



Efficient and Generic Methods to Achieve Active Security in Private Information Retrieval and More Advanced Database Search

Reo Eriguchi¹(✉), Kaoru Kurosawa^{1,2}, and Koji Nuida^{1,3}

¹ National Institute of Advanced Industrial Science and Technology, Tokyo, Japan

`eriguchi-reo@aist.go.jp`

² Research and Development Initiative, Chuo University, Tokyo, Japan

`kaoru.kurosawa.kk@vc.ibaraki.ac.jp`

³ Institute of Mathematics for Industry, Kyushu University, Fukuoka, Japan

`nuida@imi.kyushu-u.ac.jp`

Abstract. Motivated by secure database search, we present secure computation protocols for a function f in the client-servers setting, where a client can obtain $f(x)$ on a private input x by communicating with multiple servers each holding f . Specifically, we propose generic compilers from passively secure protocols, which only keep security against servers following the protocols, to actively secure protocols, which guarantee privacy and correctness even against malicious servers. Our compilers are applied to protocols computing any class of functions, and are efficient in that the overheads in communication and computational complexity are only polynomial in the number of servers, independent of the complexity of functions. We then apply our compilers to obtain concrete actively secure protocols for various functions including private information retrieval (PIR), bounded-degree multivariate polynomials and constant-depth circuits. For example, our actively secure PIR protocols achieve exponentially better computational complexity in the number of servers than the currently best-known protocols. Furthermore, our protocols for polynomials and constant-depth circuits reduce the required number of servers compared to the previous actively secure protocols. In particular, our protocol instantiated from the sparse Learning Parity with Noise (LPN) assumption is the first actively secure protocol for multivariate polynomials which has the minimum number of servers, without assuming fully homomorphic encryption.

1 Introduction

Client-server outsourcing is a central problem in secure computation. In particular, there are a wide variety of deployed systems which allow a client to search a database stored in one or more servers for desired contents. Since a client's query may contain sensitive information, it is important to realize *secure database search*, enabling a client to search a database without revealing his or

her query to the servers. A trivial solution is downloading the whole database and searching it locally. However, since the database size is typically very large, we need to construct protocols whose communication and client-side computational complexity is sublinear in the database size.

Traditionally, the problem of secure database search has been considered in two types of setting. In the *single-server* setting, there is only one server storing a database who may be corrupted; In the *multi-server* setting, there are m servers storing copies of a database and any t of them are corrupted. In this work, we focus on the multi-server setting since it is known to be impossible to efficiently achieve information-theoretic security in the single-server setting [21] and even in computational settings, the bounded collusion of servers allows better efficiency and weaker cryptographic assumptions than single-server protocols [12, 13, 26].

Private Information Retrieval (PIR) is a fundamental cryptographic primitive to realize the most basic database search. The goal of PIR is to enable an honest client to retrieve a data item a_i from a database $\mathbf{a} = (a_1, \dots, a_N)$ hiding the index i from the servers¹. To allow more complex queries such as partial match search, Barkol and Ishai [5] considered a more general setting in which a client has a private input x and servers share a function f , and the goal of the client is to obtain $f(x)$ by communicating with the servers. A rich line of works proposed secure protocols computing various classes of functions f including PIR [8–10, 14, 17, 27, 28], bounded-degree multivariate polynomials [6, 24, 36, 43], and bounded-depth circuits [5, 16, 39, 40]. Note that the communication complexity of these protocols is much smaller than the size of circuits computing f , which is the main advantages over usual multiparty computation protocols [15, 23, 34, 35].

We note that the above-mentioned protocols are *passively secure*, i.e., the privacy and correctness are guaranteed only if servers follow the protocol specifications. On the other hand, it is desirable to achieve *active security* in real-world scenarios. Namely, protocols should not only protect the privacy of queries but also guarantee the correctness of results even if some servers deviate from the protocols arbitrarily. For example, servers may try to let a client accept an incorrect result, or compute responses from an out-of-date copy of a database. This paper concerns a fundamental problem of constructing an efficient compiler from passively secure protocols to actively secure protocols. Given such a compiler, existing passively secure protocols can be directly upgraded into actively secure protocols with small overheads. Prior to our work, however, the only known passive-to-active compilers are the inefficient ones applied to PIR [11, 29], which incurs exponentially large computational overheads in the number of servers (see Sect. 1.2 for more related works on compilers in different settings including GMW-style compilers).

1.1 Our Results

In this paper, we study the problem of secure computation in the client-servers setting, where a client can obtain $f(x)$ on a private input x by communicating

¹ Throughout the paper, a client is assumed to be honest.

with multiple servers each holding f . We demonstrate theoretical feasibility of compilers that upgrade passively secure k -server protocols into actively secure m -server protocols with $m > k$. We present two such compilers: The first one upgrades a one-round passively secure protocol into a multi-round actively secure protocol. It increases the number of servers by a corruption threshold t , which seems the best possible (see Remark 1). The second one upgrades a one-round passively secure protocol into a one-round actively secure protocol while increasing the number of servers by a larger factor. More specifically,

- Our first compiler transforms a one-round passively secure k -server protocol into an $O(m^2)$ -round actively secure m -server protocol such that $m = k + t$.
- Our second compiler transforms a one-round passively secure k -server protocol into a one-round actively secure m -server protocol such that $m = O(k \log k) + 2t$.

Our compilers are generic and efficient in the sense that they are applied to protocols computing any class of functions f and the overheads in communication and computational complexity are only polynomial in m , independent of the complexity of f . Furthermore, our compilers are unconditional, i.e., requires no additional assumptions, which allows us to obtain actively secure protocols from various assumptions or even information-theoretically as shown below.

Along the way, we introduce two novel notions, *conflict-finding protocols* and *locally surjective map families*. The former is an intermediate notion between passively secure and actively secure protocols, which is used in our first compiler. The latter is a variant of perfect hash families with a stronger property, which is used in our second compiler. A key observation behind our techniques is that if a pair of servers return different answers to the same query, then a client finds that at least one of them is malicious. A difficulty is that we have to carefully design such a strategy, since just disclosing a query for one server to another may reveal his private input. See Sect. 2 for details on our techniques.

Remark 1. Our first compiler increases the number of servers by t but this seems the best possible. Indeed, the existence of a generic compiler for an actively secure protocol with $m' < t + k$ servers implies a compiler from a k -server protocol to a k' -server protocol for $k' := m' - t < k$ since an actively secure m' -server protocol implies a passively secure $(m' - t)$ -server protocol². Thus, the increase in the number of servers is optimal unless there is a generic method to reduce the number of servers. Such a method has not been found up until now.

Instantiations. Based on our compilers, we show concrete actively secure protocols for PIR, bounded-degree multivariate polynomials and constant-depth circuits. Remarkably, our protocol instantiated from the sparse LPN assumption is the first actively secure protocol for multivariate polynomials which has the minimum number of servers, without assuming fully homomorphic encryption.

² The active security enables a client to obtain a correct result by interacting with $m' - t$ servers and computing the responses of the other t servers arbitrarily.

PIR. There are compilers from a passively secure k -server PIR protocol to an actively secure m -server protocol for $m = k + 2t$ [11] and $m = k + t$ [29]. However, these compilers incur exponentially large multiplicative overheads $m^{O(t)}$ in client-side computational complexity. On the other hand, our first compiler gives an actively secure m -server protocol such that $m = k + t$ with a polynomial computational overhead $m^{O(1)}$. The only cost is that it requires $O(m^2)$ rounds of interaction between a client and servers. Our second compiler gives a one-round actively secure protocol with a polynomial computational overhead at the cost of a larger number of servers $m = O(k \log k) + 2t$. A detailed comparison is shown in Table 1.

In the information-theoretic setting, the currently most communication-efficient passively secure PIR protocol for $t \geq 2$ is the 3^t -server protocol in [10], which has sub-polynomial communication and computational complexity $N^{\omega(1)} \cdot 3^{t+\omega(t)}$ in the database size N . (Although the original protocols in [10, 14] assume non-colluding servers, i.e., $t = 1$, the corruption threshold t can be amplified by using the technique in [7] as pointed out in [32].) By applying our compilers, we obtain actively secure $3^{t+\omega(t)}$ -server PIR protocols whose computational complexity is $N^{\omega(1)} \cdot 2^{O(t)}$. It exponentially (in t) improves the complexities $N^{\omega(1)} \cdot 3^{t+\omega(t)} \cdot m^{O(t)} = N^{\omega(1)} \cdot 2^{O(t^2)}$ of actively secure protocols that are obtained from the previous compilers [11, 29]. In the computational setting, if we apply our compilers to the protocol assuming one-way functions [14], we can achieve logarithmic communication and computational complexity in N and reduce the number of servers. A detailed description is shown in Table 2.

Table 1. Comparison of passive-to-active compilers for PIR

Method	Multiplicative overhead to client-side computation	# servers	# rounds
[11]	$m^{O(t)}$	$k + 2t$	1
[29]	$m^{O(t)}\lambda$	$k + t$	1
Ours (Theorem 3)	$O(m^5\lambda)$	$k + t$	$O(m^2)$
Ours (Theorem 5)	$O(tm^2)$	$O(k \log k) + 2t$	1

k and m denote the numbers of servers in passively secure and actively secure protocols, respectively, t denotes a corruption threshold, and λ denotes a security parameter.

