12x

Table 3. Comparative performance of our protocol vs [23] based protocol for $N = T = 128$

In addition, we summarize how we compare with the protocol from [23] for relatively small N but long elements (640 – 64000 bits) in Table 3. We get a performance gain of 3 – 12x depending on the speed of the network.




Fig. 3. Total running time of Permute+Share for 72Mbps network

7.5 Choosing optimal subpermutation size T .

We choose the best value of T empirically: by fixing desired N, ℓ , enumerating over all possible T from 2 to N , and using the following formula to find the running time for each value of T (as before, d is set to be $2 \lceil \log N / \log T \rceil - 1$), which is: $(3dN \log T + dNT(2 + \ell/\lambda)) \cdot \text{TimePerAES} + (3dN\lambda \log T + (d + 1)N\ell) \cdot \text{TimePerBitSent}$. We give more details in section D.




Fig. 4. Total running time of Permute+Share for 100Mbps network

Acknowledgements

We are very grateful to the anonymous reviewers of Asiacrypt 2020, who have found inconsistencies in our experimental data, which was caused by a bug in our script. This version contains an updated estimates.

References

1. Ajtai, M., Komlós, J., Szemerédi, E.: An $o(n \log n)$ sorting network. In: Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA. pp. 1–9 (1983)
2. Benes, V.E.: Optimal rearrangeable multistage connecting networks. *The Bell System Technical Journal* **43**(4), 1641–1656 (1964)
3. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Rindal, P., Scholl, P.: Efficient two-round ot extension and silent non-interactive secure computation. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security. p. 291–308. CCS ’19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3319535.3354255>, <https://doi.org/10.1145/3319535.3354255>
4. Chida, K., Hamada, K., Ikarashi, D., Kikuchi, R., Kiribuchi, N., Pinkas, B.: An efficient secure three-party sorting protocol with an honest majority. *IACR Cryptol. ePrint Arch.* **2019**, 695 (2019), <https://eprint.iacr.org/2019/695>
5. Ciampi, M., Orlandi, C.: Combining private set-intersection with secure two-party computation. In: Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September 5-7, 2018, Proceedings. pp. 464–482 (2018)
6. Doerner, J., Shelat, A.: Scaling ORAM for secure computation. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017. pp. 523–535 (2017)
7. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. In: Blakley, G.R., Chaum, D. (eds.) *Advances in Cryptology*. pp. 10–18. Springer Berlin Heidelberg, Berlin, Heidelberg (1985)




Fig. 5. Total running time of Permute+Share for 1Gbps network

8. Foundation, O.S.: OpenSSL. <https://www.openssl.org/>
9. Garg, S., Gupta, D., Miao, P., Pandey, O.: Secure multiparty RAM computation in constant rounds. In: Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I. pp. 491–520 (2016)
10. Gentry, C., Halevi, S., Lu, S., Ostrovsky, R., Raykova, M., Wichs, D.: Garbled RAM revisited. In: Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings. pp. 405–422 (2014), https://doi.org/10.1007/978-3-642-55220-5_23
11. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. *J. ACM* **33**(4), 792–807 (Aug 1986). <https://doi.org/10.1145/6490.6503>, <http://doi.acm.org/10.1145/6490.6503>
12. Gordon, S.D., Katz, J., Kolesnikov, V., Krell, F., Malkin, T., Raykova, M., Vahlis, Y.: Secure two-party computation in sublinear (amortized) time. In: the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012. pp. 513–524 (2012)
13. Group, B.C.: Bristol Cryptography Blog. <http://bristolcrypto.blogspot.com/2017/01/rwc-2017-secure-mpc-at-google.html>
14. Guo, C., Katz, J., Wang, X., Yu, Y.: Efficient and secure multiparty computation from fixed-key block ciphers. Cryptology ePrint Archive, Report 2019/074 (2019), <https://eprint.iacr.org/2019/074>
15. Hamada, K., Ikarashi, D., Chida, K., Takahashi, K.: Oblivious radix sort: An efficient sorting algorithm for practical secure multi-party computation. *IACR Cryptology ePrint Archive* **2014**, 121 (2014), <http://eprint.iacr.org/2014/121>
16. Huang, Y., Evans, D., Katz, J.: Private set intersection: Are garbled circuits better than custom protocols? In: NDSS (2012)
17. Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending oblivious transfers efficiently. In: Boneh, D. (ed.) *Advances in Cryptology - CRYPTO 2003*. pp. 145–161. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)

