

Comparison-Based MPC in Star Topology (Full Version)

Gowri R Chandran¹, Carmit Hazay², Robin Hundt¹, Thomas Schneider¹

¹ ENCRYPTO, Technical University of Darmstadt, Germany

² Bar-Ilan University, Israel

Abstract. With the large amount of data generated nowadays, analysis of this data has become eminent. Since a vast amount of this data is private, it is also important that the analysis is done in a secure manner. Comparison-based functions are commonly used in data analysis. These functions use the comparison operation as the basis. Secure computation of such functions have been discussed for median by Aggarwal et al. (EUROCRYPT'04) and for convex hull by Shelat and Venkatasubramanian (ASIACRYPT'15).

In this paper, we present a generic protocol for the secure computation of comparison-based functions. In order to scale to a large number of participants, we propose this protocol in a star topology with an aim to reduce the communication complexity. We also present a protocol for one specific comparison-based function, the k^{th} ranked element. The construction of one of our protocols leaks some intermediate values but does not reveal information about an individual party's inputs. We demonstrate that our protocol offers better performance than the protocol for k^{th} ranked element by Tueno et. al. (FC'20) by providing an implementation.

Keywords: Secure multi-party computation, k^{th} ranked element, Scheduling problems, Optimisation problems

1 Introduction

Data is being constantly generated by organisations as well as individuals. To leverage this massive volume of unstructured data, organisations seek to use data analysis techniques. Data analysis can benefit organisations in various ways. For example, businesses can learn more about their target group by running analysis on the consumers trends, multiple companies can run analysis on their combined data to compare various data points such as salaries of employees, the key performance indicator, etc., hospitals and healthcare companies analyse their data using artificial intelligence and machine learning to obtain faster and more accurate diagnosis. In addition to using the data for analytics, these organisations, at times, wish to keep the data private as it contains sensitive information. In all the above mentioned examples the data being analysed is sensitive. In cases such as that of healthcare companies, the sharing of patient's data is forbidden by law. The analysis therefore has to be done in a manner which does not reveal the inputs. This is where secure Multiparty Computation (MPC) comes into play.

For the past couple of decades, MPC has been a prominent field of research. Starting with the seminal works of [Yao82,GMW87,BMR90] it is still a widely researched topic with recent works like [LPSY15,WRK17,CCG⁺20]. The problem of secure MPC focuses on a group of parties that do not trust each other, but still wish to compute a function f of their inputs while keeping their inputs private. Namely, it allows a set of mutually distrusting parties to securely compute a function on their joint inputs without revealing anything about their inputs except what can be inferred by the output. In the real world, there are adversaries present that may act maliciously to gain more information than they are supposed to. Semi-honest adversaries follow the protocol as it is but try to learn more information from the messages. In cases where companies or organisations run the protocol, semi-honest security is a realistic assumption, as the organisations would not deviate from the protocol for their reputation's sake. Another factor that is considered for the construction of an MPC protocol is the number of parties that are corrupted by the adversary. In our work, we consider semi-honest adversaries and a dishonest majority, i.e., $n - 1$ parties can be corrupted by the adversary.

One of the standard approaches to implement constant round MPC used in [BMR90,BNP08,KSS09,LPSY15,LSS16] is by using Garbled Circuits proposed by Yao [Yao82], by converting the function to be computed to a boolean circuit and privately evaluating the gates. Alternatively, the MPC protocol by Goldreich, Micali and Wigderson [GMW87], which also uses a boolean representation, works by secret sharing the wire values amongst the parties. This protocol has been improved and implemented in [CHK⁺12,BGIN21]. All the aforementioned protocols are generic MPC protocols which can be used to implement any function f by converting it to a boolean circuit.

Another line of works considers the development of protocols for specific functions to achieve improved efficiency by exploiting the properties of the underlying function and optimising the concrete protocols accordingly. Several

works have proposed protocols for specific functions such as private set-intersection [HV17,IOP18,PSZ18,RT21] for finding the intersection of multiple sets, secure pattern matching [HL08,HT10,YSK⁺13,FHV13] for finding matching patterns in texts and RSA key generation [FLOP18,HMR⁺19,CHI⁺21] for generation of the RSA modulus. Relevant to our work, in [GSV07,DGK07,DGK08,KSS09,Cou16], protocols for the secure comparison of integers have been proposed using various techniques.

The computation of these specific functions can be optimised by reducing the function f to the computation of smaller/easier computable primitives. One method is to reduce the secure computation of f into the secure evaluation of a boolean circuit. Another technique is to reduce the function f to multiple instances of a smaller function and perform a secure computation of this primitive. The latter reduction technique is used in [SV15] to reduce f to the comparison function, which takes two integer inputs and returns 1 if the first input is smaller than the second and 0 otherwise. Here the authors present a two-party protocol for the computation of a class of functions, where the parties only interact for implementing the comparison function. This reduction results in a much more efficient protocol, as the parties only communicate to evaluate the comparison function. The output of these functions is a tuple, hence the total communication depends on the number of elements in the output tuple, in comparison to a circuit-based approach, where the communication depends on the number of inputs of both parties. In most cases, the number of elements in the output tuple is considerably less than the number of elements in the input set.

The comparison-based functions considered in [SV15] are widely used for various data analytic purposes. They include functions such as finding the convex hull, finding the median, job scheduling problems, matroid optimisations and many more optimisation problems. One of the functions that we discuss in detail in our work is the convex hull. The convex hull of a set of points is the smallest convex set which contains all the points in that set. Another function that we consider is the median of a set, which is the element that is in the middle of an ordered set. We also discuss job scheduling, which is an optimisation algorithm where multiple parties have jobs that require the use of a common resource and these jobs are assigned to the resource at a particular time.

These functions have important real world applications. For instance, the secure computation of the convex hull of a set is useful in tracking a disease epidemic, where the extent of the spread of a disease can be monitored without revealing all the locations of the infected patients. A two-party variant of this problem is discussed in [SV15] where the authors use the *Gift Wrapping Algorithm* for computing the convex hull. Another function that is highly applicable in data analysis, especially in financial analysis, is finding the median. The secure computation of computing the median of the union of multiple sorted sets is essential in cases where the elements in the set are sensitive, for example finding the median of the salary of employees of various companies without actually revealing the salaries. Secure computation of the median has been studied intensely, some of the notable works being [AMP04,SV15]. Scheduling problems, such as job scheduling, have many applications in settings where the resources are limited and more than one user wishes to use them. Secure job scheduling can be used in applications where the details of the job (e.g. duration, amount of resource used etc.) are to be kept private, for instance in booking appointments at a doctor's, where the time taken is kept private.

A generalisation of the median functionality is finding the k^{th} ranked element of the union of multiple sorted sets. This function has similar applications to that of the median in financial and medical analysis. [AMP04] present a secure protocol for the computation of the k^{th} ranked element, where the k^{th} ranked element is computed using the binary search algorithm. Following that, a constant round protocol for this function is presented in [TKK⁺20], where the protocol is presented in a star topology with all the parties communicating only with a dedicated server.

In instances where multiple organisations wish to jointly analyse their data, often the communication occurs via WAN connections, making communication the bottleneck for running the protocol, as most organisations have high computational power. Therefore, protocols with low communication complexity are crucial. One method to achieve this is by constructing the protocols in a star network topology. The parties would then only need to communicate to one central party. The central party has its own input for the computation and also interacts with all the other parties in a series of secure two-party computations. This eliminates the need for broadcast channels, thus reducing the communication complexity. One drawback of this topology is that the central party has about n times more communication than the other parties. This overhead can be balanced by letting any party take the role of the central party when multiple instantiations of the protocol are carried out. To compensate for the additional computational overhead on the central party, it would suffice to increase its computational power.

Our Contribution and Outline

In this paper, we explore the concrete efficiency of comparison-based functions, i.e. a class of functions that can be reduced to a secure computation of comparison of integers. We present two different protocols for this class of functions. The first is a generic protocol (see §3) for secure computation of a class of functions called the Greedy

Compatible functions (see §2.3) where we extend the two-party protocol from [SV15]. We implement a reduction technique to reduce the computation of the function to a multi-party computation of the minimum of n integers and propose an efficient new method for computing the minimum. Instead of using garbled circuits to compute the entire function altogether, we use garbled circuits for implementing the comparisons. We chain the garbled circuits together to implement a series of two-party comparisons, with the final output being the minimum of n integers.

Next, we give concrete instantiations (see §3.2) of a few functions demonstrating the practical applications of our protocol. The only concrete multi-party protocol for computing the median was proposed in [AMP04]. To the best of our knowledge, there have not been any previous works on specific multi-party protocols for the other two problems, i.e. convex hull and job scheduling. We show that our multi-party protocol can be used for the computation of these functions and that it has better performance compared to specific approaches, such as [AMP04] for the median, and to generic MPC protocols that can be used for any of these functions.

Lastly, we present an alternative way of computing the k^{th} ranked element (see §4) by using reduction techniques similar to the generic protocol. Our protocol is the first that uses these reduction techniques to compute this function in a star topology. If the protocol is instantiated in a star topology, the computation of this function can be reduced to a secure computation of a summation function. Then communication only occurs during the computation of the summations, and the remaining computations are done locally by the parties and the comparisons are done locally by the central party. In particular, we reduce the computation of the k^{th} ranked element to a computation of secure summation which is instantiated using a threshold homomorphic encryption scheme and is implemented in a star topology where one of the participating parties plays the central party which interacts with the rest of the parties.

We note that our protocol for the k^{th} ranked element leaks the intermediate result of the summations to the central party. This leakage can reveal the distribution of the data in the union of sets. For some applications, revealing the distribution of data is a tolerable leakage. A potential approach for protecting this leakage can be by using differential privacy by revealing only the differentially private leakage. Combining MPC with differentially private tools has been used previously in [GRR19] to achieve cheaper private set intersection by allowing differentially private leakage. In [HMFS17] the problem of private record linkage with differentially private leakage is studied. Sometimes, allowing some leakage can result in a more efficient protocol. There have been many works that discuss this trade-off on privacy for better performance. In [CJJ⁺13,PKV⁺14] some information related to the search queries is leaked in order to achieve more efficient database search functionalities. We leave the idea of using differentially private leakage for future work.

	[FH96]	[GSV07]	[DGK07,DGK08]	[Cou16]	[ABJ ⁺ 19]	This work
Offline comm.	—	$\mathcal{O}(\kappa nd / \log \kappa)$	—	$\mathcal{O}(\kappa d / \log \kappa)$	—	$\mathcal{O}(\kappa d)$
Online comm.	$\mathcal{O}(n^2)$	$\mathcal{O}(nd)$	$\mathcal{O}(nd(d + \kappa))$	$\mathcal{O}(nd)$	$\mathcal{O}(\kappa nd \log n \log d)$	$\mathcal{O}(\kappa nd)$
Rounds	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n \log d)$	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n \log \log d)$	2	$\mathcal{O}(\log n)$
Assumption	DDH	OT	DGK	OT	LWE	OT

Table 1: Comparison of complexities for the computation of minimum function for n parties, using generic multi-party protocols ([FH96,ABJ⁺19]) or two-party comparison protocols ([GSV07,DGK07,DGK08,Cou16]) with our garbled circuit-based approach. Here, κ is the security parameter, d is the length of the input and n is the number of parties.