Bounded-Degree Multivariate Polynomials. In the information-theoretic setting, there is a passively secure protocol for polynomials in [43], which can be made actively secure by using the technique in [38]. In the computational setting, a passively secure protocol is given in [24], which can be made actively secure by the standard error correction algorithm [41]. Now, by applying our compilers, we can reduce the required number of servers of these protocols by t . Based on the passively secure protocol in [36], we can further reduce the number of servers

Table 2. Comparison of PIR protocols with sub-polynomial communication and computational complexity in the database size N for a corruption threshold $t \geq 2$

	Method	Client-side computation	# servers	# rounds	Assumption
<i>Passive</i>	[10] + [7]	$N^{o(1)}2^{O(t)}$	3^t	1	IT
	[14] + [7]	$\log N \cdot 2^{O(t)}\lambda^{O(1)}$	2^t	1	OWF
<i>Active</i>	[11] + [7] + [10]	$N^{o(1)} \cdot 2^{O(t^2)}$	$3^t + 2t$	1	IT
	[11] + [7] + [14]	$\log N \cdot 2^{O(t^2)}$	$2^t + 2t$	1	OWF
	[29] + [7] + [10]	$N^{o(1)} \cdot 2^{O(t^2)}\lambda$	$3^t + t$	1	IT
	[29] + [7] + [14]	$\log N \cdot 2^{O(t^2)}\lambda^{O(1)}$	$2^t + t$	1	OWF
	Theorem 3 + [7] + [10]	$N^{o(1)} \cdot 2^{O(t)}\lambda$	$3^t + t$	$O(3^{2t})$	IT
	Theorem 3 + [7] + [14]	$\log N \cdot 2^{O(t)}\lambda^{O(1)}$	$2^t + t$	$O(2^{2t})$	OWF
	Theorem 5 + [7] + [10]	$N^{o(1)} \cdot 2^{O(t)}$	$O(t3^t)$	1	IT
Theorem 5 + [7] + [14]	$\log N \cdot 2^{O(t)}$	$O(t2^t)$	1	OWF	

λ denotes a security parameter. “IT” stands for “information-theoretic” (i.e., no cryptographic assumption is necessary) and “OWF” stands for “one-way functions.”

by a factor of d assuming homomorphic encryption for degree- d polynomials. Notably, our protocol instantiated from [24] achieves the minimum number of servers $2t + 1$.³ A detailed comparison is shown in Table 3.

Table 3. Comparison of actively secure protocols for multivariate polynomials

Method	# servers	# rounds	Assumption
[38] + [43]	$(\frac{D}{2} + 2)t + 1$	1	IT
Theorem 3 + [43]	$(\frac{D}{2} + 1)t + 1$	$O((Dt)^2)$	IT
Theorem 3 + [36]	$(\frac{D}{d+1} + 1)t + 1$	$O((Dt)^2)$	d -HE
[41] + [24]	$3t + 1$	1	sparse LPN
Theorem 3 + [24]	$2t + 1$	$O(t^2)$	sparse LPN

D denotes the degree of the polynomials and t denotes a corruption threshold. “IT” stands for “information-theoretic” and “ d -HE” stands for homomorphic encryption for polynomials of degree d .

Constant-Depth Circuits. Barkol and Ishai [5] proposed a passively secure protocol for unbounded fan-in constant-depth circuits (i.e., the complexity class AC^0). It can be made actively secure by applying the error correction algorithm [41], and the resulting protocol needs at least $(\frac{1}{2}(\log M + O(1))^{D-1} + 2)t$ servers, where M and $D = O(1)$ are the size and depth of circuits, respectively. On the other hand, if we apply our first compiler, we need only $(\frac{1}{2}(\log M + O(1))^{D-1} + 1)t$

³ The active security is impossible if the majority of servers are corrupted, i.e., $m \leq 2t$, since there is an attack for corrupted servers to replace their input f with another function f' .

servers, which decreases the number of servers of [5] by t . For example, for the partial match problem on an M -sized database (which can be captured by depth-2 circuits of size M), our protocol requires only $(\log M + 2.5)t$ servers while the protocol obtained from [5] requires $(\log M + 3.5)t$ servers.

A beneficial consequence is that our compilers can be directly applied to future developments in passively secure protocols in the client-servers scenario and may yield new efficient constructions of actively secure protocols.

1.2 Related Work

Passive-to-Active Compilers. Within the context of PIR, there are compilers from a passively secure k -server protocol to an actively secure m -server protocol for $m = k + 2t$ [11] and $m = k + t$ [29]. As said above, however, these compilers are not only less generic in that they are applied only to PIR, but also inefficient since they incur exponential overhead $\binom{m}{t} = m^{O(t)}$ in computational complexity. There are also passive-to-active compilers in a more general multi-client setting where a private input is arbitrarily distributed among multiple clients [15, 23, 34, 35]. However, in actively secure protocols resulting from these compilers, servers need to interactively evaluate a circuit gate by gate. Consequently, protocols require communication and computational complexity that is proportional to the size of a function, and do not work efficiently if the function encodes a large database.

PIR. There are direct constructions of m -server PIR protocols in a malicious setting [3, 25, 33, 38, 47]. However, the communication complexities of [3, 25, 33, 38] are all polynomial $N^{O(t/m)}$ in the database size N , while those of ours are $N^{o(1)}$, i.e., smaller than any polynomial function in N . The protocol in [47] does not guarantee privacy if malicious servers collude, and thus does not satisfy active t -security for $t > 1$ in our sense⁴. There are also constructions of PIR with a weaker security guarantee [19, 22], which can only tell a client the existence of malicious servers. Actively secure PIR is also considered in a special setting where the length of each entry of a database is sufficiently large (e.g., [4, 44] and references therein). The protocols in [4, 44] assume that the length of each entry of a database is at least exponential in N and hence result in exponentially large communication complexity in N .

Protocols in the Single-Server Setting. Generally, if we have a passively secure single-server protocol, then we can obtain an actively secure protocol with the minimum number of servers $2t + 1$ since a client just runs the passively secure protocol with each server and computes the majority of $2t + 1$ outputs. However, there is an impossibility result on efficient single-server protocols for PIR in the information-theoretic setting [21]. Even if we go for computational security, it seems to be impossible to construct single-server PIR protocols from the minimal assumption of one-way functions [26], and for a general function, we currently

⁴ We mean by t -security that at most t servers are corrupted.

need to assume fully homomorphic encryption, which is only instantiated from a narrow class of assumptions [31, 42].

Verifiable Computation. The problem of dealing with malicious servers has also been considered within the context of verifiable computation in the single-server setting [2, 20, 30] and in the multi-server setting [1, 45, 46]. However, these verifiable computation protocols only detect malicious behavior of servers and cannot achieve active security in our sense. The protocol in [18] uses a similar idea that a client compares answers from one server with those from another. However, it does not consider the setting where a client’s input x should be private and also assumes that all parties agree on a function f in advance, while in our setting the client does not know f since it corresponds to an unknown database.

2 Technical Overview

In this section, we provide an overview of our compilers to construct an actively t -secure m -server protocol Π' from any one-round passively t -secure k -server protocol Π such that $k < m$. Let $V = \{S_1, \dots, S_m\}$ be the set of m servers of Π' . A key observation behind our constructions is that if a pair of servers in V return different answers to the same query of Π , then at least one of them is malicious⁵. We call such two servers a *conflicting pair*. The client continues to remove conflicting pairs from V in an appropriate way. Finally, the client executes a protocol only with remaining honest servers and obtains a correct result. For ease of exposition, we first explain our non-interactive actively secure protocol and then explain our interactive protocol with fewer servers.

2.1 Non-interactive Actively Secure Protocols

As a first attempt, we consider the following basic construction.

1. A client C partitions $V = \{S_1, \dots, S_m\}$ into k groups $V = G_1 \cup \dots \cup G_k$ in such a way that each G_j contains at least one honest server.
2. C computes k queries of Π on his private input and sends the j -th query to all servers in the j -th group G_j .

If every group contains no conflicting pair (i.e., all servers in each G_j return the same answer), then C can compute the correct result from the k answers of the k groups. Otherwise, C removes a conflicting pair from V , and repeats the above process at most t times to remove all malicious servers. This method, however, requires a large number of servers $m = \Omega(kt)$ since the size of each group G_j needs to be larger than t .

We reduce the number of servers to $m \approx 2t + k$ by introducing a novel notion of locally surjective map families. Technically, we consider a family \mathcal{F} of maps from the set $V = \{S_1, \dots, S_m\}$ of m servers to $[k] := \{1, 2, \dots, k\}$. Each map

⁵ Here, we assume that the server-side computation is deterministic.