18. Keller, M., Scholl, P.: Efficient, oblivious data structures for MPC. In: *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security*, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II. pp. 506–525 (2014)
19. Liu, C., Huang, Y., Shi, E., Katz, J., Hicks, M.W.: Automating efficient ram-model secure computation. In: *2014 IEEE Symposium on Security and Privacy*, SP 2014, Berkeley, CA, USA, May 18-21, 2014. pp. 623–638 (2014)
20. Liu, C., Wang, X.S., Nayak, K., Huang, Y., Shi, E.: Oblivm: A programming framework for secure computation. In: *2015 IEEE Symposium on Security and Privacy*, SP 2015, San Jose, CA, USA, May 17-21, 2015. pp. 359–376 (2015)
21. Mohassel, P., Zhang, Y.: Secureml: A system for scalable privacy-preserving machine learning. In: *2017 IEEE Symposium on Security and Privacy (SP)*. pp. 19–38 (May 2017). <https://doi.org/10.1109/SP.2017.12>
22. Mohassel, P., Rindal, P., Rosulek, M.: Fast database joins for secret shared data. *Cryptology ePrint Archive*, Report 2019/518 (2019), <https://eprint.iacr.org/2019/518>
23. Mohassel, P., Sadeghian, S.: How to hide circuits in mpc an efficient framework for private function evaluation. In: Johansson, T., Nguyen, P.Q. (eds.) *Advances in Cryptology – EUROCRYPT 2013*. pp. 557–574. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
24. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) *Advances in Cryptology — EUROCRYPT '99*. pp. 223–238. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)
25. Pinkas, B., Schneider, T., Tkachenko, O., Yanai, A.: Efficient circuit-based psi with linear communication. *Cryptology ePrint Archive*, Report 2019/241 (2019), <https://eprint.iacr.org/2019/241>
26. Pinkas, B., Schneider, T., Weinert, C., Wieder, U.: Efficient circuit-based PSI via cuckoo hashing. In: *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III. pp. 125–157 (2018). https://doi.org/10.1007/978-3-319-78372-7_5, https://doi.org/10.1007/978-3-319-78372-7_5
27. Schoppmann, P., Gascón, A., Reichert, L., Raykova, M.: Distributed vector-ole: Improved constructions and implementation. In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. p. 1055–1072. CCS '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3319535.3363228>, <https://doi.org/10.1145/3319535.3363228>
28. Wang, X., Gordon, S.D., McIntosh, A., Katz, J.: Secure computation of MIPS machine code. In: *Computer Security - ESORICS 2016 - 21st European Symposium on Research in Computer Security*, Heraklion, Greece, September 26-30, 2016, Proceedings, Part II. pp. 99–117 (2016)
29. Wang, X., Malozemoff, A.J., Katz, J.: EMP-toolkit: Efficient MultiParty computation toolkit. <https://github.com/emp-toolkit> (2016)
30. Zahur, S., Wang, X., Raykova, M., Gascón, A., Doerner, J., Evans, D., Katz, J.: Revisiting square-root ORAM: efficient random access in multi-party computation. In: *IEEE Symposium on Security and Privacy*, SP 2016, San Jose, CA, USA, May 22-26, 2016. pp. 218–234 (2016)

A Oblivious Punctured Vector (OPV)

In this section we present communication-efficient protocol for OPV. A very similar protocol appears in concurrent and independent work of [27, 3].

A.1 Construction

To implement a OPV protocol for a domain \mathbb{D} , we first define two algorithms as follows.

Setup($1^\lambda, n, i$) \rightarrow ($\mathbf{s}_0, \mathbf{s}_1$): Setup is a PPT algorithm, that, given a security parameter λ , a vector length n and an index $i \in [n]$, outputs a pair of seeds ($\mathbf{s}_0, \mathbf{s}_1$), where $\mathbf{s}_0, \mathbf{s}_1 \in \{0, 1\}^{\text{poly}(\lambda)}$ and \mathbf{s}_1 includes i .

Expand(b, \mathbf{s}_b) \rightarrow (\mathbf{v}_b): Expand is a polynomial time algorithm that, given a party index $b \in \{0, 1\}$ and a seed \mathbf{s}_b , outputs a vector \mathbf{v}_b of length n , $\mathbf{v}_b \in [\mathbb{D}]^n$.

We implement the algorithms as follows. First we give a 2 party protocol **OblivSetup** that realizes the functionality $\mathcal{F}_{[\mathbb{D}]((\lambda, n), (\lambda, n, i))} = \text{Setup}(1^\lambda, n, i)$. We fix our domain \mathbb{D} to strings of length λ , i.e., $\{0, 1\}^\lambda$. Then we give the construction for **Expand** which P_0 and P_1 run non-interactively.