Related Work

Here, we mention several closely related works. The development of general-purpose MPC started with [GMW87] and is still a major area of research with seminal works like [FH96,CDN01,BLO16,LSS16,ABJ⁺19]. These protocols can be used to instantiate the function for finding the minimum integer, which is one of the major underlying computations of our generic protocol. If we use the multi-party protocol in [ABJ⁺19] to instantiate the computation of minimum function, the resulting communication complexity is $\mathcal{O}(\kappa nd \log n \log d)$ with 2 rounds, where κ is the security parameter, n is the number of parties, and d is the input size. This protocol uses functional encryption combiners to achieve a constant round multi-party computation protocol, and although introducing good asymptotic results, it is not practical enough for an implementation.

	[AMP04]	[TKK ⁺ 20]			This work
		YGC	AHE1	AHE2	
Comm.	$\mathcal{O}(n^2 \log S)$	$\mathcal{O}(\kappa n^2)$	$\mathcal{O}(\kappa n^2 dt)$	$\mathcal{O}(\kappa n^2 dt)$	$\mathcal{O}(n \log S)$
Rounds	$\log S$	4	4	4	$\log S$

Table 2: Comparison of protocols for secure computation of the k^{th} ranked element. We compare our protocol with the one in [AMP04] and the three protocols presented in [TKK⁺20] that are based on Yao’s garbled circuit (YGC) and additively homomorphic encryption (AHE). n is the number of parties, κ is the security parameter, d is the bit-length of the input, t is the threshold of the additive homomorphic scheme, and S is the range of elements in the database.

Often specific purpose protocols are developed to replace the use of generic MPC and improve the performance. There has been abundant work done to develop protocols specifically for the secure comparison of two integers. In [DGK07, DGK08], homomorphic encryption is used to build a protocol for the two-party comparison of integers. Using either of these protocols for computing the minimum, a communication complexity of $\mathcal{O}(nd(d + \kappa))$ is achieved with $\mathcal{O}(\log n)$ rounds. In [GSV07], a two-party protocol for the comparison of integers is presented using the encryption scheme of [CDN01]. Using this protocol to implement the minimum function the online communication complexity is $\mathcal{O}(nd)$ with a round complexity of $\mathcal{O}(\log n \log d)$. In [Cou16], the authors use a block decomposition technique to compare the blocks of the integer and execute the comparison with Oblivious Transfer as a building block. By implementing the minimum function using [Cou16], an online communication complexity of $\mathcal{O}(nd)$ is achieved with a round complexity of $\mathcal{O}(\log n \log \log d)$. We compare the efficiency of our implementation of the minimum function with that of instantiating it with any of the above protocols in Tab. 1.

In [TKK⁺20], a constant round protocol for the computation of the k^{th} ranked element is presented. The protocol is presented in a star network topology where the clients interact with a server. Unlike this setup, the central party in our protocol is one of the parties participating in the protocol and also provides input. Tueno et al. present several protocols in [TKK⁺20], using different building blocks. The protocol using the garbled circuit approach has a total communication complexity of $\mathcal{O}(\kappa n^2)$, while the protocol using an additively homomorphic encryption scheme has a communication complexity of $\mathcal{O}(\kappa n^2 dt)$, where t is the threshold of the homomorphic encryption scheme. Our work achieves a communication complexity of $\mathcal{O}(\kappa n \log S)$, where S is the range of elements in the database, and a round complexity of $\mathcal{O}(\log S)$ rounds. We provide a comparison of our work and the previous works on the k^{th} ranked element in Tab. 2.

2 Preliminaries

In this section we define some basic cryptographic primitives that are used in our protocols.

Basic notations. We denote a security parameter by κ . We say that a function $\mu : \mathbb{N} \rightarrow \mathbb{N}$ is *negligible* if for every positive polynomial $p(\cdot)$ and all sufficiently large κ ’s it holds that $\mu(\kappa) < \frac{1}{p(\kappa)}$. We use the abbreviation PPT to denote probabilistic polynomial-time and denote by $[n]$ the set of elements $\{1, \dots, n\}$ for some $n \in \mathbb{N}$.

2.1 Computational Indistinguishability

We specify next the definition of computational indistinguishability.

Definition 1 Let $X = \{X(a, \kappa)\}_{a \in \{0,1\}^*, \kappa \in \mathbb{N}}$ and $Y = \{Y(a, \kappa)\}_{a \in \{0,1\}^*, \kappa \in \mathbb{N}}$ be two distribution ensembles. We say that X and Y are computationally indistinguishable, denoted $X \stackrel{c}{\approx} Y$, if for every PPT machine \mathcal{D} , there exists a negligible function negl such that:

$$|\Pr[\mathcal{D}(X(a, \kappa), 1^\kappa) = 1] - \Pr[\mathcal{D}(Y(a, \kappa), 1^\kappa) = 1]| \leq \text{negl}(\kappa).$$

2.2 Decisional Diffie-Hellman Problem

Let \mathcal{G} be a group generation algorithm, which outputs $(p, \mathbb{G}, \mathbb{G}_1, g)$ given 1^κ , where \mathbb{G}, \mathbb{G}_1 are descriptions of groups of prime order p and g is a generator.

Definition 2 (DDH) We say that the decisional Diffie-Hellman (DDH) problem is hard relative to \mathcal{G} if for any PPT distinguisher \mathcal{D} , there exists a negligible function negl such that:

$$|\Pr[\mathcal{D}(\mathbb{G}, p, g, g^x, g^y, g^z) = 1] - \Pr[\mathcal{D}(\mathbb{G}, p, g, g^x, g^y, g^{xy}) = 1]| \leq \text{negl}(\kappa)$$

where $(\mathbb{G}, p, g) \leftarrow \mathcal{G}(1^\kappa)$.

2.3 Greedy Compatible Functions

A function f is said to be *Greedy Compatible* [SV15], if f on the union of given sets can be defined using two functions \mathcal{F}_{MIN} and \mathcal{F}_{UPT} , such that these functions have a few specific properties as specified in definition 3.

Definition 3 (Greedy Compatible Functions) The necessary and sufficient conditions for a function f to be Greedy Compatible are:

1. *Unique solution:* Given inputs $X_i, i = [1, n]$ there is a unique solution.
2. *Unique order:* The output (c_1, \dots, c_l) is released in a unique order, i.e.,

$$f(X_1, \dots, X_n) = (c_1, \dots, c_l)$$

where $c_1 = \mathcal{F}_{\text{UPT}}(\perp, \bigcup X_i)$ and $c_{i+1} = \mathcal{F}_{\text{UPT}}((c_1, \dots, c_j), \bigcup X_i)$ for $i = [1, l-1]$.

3. *Local updatability:* The function \mathcal{F}_{UPT} on the union of all sets can be computed by computing the function \mathcal{F}_{UPT} on each set individually and then computing \mathcal{F}_{MIN} on its result. Namely,

$$\mathcal{F}_{\text{UPT}}((c_1, \dots, c_j), \bigcup X_i) = \mathcal{F}_{\text{MIN}}(\mathcal{F}_{\text{UPT}}((c_1, \dots, c_j), X_1), \dots, \mathcal{F}_{\text{UPT}}((c_1, \dots, c_j), X_n)).$$

2.4 Oblivious Transfer

1-out-of-2 Oblivious Transfer (OT) is a two-party protocol run between a sender \mathcal{S} and a receiver \mathcal{R} . The sender \mathcal{S} inputs a pair of l -bit strings $s_0, s_1 \in \{0, 1\}^l$ and \mathcal{R} inputs a choice bit $b \in \{0, 1\}$. At the end of the protocol, \mathcal{R} learns the chosen string s_b , but nothing about the unchosen string s_{1-b} , whereas \mathcal{S} learns nothing about the choice bit b .

Oblivious Transfer Extension. OT protocols require costly public-key cryptography, but their performance can be improved using OT extension [IKNP03, ALSZ13]. OT extension allows extending a few public key-based base OTs using only symmetric cryptography and a constant number of rounds.

2.5 Garbled Circuits

An efficient way to evaluate a boolean circuit C in a constant number of rounds is Yao's garbled circuit [Yao86, LP04]. In this approach, the circuit constructor \mathcal{S} creates a garbled circuit \tilde{C} as follows: for each wire W_i of the circuit, \mathcal{S} randomly chooses two garbled values $\tilde{w}_i^0, \tilde{w}_i^1$, where \tilde{w}_i^j represents the value j of W_i . Further, for each gate G_i , \mathcal{S} creates a garbled table \tilde{T}_i with the following property: given a set of garbled values of G_i 's inputs, \tilde{T}_i allows to recover the garbled value of the corresponding G_i 's output, but nothing else. \mathcal{S} sends these garbled tables, called garbled circuit \tilde{C} to the evaluator \mathcal{C} . Additionally, \mathcal{C} obliviously (via OT) obtains the garbled inputs \tilde{w}_i corresponding to the inputs of both parties. Now \mathcal{C} can evaluate the garbled circuit by evaluating \tilde{C} gate by gate, using the garbled tables \tilde{T}_i . Finally, \mathcal{C} translates the garbled output into the output values given for the respective parties.

2.6 Secure Multi-Party Computation

The aim of MPC is to compute an agreed upon function correctly on the private inputs and not reveal anything beyond the result. The security of MPC is defined using the *real/ideal paradigm* which is the conceptual core of the definition.

The Real/Ideal Paradigm The real/ideal paradigm introduces an *ideal world* that implicitly captures all security properties and security is defined with respect to this ideal world. This avoids the need to list out all the attacks that constitute violations of security, which is tedious and error-prone.

Ideal World. In the ideal world, the parties securely compute a function \mathcal{F} by sending their inputs privately to a trusted third party \mathcal{T} . Each party P_i sends its input x_i to \mathcal{T} . \mathcal{T} simply computes $\mathcal{F}(x_1, \dots, x_n)$ and returns the output to all parties. An adversary in the ideal world can corrupt any party P_i but not \mathcal{T} . The adversary learns nothing other than $\mathcal{F}(x_1, \dots, x_n)$ since that is the only message it receives. The ideal world is used as a security benchmark compared to an actual protocol. Specifically, an adversary in the real world should not be able to achieve anything more than it does in the ideal world.

Real World. In the real world, there is no trusted party. All parties use a protocol to communicate with each other. A real world protocol π is considered secure if anything an adversary can achieve in the real world can also be achieved by an adversary in the ideal world.

Semi-Honest Security A *semi-honest* adversary never deviates from the protocol but tries to gain information about the honest parties' inputs by observing the execution of the protocol.