$f \in \mathcal{F}$ defines a partition $V = G_{f,1} \cup \dots \cup G_{f,k}$, where $G_{f,j} = \{S_i : f(S_i) = j\}$. For each map $f \in \mathcal{F}$, the client C computes k queries of Π on his private input and sends the j -th query to all servers in the j -th group $G_{f,j}$. Our strategy is that C proceeds in t steps to detect and remove at least one new malicious server per step. In each step,

- If for every (f, j) , all the remaining servers in $G_{f,j}$ return the same answer $\text{ans}_{f,j}$, then C computes an output x_f of Π from $(\text{ans}_{f,1}, \dots, \text{ans}_{f,k})$ for each f and decides the final output by the majority vote over the x_f 's;
- Otherwise, i.e., if two remaining servers in some $G_{f,j}$ give different answers, then C removes this conflicting pair and proceeds to the next step.

Observe that in the latter case, at least one of the two servers is malicious and hence at least one malicious server is always removed. The requirement for C to succeed is that in the former case, more than half of the x_f 's are correct. A sufficient condition is that for more than half of the f 's, there remains at least one honest server in each of $G_{f,1}, \dots, G_{f,k}$. Indeed, for such f 's, C receives the correct answer from servers in each of $G_{f,1}, \dots, G_{f,k}$, or proceeds to the latter case and removes a conflicting pair. Since there remains at least $m - 2t$ honest servers at every step, the condition can be formulated as the family \mathcal{F} of maps satisfying that for any subset $H \subseteq V$ of size $m - 2t$, there exist more than half of the f 's such that $f(H) = [k]$. We name such a family as a *locally surjective map family*.

We can prove by a probabilistic argument the existence of a locally surjective map family \mathcal{F} of size $O(m)$ if $k = O((m - 2t)/\log(m - 2t))$. Therefore, we can obtain an actively t -secure m -server protocol Π' from a passively t -secure k -server protocol Π if $m = O(k \log k) + 2t$. Since the client can run all instances of Π in parallel, the resulting protocol Π' is one-round and only incurs a $O(tm|\mathcal{F}|) = O(tm^2)$ multiplicative overhead to communication and computational complexity.

2.2 Interactive Actively Secure Protocols

We further reduce the number of servers from $m = O(k \log k) + 2t$ to $m = t + k$. In our first construction, if a client C finds a conflicting pair of servers, then he removes both servers from the set V . After eliminating all t malicious servers, the number of remaining servers is reduced to $m - 2t$ in the worst case. Therefore, as long as this approach is used, the number of servers must be $m \geq 2t + k$ since it should hold that $m - 2t \geq k$.

Our second construction reduces the required number of servers to $m = t + k$ by introducing a notion of *t-conflict-finding* protocols, which is an intermediate notion between passively t -secure protocols and actively t -secure ones. Intuitively, in a conflict-finding protocol, a client C obtains a correct result or a non-trivial partition (G_0, G_1) of the set V of servers such that all honest servers are included in G_0 or G_1 (and hence the other group consists of malicious servers

only)⁶. A pair of servers crossing the partition (G_0, G_1) is supposed to be conflicting.

More concretely, we consider a graph \mathcal{G} with m vertices each of which represents a server. Our protocol starts with \mathcal{G} being a complete graph, and repeats the following steps:

1. The client C executes a conflict-finding protocol Π_{CF} with some subset $V' \subseteq V$ which forms a connected subgraph of size $k = m - t$ in \mathcal{G} (which can be efficiently found).
2. If all servers in V' behave honestly, then C obtains the correct output.
3. Otherwise C can find a partition (G_0, G_1) of V' thanks to the conflict-finding property of Π_{CF} . Note that there is always an edge $e = (S_i, S_j)$ between G_0 and G_1 since $G_0 \cup G_1 = V'$ is connected. Furthermore, since all honest servers in V' are included in G_0 or G_1 , at least one of S_i and S_j is malicious. Now, C removes the edge e from \mathcal{G} instead of eliminating the two servers, and goes back to the first step.

Since all edges among honest servers remain unremoved (and hence the set of all honest servers remains connected), C can successfully find a set of $k = m - t$ honest servers within $O(m^2)$ rounds. Note that in the above construction, C chooses a set of servers with which he executes Π_{CF} , depending on the answers that are maliciously computed in the previous rounds. Thus servers may learn some information on the client input x by seeing which servers C removes. To address this problem, we impose an additional property that the distribution of the partition (G_0, G_1) is independent of x regardless of how malicious servers behave. Then, an edge removed in each round leaks no information on x and hence the privacy of x is preserved.

Two-Round Conflict-Finding Protocols. The remaining problem is how to construct conflict-finding protocols. We show a construction of a two-round t -conflict-finding k -server protocol Π_{CF} from a passively t -secure k -server one Π . For simplicity, let $V' = (S_1, \dots, S_k)$. In the first round, a client computes real queries $(\text{que}_i)_{i \in [k]}$ on his private input x according to Π as usual. He also computes dummy queries $(\text{que}'_i)_{i \in [k]}$ on a default input x_{def} which is independent of x . He then sends a random permutation of $(\text{que}_i, \text{que}'_i)$ to each server S_i . Note that the privacy of Π and the random permutation ensure that servers cannot distinguish which queries are computed on x or x_{def} . Each server S_i returns answers $(\text{ans}_i, \text{ans}'_i)$ to the two queries as usual. In the second round, the client sends all the dummy queries $(\text{que}'_i)_{i \in [k]}$ on x_{def} to all servers in V' , which does not affect privacy since x_{def} is independent of x . In response, each server S_j returns $v_j := (\text{ans}'_i(j))_{i \in [k]}$ to $(\text{que}'_i)_{i \in [k]}$, where $\text{ans}'_i(j)$ is the answer which S_i would compute to the dummy query que'_i .

For simplicity, suppose that S_1 is the only malicious server. If S_1 behaved honestly in the first round, it holds that $\text{ans}'_1 = \text{ans}'_1(2) = \dots = \text{ans}'_1(k)$.

⁶ We say that a partition (G_0, G_1) is non-trivial if neither of G_0 nor G_1 is empty.

If S_1 returned an incorrect answer to que'_1 , it is different from any of $\text{ans}'_1(2), \dots, \text{ans}'_1(k)$. From this observation, the client C trusts the answer ans_1 of S_1 to the main query que_1 in the first round if and only if $\text{ans}'_1 = \text{ans}'_1(2) = \dots = \text{ans}'_1(k)$. Generalizing this, we let C compute an output based on $(\text{ans}_1, \text{ans}_2, \dots, \text{ans}_k)$ if all the v_j 's take the same value. Otherwise, he partitions the set of servers into equivalence classes by placing S_i and S_j into the same class if and only if $v_i = v_j$, and outputs a non-trivial partition (G_0, G_1) in some way. Note that any pair of honest servers S_i, S_j return the same answer in the second round, i.e., $v_i = v_j$, and hence they are placed in the same class. A malicious server successfully submits an incorrect answer without being detected only if it guesses correctly which query encodes the client's true input x . As we said above, it happens with probability $1/2$. More generally, if the client prepares $M - 1$ sets of dummy queries, the cheating probability of malicious servers can be reduced to $1/M$. This can be made even negligible by executing sufficiently many (say, κ) instances in parallel. If a conflict is found in some instance, C outputs a non-trivial partition (G_0, G_1) obtained in that instance. Otherwise, he outputs the majority of the κ outputs if it exists. To let this protocol fail, malicious servers need to let the client output valid but incorrect outputs in at least $\kappa/2$ instances. The cheating probability is thus $O(M^{-\kappa/2})$, which is negligible.

To see that the partition (G_0, G_1) leaks no information on the client's input x , observe that (G_0, G_1) is determined by answers $(v_j)_{j \in [k]}$. These answers are independent of x since they can be simulated from dummy queries and t queries for x that malicious servers see. The former is independent of x in the first place and the latter leaks no information on x due to the privacy of Π .

To summarize, we obtain an $O(m^2)$ -round actively t -secure m -server protocol from a passively t -secure k -server protocol if $k \leq m - t$. The communication and computational overhead is a multiplicative polynomial factor in m .

3 Preliminaries

Notations. For $m \in \mathbb{N}$, define $[m] = \{1, 2, \dots, m\}$. Let X, Y be sets. If $X \subseteq Y$, we define $Y \setminus X = \{y \in Y : y \notin X\}$ and simply denote it by \bar{X} if Y is clear from the context. We write $u \leftarrow_s X$ if u is chosen uniformly at random from X . Define $\binom{X}{k}$ as the set of all subsets of X of size k . Define $\text{Map}(X, Y)$ as the set of all maps from X to Y . If $X = [m]$ and $Y = [k]$, we simply denote it by $\text{Map}(m, k)$. Let $\log x$ denote the base-2 logarithm of x and $\ln x$ denote the base-e logarithm of x , where e is the Napier's constant. We call a function $f : \mathbb{N} \ni \lambda \mapsto f(\lambda) \in \mathbb{R}$ *negligible* if for any $c > 0$, there exists $\lambda_0 \in \mathbb{N}$ such that $0 \leq f(\lambda) < \lambda^{-c}$ for any $\lambda > \lambda_0$. We call f *polynomial* there exists $c > 0$ such that $0 \leq f(\lambda) < \lambda^c$ for all λ . Throughout the paper, we use the following notations:

- m denotes the total number of servers, which is polynomial in a security parameter λ .
- t denotes the number of corrupted servers.
- C denotes a client and S_i denotes the i -th server.