Given an OPV for \mathbb{D} of strings of length λ , we can build an OPV for domain \mathbb{D}' , where \mathbb{D}' is strings of length $l \geq \lambda$, in a blackbox way. We give this construction in Section 3.2.

G	Length doubling PRG
$i = \sigma_1 \sigma_2 \dots \sigma_{\log n}$	binary representation of input index i
$l = k_1 k_2 \dots k_j$	j -bit binary representation of l
$x^{j,l}$	l^{th} node from the left at level j in the tree, where $l \in [0, 2^j - 1], j \in [1, \log n]$

Table 4. Notations

Setup:

- Pick $\text{seed}_\epsilon \leftarrow \{0, 1\}^m$. Let $\text{seed}_0 \circ \text{seed}_1 \leftarrow G(\text{seed}_\epsilon)$.
- For $l = 1, \dots, \log n - 1$: $\text{seed}_{\sigma_1 \dots \sigma_l 0} \circ \text{seed}_{\sigma_1 \dots \sigma_l 1} \leftarrow G(\text{seed}_{\sigma_1 \dots \sigma_l})$.
- Set $\mathbf{s}_0 := (n, \text{seed}_\epsilon), \mathbf{s}_1 := (n, i, \text{seed}_{\overline{\sigma_1}}, \text{seed}_{\sigma_1 \overline{\sigma_2}}, \dots, \text{seed}_{\sigma_1 \sigma_2 \dots \overline{\sigma_{\log n}}})$ and output $(\mathbf{s}_0, \mathbf{s}_1)$.

OblivSetup: Let us assume both parties hold an implicit full binary tree and the levels of the tree are numbered as follows: root is at level 0 and leaves are at level $\log n$. The protocol proceeds as follows:

1. Party P_0 picks $\text{seed}_\epsilon \leftarrow \{0, 1\}^\lambda$.
2. For $j = 1, \dots, \log n$: do the following:

$$\{x^{j,2l} \circ x^{j,2l+1}\}_{l \in [0, 2^{j-1}-1]} \leftarrow \{G(x^{j-1,l})\}_{l \in [0, 2^{j-1}-1]}$$

$$\text{str}^{j,0} \leftarrow \bigoplus_{l \in [0, 2^{j-1}-1]} x^{j,2l}, \quad \text{str}^{j,1} \leftarrow \bigoplus_{l \in [1, 2^{j-1}-1]} x^{j,2l+1}$$

Note that $x^{j,l} = \text{seed}_{k_1 k_2 \dots k_j}$.

3. For $j = 1, \dots, \log n$: P_0 and P_1 run OT : $((\text{str}^{j,0}, \text{str}^{j,1}), \sigma_j) = (\perp, \text{str}^{j, \overline{\sigma_j}})$.
4. At the end of the OT phase P_1 locally expands the strings it received through OT to compute $\text{seed}_{\overline{\sigma_1}}, \text{seed}_{\sigma_1 \overline{\sigma_2}}, \dots, \text{seed}_{\sigma_1 \sigma_2 \dots \overline{\sigma_{\log n}}}$. The expansion works as follow. For $j = 1, \dots, \log n$: P_1 has received, through the OT, $\text{str}^{j, \overline{\sigma_j}}$. Note that $\text{str}^{j, \overline{\sigma_j}}$ contains $\text{seed}_{\sigma_1 \sigma_2 \dots \overline{\sigma_{\log j}}}$ and P_1 can take off the extra terms by expanding the $2^{j-1} - 1$ seeds from the previous levels. More concretely,

$$\bigoplus_{k_1, k_2, \dots, k_{j-1} \in \{0, 1\}, k_j = \overline{\sigma_j} \wedge k_1 k_2 \dots k_{j-1} \neq \sigma_1 \sigma_2 \dots \sigma_{j-1}} \text{seed}_{k_1 k_2 \dots k_j} \leftarrow \text{str}^{j, \overline{\sigma_j}}$$

5. At the end of this step, P_1 outputs $\mathbf{s}_0 := (n, \text{seed}_\epsilon)$ and P_1 outputs $\mathbf{s}_1 := (n, i, \text{seed}_{\overline{\sigma_1}}, \text{seed}_{\sigma_1 \overline{\sigma_2}}, \dots, \text{seed}_{\sigma_1 \sigma_2 \dots \overline{\sigma_{\log n}}})$.