A party's view consists of its inputs, its random tape and the list of messages that it receives during the protocol execution. An adversary's view consists of the views of all corrupted parties. A protocol is said to be secure against semi-honest adversaries if the corrupted parties in the real world have views that are indistinguishable to their views in the ideal world. Such an adversary in an ideal world is called a simulator. The existence of such a simulator that can simulate the view of an adversary proves that an adversary cannot learn any additional information in the real world that cannot be achieved in an ideal world.

More formally, consider a protocol π and a functionality \mathcal{F} . We denote by C the set of corrupted parties and by Sim a simulator algorithm. Then, two distributions of random variables are defined as:

- **REAL $_{\pi}(\kappa, \mathcal{J}; x_1, \dots, x_n)$:** each party P_i runs the protocol honestly using its private input x_i and security parameter κ . Let V_i be the final view of party P_i , and let y_i be its final output.
Output $\{V_i | i \in \mathcal{J}\}, (y_1, \dots, y_n)$.
- **IDEAL $_{\mathcal{F}, \text{Sim}}(\kappa, \mathcal{J}; x_1, \dots, x_n)$:**
Compute $(y_1, \dots, y_n) \leftarrow \mathcal{F}(x_1, \dots, x_n)$.
Output $\text{Sim}(C, \{(x_i, y_i) | i \in \mathcal{J}\}), \{y_i | i \notin \mathcal{J}\}$.

A protocol is said to be secure against semi-honest adversaries if the adversary's view in the real world is indistinguishable from its view in the ideal world.

Definition 4 A protocol π securely realises a functionality \mathcal{F} in the presence of semi-honest adversaries if there exists a simulator Sim such that for every subset of C and for all inputs x_1, \dots, x_n , the two distributions are computationally indistinguishable, i.e.,

$$\mathbf{REAL}_{\pi}(\kappa, C; x_1, \dots, x_n) \stackrel{c}{\approx} \mathbf{IDEAL}_{\mathcal{F}, \text{Sim}}(\kappa, \mathcal{J}; x_1, \dots, x_n).$$

2.7 Gift-Wrapping Algorithm

The Gift Wrapping algorithm [Jar73] for finding the convex hull of a set works as follows: the first point in the convex hull is the leftmost point of the set. From this point a vertical line is considered. Then this line is rotated in a clockwise direction until it touches another point in the set. The first point that touches this line is the second point of the convex hull. Then a vertical line is considered from this point and again rotated in a clockwise direction. This process continues till the last point that falls on the rotation line is the first point of the convex hull.

3 Comparison-Based Functions

Here we propose an extension to the two-party protocol in [SV15], where a secure protocol to compute a class of functions called Greedy Compatible functions (cf. §2.3) is discussed. We extend their protocol to the multi-party setting and propose optimisations for the multi-party computation of the minimum function.

The Greedy Compatible functions can be defined in an iterative manner with all the computations done locally by each party except for computing the minimum. After each iteration, the output is slowly released so that the final output of the computation is the tuple of outputs from each iteration. Furthermore, the output of each iteration is

given as the input for the next iteration. The only step where the parties interact with each other is for computing the minimum. We propose to reduce the communication between the parties by constructing the protocol in a star topology. Thus the parties only communicate with one central party which eliminates the need for broadcast channels. As discussed above, the parties first compute a part of the function locally and then the minimum together. Consequently, the protocol π_{GP} (Fig. 3) for computing the function f involves instantiating two sub-functionalities: first, the local update function \mathcal{F}_{UPT} (Fig. 1) and second, the minimum function \mathcal{F}_{MIN} (Fig. 2). The functionality \mathcal{F}_{UPT} updates the input of each party for the computation of the minimum functionality. The output of \mathcal{F}_{UPT} is the input of \mathcal{F}_{MIN} for the next iteration of the protocol. \mathcal{F}_{UPT} is computed locally by each party, where its inputs are the party P_i 's set of elements X_i and the output of \mathcal{F}_{MIN} . The output of \mathcal{F}_{UPT} is the pair (x_i, δ_i) . \mathcal{F}_{MIN} takes (x_i, δ_i) as input and computes the minimum of all values δ_i and then returns the x_i corresponding to the smallest δ_i . Now, combining these two functionalities we describe the protocol π_{GP} for securely computing f . In the initialisation step, the parties locally compute their first input pair by calling the functionality \mathcal{F}_{UPT} on the set X_i . Then the iteration begins; for the first iteration, the parties send their input pair (x_i^1, δ_i^1) to the functionality \mathcal{F}_{MIN} , which computes $\min\{\delta_1^1, \delta_2^1, \dots, \delta_n^1\}$ and returns some $c_1 = x_t^1$ corresponding to the smallest δ_t^1 . Then the parties update their inputs for the next iteration by calling \mathcal{F}_{UPT} on c_1 and X_i . The protocol runs for $j = 1, \dots, l$ iterations, with l depending on the computed function. As mentioned earlier, the protocol releases the output slowly, i.e., after each iteration, the parties receive c_j , which is a part of the final output (c_1, c_2, \dots, c_l) .

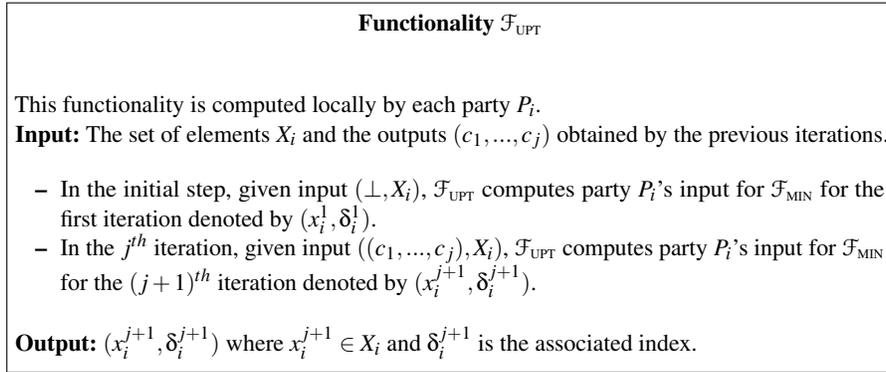


Fig. 1: Local update function.

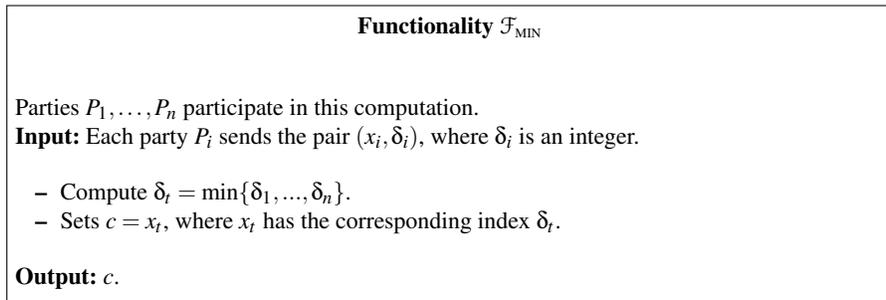


Fig. 2: Minimum function.

Security. We state and prove the security of this protocol next.

Theorem 1 *The class of Greedy Compatible functions (cf. definition 3) is securely computed by protocol π_{GP} (Fig. 3) in the presence of semi-honest adversaries for $n \geq 2$ in the \mathcal{F}_{MIN} -hybrid.*

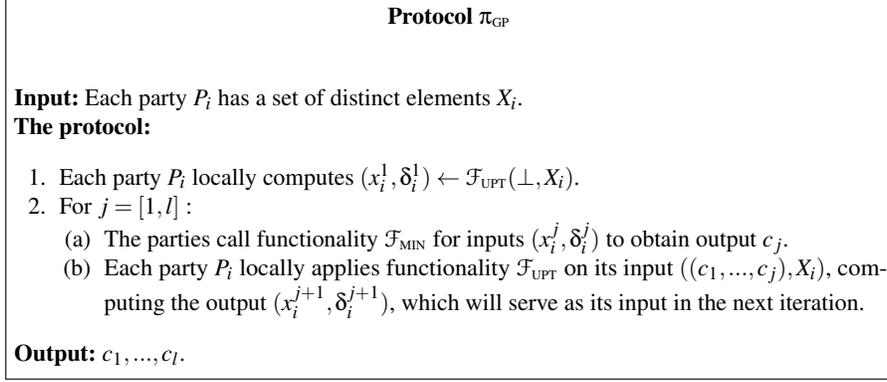


Fig. 3: Semi-honest protocol for a *greedy compatible* function f .

Proof. We prove the security of the protocol in a hybrid model, where the function \mathcal{F}_{MIN} is computed by a trusted third party. Consider \mathcal{A} to be an adversary that corrupts a subset \mathcal{J} of the parties. Let (c_1, \dots, c_l) be the final output of the computation. We construct a simulator S that generates the view of P_i , $i \in \mathcal{J}$. S is given P_i 's input X_i and the output (c_1, \dots, c_l) , then S works as follows:

1. Given X_i , $i \in \mathcal{J}$ and (c_1, \dots, c_l) , the simulator S invokes the corrupted parties on their corresponding inputs.
2. S plays the honest parties' role against the corrupted parties on arbitrary sets of inputs.
3. In the j^{th} iteration, given the inputs $\{(x_i^j, \delta_i^j)\}_{i \in \mathcal{J}}$ to \mathcal{F}_{MIN} , S simulates c_j as the output of \mathcal{F}_{MIN} .

In this case, the view of the corrupted party in the simulation is identical to that in the real execution of the protocol. From the unique ordering property of the solution, the two views are identical. Hence, the protocol π_{GP} securely computes any function f in the presence of semi-honest adversaries.

Complexity. In protocol π_{GP} , communication occurs only during the execution of \mathcal{F}_{MIN} . Let $\mathcal{O}(C)$ be the communication complexity of \mathcal{F}_{MIN} and l be the number of rounds of the protocol. Since the parties execute \mathcal{F}_{MIN} once per round, the total communication complexity is $\mathcal{O}(lC)$. We discuss the total complexity of the protocol in §3.1.

3.1 Realising \mathcal{F}_{MIN}

Recall that during the execution of π_{GP} , communication between the parties only occurs during the instantiation of \mathcal{F}_{MIN} . To reduce this communication, we propose an efficient way of computing the minimum function by splitting the multi-party computation of \mathcal{F}_{MIN} into a series of two-party computations. We achieve this by performing the comparisons pairwise. Thus, we implement \mathcal{F}_{MIN} in a star network topology where all the parties interact with one central party (say P_1) and not with any other party. This implies that the protocol π_{GP} can be implemented in a star topology as well. We define the pairwise computation of \mathcal{F}_{MIN} in Fig. 4.

Correctness. The correctness follows directly from a linear search algorithm. Nevertheless, we prove that the protocol $\mathcal{F}_{\text{MIN}}^2$ (Fig. 4) correctly computes the functionality \mathcal{F}_{MIN} .