The notation $\tilde{O}(\cdot)$ hides a polylogarithmic factor in a security parameter λ .

3.1 Secure Computation in the Client-Servers Setting

We follow the client-servers model used in [5]. In this model, there is an honest client C who holds a private input x , and m servers S_1, \dots, S_m who all hold the same input p . The goal is:

Byzantine-Robustness. The client learns the value $F(p, x)$ for a publicly known function F even if t servers behave maliciously;

Privacy. The client keeps his input x hidden from any collusion of t servers.

We do not assume any interaction between servers. We call a message from a client to servers a *query* and a message from servers to the client an *answer*.

In the above setting, we assume that the function F takes a common input p from servers. Typically, the input p will be a description of a function f applied to the input x of a client (e.g., a description of a circuit or a polynomial) and F is the universal function defined by $F(p, x) = f(x)$.

If $m \geq 2t + 1$, there is a trivial 1-round protocol achieving the above goal: C downloads p from all servers, finds the correct p by the majority vote, and computes $F(p, x)$ by himself. However, this protocol results in large communication and client-side computational complexity that is linear in the description length of p . In applications to database search, p encodes a large database and its size is proportional to the database size. From this point of view, we say that a protocol is *efficient* if its communication and client-side computational complexity is sublinear in the description length of p and linear in that of x .

More formally, we define a secure computation protocol in the client-servers setting as an abstract primitive. First, we show the syntax and correctness.

Definition 1. Let $\mathcal{P} = (P_\lambda)_{\lambda \in \mathbb{N}}$, $\mathcal{X} = (X_\lambda)_{\lambda \in \mathbb{N}}$, and $\mathcal{Y} = (Y_\lambda)_{\lambda \in \mathbb{N}}$ be sequences of sets with polynomial-size descriptions and $\mathcal{F} = (F_\lambda : P_\lambda \times X_\lambda \rightarrow Y_\lambda)_{\lambda \in \mathbb{N}}$ be sequences of functions with polynomial-size descriptions. An ℓ -round m -server protocol for \mathcal{F} is a tuple of three polynomial-time algorithms $\Pi = (\text{Query}, \text{Answer}, \text{Output})$, where:

- $\text{Query}(1^\lambda, x, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]}) \rightarrow ((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)})$: **Query** is a possibly randomized algorithm that takes $x \in X_\lambda$, a state $\text{st}^{(j-1)}$ and answers $(\text{ans}_i^{(j-1)})_{i \in [m]}$ in round $j - 1$ as input, and outputs queries $(\text{que}_i^{(j)})_{i \in [m]}$ and a state $\text{st}^{(j)}$ in round j , where we define $\text{st}^{(0)}$, $\text{ans}_i^{(0)}$ as the empty string;
- $\text{Answer}(1^\lambda, p, \text{que}_i^{(j)}) \rightarrow \text{ans}_i^{(j)}$: **Answer** is a deterministic algorithm that takes $p \in P_\lambda$ and a query $\text{que}_i^{(j)}$ in round j as input, and outputs an answer $\text{ans}_i^{(j)}$ in round j ;
- $\text{Output}(1^\lambda, \text{st}^{(\ell)}, (\text{ans}_i^{(\ell)})_{i \in [m]}) \rightarrow y$: **Output** is a possibly randomized algorithm that takes a state $\text{st}^{(\ell)}$ and answers $(\text{ans}_i^{(\ell)})_{i \in [m]}$ in round ℓ as input, and outputs $y \in Y_\lambda$;

satisfying the following property:

Correctness. *There exists a negligible function $\text{negl}(\lambda)$ such that for any $\lambda \in \mathbb{N}$ and any $(p, x) \in P_\lambda \times X_\lambda$,*

$$\Pr \left[\text{Output}(1^\lambda, \text{st}^{(\ell)}, (\text{ans}_i^{(\ell)})_{i \in [m]}) = F_\lambda(p, x) \right] \geq 1 - \text{negl}(\lambda),$$

where

$$\begin{aligned} ((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)}) &\leftarrow \text{Query}(1^\lambda, x, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]}), \\ \text{ans}_i^{(j)} &\leftarrow \text{Answer}(1^\lambda, p, \text{que}_i^{(j)}) \end{aligned}$$

for all $j \in [\ell]$ and $i \in [m]$.

We note that an answer algorithm **Answer** is not always defined to be deterministic in the literature but all the instantiations considered in this paper actually have deterministic answer algorithms. We omit a security parameter 1^λ from inputs if it is clear from the context.

An abstract primitive $\Pi = (\text{Query}, \text{Answer}, \text{Output})$ immediately implies an ℓ -round protocol in the above client-servers setting. Indeed, a client has a private input x and m servers have a common input p . In each round, the client runs **Query**, sends queries to servers and stores a state in his memory. In response, servers run **Answer** on the queries that they receive, and send answers back to the client. In the final round, the client runs **Output** on his state and servers' answers, and obtains $y = F_\lambda(p, x)$. Due to this correspondence, we will use the terminologies interchangeably for the sake of readability.

The above-mentioned trivial 1-round protocol corresponds to the scheme in which **Query** outputs nothing, **Answer** outputs p and then **Output** computes $y = F_\lambda(p, x)$. To rule out this, we define the efficiency measures of Π as follows. Let $\text{que}_i^{(j)}$ and $\text{ans}_i^{(j)}$ be queries and answers computed by Π and denote their bit-lengths by $|\text{que}_i^{(j)}|$ and $|\text{ans}_i^{(j)}|$, respectively. Define the *communication complexity* $\text{Comm}_\lambda(\Pi)$ as

$$\text{Comm}_\lambda(\Pi) = \sup_{(p, x) \in P_\lambda \times X_\lambda} \sum_{i \in [m], j \in [\ell]} (|\text{que}_i^{(j)}| + |\text{ans}_i^{(j)}|).$$

Define the *client-side computational complexity* $\text{c-Comp}_\lambda(\Pi)$ as the sum of the running time of $\text{Query}(1^\lambda, x, \cdot, \cdot)$ and $\text{Output}(1^\lambda, \cdot, \cdot)$ with worst-case inputs $(p, x) \in P_\lambda \times X_\lambda$. Let $\text{Comm}(\Pi) = (\text{Comm}_\lambda(\Pi))_{\lambda \in \mathbb{N}}$ and $\text{c-Comp}(\Pi) = (\text{c-Comp}_\lambda(\Pi))_{\lambda \in \mathbb{N}}$. We say that Π is *efficient* if there exists a sublinear function $g(\ell) = o(\ell)$ such that

$$\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} \in g(|p|) \cdot |x| \cdot \text{poly}(m, \lambda), \quad (1)$$

where $|p|$ and $|x|$ are the description lengths of elements of P_λ and X_λ , respectively. One can also define the *server-side computational complexity* $\text{s-Comp}_\lambda(\Pi)$ as the running time of $\text{Answer}(1^\lambda, p, \cdot)$, and define $\text{s-Comp}(\Pi) = (\text{s-Comp}_\lambda(\Pi))_{\lambda \in \mathbb{N}}$. We see the communication and client-side complexity as a primary efficiency measure and the server-side computational complexity as a secondary measure.

Next, we show the security requirements.

Definition 2. Let $\Pi = (\text{Query}, \text{Answer}, \text{Output})$ be an ℓ -round m -server protocol for $\mathcal{F} = (F_\lambda : P_\lambda \times X_\lambda \rightarrow Y_\lambda)_{\lambda \in \mathbb{N}}$. We say that Π is actively t -secure if it satisfies the following requirements:

Privacy. There exists a negligible function $\text{negl}(\lambda)$ such that for any stateful algorithm \mathcal{A} and any $\lambda \in \mathbb{N}$,

$$\text{Adv}_{\Pi, \mathcal{A}}(\lambda) := |\Pr[\text{Priv}_{\Pi, \mathcal{A}}^0(\lambda) = 0] - \Pr[\text{Priv}_{\Pi, \mathcal{A}}^1(\lambda) = 0]| < \text{negl}(\lambda),$$

where for $b \in \{0, 1\}$, $\text{Priv}_{\Pi, \mathcal{A}}^b(\lambda)$ is the output b' of \mathcal{A} in the following experiment:

1. $(x_0, x_1, p, T) \leftarrow \mathcal{A}(1^\lambda)$, where $x_0, x_1 \in X_\lambda$, $p \in P_\lambda$ and $T \subseteq [m]$ is of size at most t .
2. For each $j = 1, 2, \dots, \ell$,
 - (a) Let $((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)}) \leftarrow \text{Query}(1^\lambda, x_b, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]})$ and give $(\text{que}_i^{(j)})_{i \in T}$ to \mathcal{A} .
 - (b) If $j < \ell$, \mathcal{A} outputs $(\text{ans}_i^{(j)})_{i \in T}$. If $j = \ell$, \mathcal{A} outputs a bit $b' \in \{0, 1\}$.