Expand: For party b , construct \mathbf{v}_b as follows.

- $b = 0$: Parse \mathbf{s}_0 as $(n, \text{seed}_\epsilon)$. Compute $\text{seed}_0 \circ \text{seed}_1 \leftarrow G(\text{seed}_\epsilon)$.
 For $j = 1, \dots, \log n$: do the following: $\text{seed}_{k_1 k_2 \dots k_{j-1} k_j} \circ \text{seed}_{k_1 k_2 \dots k_{j-1} \overline{k_j}} \leftarrow G(\text{seed}_{k_1 k_2 \dots k_{j-1}})$ for $k_1, \dots, k_j \in \{0, 1\}$.
 For $t \in [1, n]$, set $\mathbf{v}_0[t] := \text{seed}_{k_1 k_2 \dots k_{\log n}}$ where $t = k_1 k_2 \dots k_{\log n}$, i.e., the binary representation of t .
- $b = 1$: Parse \mathbf{s}_1 as $(n, i = \sigma_1 \dots \sigma_{\log n}, \text{seed}_{\overline{\sigma_1}}, \text{seed}_{\sigma_1 \overline{\sigma_2}}, \dots, \text{seed}_{\sigma_1 \sigma_2 \dots \overline{\sigma_{\log n}}})$.
 For $j = 2 \dots, \log n$, expand each of the seeds as follows:
 $\text{seed}_{k_1 k_2 \dots k_{j-1} k_j} \circ \text{seed}_{k_1 k_2 \dots k_{j-1} \overline{k_j}} \leftarrow G(\text{seed}_{k_1 k_2 \dots k_{j-1}})$ for $k_1, \dots, k_j \in \{0, 1\} \wedge k_1 k_2 \dots k_{j-1} \neq \sigma_1 \sigma_2 \dots \sigma_{j-1}$.
 For $t \in [1, n] \wedge t \neq i$, set $\mathbf{v}_1[t] := \text{seed}_{k_1 k_2 \dots k_{\log n}}$ where $t = k_1 k_2 \dots k_{\log n}$, i.e., the binary representation of t . Set $\mathbf{v}_1[i] = \perp$.

Security Proof Correctness OPV according to Def 3.1 follows from the correctness of OT protocols. Now we will prove that our construction satisfies both position and value hiding. In order to prove that, we first prove some helper theorems.

Theorem 6. *OblivSetup securely realizes the ideal functionality $\mathcal{F}_{\mathbb{D}}((n), (n, i)) = \text{Setup}(1^\lambda, n, i) = (s_0, s_1)$ as per Definition 1.*

Proof. We first construct a simulator that works as follows:

If $b = 0$ (i.e. P_0 is corrupt): $\text{sim}(1^\lambda, 0, n, s_0)$ will parse s_0 as n, seed_ϵ . Then it will run the protocol steps to generate $\text{str}^{j,0}, \text{str}^{j,1}$ for $j = 1, \dots, \log n$ and simulate the view from the OTs with $\text{sim}^{\text{OT}}(1^\lambda, 0, (\text{str}^{j,0}, \text{str}^{j,1}), \perp)$.

If $b = 1$ (i.e. P_1 is corrupt): $\text{sim}(1^\lambda, 1, (n, i), s_1)$ will parse i as $i = \sigma_1 \sigma_2 \dots \sigma_{\log n}$ and s_1 as $(n, i, \text{seed}_{\overline{\sigma_1}}, \text{seed}_{\sigma_1 \overline{\sigma_2}}, \dots, \text{seed}_{\sigma_1 \sigma_2 \dots \overline{\sigma_{\log n}}})$. It will simulate the view from the OTs with $\text{sim}^{\text{OT}}(1^\lambda, 1, \sigma_j, \text{str}^{j, \overline{\sigma_j}})$, where it generates $\text{str}^{j, \overline{\sigma_j}}$ as follows:

$$\text{seed}_{k_1 k_2 \dots k_{j-1} k_j} \circ \text{seed}_{k_1 k_2 \dots k_{j-1} \overline{k_j}} \leftarrow G(\text{seed}_{k_1 k_2 \dots k_{j-1}})$$

for $k_1, \dots, k_j \in \{0, 1\} \wedge k_1 k_2 \dots k_{j-1} \neq \sigma_1 \sigma_2 \dots \sigma_{j-1}$.