Assume that δ_t is the smallest element in $\{\delta_1, \delta_2, \dots, \delta_n\}$ and $t \in [1, n]$. Now, if $\delta_t \in \{\delta_1, \delta_2\}$ then $(a_1, b_1) = (x_t, \delta_t)$ and at each round $i = 2, \dots, n-1$, the variables (a_i, b_i) will be set to (x_t, δ_t) . If $\delta_t \in \{\delta_3, \dots, \delta_n\}$, then at round $i = t-1$, the variables (a_i, b_i) are set to be (x_t, δ_t) and for all $i \geq t$, $(a_i, b_i) = (x_t, \delta_t)$. Hence, the output is $a_{n-1} = x_t$.

Conversely, let the output be $c = x_t$, i.e. $(a_{n-1}, b_{n-1}) = (x_t, \delta_t)$. Then two cases arise: either $\delta_t = \delta_n$ or $\delta_t = b_{n-2}$. If $\delta_t = \delta_n$, then we get that δ_t is the smallest element in $\{\delta_1, \dots, \delta_n\}$, because $\delta_n < b_{n-2}$ and $b_{n-2} = \min\{\delta_1, \dots, \delta_{n-1}\}$. If $\delta_t = b_{n-2}$, then we know that $b_{n-2} = \min\{\delta_1, \dots, \delta_{n-1}\}$ and $\delta_t < \delta_n$, therefore, δ_t is the smallest element in $\{\delta_1, \dots, \delta_n\}$. Hence, the protocol $\mathcal{F}_{\text{MIN}}^2$ correctly computes \mathcal{F}_{MIN} .

Instantiation. We instantiate the pairwise comparisons using garbled circuits and construct the protocol in a star topology. Therefore, all communications between the parties happen via party P_1 .

Our protocol is based on the idea of mobile agents [CCKM00] that chains together multiple garbled circuit computations. We consider party P_1 to be the originator of the mobile agents protocol. Then the remaining parties are the hosts. The protocol works as follows: P_1 generates the message for $n-1$ parallel OTs. The next party, i.e., P_2 constructs a

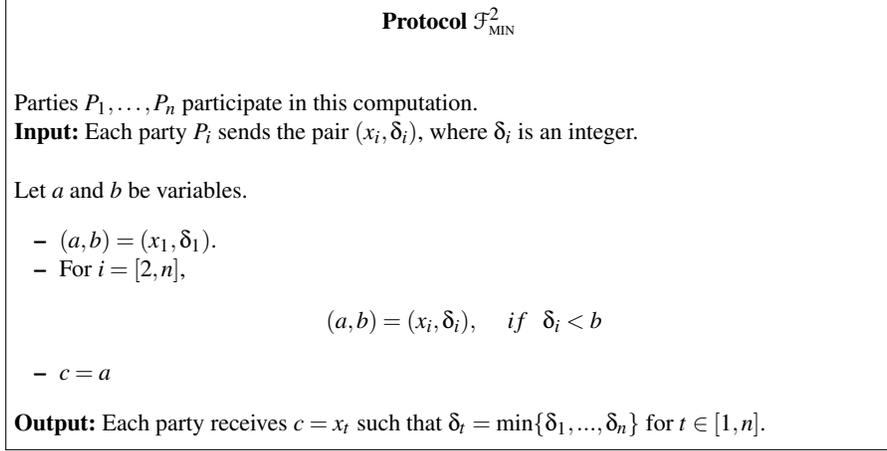


Fig. 4: Protocol for minimum function using pairwise comparisons

circuit using the output of P_1 and sends it to P_3 via P_1 . P_3 then generates a circuit using P_2 's output and sends it to P_4 via P_1 . The parties continue similarly till the last party P_n sends the circuit to P_1 and P_1 evaluates the final circuit to obtain the final output. This construction can be modified to a binary tree based structure, where parties P_{2^i} , for $i = 1, \dots, \lfloor n/2 \rfloor$ can parallelly run the second step and send the circuits to $P_{2^{i+1}}$, for $i = 1, \dots, \lfloor (n-1)/2 \rfloor$. In the third step, parties $P_{2^{i+1}}$ construct their circuits based on the received circuits and send it to party $P_{4^{i+1}}$, for $i = 1, \dots, \lfloor n/4 \rfloor$. We can also view this communication pattern as an evaluation of a hypercube, as presented in [IOP18], where each party represents a vertex of the cube.

Complexity. We first discuss the complexity for computing \mathcal{F}_{MIN} and then compute the total complexity of the protocol π_{GP} .

The number of rounds required for the computation of the minimum is $\mathcal{O}(\log n)$. To increase the computation efficiency, we use OT extensions instead of plain OTs. As the precomputation for the OT extensions can be done in parallel, the communication complexity of the precomputation is $\mathcal{O}(\kappa d)$, where d is the size of the input. The total number of comparisons performed for each minimum function is $n - 1$, hence the total communication complexity is $\mathcal{O}(\kappa n d)$.

We can instantiate the minimum functionality using either specific two-party protocols, or specific multi-party protocol or even using generic multi-party protocols. We now compare our results with those of using existing protocols. If we instantiate the pairwise comparisons using a two-party protocol for integer comparison [Cou16], it would require $\mathcal{O}(\log n \log \log d)$ rounds and would have a communication complexity of $\mathcal{O}(nd)$. Using a generic MPC protocol [ABJ⁺19], we can compute the minimum value in a constant number of rounds, with a communication complexity of $\mathcal{O}(\kappa n d \log n \log d)$. In Tab. 1, we give a detailed comparison of the protocols.

Thus, we see that our garbled circuit-based approach gives the most efficient instantiation of $\mathcal{F}_{\text{MIN}}^2$ in terms of communication complexity.

3.2 Concrete Instantiations of f

As discussed in §3, the protocol π_{GP} computes a class of functions f . The main challenge in implementing π_{GP} is to define the functionalities \mathcal{F}_{UPT} and \mathcal{F}_{MIN} , such that correctness still holds. For each function f , the definition of these functionalities changes according to f . Now we discuss some examples of f in detail and show how we can define \mathcal{F}_{UPT} and \mathcal{F}_{MIN} to realise the functions. Specifically, we consider the following functions: median of a set of elements, convex hull of a set and job scheduling. Two-party protocols for computation of median and convex hull have been given in [SV15]. We give the multi-party protocols for these functions and also present a new protocol for secure job scheduling.

Median The median of a set of N elements is the element that is in the middle of the ordered set. Consider n parties where each party P_i has a set of elements X_i . The objective is to find the median of the union of all the sets. We

assume that the number of elements in the union of the sets is public, let it be N . Then if N is odd, the median is at position $M = (N + 1)/2$ and if N is even, the median is at position $M = N/2 + 1$. The range of the elements in the union is also known and is denoted by $[a, b]$. We use the binary search algorithm to compute the median. In the j^{th} iteration, let s^j be the midpoint of the current range, l_i^j be the number of elements in X_i less than s^j , g_i^j be the number of elements in X_i greater than s^j , and m_i^j be the element in X_i which is nearest to s^j . Let the associated index be defined as $\delta_i^j = |s^j - m_i^j|$. Then, \mathcal{F}_{MIN} is defined as follows: $\mathcal{F}_{\text{MIN}}((l_i^j, g_i^j, m_i^j), \delta_i^j) = c_j$, where $c_j = (L^j, G^j, m_t^j)$ such that $t = \text{argmin}\{\delta_1^j, \dots, \delta_n^j\}$, $L^j = \sum_{i=1}^n l_i^j$ and $G^j = \sum_{i=1}^n g_i^j$.

The function \mathcal{F}_{UPT} is defined as follows:

- For $j = 1$, $\mathcal{F}_{\text{UPT}}(\perp, X_i) = ((l_i^1, g_i^1, m_i^1), \delta_i^1)$, where $l_i^1, g_i^1, m_i^1, \delta_i^1$ are all computed with respect to $s^1 = [b + a/2]$.
- For $j > 1$, if $L^j = G^j$, then the parties send terminate to \mathcal{F}_{MIN} . Else, if $L^j < G^j$, then \mathcal{F}_{UPT} sets $b = s^{j-1}$. If $L^j > G^j$, then \mathcal{F}_{UPT} sets $a = s^{j-1}$. $\mathcal{F}_{\text{UPT}}((c_1, \dots, c_j), X_i) = ((l_i^j, g_i^j, m_i^j), \delta_i^j)$, where $s^j = [a + b/2]$.

The median will be m_t^j , $t \in [1, n]$. The final output of the computation is $c_j = (L^j, G^j, m_t^j)$. The median of the union of sets is the output of the last iteration, i.e., c_j . The maximum number of iterations of the protocol is $\log N$.

Convex Hull The convex hull of a set of points is the smallest convex set, which contains all the points in that set. Suppose there are n parties, each having a set of points. Then the convex hull of the union of these sets is the smallest convex set that contains all the points of all the sets. There are various algorithms that are used to find the convex hull of a set. Here we consider the Gift Wrapping algorithm [Jar73] (see §2.7) which is the most efficient algorithm for this function.

Now, consider n parties and suppose each party P_i has a set X_i , then our objective is to compute the convex hull of the union of these sets. Each element in X_i is a point on a plane which is represented as $p_i = (x_i, y_i)$ where x_i and y_i are the X and Y coordinate of the point, respectively. Now we apply the Gift Wrapping Algorithm and define the functionality \mathcal{F}_{UPT} .

- For $j = 1$, $\mathcal{F}_{\text{UPT}}(\perp, X_i) = (p_i^1, \delta_i^1)$, where p_i^1 is the leftmost point in the set X_i (i.e., the point with the smallest X -coordinate) and δ_i^1 is the X -coordinate of p_i^1 .
- For $j > 1$, if $p_i \in \{c_1, \dots, c_{j-1}\}$, then P_i sends terminate to \mathcal{F}_{MIN} . Else, $\mathcal{F}_{\text{UPT}}((c_1, \dots, c_{j-1}), X_i) = (p_i^j, \delta_i^j)$, where p_i^j is the point that makes the smallest clockwise angle (larger than zero) with the vertical dropped from c_{j-1} and δ_i^j is the magnitude of the angle between the line joining c_{j-1} and p_i^j and the vertical from c_{j-1} . (The next point in the convex hull is the point that makes the least clockwise angle with the point p_{j-1} . Hence, each party sets its input for the next iteration by comparing the angles.)

The functionality \mathcal{F}_{MIN} is defined exactly as in Fig. 2. Thus, the final output is (c_1, \dots, c_f) , where each c_j is a point of the convex hull. The number of iterations of the protocol is equal to the number of points on the convex hull.

If there are three collinear points in $\bigcup X_i$, the rotation line will touch two points at the same time, which will give the same angle δ_i for two points p_i . This will create a conflict in the computation as the outputs may not be in unique order, which does not satisfy the properties of f in definition 3. Hence, we assume that no three points in the set are collinear.