Byzantine-robustness. There exists a negligible function $\text{negl}(\lambda)$ such that for any stateful algorithm \mathcal{A} and any $\lambda \in \mathbb{N}$,

$$\Pr[\text{BR}_{\Pi, \mathcal{A}}(\lambda) = 1] < \text{negl}(\lambda),$$

where $\text{BR}_{\Pi, \mathcal{A}}(\lambda)$ is the output of the following experiment:

1. $(x, p, T) \leftarrow \mathcal{A}(1^\lambda)$, where $x \in X_\lambda$, $p \in P_\lambda$ and $T \subseteq [m]$ is of size at most t .
2. For each $j = 1, 2, \dots, \ell$,
 - (a) Let $((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)}) \leftarrow \text{Query}(1^\lambda, x_b, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]})$ and give $(\text{que}_i^{(j)})_{i \in T}$ to \mathcal{A} .
 - (b) \mathcal{A} outputs $(\text{ans}_i^{(j)})_{i \in T}$.
3. Return 1 if $\text{Output}(1^\lambda, \text{st}^{(\ell)}, (\text{ans}_i^{(\ell)})_{i \in [m]}) \neq F_\lambda(p, x)$, and otherwise return 0.

We say that Π is *passively* t -secure if it satisfies the above requirements for *semi-honest* adversaries \mathcal{A} , i.e., those following the instructions of Π . Note that for semi-honest adversaries, the Byzantine-robustness of Π immediately follows from the correctness of Π . We say that Π is *computationally actively* t -secure (resp. *computationally passively* t -secure) if it satisfies the above requirements for probabilistic polynomial-time (PPT) adversaries \mathcal{A} (resp. semi-honest PPT adversaries \mathcal{A}).

3.2 Existing Passively Secure Protocols

Private Information Retrieval. Let $N = N(\lambda)$ be a polynomial function. Define $\text{INDEX}_N = (F_\lambda : \{0, 1\}^N \times [N] \rightarrow \{0, 1\})_{\lambda \in \mathbb{N}}$ as a sequence of functions such that for each $\lambda \in \mathbb{N}$,

$$F_\lambda((a_1, \dots, a_N), x) = a_x, \quad \forall (a_1, \dots, a_N) \in \{0, 1\}^N, \forall x \in [N].$$

An m -server protocol for INDEX_N is called an m -server private information retrieval (PIR) protocol for N -sized databases. In the information-theoretic setting, the most communication-efficient passively secure 3-server PIR protocol was given by [10] and in the computational setting, the passively secure 2-server PIR protocol was given by [14] assuming the existence of one-way functions. Although the original protocols in [10, 14] assume $t = 1$, the corruption threshold t can be amplified by using the technique in [7] as pointed out in [32]. More specifically, the following propositions hold.

Proposition 1. *There exists a passively t -secure 1-round 3^t -server protocol Π for INDEX_N such that*

- $\text{Comm}(\Pi) = \exp(O(\sqrt{\log N \log \log N})) \cdot t3^t = N^{o(1)} \cdot 2^{O(t)}$;
- $\text{c-Comp}(\Pi) = \exp(O(\sqrt{\log N \log \log N})) \cdot t3^t = N^{o(1)} \cdot 2^{O(t)}$;
- $\text{s-Comp}(\Pi) = N^2 \cdot \exp(O(\sqrt{\log N \log \log N})) \cdot 2^t = N^{2+o(1)} \cdot 2^t$.

Note that the above protocol satisfies the efficiency requirement (1) since $\text{Comm}(\Pi)$ and $\text{c-Comp}(\Pi)$ are sub-polynomial (i.e., less than any polynomial) in the description length N of elements of $P_\lambda = \{0, 1\}^N$.

Proposition 2. *Assume a pseudorandom generator $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2(\lambda+1)}$. There exists a computationally passively t -secure 1-round 2^t -server protocol Π for INDEX_N such that*

- $\text{Comm}(\Pi) = O(\log N \cdot \lambda \cdot t2^t)$;
- $\text{c-Comp}(\Pi)$ is $O(\log N \cdot t2^t)$ invocations of G ;
- $\text{s-Comp}(\Pi)$ is $O(N^2 \log N \cdot t)$ invocations of G .

Remark 2. Dvir and Gopi [27] devised a technique to optimize the 3-server protocol in [28] and obtained a 2-server PIR protocol with $N^{o(1)}$ communication. However, since the answer length is not constant, the passively t -secure protocol obtained by applying the amplification technique of [7] has larger communication complexity $\exp(O(t\sqrt{\log N \log \log N}))$ and does not satisfy the efficiency requirement (1).

Bounded-Degree Polynomials. Let $N = N(\lambda)$, $D = D(\lambda)$ and $M = M(\lambda)$ be polynomial functions. We define $\text{POLY}_{N,D,M}(R) = (F_\lambda)_{\lambda \in \mathbb{N}}$ as a sequence of functions such that $F_\lambda(p, \mathbf{x}) = p(\mathbf{x})$ for any N -variate polynomial p over a ring R with degree D and number of monomials M , and for any $\mathbf{x} \in R^N$. The following is implicit in [43].

Proposition 3. *Let $N, D, M \in \text{poly}(\lambda)$. Let R be a ring such that for any $a \in \{1, 2, \dots, m-1\}$, an element $a \cdot 1_R$ has an inverse in R , where 1_R is the multiplicative identity of R . Suppose that $m > Dt/2$. Then, there exists a passively t -secure 1-round m -server protocol Π for $\text{POLY}_{N,D,M}(R)$ such that*

- $\text{Comm}(\Pi)$ is $O(Nm)$ ring elements;
- $\text{c-Comp}(\Pi)$ is $O(Ntm)$ ring operations;
- $\text{s-Comp}(\Pi)$ is $O(NMD)$ ring operations.

Since the description length of a polynomial with M monomials is $\tilde{O}(MD \log |R|)$, the above protocol satisfies the efficiency requirement (1) if $MD = \omega(N)$.

Ishai, Lai and Malavolta [36] showed that assuming homomorphic encryption for degree- d polynomials, the number of servers in Proposition 3 can be decreased by a factor of d .

Proposition 4. *Let $d = O(1)$ and R be a ring such that for any $a \in \{1, 2, \dots, \max\{d, m - 1\}\}$, an element $a \cdot 1_R$ has an inverse in R , where 1_R is the multiplicative identity of R . Assume a homomorphic encryption scheme HE for degree- d polynomials over R . Let $M, N \in \text{poly}(\lambda)$ and $D = O(1)$. Suppose that $m > Dt/(d+1)$. Then there exists a computationally passively t -secure 1-round m -server protocol Π for $\text{POLY}_{N,D,M}(R)$ such that*

- $\text{Comm}(\Pi) = O(Nm \cdot \ell_{\text{ct}})$, where ℓ_{ct} is the description length of ciphertexts of HE;
- $\text{c-Comp}(\Pi) = O((Nt \cdot \tau_{\text{Enc}} + \tau_{\text{Dec}})m)$, where τ_{Enc} and τ_{Dec} are the running time of the encryption and decryption algorithms of HE, respectively;
- $\text{s-Comp}(\Pi) = O(MN \cdot \tau_{\text{Eval}})$, where τ_{Eval} is the running time of the evaluation algorithm of HE per operation.

Note that we have $\max\{d, m - 1\} = \text{poly}(\lambda)$ since $d = O(1)$ and $m = \text{poly}(\lambda)$. On the other hand, homomorphic encryption schemes mentioned in [36] assume that R is a prime field of size q or a ring of integers modulo $n = q_1 q_2$ for exponentially large primes q, q_1, q_2 . In these cases, $a \cdot 1_R$ has an inverse in R if $a \in \{1, 2, \dots, \max\{d, m - 1\}\}$.

Under the sparse Learning Parity with Noise (LPN) assumption over a field \mathbb{F}_q , Dao et al. [24] proposed a passively t -secure $(t+1)$ -server protocol for polynomials of degree $D = O(\log \lambda / \log \log \lambda)$. Although the original protocol does not have sublinear-size upload cost when evaluating a single polynomial, it can be seen that the upload cost is amortized if sufficiently many polynomials are evaluated on the same input. Specifically, let $N = N(\lambda)$, $D = D(\lambda)$, $M = M(\lambda)$, and $L = L(\lambda)$ be polynomial functions. We define $\text{POLY}_{N,D,M}^L(R) = (F_\lambda)_{\lambda \in \mathbb{N}}$ as a sequence of functions such that $F_\lambda((p_1, \dots, p_L), \mathbf{x}) = (p_1(\mathbf{x}), \dots, p_L(\mathbf{x}))$ for any N -variate polynomials p_1, \dots, p_L over a ring R with degree D and number of monomials M , and for any $\mathbf{x} \in R^N$.