$$\text{str}^{j, \overline{\sigma_j}} \leftarrow \text{seed}_{\sigma_1 \sigma_2 \dots \overline{\sigma_j}}$$

$$\bigoplus_{k_1, k_2, \dots, k_{j-1} \in \{0, 1\}, k_j = \overline{\sigma_j} \wedge k_1 k_2 \dots k_{j-1} \neq \sigma_1 \sigma_2 \dots \sigma_{j-1}} \text{seed}_{k_1 k_2 \dots k_j}$$

We show that this simulator produces an ideal experiment that is indistinguishable from the real experiment. We start with the case where $b = 0$ and show this through a series of games:

We define Game k as the following: In Game k , run the OT simulator $\text{sim}^{\text{OT}}(1^\lambda, 0, (\text{str}^{j,0}, \text{str}^{j,1}), \perp)$ for $j = 0, \dots, k$ and for $j = k+1, \dots, \log n$, run the OT protocol. Notice that Game 0 is identical to the real experiment and Game $\log n$ is identical to the ideal experiment. Now, Games k and $k+1$ are computationally indistinguishable by the security of the OT protocol. Therefore for $b = 0$ the simulator produces an ideal experiment that is computationally indistinguishable from the real experiment.

Now we look at the case where $b = 1$ and proceed through a series of games as before. In Game k , run the OT simulator $\text{sim}^{\text{OT}}(1^\lambda, 1, \sigma_j, \text{str}^{j, \overline{\sigma_j}})$ for $j = 0, \dots, k$ and for $j = k+1, \dots, \log n$, run the OT protocol. Notice that Game 0 is identical to the real experiment and Game $\log n$ is identical to the ideal experiment. Games k and $k+1$ are computationally indistinguishable by the security of the OT protocol. Therefore for $b = 1$ the simulator produces an ideal experiment that is computationally indistinguishable from the real experiment.

Theorem 7. *Our scheme satisfies the following property: for any $n \in \mathbb{N}, i, i' \in [n]$, the following distributions are computationally indistinguishable:*

$$\begin{aligned} \mathcal{D}_1 &= \{(s_0, s_1) \leftarrow \text{Setup}(1^\lambda, n, i) : (1^\lambda, n, i, i', s_0)\} \\ \mathcal{D}_2 &= \{(s_0, s_1) \leftarrow \text{Setup}(1^\lambda, n, i') : (1^\lambda, n, i, i', s_0)\} \end{aligned}$$

Proof. Since the seed $s_0 = (n, \text{seed}_\epsilon \leftarrow^{\$} \{0, 1\}^\lambda)$, it does not depend on i . Hence the two distributions are identical.

Theorem 8. *Our construction satisfies the following property: for any $n \in \mathbb{N}, i \in [n]$, the following distributions are computationally indistinguishable:*

$$\begin{aligned} \mathcal{D}_1 &= \{(s_0, s_1) \leftarrow \text{Setup}(1^\lambda, n, i), v_0 \leftarrow \text{Expand}(0, s_0) : (1^\lambda, n, i, v_0, s_1)\} \\ \mathcal{D}_2 &= \{(s_0, s_1) \leftarrow \text{Setup}(1^\lambda, n, i), v_1 \leftarrow \text{Expand}(1, s_1), \\ &v_0[j] := v_1[j] \forall j \neq i, v_0[i] := r \text{ where } r \leftarrow^{\$} \mathbb{D} : (1^\lambda, n, i, v_0, s_1)\} \end{aligned}$$

Proof. We show that the two distributions are computationally indistinguishable through a series of distributions defined as follows:

- H_0 : $\mathcal{D}_1 = \{(s_0, s_1) \leftarrow \text{Setup}(1^\lambda, n, i), v_0 \leftarrow \text{Expand}(0, s_0) : (1^\lambda, n, i, v_0, s_1)\}$
 H_1 : Identical to the previous distribution except the following: In **Setup**, instead of generating seed_ϵ , set $\text{seed}_{\sigma_1}, \text{seed}_{\overline{\sigma_1}} \leftarrow^{\$} \{0, 1\}^\lambda$. Run the rest of the protocol steps to generate all the leaves, set v_0 and s_1 .
 H_k : Identical to the previous distribution except the following: In setup, set $\text{seed}_{\sigma_1 \dots \sigma_k}, \text{seed}_{\sigma_1 \dots \overline{\sigma_k}} \leftarrow^{\$} \{0, 1\}^\lambda$ for $k = 2, \dots, \log n$. Run the rest of the protocol steps to generate all the leaves, set v_0 and s_1 .
 $H'_{\log n}$: Identical to $H_{\log n}$ except the following. Instead of generating v_0 , run $\text{Expand}(1, s_1)$ to generate v_1 , set $v_0[i] \leftarrow^{\$} \{0, 1\}^\lambda$.