Job Scheduling Job scheduling is an optimisation algorithm where multiple parties have jobs that require the use of a common resource and these jobs are assigned to the resource at a particular time. Secure job scheduling can be used in applications where the details of the job (e.g., duration, amount of resource used, etc.) are to be kept private, for example in a car-sharing service where the clients would like to protect the information about the usage of the car. Here we consider a job scheduling problem with one shared resource and multiple jobs. Consider n parties, each party P_i having a set of jobs $J_i = \{b_i^1, \dots, b_i^k\}$. The goal is to find the order in which to assign these jobs to a common resource R . We consider the Shortest Job First (SJF) algorithm for the scheduling as this algorithm gives the best average waiting time for the parties. We instantiate the generic protocol in Fig. 3 to realise the function. The functionality \mathcal{F}_{MIN} runs exactly like in Fig. 2 and we define the functionality \mathcal{F}_{UPT} as follows:

- For $j = 1$, $\mathcal{F}_{\text{UPT}}(\perp, X_i) = (b_i, \delta_i)$, where b_i is the shortest job in J_i and δ_i is the completion time for b_i .
- For $j > 1$, if $J_i \subseteq \{c_1, \dots, c_{j-1}\}$, then $\mathcal{F}_{\text{UPT}}((c_1, \dots, c_{j-1}), J_i) = (\perp, \delta_i)$ where $\delta_i = \infty$. If $b_i^j \in \{c_1, \dots, c_{j-1}\}$, then $\mathcal{F}_{\text{UPT}}((c_1, \dots, c_{j-1}), J_i) = (b_i, \delta_i)$ where $b_i \in J_i \setminus \{b_i^j\}$ is the smallest job with completion time δ_i . Else if $jb_i^j \in \{c_1, \dots, c_{j-1}\}$, then $\mathcal{F}_{\text{UPT}}((c_1, \dots, c_{j-1}), J_i) = (b_i, \delta_i)$ where $b_i \in J_i$ is the shortest job and its completion time is δ_i .

The final output is (c_1, \dots, c_j) , which gives the order in which to assign the jobs to the resource R . The protocol runs for $N = \sum_i |J_i|$ iterations.

4 Leaky k^{th} Ranked Element

Now we focus on one specific comparison-based function, namely finding the k^{th} ranked element of a set. Recalling that protocol π_{GP} (in §3) computes comparison-based functions that possess the properties specified in definition 3, it can therefore be implemented for computing the k^{th} ranked element of a union of sets. In this section, we construct a special protocol for the computation of the k^{th} ranked element. The protocol discussed here has a smaller communication complexity than a generic protocol computing this function, as well as previous protocols for the computation of this function. We also see that using this specific protocol is more efficient than π_{GP} (from §3) for computing the k^{th} ranked element. The functionality for finding the k^{th} ranked element for N parties is defined by $\mathcal{F}_k(D_1, \dots, D_N) \mapsto (x_k, \dots, x_k)$, where x_k is the k^{th} element of the union over all input sets D_i . We present a Leaky k^{th} Ranked Element protocol π_k^{Leaky} (Fig. 7) that securely realises functionality $\mathcal{F}_k^{\text{Leaky}}$ (Fig. 6). We use the binary search algorithm for finding the k^{th} element. This protocol works iteratively and requires two summations and two comparisons per iteration. We construct the protocol in a star network topology, hence splitting the computations into a series of secure two-party computations. By doing this, we reduce the communication complexity, but leak the result of the summations to the parties. This leakage is discussed in detail in §4.1.

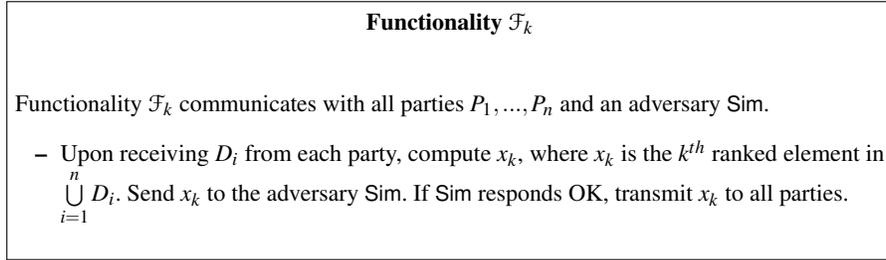


Fig. 5: Functionality for computing the k^{th} ranked element of a union of n sets.

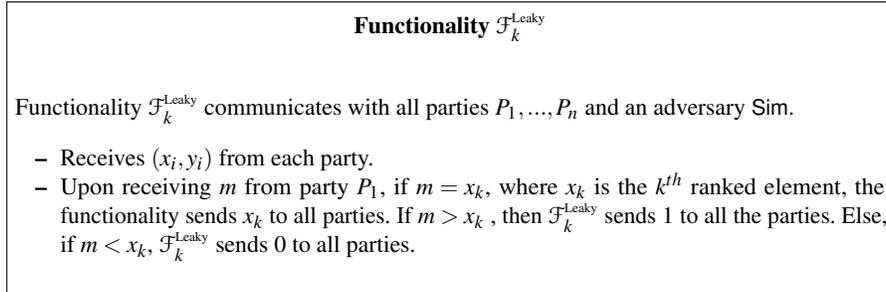


Fig. 6: Leaky functionality for computing the k^{th} ranked element of a union of n sets.

Now we present the Leaky protocol in detail. The functionality $\mathcal{F}_k^{\text{Leaky}}$ (Fig. 6) computes the k^{th} ranked element of the union of N sets. It takes inputs from the parties for computing the sum. The result of this addition is then returned to all the parties. Next, the functionality receives an input from party P_1 . If this input is equal to the k^{th} ranked element, then $\mathcal{F}_k^{\text{Leaky}}$ sends x_k (the k^{th} ranked element) to all the parties. Otherwise, if the input is smaller than x_k , $\mathcal{F}_k^{\text{Leaky}}$ sends 1 to all the parties and if the input is greater, it sends 0. The functionality is called leaky as it reveals the intermediate sum in each iteration to all the parties. This provides additional information to the parties which otherwise an adversary could not compute from the output alone.

The protocol π_k^{Leaky} , which realises functionality $\mathcal{F}_k^{\text{Leaky}}$, works in iterations and as follows: in the precomputation phase, each party computes $m = \lceil a + b/2 \rceil$ and counts the number of elements in its database that are smaller than m

(and larger than m). The parties encrypt inputs using a threshold homomorphic encryption scheme and send them to P_1 . Party P_1 computes the sum of these encrypted inputs and all the parties jointly decrypt the result and obtain the plaintext. Next, P_1 compares the result of the sum to conclude whether the k^{th} ranked element is smaller than m (or larger than m). Based on the result of the comparison, the value of m is updated for the next iteration. Note that each party only communicates with party P_1 throughout the execution. This reduces the communication cost as compared to the multi-party protocol given in [AMP04] that computes the same functionality.

Leaky k^{th} Ranked Element (π_k^{Leaky})

Input: Each party P_i has a database D_i . The rank k , the range of elements in the union of databases $([a, b])$ and the size of each database $(|D_i|)$ are public.

Primitives: A homomorphic encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ having a key generation protocol π_{Gen} .

Initial phase: Each party P_i ranks its elements in ascending order. $N = \sum_i |D_i|$ is the total number of elements in $\cup D_i$.

Key generation phase: The parties engage in a semi-honestly secure protocol π_{Gen} to generate a public key pk and their respective shares of secret key sk_i of sk .

Local computation phase: Each party P_i does the following:

1. Computes $m = \lfloor (a + b) / 2 \rfloor$.
2. Computes the number of elements (l_i) less than m and the number of elements (g_i) greater than m .

Multi-party phase:

4. Each party P_i encrypts its masked inputs, $c_i = \text{Enc}_{\text{pk}}(l_i)$ and $c'_i = \text{Enc}_{\text{pk}}(g_i)$, and sends the ciphertexts to P_1 .
5. P_1 computes $[L] = \sum_{i=1}^n c_i$ and $[G] = \sum_{i=1}^n c'_i$.
6. The parties jointly decrypt $[L]$ and $[G]$ to obtain the sums L and G , respectively. Party P_1 receives the decrypted values L and G .
7. P_1 does the following comparisons
 - (a) If $L < k$ and $G \leq N - k$, then m is the k^{th} ranked element and P_1 sends Foundk to all the parties.
 - (b) If $L \geq k$, then P_1 sends 1 to all parties. Then each party sets $b = m - 1$ and repeats the protocol from the local computation phase.
 - (c) If $G > N - k$, then P_1 sends 0 to all parties and each party sets $a = m + 1$ and repeats the protocol from the local computation phase.

Output: x_k (the k^{th} element of $\cup D_i$).

Fig. 7: Leaky protocol for passively secure computation of the k^{th} ranked element

Correctness. We prove the correctness of the protocol in the following argument. Let $[a, b]$ be the range of elements in the union of all the sets and N be the number of elements in the union. Let x_k be the element at the k^{th} position in the union of the sets where the elements are arranged in ascending order. Let m_i be the value of m in the i^{th} iteration. In the i^{th} iteration, each party counts the number of elements smaller than m_i and the number of elements greater than m_i and the sum of these values from all the parties is computed respectively. Let L_i and G_i be the total number of elements smaller than and larger than m_i respectively, in the i^{th} iteration. Then three cases arise.

- If $x_k < m_i$, then $x_k \in [a, m_i - 1]$. Then the number of elements smaller than m_i is greater than the number of elements smaller than x_k , i.e., $L_i > k - 1$. Thus, m_{i+1} will be computed as $\left\lfloor \frac{a + m_i - 1}{2} \right\rfloor$ and the procedure is repeated.

- If $x_k > m_i$, then $x_k \in [m_i + 1, b]$. Then, the number of elements larger than m_i will be greater than the number of elements larger than x_k , i.e., $G_i > N - k$. Thus, m_{i+1} is computed as $\left\lceil \frac{m_i + 1 + b}{2} \right\rceil$ and the procedure is repeated with m_{i+1} .
 - Now if $L_i < k$ and $G_i \leq N - k$, then $x_k \notin [a, m_i - 1]$ and $x_k \notin [m_i + 1, b]$, which implies $x_k = m_i$.
- Hence, the protocol correctly computes the k^{th} element.

Security. The protocol π_k^{Leaky} securely realises $\mathcal{F}_k^{\text{Leaky}}$ in the presence of semi-honest adversaries for $n \geq 2$. We discuss the proof of security in detail in §B.

Instantiations of the threshold PKE. The threshold homomorphic encryption in protocol π_k^{Leaky} can be instantiated using any homomorphic encryption schemes (see §A). In the implementation of our protocol, described in §4.3, we use the threshold Paillier PKE [Pai99]. The threshold Paillier can be implemented with distributed RSA modulus generation, as discussed in [HMR⁺19].