Proposition 5. *Assume that the (δ, q) -sLPN assumption holds for a constant $0 \leq \delta \leq 1$ and a sequence $q = (q(\lambda))_{\lambda \in \mathbb{N}}$ of prime powers that are computable in polynomial time in λ . Let $L, M, N \in \text{poly}(\lambda)$ and $D = O(\log \lambda / \log \log \lambda)$. Then, there exists a computationally passively t -secure 1-round $(t+1)$ -server protocol Π for $\text{POLY}_{N,D,M}^L(\mathbb{F}_q)$ such that*

- $\text{Comm}(\Pi) = \tilde{O}((M^{2/\delta}N + L)(\log q)m\lambda)$;
- $\text{c-Comp}(\Pi) = \tilde{O}((M^{2/\delta}N + L)(\log q)m\lambda)$;
- $\text{s-Comp}(\Pi) = \tilde{O}(M^{1/\delta+1}L(\log q)\lambda)$.

Note that the description length of L polynomials each with M monomials is $\tilde{O}(ML \log q)$ if the degree is $D = o(\log \lambda)$. Thus, if $L = \omega(M^{2/\delta-1})$, the above protocol satisfies the efficiency requirement (1). See [24] for the details including the definition of the sparse LPN assumption.

Constant-Depth Circuits. We consider Boolean circuits of constant depth with unbounded fan-in and fan-out. Formally, a Boolean circuit C is a labelled directed acyclic graph. The nodes with no incoming edges are labelled with input variables, their negations, or constants. The other nodes are called gates and are labelled with one of operators in $\{\text{AND}, \text{OR}, \text{NOT}\}$. Nodes with no outgoing edges are called output nodes. We only consider a circuit with a single output node. The size of a circuit is the number of edges and its depth is the length of the longest path from an input node to the output node. We define the output of C on input x , which we denote by $C(x)$, as the value of the output node after input values proceed through a sequence of gates.

Let $N = N(\lambda)$, $D = D(\lambda)$ and $M = M(\lambda)$ be polynomial functions. We define $\text{CIRC}_{N,D,M} = (F_\lambda)_{\lambda \in \mathbb{N}}$ as a sequence of functions such that $F_\lambda(C, x) = C(x)$ for any Boolean circuit C with N input variables, depth D and size M , and for any N -bit string x .

Proposition 6. *Let $N, M \in \text{poly}(\lambda)$ and $D = O(1)$. Suppose that $m \geq (\log M + 3)^{D-1}t/2$. Then, there exists a passively t -secure 1-round m -server protocol Π for $\text{CIRC}_{N,D,M}$ such that*

- $\text{Comm}(\Pi) = O((\log M)^{D-1}N(\log N)\lambda m)$;
- $\text{c-Comp}(\Pi) = O((\log M)^{D-1}N(\log N)tm + (\log M)^2\lambda m)$;
- $\text{s-Comp}(\Pi) = O(M(\log M)N\lambda)$.

The protocol is efficient since $\text{Comm}(\Pi)$ and $\text{c-Comp}(\Pi)$ are linear in N and polylogarithmic in the size M of circuits, omitting factors in λ and m .

4 Interactive Actively Secure Protocols

In this section, we show our compiler from one-round passively t -secure k -server protocols to $O(m^2)$ -round actively t -secure m -server protocols such that $m \geq k + t$. To this end, we introduce a notion of *conflict-finding* protocols, which is an intermediate notion between passively secure and actively secure protocols. We show a generic compiler from conflict-finding to actively secure protocols in Sect. 4.3 and then show a generic compiler from passively secure to conflict-finding protocols.

4.1 Graph Theory

To begin with, we recall the standard terminology of graph theory (see [37, Chapter 2] for instance). A (simple and undirected) graph \mathcal{G} is a pair (V, E) , where V is a set of vertices and E is a set of edges $(i, j) \in V \times V$. Throughout the paper, we only consider the cases where V is either $[m]$ or a subset of $[m]$. Thus we may assume that V is a totally ordered set. The total order on V naturally induces a lexicographic order on E , which is also a total order on E . A graph \mathcal{G} is called *connected* if there is a path between each pair of vertices. It is a standard result that there is a deterministic algorithm \mathcal{D} which decomposes \mathcal{G} into connected components in time $O(|V| + |E|)$ [37]. For $S \subseteq V$, we denote by $\mathcal{G}[S]$ the induced subgraph, i.e., the graph whose vertex set is S and whose edge set consists of the edges in E that have both endpoints in S .

We show a deterministic algorithm \mathcal{C}'_k such that for any connected graph $\mathcal{G} = (V, E)$ with at least k vertices, $\mathcal{C}'_k(\mathcal{G})$ outputs a subset $S \subseteq V$ such that $|S| = k$ and $\mathcal{G}[S]$ is connected. First, \mathcal{C}'_k chooses the minimum node s of V with respect to the total order on V . Secondly, \mathcal{C}'_k runs the “textbook” depth-first search algorithm [37] starting at the vertex s , except that it stops searching if it visits k vertices. Finally, \mathcal{C}'_k outputs the set S of all vertices it visited so far. By definition, S is of size k . Since any pair of vertices in S are connected via s , $\mathcal{G}[S]$ is connected. The running time of \mathcal{C}'_k is $O(|V| + |E|)$.

Next, we show a deterministic algorithm \mathcal{C}_k such that for any graph $\mathcal{G} = (V, E)$, if \mathcal{G} contains a connected component of size at least k , $\mathcal{C}_k(\mathcal{G})$ outputs a subset $S \subseteq V$ of size k such that $\mathcal{G}[S]$ is connected, and otherwise, it outputs the empty set \emptyset . First, \mathcal{C}_k lists all the connected components of \mathcal{G} , $(\mathcal{G}_1, \dots, \mathcal{G}_q) \leftarrow \mathcal{D}(\mathcal{G})$. Secondly, \mathcal{C}_k lets q_{\min} be the minimum index q such that \mathcal{G}_q has at least k vertices. If no component has k vertices, \mathcal{C}_k outputs \emptyset . Otherwise, \mathcal{C}_k outputs $S \leftarrow \mathcal{C}'_k(\mathcal{G}_{q_{\min}})$. The correctness of \mathcal{C}_k immediately follows from those of \mathcal{D} and \mathcal{C}'_k . The running time of \mathcal{C}_k is $O(|V| + |E|)$.

Finally, we show a trivial but frequently-used algorithm \mathcal{E} , which takes as input a graph $\mathcal{G} = (V, E)$ and a pair of disjoint non-empty subsets $G_0, G_1 \subseteq V$, and outputs the minimum edge $e = (i, j) \in E$ (with respect to the total order on E) such that $i \in G_0$ and $j \in G_1$, or $j \in G_0$ and $i \in G_1$. The running time of \mathcal{E} is $O(|E|)$.

4.2 Formalization of Conflict-Finding Protocols

Roughly speaking, in a conflict-finding protocol, a client obtains (y, z) , where y is the main output (supposed to be $F(p, x)$) and z is an auxiliary string. The string z is either **output**, **failure**, or a non-trivial partition (G_0, G_1) of the set of servers⁷. The security requirements are:

Soundness. The probability that $z = \text{output}$ and $y \neq F(p, x)$ is negligible, and the probability that the protocol outputs $z = \text{failure}$ is also negligible;

⁷ We say that a partition (G_0, G_1) is *non-trivial* if $G_0 \neq \emptyset$ and $G_1 \neq \emptyset$.

Conflict-Finding. If z is a non-trivial partition (G_0, G_1) of the set of servers, then one of G_0 or G_1 contains all honest servers (and hence the other group consists of malicious servers only);

Privacy. An adversary should not learn a client's input x even if she knows z .

Intuitively, the conflict-finding property ensures that a client learns a subset of malicious servers only, which allows him to find a pair of servers such that at least one of them is malicious. We require the privacy should hold even if z is leaked, in order for an adversary not to learn additional information from a set of servers the client removes. Below, we show formal definitions.

Definition 3. We say that $\Pi = (\text{Query}, \text{Answer}, \text{Output})$ is an ℓ -round t -conflict-finding m -server protocol for $\mathcal{F} = (F_\lambda : P_\lambda \times X_\lambda \rightarrow Y_\lambda)_{\lambda \in \mathbb{N}}$ if it satisfies the following properties:

Syntax. The syntax of **Query** and **Answer** is the same as that of Π as an ℓ -round m -server protocol for \mathcal{F} (Definition 1). The algorithm **Output** takes a state $\text{st}^{(\ell)}$ and answer $(\text{ans}_i^{(\ell)})_{i \in [m]}$ in round ℓ as input, and outputs (y, z) such that (1) $y \in Y_\lambda$ and $z = \text{output}$, (2) $y = \perp$ and $z = (G_0, G_1)$, which is a non-trivial partition of $[m]$, or (3) $y = \perp$ and $z = \text{failure}$. We call the first (resp. second) component of the output of **Output** the y -output (resp. z -output).