By the security of PRG, distributions H_k, H_{k+1} are identical for $k = 1, \dots, \log n$. Finally, distributions $H_{\log n}$ and $H'_{\log n}$ are identical.

Now we define another series of hybrid distributions as follows:

- $G_{\log n}$: This distribution is identical to $H'_{\log n}$ except the following: compute $\text{seed}_{\sigma_1 \dots \sigma_{\log n-1}} \leftarrow^{\$} \{0, 1\}^\lambda$

$$\text{seed}_{\sigma_1 \dots \sigma_{\log n}} \circ \text{seed}_{\sigma_1 \dots \overline{\log n}} \leftarrow G(\text{seed}_{\sigma_1 \dots \sigma_{\log n-1}})$$

. Then replace $\text{seed}_{\sigma_1 \dots \sigma_{\log n}} \leftarrow^{\$} \{0, 1\}^\lambda$.

- G_k : This distribution is identical to the previous, except the following: For $k = \log n - 1, \dots, 1$: Instead of setting $\text{seed}_{\sigma_1 \dots \sigma_k}, \text{seed}_{\sigma_1 \dots \overline{k}} \leftarrow^{\$} \{0, 1\}^\lambda$, compute $\text{seed}_{\sigma_1 \dots \sigma_{k-1}} \leftarrow^{\$} \{0, 1\}^\lambda$

$$\text{seed}_{\sigma_1 \dots \sigma_k} \circ \text{seed}_{\sigma_1 \dots \overline{k}} \leftarrow G(\text{seed}_{\sigma_1 \dots \sigma_{k-1}})$$

- G_0 : This distribution is identical to the previous, except the following: Instead of generating $\text{seed}_{\sigma_1}, \text{seed}_{\overline{\sigma_1}} \leftarrow^{\$} \{0, 1\}^\lambda$, generate $\text{seed}_\epsilon \leftarrow^{\$} \{0, 1\}^\lambda$ and set

$$\text{seed}_{\sigma_1} \circ \text{seed}_{\overline{\sigma_1}} \leftarrow G(\text{seed}_\epsilon)$$

This distribution is identical to \mathcal{D}_2 .

For $k = 0, \dots, \log n$, distributions G_k and G_{k+1} are computationally indistinguishable from the security of PRG. It remains to show that $H'_{\log n}$ and $G_{\log n}$ are computationally indistinguishable as well.

To show this, we show that if there is a PPT distinguisher D that distinguishes $H'_{\log n}$ and $G_{\log n}$ with non-negligible probability, then we can use D to build a PPT distinguisher \mathcal{A} that breaks PRG security with same advantage. \mathcal{A} does the following: on input $w_1 \circ w_2 \in \{0, 1\}^{2\lambda}$, chooses $w'_1 \xleftarrow{\$} \{0, 1\}^\lambda$ and runs D with $w'_1 \circ w_2$. If $w_1 \circ w_2 \xleftarrow{\$} \{0, 1\}^{2\lambda}$, then D exactly simulates game $H'_{\log n}$, otherwise it simulates game $G_{\log n}$. Now if D can distinguish $H'_{\log n}$ and $G_{\log n}$, then \mathcal{A} can distinguish whether $w_1 \circ w_2$ is the output of a PRG or a truly random string immediately with the same advantage as D . Hence, $H'_{\log n}$ and $G_{\log n}$ are computationally indistinguishable.

Now we are ready to prove the main theorem.

Proof of Theorem 1

Proof. Since our protocol satisfies Theorem 6 and Theorem 7, it implies that our construction satisfies position hiding. Since our protocol satisfies Theorem 6 and Theorem 8, it implies that our construction satisfies value hiding. The round complexity, communication and computation costs are evident from the construction.

B Fixed-key blockcipher

In this section we give more details about the definition and security of primitives from section 7.1.

Definition 2. [14] Let $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ be a function and $R \in \{0, 1\}^\lambda$. Define $\mathcal{O}_R(x) = H(x \oplus R)$. Let \mathbb{F}_λ denote the set of all functions from $\{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ and f be randomly picked from \mathbb{F}_λ . For a distinguisher D and for any sufficiently large $\lambda \in \mathbb{N}$, let

$$\text{Adv}_{H, \mathcal{R}}(D) = |\text{Pr}_{R \leftarrow \{0, 1\}^\lambda} [D^{\mathcal{O}_R(\cdot)}(1^\lambda) = 1] - \text{Pr}_{f \leftarrow \mathbb{F}_\lambda} [D^{f(\cdot)}(1^\lambda) = 1]|$$

H is CRH if, for any PPT D making at most q queries to $\mathcal{O}_R(\cdot)$, there exists a negligible function negl such that $\text{Adv}_{H, \mathcal{R}}(D) \leq \text{negl}(\lambda)$ where q is polynomial in λ .