4.1 Leakage

Now we discuss the leakage mentioned above in protocol π_k^{Leaky} . In each iteration of π_k^{Leaky} , the sum of the number of elements smaller than m (L) and the sum of the number of elements greater than m (G) is leaked to party P_1 . Therefore for each m , the number of elements greater or smaller than m in the union of all databases, i.e., $\bigcup D_i$, is revealed. Using this leakage from each iteration, an adversary can calculate the number of elements that lie between two values of m . This shows the distribution of the elements in $\bigcup D_i$ and $\bigcup D_j$, for $j = [2, n]$. However, this leak only reveals a collective information about the union of the databases, and the distribution of elements in each individual set D_j cannot be computed from this leakage.

In some applications certain leakage may be tolerable as a tradeoff between privacy and efficiency. This can be demonstrated by several works like [CJJ⁺13,PKV⁺14,KMRR15,SGB18] that leak some information in order to obtain more efficient protocols. For instance, in [CJJ⁺13,PKV⁺14,SGB18] DBMS search protocols are presented that allow leakage of some information to improve the efficiency of the search. [KMRR15] studies the dual-execution paradigm [MF06] where the efficiency of two-party computation is improved by revealing a single bit of the honest party's input to the adversary. These works demonstrate tradeoffs between privacy and efficiency, where some leakage may be accepted in order to achieve higher efficiency. Moreover, if the leakage is to be reduced, a potential solution may be to use differential privacy. Then the leakage in the protocols would be the differentially private leakage as demonstrated in [GRR19].

4.2 Complexity

Let $S = b - a + 1$, where $[a, b]$ is the range of elements in $\bigcup D_i$. Then, the maximum number of rounds is $\log S$. The communication occurs at the setup phase for generating correlated randomness and the multi-party phase for finding the k^{th} element. The communication complexity of the key generation phase is $\mathcal{O}(\kappa \cdot n^2)$ where n is the number of parties participating. In the multi-party phase, in each round, the communication occurs for: n encryptions, 1 decryption and 1 broadcast by P_1 . Hence, the communication complexity of the online phase protocol is $\mathcal{O}(\kappa n d \log S)$, where d is the length of the inputs.

The protocol in [AMP04] also uses the binary search algorithm and requires $\log S$ rounds. The circuit consists of two summations and two integer comparisons, therefore the circuit size is $\mathcal{O}(n \log S)$. Hence, using an efficient MPC protocol [ABJ⁺19] for the computation of the circuit, the total complexity for the protocol becomes $\mathcal{O}(\kappa n d \log n \log d \log S)$.

The protocol in [TKK⁺20] uses a star topology to achieve a constant round protocol. Using the additively homomorphic encryption as the basis, they obtain a 4 round protocol. They compute the rank of the element by comparing each element with every other element. The communication complexity of their protocol is therefore quadratic in the number of participating parties, i.e., $\mathcal{O}(\kappa n^2 d)$. We give the comparison of the complexities of the above protocols with our work in Tab. 2.

4.3 Implementation

We implemented the protocol π_k^{Leaky} in the Rust programming language and instantiated the threshold homomorphic encryption with the threshold Paillier PKE using 2048 and 3072 bit modulus N . The implementation of the threshold encryption is based on the C library *libhcs* [Tie18] and the Java library *Paillier Threshold Encryption Toolbox* [UTD10].

4.4 Benchmarks

Benchmark environment: We ran our benchmarks on a server with $2 \times$ Intel Xeon Gold 6144 @ 3.5 GHz (8 physical cores) and $16 \times 32 = 512$ GB DDR4 RAM. We created 20 containers, each with 32 GB RAM and one core. The containers, running Arch Linux, were connected via a simulated WAN with a bandwidth of 100 Mbits and latency of 100 ms.

For our benchmarks, we consider the worst-case scenario for our protocol, i.e., when $k = 1$. We assume that each party holds a set of elements and the range of the elements in the union of these sets is S . Therefore, the protocol runs for $\log S$ rounds. We benchmark values of S from 10^4 to 10^{14} . We benchmark the total communication of the protocol execution for 3 to 19 parties, where each party's database has a size of 1 GB (Figs. 8c and 8d). To emulate the setting of [TKK⁺20], we also evaluate the runtime for 20 to 200 parties with a single element each (Figs. 8a and 8b). The results are averaged over 10 runs for each configuration.

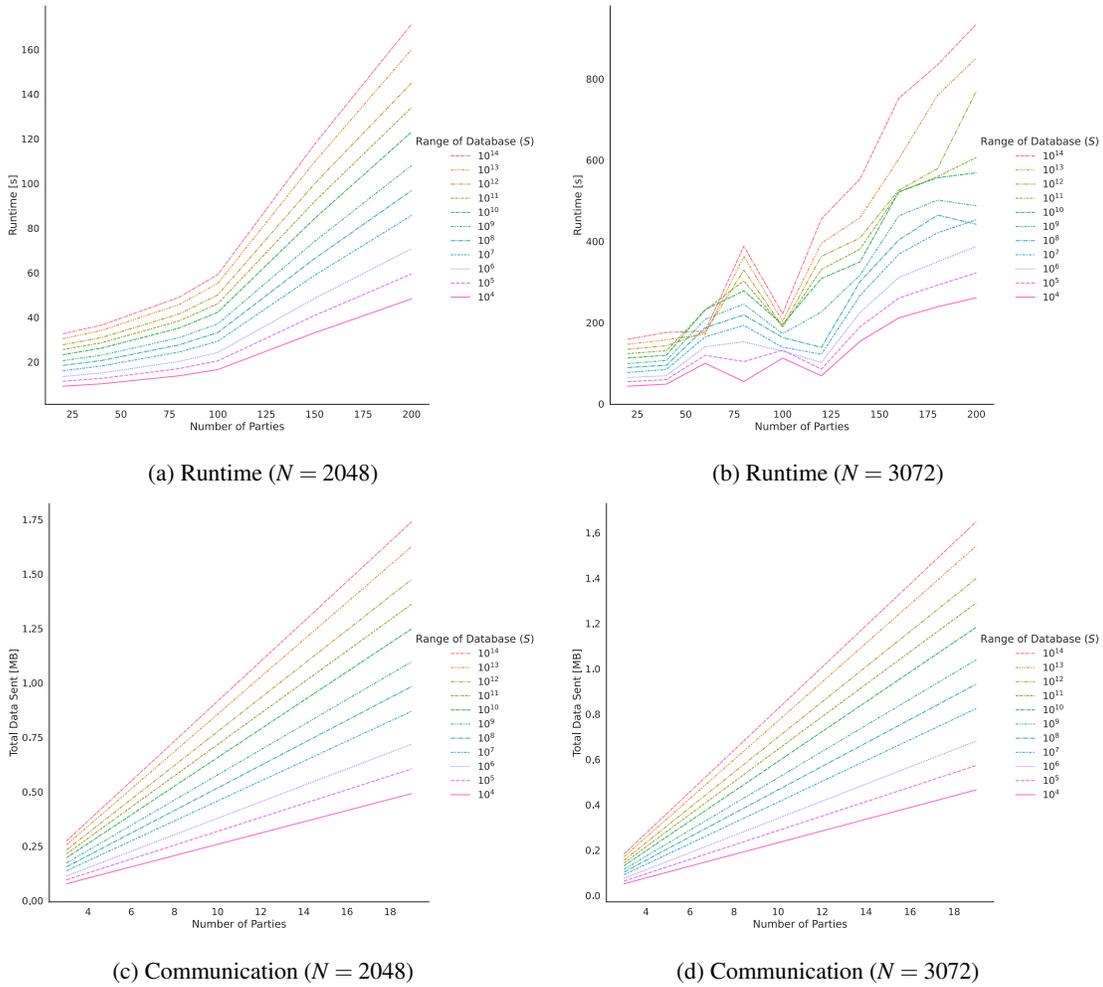


Fig. 8: Experimental analysis of protocol π_k^{Leaky} (Fig. 7), which computes the k^{th} ranked element of the union of n sets, for varying number of parties n . The plots show the results for different ranges of values in the database. N is the RSA modulus for Paillier threshold encryption.

	[TKK ⁺ 20] (no leakage)			This work (leaky)			
	YGC	AHE1	AHE2	$(N = 2048)$		$(N = 3072)$	
				$S = 10^4$	$S = 10^{14}$	$S = 10^4$	$S = 10^{14}$
Time (s)	197	1749	441	16.6	59.19	112.6	222.2
Comm. (MB)	0.31	1.11	0.32	0.027	0.096	0.040	0.143

Table 3: Performance comparison of our protocol π_k^{Leaky} (Fig. 7) for computing the k^{th} ranked element among 100 parties connected via WAN with the numbers reported in [TKK⁺20] based on Yao’s garbled circuit (YGC) and additively homomorphic encryption (AHE). Comm. is the communication from the client to the server in MB. S is the range of elements in the database. N is the RSA modulus for Paillier threshold encryption.

Results Our experimental results match with the expected asymptotic complexities. We see that the total communication scales linearly with the number of parties, and increasing the number of parties does not affect the individual communication. Moreover, the communication increases logarithmically with the range of the database.

The outliers in Fig. 8b are due to thread scheduling issues on the container executing party P_1 .

Comparison: We compare our results with the experimental results of the previous work in [TKK⁺20]. The comparison of results for 100 parties with a database range of either 10^4 or 10^{14} is given in Tab. 3. In the case of $S = 10^{14}$, our protocol reduces the client communication by more than a factor of two compared to the best protocol of [TKK⁺20].

5 Conclusion and Future Work

In this paper, we have presented two multi-party protocols, one for computing a class of comparison-based functions and the second for computing the k^{th} ranked element. The protocols that we constructed have better communication complexities as compared to the previous works on these specific functions, and to generic multi-party protocols that can be used for these functions. We reduce the functions to a computation of some high-level primitives and perform the computations in a star network topology, where one party communicates with every other party to execute a series of two-party computations. We show that such a design improves the efficiency of the protocols as the communication between parties during the execution are reduced to a minimum. Our protocols have communication complexities linear in the number of parties, which makes it easily scalable.

An interesting future direction would be to hide the leakage in the second protocol without losing the efficiency. A potential technique to reduce the leakage is to use differential privacy where the differentially private leakage will be revealed.

Another future direction would be to extend these protocols to a malicious setting by adding additional consistency check without compromising much on the efficiency.

Acknowledgments

This project received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 850990 PSOTI). It was co-funded by the Deutsche Forschungsgemeinschaft (DFG) within SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251805230, and by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within ATHENE.

This project was also supported by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Minister’s Office, and by ISF grant No. 1316/18.