Correctness. There exists a negligible function $\text{negl}(\lambda)$ such that for any $\lambda \in \mathbb{N}$ and any $(p, x) \in P_\lambda \times X_\lambda$, it holds that

$$\Pr \left[(y, z) \leftarrow \text{Output}(1^\lambda, \text{st}^{(\ell)}, (\text{ans}_i^{(\ell)})_{i \in [m]}) : y = F_\lambda(p, x) \right] \geq 1 - \text{negl}(\lambda),$$

where

$$\begin{aligned} ((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)}) &\leftarrow \text{Query}(1^\lambda, x, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]}), \\ \text{ans}_i^{(j)} &\leftarrow \text{Answer}(1^\lambda, p, \text{que}_i^{(j)}) \end{aligned}$$

for all $j \in [\ell]$ and $i \in [m]$.

Soundness. There exists a negligible function $\text{negl}(\lambda)$ such that for any stateful algorithm \mathcal{A} and any $\lambda \in \mathbb{N}$,

$$\Pr [\text{Sound}_{\Pi, \mathcal{A}}(\lambda) = 1] < \text{negl}(\lambda), \quad (2)$$

where $\text{Sound}_{\Pi, \mathcal{A}}(\lambda)$ is the output of the following experiment:

1. $(x, p, T) \leftarrow \mathcal{A}(1^\lambda)$, where $x \in X_\lambda$, $p \in P_\lambda$ and $T \subseteq [m]$ is of size at most t .
2. For each $j = 1, 2, \dots, \ell$,
 - (a) Let $((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)}) \leftarrow \text{Query}(1^\lambda, x, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]})$ and give $(\text{que}_i^{(j)})_{i \in T}$ to \mathcal{A} .
 - (b) \mathcal{A} outputs $(\text{ans}_i^{(j)})_{i \in T}$.
3. Let $(y, z) \leftarrow \text{Output}(1^\lambda, \text{st}^{(\ell)}, (\text{ans}_i^{(\ell)})_{i \in [m]})$.
4. Return 1 if $y \in Y_\lambda \setminus \{F_\lambda(p, x)\}$ and $z = \text{output}$, or $y = \perp$ and $z = \text{failure}$. Otherwise return 0.

Conflict-Finding. For any stateful algorithm \mathcal{A} and any $\lambda \in \mathbb{N}$,

$$\Pr[\text{CF}_{\Pi, \mathcal{A}}(\lambda) = 1] = 0,$$

where $\text{CF}_{\Pi, \mathcal{A}}(\lambda)$ is the output of the following experiment:

1. $(x, p, T) \leftarrow \mathcal{A}(1^\lambda)$, where $x \in X_\lambda$, $p \in P_\lambda$ and $T \subseteq [m]$ is of size at most t .
2. For each $j = 1, 2, \dots, \ell$,
 - (a) Let $((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)}) \leftarrow \text{Query}(1^\lambda, x, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]})$ and give $(\text{que}_i^{(j)})_{i \in T}$ to \mathcal{A} .
 - (b) \mathcal{A} outputs $(\text{ans}_i^{(j)})_{i \in T}$.
3. Let $(y, z) \leftarrow \text{Output}(1^\lambda, \text{st}^{(\ell)}, (\text{ans}_i^{(\ell)})_{i \in [m]})$.
4. Return 1 if $z = (G_0, G_1)$, $G_0 \not\subseteq T$ and $G_1 \not\subseteq T$. Otherwise return 0.

Privacy. There exists a negligible function $\text{negl}(\lambda)$ such that for any stateful algorithm \mathcal{A} and any $\lambda \in \mathbb{N}$,

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{CF}}(\lambda) := \left| \Pr \left[\text{Priv}_{\Pi, \mathcal{A}}^{\text{CF}, 0}(\lambda) = 0 \right] - \Pr \left[\text{Priv}_{\Pi, \mathcal{A}}^{\text{CF}, 1}(\lambda) = 0 \right] \right| < \text{negl}(\lambda),$$

where for $b \in \{0, 1\}$, $\text{Priv}_{\Pi, \mathcal{A}}^{\text{CF}, b}(\lambda)$ is the output b' of \mathcal{A} in the following experiment:

1. $(x_0, x_1, p, T) \leftarrow \mathcal{A}(1^\lambda)$, where $x_0, x_1 \in X_\lambda$, $p \in P_\lambda$ and $T \subseteq [m]$ is of size at most t .
2. For each $j = 1, 2, \dots, \ell$,
 - (a) Let $((\text{que}_i^{(j)})_{i \in [m]}, \text{st}^{(j)}) \leftarrow \text{Query}(1^\lambda, x_b, \text{st}^{(j-1)}, (\text{ans}_i^{(j-1)})_{i \in [m]})$ and give $(\text{que}_i^{(j)})_{i \in T}$ to \mathcal{A} .
 - (b) \mathcal{A} outputs $(\text{ans}_i^{(j)})_{i \in T}$.
3. Let $(y, z) \leftarrow \text{Output}(1^\lambda, \text{st}^{(\ell)}, (\text{ans}_i^{(\ell)})_{i \in [m]})$ and give z to \mathcal{A} .
4. \mathcal{A} outputs a bit $b' \in \{0, 1\}$.

For a (possibly non-negligible) function $\epsilon(\lambda)$, we define a weaker notion of a ϵ -sound t -conflict-finding protocol Π as the one satisfying the requirements in Definition 3 except that the condition (2) is replaced with

$$\Pr[\text{Sound}_{\Pi, \mathcal{A}}(\lambda) = 1] < \epsilon.$$

We say that Π is *computationally t -conflict-finding* if it satisfies the above requirements for PPT adversaries \mathcal{A} .

4.3 Compiler from Conflict-Finding to Actively Secure Protocols

We construct an actively t -secure m -server protocol from a t -conflict-finding $(m - t)$ -server protocol. We give a sketch here and defer the formal proof to the full version.

Theorem 1. *Suppose that there exists an ℓ -round (resp. computationally) t -conflict-finding k -server protocol Π_{CF} for $\mathcal{F} = (F_\lambda : P_\lambda \times X_\lambda \rightarrow Y_\lambda)_{\lambda \in \mathbb{N}}$. If $m \geq t + k$, there exists an $O(\ell m^2)$ -round (resp. computationally) actively t -secure m -server protocol Π for \mathcal{F} such that*

- $\text{Comm}(\Pi) = O(m^2 \cdot \text{Comm}(\Pi_{\text{CF}}))$;
- $\text{c-Comp}(\Pi) = O(m^2 \cdot \text{c-Comp}(\Pi_{\text{CF}}) + m^4)$;
- $\text{s-Comp}(\Pi) = O(m^2 \cdot \text{s-Comp}(\Pi_{\text{CF}}))$.

Proof (sketch). Define $N := \binom{m}{2} - \binom{m-t}{2} + 1 = O(m^2)$. Let V be the set of all m servers and $\mathcal{G}^{(1)}$ be the complete graph on V . Consider the following protocol Π : For each $j = 1, 2, \dots, N$,

1. The client C finds a k -sized subset $S^{(j)}$ of V such that $\mathcal{G}^{(j)}[S^{(j)}]$ is connected, based on the algorithm \mathcal{C}_k in Sect. 4.1.
2. C executes the conflict-finding protocol Π_{CF} with k servers in $S^{(j)}$, and obtain an output $(y^{(j)}, z^{(j)})$.
3. If $z^{(j)} = \text{output}$, then C outputs the y -output $y^{(j)}$.
4. If $z^{(j)} = \text{failure}$, then C outputs any default value y_0 .
5. If $z^{(j)}$ is a non-trivial partition $(G_0^{(j)}, G_1^{(j)})$ of $S^{(j)}$, then C does the following:
 - (a) Find an edge $e^{(j)}$ of $\mathcal{G}^{(j)}$ crossing the partition $(G_0^{(j)}, G_1^{(j)})$ based on the algorithm \mathcal{E} in Sect. 4.1. Such an edge exists since $\mathcal{G}^{(j)}[S^{(j)}]$ is connected.
 - (b) Let $\mathcal{G}^{(j+1)}$ be a graph obtained by removing $e^{(j)}$ from $\mathcal{G}^{(j)}$.
 - (c) Go back to Step 1.

Privacy. An adversary corrupting a set T of at most t servers cannot learn a client's input from interaction at Step 2 due to the fact that $|T \cap S^{(j)}| \leq |T| \leq t$ and the privacy of Π_{CF} . The adversary can also see a sequence of graphs $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \dots, \mathcal{G}^{(N)}$ but as shown at Step 5, the sequence is determined only by a sequence of z -outputs $z^{(1)}, z^{(2)}, \dots, z^{(N)}$. Since Π_{CF} guarantees privacy even if z -outputs are leaked, she learns no additional information.

Byzantine-Robustness. The client C outputs an incorrect result only if one of the following events occurs: (1) $z^{(j)} = \text{output}$ and $y^{(j)}$ is an incorrect result for some $j \in [N]$, (2) $z^{(j)} = \text{failure}$ for some $j \in [N]$, or (3) $z^{(j)}$ is a non-trivial partition for all $j \in [N]$. It follows from the soundness of Π_{CF} that the first and second cases occur only with negligible probability.