We note that, [14] defined a more general definition where R is picked from a distribution with sufficient min-entropy (at least λ), but this definition suffices for our purpose. Now, we are ready to prove the following theorem.

Theorem 9. *if H is a correlation-robust hash function (CRH,) then $G(x)$ defined as $G(x) = H(1 \oplus x) \circ H(2 \oplus x)$ is a length doubling PRG.*

Proof. For the sake of contradiction, suppose not. Then there exists a PPT distinguisher D that can break the PRG security game with overwhelming advantage. We will use D to construct a distinguisher D' that can win the CRH game in Definition 2 with the same advantage. D' functions as follows. It invokes its own oracle with messages 1 and 2 to get strings w_1, w_2 respectively. Then it constructs $w_1 \circ w_2$ and sends it to D as the PRG challenge. It outputs D 's guess bit as its output, thereby inheriting its success probability. Note that, in this reduction, D' implicitly uses the fixed R as the PRG seed, even though it does not know it. This concludes the proof. \square

C Public Key Encryption (PKE) based solution

Permute + Share can be implemented using an additively homomorphic public key encryption scheme such as Paillier [24]. Another alternative to using Paillier encryption is to use El Gamal encryption [7] which provides multiplicative homomorphism. But using El Gamal encryption will result in multiplicative shares instead of additive, and converting them to additive shares introduces huge overhead and quickly makes the scheme infeasible. We discuss both approaches here.

Pailler-encryption-based solution: Before getting into the the Permute + Share construction, let us recall Pailler = (Gen, Enc, Dec)[24].

Key Generation: This algorithm consists of the following:

1. $n = pq$ where p, q are two large primes of equal length.
2. Define $\phi(n) = (p-1)(q-1)$.
3. Set $g = n + 1$ and $\mu = \phi(n)^{-1} \pmod n$
4. Set $\text{sk} = (p, q)$ and $\text{pk} = (n, g)$.

Encryption: Let m be a message to be encrypted where $0 \leq m < n$. Select a random r where $0 < r < n$ and $r \in \mathbb{Z}_{n^2}^*$. Compute ciphertext $c \leftarrow g^m r^n \pmod{n^2}$. Let us denote this as $c = [m]$.

Decryption: Given a ciphertext $c < n^2$, compute

$$m \leftarrow (L(c^{\phi(n)} \pmod{n^2}) \cdot \mu \pmod n) \text{ where } L(u) = \frac{u-1}{n} \text{ for } u = 1 \pmod n.$$

We will be using the following properties of Pailler in our construction:

Homomorphism: The product of a ciphertext c with a plaintext m' raising g will decrypt to the sum of the corresponding plaintexts: $\text{Decrypt}([m] \cdot g^{m'} \pmod{n^2}) = m + m' \pmod n$.

Ciphertext Randomization: To randomize a ciphertext c , pick a random r' where $0 < r' < n$ and compute $c \cdot r'^n \pmod{n^2}$.

Now let us define the Permute + Share protocol using Pailler. Let (sk, pk) be P_1 's encryption keys. In the following we denote the component-wise Hadamard product of two vectors a, b by $a \odot b$.

1. P_1 sends encrypted vector \mathbf{x} , denoted as $\mathbf{c} = [\mathbf{x}]$ to P_0 .
2. P_0 picks a vector of random elements \mathbf{r}_1 where each element $e \in \mathbb{Z}_{n^2}^*$ and $0 < e < n$ and randomizes the ciphertexts $\mathbf{c}' \leftarrow \mathbf{c} \odot \mathbf{r}_1^n \pmod{n^2}$.

3. P_0 permutes \mathbf{c}' to obtain $\mathbf{b} \leftarrow [\pi(\mathbf{x})]$
4. Then P_0 picks another vector of random elements \mathbf{r}_2 where each element $e \in \mathbb{Z}_{n^2}^*$ is in $0 < e < n$ and computes $\mathbf{b} \cdot \mathbf{g}^{-r_2} \bmod n^2$ and sends it back to P_1 .
5. P_0 's share is \mathbf{r}_2 .
6. P_1 decrypts $[\pi(\mathbf{x}) - \mathbf{r}_2]$ to receive $\pi(\mathbf{x}) - \mathbf{r}_2$.