References

- ABJ⁺19. Prabhanjan Ananth, Saikrishna Badrinarayanan, Aayush Jain, Nathan Manohar, and Amit Sahai. From FE combiners to secure MPC and back. In *TCC (1)*, 2019.
- ALSZ13. Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. More efficient oblivious transfer and extensions for faster secure computation. In *CCS*, 2013.
- AMP04. Gagan Aggarwal, Nina Mishra, and Benny Pinkas. Secure computation of the k th-ranked element. In *EUROCRYPT*, 2004.
- BGIN21. Elette Boyle, Niv Gilboa, Yuval Ishai, and Ariel Nof. Sublinear GMW-style compiler for MPC with preprocessing. In *CRYPTO (2)*, 2021.
- BLO16. Aner Ben-Efraim, Yehuda Lindell, and Eran Omri. Optimizing semi-honest secure multiparty computation for the internet. In *CCS*, 2016.
- BMR90. Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols (extended abstract). In *STOC*, 1990.
- BNP08. Assaf Ben-David, Noam Nisan, and Benny Pinkas. FairplayMP: a system for secure multi-party computation. In *CCS*, 2008.
- BV14. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM J. Comput.*, 2014.
- CCG⁺20. Arka Rai Choudhuri, Michele Ciampi, Vipul Goyal, Abhishek Jain, and Rafail Ostrovsky. Round optimal secure multiparty computation from minimal assumptions. In *TCC (2)*, 2020.
- CCKM00. Christian Cachin, Jan Camenisch, Joe Kilian, and Joy Müller. One-round secure computation and secure autonomous mobile agents. In *ICALP*, 2000.
- CDN01. Ronald Cramer, Ivan Damgård, and Jesper Buus Nielsen. Multiparty computation from threshold homomorphic encryption. In *EUROCRYPT*, 2001.
- CHI⁺21. Megan Chen, Carmit Hazay, Yuval Ishai, Yuriy Kashnikov, Daniele Micciancio, Tarik Riviere, Abhi Shelat, Muthu Venkatasubramanian, and Ruihan Wang. Diogenes: Lightweight scalable RSA modulus generation with a dishonest majority. In *S&P*, 2021.
- CHK⁺12. Seung Geol Choi, Kyung-Wook Hwang, Jonathan Katz, Tal Malkin, and Dan Rubenstein. Secure multiparty computation of boolean circuits with applications to privacy in on-line marketplaces. In *CT-RSA*, 2012.
- CJJ⁺13. David Cash, Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Catalin Rosu, and Michael Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *CRYPTO (1)*, 2013.
- Cou16. Geoffroy Couteau. Efficient secure comparison protocols. *IACR Cryptol. ePrint Arch.*, page 544, 2016.
- DGK07. Ivan Damgård, Martin Geisler, and Mikkel Krøigaard. Efficient and secure comparison for on-line auctions. In *ACISP*, 2007.
- DGK08. Ivan Damgård, Martin Geisler, and Mikkel Krøigaard. Homomorphic encryption and secure comparison. *Int. J. Appl. Cryptogr.*, 2008.
- DJ01. Ivan Damgård and Mads Jurik. A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In *PKC*, 2001.
- DJN10. Ivan Damgård, Mads Jurik, and Jesper Buus Nielsen. A generalization of Paillier’s public-key system with applications to electronic voting. *Int. J. Inf. Sec.*, 2010.
- FH96. Matthew K. Franklin and Stuart Haber. Joint encryption and message-efficient secure computation. *J. Cryptol.*, 1996.
- FHV13. Sebastian Faust, Carmit Hazay, and Daniele Venturi. Outsourced pattern matching. In *ICALP (2)*, 2013.
- FLOP18. Tore Kasper Frederiksen, Yehuda Lindell, Valery Osheter, and Benny Pinkas. Fast distributed RSA key generation for semi-honest and malicious adversaries. In *CRYPTO (2)*, 2018.
- Gam85. Taher El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inf. Theory*, 1985.
- Gil99. Niv Gilboa. Two party RSA key generation. In *CRYPTO*, 1999.
- GMW87. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *STOC*, 1987.
- GRR19. Adam Groce, Peter Rindal, and Mike Rosulek. Cheaper private set intersection via differentially private leakage. *PoPETs*, 2019.
- GSV07. Juan A. Garay, Berry Schoenmakers, and José Villegas. Practical and secure solutions for integer comparison. In *PKC*, 2007.

- HL08. Carmit Hazay and Yehuda Lindell. Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries. In *TCC*, 2008.
- HMFS17. Xi He, Ashwin Machanavajjhala, Cheryl J. Flynn, and Divesh Srivastava. Composing differential privacy and secure computation: A case study on scaling private record linkage. In *CCS*, 2017.
- HMR⁺19. Carmit Hazay, Gert Læssøe Mikkelsen, Tal Rabin, Tomas Toft, and Angelo Agatino Nicolosi. Efficient RSA key generation and threshold Paillier in the two-party setting. *J. Cryptol.*, 2019.
- HT10. Carmit Hazay and Tomas Toft. Computationally secure pattern matching in the presence of malicious adversaries. In *ASIACRYPT*, 2010.
- HV17. Carmit Hazay and Muthuramakrishnan Venkitasubramaniam. Scalable multi-party private set-intersection. In *PKC (1)*, 2017.
- IKNP03. Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In *CRYPTO*, 2003.
- IOP18. Roi Inbar, Eran Omri, and Benny Pinkas. Efficient scalable multiparty private set-intersection via garbled Bloom filters. In *SCN*, 2018.
- Jar73. Ray A. Jarvis. On the identification of the convex hull of a finite set of points in the plane. *Inf. Process. Lett.*, 1973.
- KMRR15. Vladimir Kolesnikov, Payman Mohassel, Ben Riva, and Mike Rosulek. Richer efficiency/security trade-offs in 2PC. In *TCC (1)*, 2015.
- KSS09. Vladimir Kolesnikov, Ahmad-Reza Sadeghi, and Thomas Schneider. Improved garbled circuit building blocks and applications to auctions and computing minima. In *CANS*, 2009.
- LP04. Yehuda Lindell and Benny Pinkas. A proof of Yao’s protocol for secure two-party computation. *Electron. Colloquium Comput. Complex.*, 2004.
- LPSY15. Yehuda Lindell, Benny Pinkas, Nigel P. Smart, and Avishay Yanai. Efficient constant round multi-party computation combining BMR and SPDZ. In *CRYPTO (2)*, 2015.
- LSS16. Yehuda Lindell, Nigel P. Smart, and Eduardo Soria-Vazquez. More efficient constant-round multi-party computation from BMR and SHE. In *TCC (B1)*, 2016.
- MF06. Payman Mohassel and Matthew K. Franklin. Efficiency tradeoffs for malicious two-party computation. In *PKC*, 2006.
- Pai99. Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, 1999.
- PKV⁺14. Vasilis Pappas, Fernando Krell, Binh Vo, Vladimir Kolesnikov, Tal Malkin, Seung Geol Choi, Wesley George, Angelos D. Keromytis, and Steven M. Bellovin. Blind Seer: A scalable private DBMS. In *S&P*, 2014.
- PSZ18. Benny Pinkas, Thomas Schneider, and Michael Zohner. Scalable private set intersection based on OT extension. *ACM Trans. Priv. Secur.*, 2018.
- RT21. Mike Rosulek and Ni Trieu. Compact and malicious private set intersection for small sets. In *CCS*, 2021.
- SGB18. Phillipp Schoppmann, Adrià Gascón, and Borja Balle. Private nearest neighbors classification in federated databases. *IACR Cryptol. ePrint Arch.*, 2018.
- SV15. Abhi Shelat and Muthuramakrishnan Venkitasubramaniam. Secure computation from millionaire. In *ASIACRYPT (1)*, 2015.
- Tie18. M. Tiehuis. libhcs. <https://github.com/tiehuis/libhcs>, 2018. Accessed: 29.11.2021.
- TKK⁺20. Anselme Tueno, Florian Kerschbaum, Stefan Katzenbeisser, Yordan Boev, and Mubashir Qureshi. Secure computation of the k^{th} -ranked element in a star network. In *FC*, 2020.
- UTD10. UTD Data and Privacy Lab. Paillier threshold encryption toolbox. <http://cs.utdallas.edu/dspl/cgi-bin/pailliertoolbox/index.php>, 2010. Accessed: 29.11.2021.
- WRK17. Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In *CCS*, 2017.
- Yao82. Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *FOCS*, 1982.
- Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, 1986.
- YSK⁺13. Masaya Yasuda, Takeshi Shimoyama, Jun Kogure, Kazuhiro Yokoyama, and Takeshi Koshihara. Secure pattern matching using somewhat homomorphic encryption. In *CCSW*, 2013.

A Additively Homomorphic PKE

A public key encryption scheme is said to be additively homomorphic if for two ciphertexts $c_1 = \text{Enc}_{\text{pk}}(m_1; r_1)$ and $c_2 = \text{Enc}_{\text{pk}}(m_2; r_2)$, we can efficiently compute $\text{Enc}_{\text{pk}}(m_1 + m_2; r)$ with an independent r and without the knowledge of the secret key sk .

To formalize the definition, we assume that both the plaintext space(\mathcal{M}) and the ciphertext space(\mathcal{C}) are groups (with operations $+$ and \cdot , respectively). We use the notation $\text{Enc}_{\text{pk}}(m)$ instead of $\text{Enc}_{\text{pk}}(m; r)$ to denote the random variable induced by the latter, where r is chosen uniformly at random. Then, we have the following definition:

Definition 5 (Homomorphic PKE) A public key encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ is said to be homomorphic if $\forall \kappa$ and all (pk, sk) output by $\text{Gen}(1^\kappa)$, it is possible to define two groups \mathcal{M}, \mathcal{C} such that:

- For all $m \in \mathcal{M}$, the value $\text{Enc}_{\text{pk}}(m)$ is an element of \mathcal{C} .³
- For every $m_1, m_2 \in \mathcal{M}$ it holds that

$$\{\text{pk}, c_1 = \text{Enc}_{\text{pk}}(m_1), c_1 \cdot \text{Enc}_{\text{pk}}(m_2)\} \equiv \{\text{pk}, \text{Enc}_{\text{pk}}(m_1), \text{Enc}_{\text{pk}}(m_1 + m_2)\} \quad (1)$$

where the group operations are carried out in the respective groups i.e., \mathcal{C} and \mathcal{M} , and the randomness for the distinct ciphertexts are independent.

Note that multiplication of a plaintext by a scalar is supported by any such scheme. We implicitly assume that each homomorphic operation on a set of ciphertexts is concluded with a refresh operation, where the party multiplies the result ciphertext with an independently generated ciphertext that encrypts zero. This is required in order to ensure that the randomness of the outcome ciphertext is not related to the randomness of the original set of ciphertexts.

A.1 Threshold PKE

In a distributed scheme, shares of the secret key are held by the parties so that the combined key remains secret. In order to decrypt, the parties use their shares to compute intermediate values, which are combined eventually to form the decrypted plaintext. To formalize this notion, we consider two multi-party functionalities: \mathcal{F}_{GEN} securely generates secret key shares for all parties (Fig. 9), and \mathcal{F}_{DEC} jointly decrypts a given ciphertext (Fig. 10).

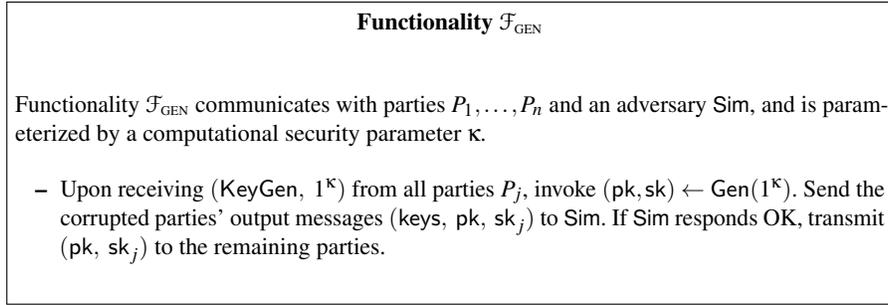


Fig. 9: Threshold key generation functionality.

A.2 The Paillier PKE

A popular instantiation of the homomorphic encryption is the Paillier encryption scheme [Pai99]. The key generation algorithm chooses two equal length primes p and q and computes $N = pq$. Then it selects an element $g \in \mathbb{Z}_{N^{s+1}}^*$ such that $g = (1+N)^j r^N \pmod{N^{s+1}}$ for a known j relatively prime to N and r^N . Let λ be the least common multiple of $p-1$ and $q-1$, then the algorithm chooses $d \pmod{N} \in \mathbb{Z}_{N^*}$ and $d = 0 \pmod{\lambda}$. The public key is (N, g) and the secret key is d . Then, encryption of plaintext $m \in \mathbb{Z}_{N^s}$ is done by computing $g^m r^{N^s} \pmod{N^{s+1}}$. And the decryption of a ciphertext c is done by first computing $c^d \pmod{N^{s+1}}$ which gives $(1+N)^{jmd} \pmod{N^s}$ and then compute discrete logarithm of the result relative to $(1+N)$. The security of the Paillier scheme is implied by the Decisional Composite Residuosity (DCR) hardness assumption.

Damgård-Jurik PKE. Damgård and Jurik [DJ01] propose a generalization of the Paillier scheme to groups of the form $\mathbb{Z}_{n^{s+1}}^*$ for $s > 0$. In this scheme, to encrypt a message, $m \in \mathbb{Z}_n^*$, a random $r \in \mathbb{Z}_n^*$ is chosen and $g^m r^{n^s} \in \mathbb{Z}_{n^{s+1}}^*$ is computed.

³ The plaintext and ciphertext spaces may depend on pk ; we leave this implicit.

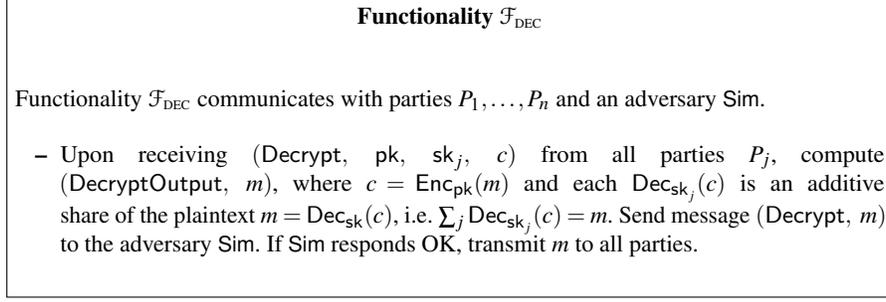


Fig. 10: Threshold decryption functionality.

Threshold Paillier with semi-honest security. The threshold variant of Paillier PKE in the passive setting is given in [Gil99], in which the parties generate an RSA modulus N mutually. Damgård, Jurik and Nielsen [DJN10] also give a threshold variant for the extension of Paillier PKE.

A.3 El Gamal PKE

Another instantiation of homomorphic encryption is the El Gamal scheme [Gam85]. It has two variants of additive and multiplicative definitions. In this work we use the additive variant. Let \mathbb{G} be a group of prime order p in which DDH is hard. Then the public key is a tuple $\text{pk} = \langle \mathbb{G}, p, g, h \rangle$ and the corresponding secret key is $\text{sk} = s$, such that $g^s = h$. Encryption is performed by choosing $r \leftarrow \mathbb{Z}_p$ and computing $\text{Enc}_{\text{pk}}(m; r) = \langle g^r, h^r \cdot g^m \rangle$. The decryption of ciphertext $c = \langle \alpha, \beta \rangle$ is performed by computing $g^m = \beta \cdot \alpha^{-s}$ and then finding m by running an exhaustive search.

Threshold El Gamal. The parties agree on a group \mathbb{G} of prime order p and a generator g . Then each party P_i selects $s_i \leftarrow \mathbb{Z}_p$ and sends $h_i = g^{s_i}$ to the other parties. Lastly, the parties compute $h = \prod_{i=1}^n h_i$ and set $\text{pk} = \langle \mathbb{G}, pg, h \rangle$. Hence, the secret key $s = \sum_{i=1}^n s_i$ associated with this public key is correctly shared between the parties. The decryption of a ciphertext $c = \langle c_1, c_2 \rangle$ works by computing $c_2 \cdot (\prod_{i=1}^n c_1^{s_i})^{-1}$, where each party sends c_1 to the power of its share s_i .

A.4 LWE-Based PKE

The standard LWE-based homomorphic encryption scheme [BV14] is constructed based on two major parameters: the dimension n and the modulus q . First, a Secret vector \mathbf{S} of n integers is chosen and then a set of public keys (\mathbf{A}_i, b) , denoted as $\text{pk} = \{\text{pk}_1, \text{pk}_2, \dots\}$, is computed using the key generation function, where \mathbf{A}_i is an arbitrary vector of n integers, and b is an integer computed as $\langle \mathbf{A}_i, \mathbf{S} \rangle + 2e_i$; where e_i is a small randomly chosen error. Now, for every j^{th} public key generation, a random arbitrary vector \mathbf{A}_j is chosen, and the public key pk_j is computed from secret key \mathbf{S} as follows: $\text{pk}_j = (\mathbf{A}_j, b_j) = (\mathbf{A}_j, \langle \mathbf{A}_j, \mathbf{S} \rangle + 2e_j)$.

Given the plaintext message bit m_i , the encryption algorithm computes the ciphertext $\mathbf{C}_i = (\mathbf{A}_i, v_i)$ using any random j^{th} public key $\text{pk}_j = (\mathbf{A}_j, b_j)$ where $v_i = b_j + m_i \pmod{q} = (\langle \mathbf{A}_j, \mathbf{S} \rangle + 2e_j + m_i \pmod{q})$ and $\mathbf{A}_i = \mathbf{A}_j$.

Decryption takes the ciphertext $\mathbf{C}_i = (\mathbf{A}_i, v_i)$ and computes the plaintext message bit m_i using the secret key \mathbf{S} . The decryption process is given as follows: $m_i = (v_i - \langle \mathbf{A}_i, \mathbf{S} \rangle) \pmod{2}$. The decryption eliminates two masks and leaves the message bit as output.

B Security of π_k^{Leaky}

Theorem 2 (Security of π_k^{Leaky}). Assume that (Gen, Enc, Dec) is an IND-CPA secure threshold homomorphic encryption scheme. Then the protocol (Fig. 7) securely realises $\mathcal{F}_k^{\text{Leaky}}$ in the presence of semi-honest adversaries for $n \geq 2$.

Proof. We construct a simulator S , where S is provided with the inputs of the corrupted parties and the output of the computation. Let \mathcal{A} be an adversary corrupting a subset of the parties, then two cases arise:

Case 1: \mathcal{A} corrupts a strict subset \mathcal{J} of the parties excluding P_1 . Let x_k be the k^{th} ranked element in the union of all the sets. The simulator S has input D_i , $i \in \mathcal{J}$ and x_k , and is defined as follows:

1. Given D_i , $i \in \mathcal{J}$ and x_k , the corrupted parties are invoked by S on their corresponding inputs.
2. S generates $(pk, sk) \leftarrow \text{Gen}(1^\kappa)$ and invokes S_{Gen} for the key generation phase.
3. S plays the role of P_1 on an arbitrary set of inputs, against the corrupted parties.
4. For each iteration j , let $m_j = \lceil (a+b)/2 \rceil$ be the median of the range as computed by each party. If $x_k < m_j$, then S returns 1 to all parties. If $x_k > m_j$, S returns 0 to the parties.

In this case, the view generated by the simulator is identical to the view of the honest parties in the real protocol. In the case when $x_k < m$, it implies that L (number of elements less than m) $\geq k$, then in the real execution as well as in the simulation, $b = m - 1$. When $x_k > m$, it implies that G (the number of elements greater than m) $\geq N - k + 1$, then in the real execution and in the simulation, $a = m + 1$. When $x_k = m$, L (number of elements less than m) $\leq k - 1$ and G (number of elements greater than m) $\leq N - k$, then in both executions m is the k^{th} element. Hence, the view of the honest parties in both executions are identical.

Case 2: \mathcal{A} corrupts a strict subset \mathcal{J} of parties including P_1 . Let x_k be the k^{th} ranked element of the union of all databases D_j . Here the simulator S has input sets D_i , $i \in \mathcal{J}$ and x_k , and is defined as follows:

1. Given D_i , $i \in \mathcal{J}$ and x_k , the corrupted parties are invoked by the simulator on their corresponding inputs.
2. S generates $(pk, sk) \leftarrow \text{Gen}(1^\kappa)$ and invokes the simulator $S_{\text{Gen}}(pk)$ for $\pi_{\text{Gen}}^{\text{SH}}$ in the key generation phase.
3. S plays the honest parties' role against P_1 in the protocol. For the j^{th} iteration, let $m_j = \lceil (a+b)/2 \rceil$ be computed by S . The simulator sends encryptions of two arbitrary inputs to P_1 and P_1 computes two sets of sums
 - (a) If $m_j < x_k$, the simulator invokes $S_{\text{Dec}}(r_1)$ for the decryption of C and $S_{\text{Dec}}(N - k + r_2)$ for the decryption of C' , where $r_1, r_2 \in \mathbb{Z}_k$. Then P_1 returns 0 and the simulator sets $a = m_j + 1$.
 - (b) If $m_j > x_k$, the simulator invokes $S_{\text{Dec}}(k + r_1)$ for the decryption of C and $S_{\text{Dec}}(r_2)$ for the decryption of C' , where $r_1, r_2 \in \mathbb{Z}_k$. Then P_1 returns 1 and the simulator sets $b = m_j - 1$.
 - (c) If $m_j = x_k$, then m_j is the k^{th} element.

In this case the difference lies in the encryptions sent to P_1 . In the real execution, the parties send encryptions of their input to P_1 whereas the simulator sends encryptions of arbitrary inputs to P_1 . Hence, the indistinguishability of the two views follows from the privacy of the threshold homomorphic encryption scheme.

Thus, the above protocol securely computes the k^{th} ranked element of the union of n sets.