Cost: In this protocol, since every element of \mathbf{x} has to be encrypted and the encryption message space is defined to be \mathbb{Z}_n , each element has to be broken into blocks of size n for this protocol. This means that P_0 computes $N * \lceil \ell/n \rceil$ encryptions and P_1 computes $N * \lceil \ell/n \rceil$ ciphertext randomizations and ciphertext with plaintext multiplications. The communication for this protocol is $N * \lceil \ell/n \rceil * 2n$ bits. The protocol takes 1 round.

Feasibility of an El Gamal-based solution: Typically, Pailler requires 4096 bit primes for the modern standard of security, which is expensive. An alternate solution will be to use El Gamal Encryption [7] which provides multiplicative homomorphism. So, if we were to implement the above scheme using El Gamal encryption, P_0, P_1 will end up with multiplicative shares and will need to run a secure protocol (using Garbled Circuits) to convert from multiplicative to arithmetic shares. El Gamal can be implemented on Elliptic Curves with small parameters, typically, 256 bits. But this means that the multiplicative shares are Elliptic Curve point shares; converting EC point shares to arithmetic shares inside a GC is prohibitively expensive. Yet another solution will be to avoid using Elliptic Curves and use large finite fields, but this will require large parameters, ~ 2000 bits or more, which will result in multiplicative shares over large finite fields. Converting to arithmetic shares using GC is also prohibitively expensive.

Alternatively, we could use El Gamal in an EC group to encrypt only short messages in the exponent, and then decrypt by computing the discrete log or simply doing a table lookup. Suppose we use 40 bit messages for and 512 bits for the ciphertext size (which requires either 2^{40} exponentiations per decryption or a roughly 70TB lookup table, either of which is already fairly unreasonable). On a database of size N where each element is ℓ bits, then the ciphertext expansion using the ElGamal variant is roughly $10N\ell$, whereas with our hybrid protocol (with $T = \sqrt{N}$) is $4N\ell$. Assuming each exponentiation in the ECC group takes around 100 microseconds, the computation cost of our hybrid protocol is almost 2 orders of magnitude faster than this ElGamal approach. Thus, this approach is not competitive and we omit it from our comparisons below.

D The Choice of Optimal Subpermutation Size T

We choose the best value of T empirically, by fixing desired N, ℓ , and using the formulas to find the running time by iterating over $T = 2, \dots, N$.

However, in this section we describe some properties of the running time as a function of T , which give a better understanding of how the running time depends on T .

Recall that the number of AES calls is $3dN \log T + dNT(2 + \ell/k)$, and the number of bits sent is $3dNk \log T + (d + 1)N\ell$. Consider the following extreme cases:

- Very large values of T , e.g. $T = N$, $d = O(1)$. Omitting constants, the computation cost becomes proportional to $N \log N + N^2(2 + \ell/k)$, and the communication cost becomes proportional to $kN \log N + N\ell$. For large enough values of N , the term $N^2(2 + \ell/k)$ (and thus *the computation cost*) dominates, and the total cost can be reduced by decreasing T .
- Very small values of T , e.g. $T = 2$, $d \approx 2 \log N$. Omitting constants and lower order terms, the computation cost becomes proportional to $N \log N + N \log N(2 + \ell/k)$, and the communication cost becomes proportional to $kN \log N + N \log N\ell$. For large enough values of N , the total cost is dominated by the term $kN \log N$ or $N \log N\ell$ (and thus by *the communication cost*). Thus, the total cost can be reduced by decreasing d , i.e. by increasing T .

From this, one can derive several useful observations about the optimal value of T :

- Since computing d from T and N involves rounding down $\frac{\log N}{\log T}$, there are several values of T which result in the same d (for a fixed N). For any fixed d , among all the values of T resulting in this d , the smallest T gives the smallest total running time.
- For large enough N , one can expect the optimal value of T to be in the lower region of $[1, \dots, N]$. This is because for large $T \approx N$, there will be a quadratic term in N , but for very small T , the dominating term is only proportional to $N \log N$.
- The computation cost is minimized by smaller values of T , but the communication cost is minimized by larger values of T . Thus, if one increases the “relative cost” of a single AES computation over the network bandwidth - e.g. by running the protocol on a slower machine, or by using faster internet - one can expect the optimal value of T to become smaller. Further, if after estimating the computation and communication costs separately, it appears that e.g. computation cost dominates, it means that T must be decreased, and vice versa.