We argue that the third case never occurs. Assume otherwise, then for all j , the z -output $z^{(j)}$ of the j -th iteration is a non-trivial partition $(G_0^{(j)}, G_1^{(j)})$ of $S^{(j)}$. Since the conflict-finding property of Π_{CF} ensures that either $G_0^{(j)}$ or $G_1^{(j)}$ includes the set of honest servers $H := [m] \setminus T$, the removed edge $e^{(j)} = (i_1, i_2)$ satisfies $i_1 \in T$ or $i_2 \in T$ and hence the subgraph $\mathcal{G}^{(j)}[H]$ is a complete graph for all j . Since N is larger than the total number $N' = \binom{m}{2} - \binom{m-t}{2}$ of unordered pairs (i_1, i_2) such that $i_1 \in T$ or $i_2 \in T$, $\mathcal{G}^{(N')}$ has no edge $e = (i_1, i_2)$ such that $i_1 \in T$ or $i_2 \in T$. Therefore, a set of servers $S^{(N')}$ involved in the N' -th iteration is a subset of H since $k \leq m - t \leq |H|$. We have assumed that $z^{(N')}$ is a non-trivial partition $(G_0^{(N')}, G_1^{(N')})$ of $S^{(N')}$ but the conflict-finding property ensures that $H \subseteq G_0^{(N')}$ or $H \subseteq G_1^{(N')}$, which is contradiction. \square

4.4 Compiler from Passively Secure to Conflict-Finding Protocols

First, we show a basic construction of ϵ -sound conflict-finding protocols for non-negligible ϵ . We give a sketch here and defer the formal proof to the full version.

Proposition 7. *Let Π be a 1-round (resp. computationally) passively t -secure m -server protocol for $\mathcal{F} = (F_\lambda : P_\lambda \times X_\lambda \rightarrow Y_\lambda)_{\lambda \in \mathbb{N}}$. Let $M = \text{poly}(\lambda)$. Then, there exists a 2-round (resp. computationally) ϵ -sound t -conflict-finding m -server protocol Π' for \mathcal{F} such that*

- $\text{Comm}(\Pi') = O(mM \cdot \text{Comm}(\Pi))$;
- $\text{c-Comp}(\Pi') = O(m^2M \cdot \text{c-Comp}(\Pi))$;
- $\text{s-Comp}(\Pi') = O(mM \cdot \text{s-Comp}(\Pi))$;

where $\epsilon = m/M + \text{negl}(\lambda)$ for some negligible function $\text{negl}(\lambda)$.

Proof (sketch). Consider the following protocol Π' :

First Round

1. The client C chooses μ_* uniformly at random from $[M]$.
2. For all $\mu \in [M]$, C computes queries $(\text{que}_1^{(\mu)}, \dots, \text{que}_m^{(\mu)})$ of Π on his true input x if $\mu = \mu_*$, and on a default input x_{def} otherwise.
3. C sends the queries $(\text{que}_i^{(\mu)})_{\mu \in [M]}$ to each server S_i as usual, who returns answers $(\text{ans}_i^{(\mu)})_{\mu \in [M]}$ to them.

Second Round

1. C sends all the queries $(\text{que}_k^{(\mu)})_{k \in [m], \mu \neq \mu_*}$ for the default input x_{def} to all servers.
2. For all $k \in [m]$ and $\mu \in [M] \setminus \{\mu_*\}$, each server S_i returns an answer $\text{ans}_k^{(\mu)}(i)$ as S_k would answer to $\text{que}_k^{(\mu)}$.

To obtain an output, C defines $v_i = (\text{ans}_k^{(\mu)}(i))_{k \in [m], \mu \neq \mu_*}$ for all $i \in [m]$. For simplicity, we here assume that $\text{ans}_i^{(\mu)}(i) = \text{ans}_i^{(\mu)}$ for all $i \in [m]$. This is because otherwise, it means that a server S_i returns different answers in the first and second rounds and hence S_i is immediately found malicious. The client C partitions the set of servers into equivalence classes G'_0, \dots, G'_ℓ under the equivalence relation defined as: $i \sim j \stackrel{\text{def}}{\iff} v_i = v_j$. If $\ell = 0$ (i.e., all servers belong to the same equivalence class), then he runs the output algorithm of Π on the answers $(\text{ans}_1^{(\mu_*)}, \dots, \text{ans}_m^{(\mu_*)})$ to the queries for his true input. He then outputs the result y along with $z = \text{output}$. If $\ell \geq 1$, then he outputs $y = \perp$ and $z = (G_0, G_1)$, where $G_0 = G'_0$ and $G_1 = G'_1 \cup \dots \cup G'_\ell$.

Conflict-Finding. Let T be a set of corrupted servers. Since honest servers $i, j \notin T$ always return the same answer to the same query, we have that $\text{ans}_k^{(\mu)}(i) = \text{ans}_k^{(\mu)}(j)$ for all $k \in [m]$ and $\mu \in [M] \setminus \{\mu_*\}$, and hence $v_i = v_j$. Therefore, the set of honest servers is contained in an equivalence class and it holds that $\overline{T} \subseteq G_0 = \overline{G_1}$ or $\overline{T} \subseteq G_1 = \overline{G_0}$.

Soundness. In the first place, the protocol Π' never outputs $z = \text{failure}$. Assume that Π' outputs $z = \text{output}$. Then, all servers belong to the same equivalence class, which implies that $v_i = v_j$ for any $i, j \in [m]$. To let the client accept an incorrect result, an adversary needs to let at least one corrupted server S_i submit an incorrect answer exactly to the query $\text{que}_i^{\langle \mu_* \rangle}$ for the client's true input. (This is because if a corrupted server submits incorrect $\text{ans}_i^{\langle \mu \rangle}$ for some $\mu \neq \mu_*$, then it is detected when compared with an answer $\text{ans}_i^{\langle \mu \rangle}(j)$ from an honest server $j \notin T$.) However, the adversary cannot learn which query encodes the client's true input due to the privacy of Π . Therefore, her best possible strategy is to guess μ_* uniformly at random, which succeeds only with probability $1/M$. The union bound implies that the error probability is at most m/M .

Privacy. Since M queries are generated independently, an adversary learns no information on the client's input x in the first round. The queries revealed in the second round are the ones for a default input x_{def} , which is independent of x , and hence the adversary learns no additional information. The privacy holds even if the z -output z is leaked, since z is determined only by $(v_i)_{i \in [m]}$, which can be simulated from information that the adversary learns up to the second round. \square

Next, we show that the error probability of the basic construction can be made negligible by parallel execution. The proof is deferred to the full version.

Theorem 2. *Let Π be a 1-round (resp. computationally) passively t -secure m -server protocol for $\mathcal{F} = (F_\lambda : P_\lambda \times X_\lambda \rightarrow Y_\lambda)_{\lambda \in \mathbb{N}}$. Then there exists a 2-round (resp. computationally) t -conflict-finding m -server protocol Π_{CF} for \mathcal{F} such that*

- $\text{Comm}(\Pi_{\text{CF}}) = O(m^2 \lambda \cdot \text{Comm}(\Pi));$
- $\text{c-Comp}(\Pi_{\text{CF}}) = O(m^3 \lambda \cdot \text{c-Comp}(\Pi));$
- $\text{s-Comp}(\Pi_{\text{CF}}) = O(m^2 \lambda \cdot \text{s-Comp}(\Pi)).$

Finally, by combining Theorems 1 and 2, we obtain our generic construction of an $O(m^2)$ -round actively t -secure m -server protocol from any 1-round passively t -secure k -server protocol for $k \leq m - t$.

Theorem 3. *Suppose that $m > 2t$. Let $k \leq m - t$ and Π be a 1-round (resp. computationally) passively t -secure k -server protocol for \mathcal{F} . Then there exists an $O(m^2)$ -round (resp. computationally) actively t -secure m -server protocol Π' for \mathcal{F} such that*

- $\text{Comm}(\Pi') = O(m^4 \lambda \cdot \text{Comm}(\Pi));$
- $\text{c-Comp}(\Pi') = O(m^5 \lambda \cdot \text{c-Comp}(\Pi));$
- $\text{s-Comp}(\Pi') = O(m^4 \lambda \cdot \text{s-Comp}(\Pi)).$

4.5 Instantiations

By applying our compiler in Theorem 3 to the passively secure protocols in Propositions 1 and 2, we obtain actively secure protocols for INDEX_N .

Corollary 1. *Suppose that $m \geq 3^t + t$. Let $N \in \text{poly}(\lambda)$. Then, there exists an actively t -secure $O(m^2)$ -round m -server protocol Π for INDEX_N such that*

- $\text{Comm}(\Pi) = \exp(O(\sqrt{\log N \log \log N})) \cdot t3^t m^4 \lambda$;
- $\text{c-Comp}(\Pi) = \exp(O(\sqrt{\log N \log \log N})) \cdot t3^t m^5 \lambda$;
- $\text{s-Comp}(\Pi) = N^2 \cdot \exp(O(\sqrt{\log N \log \log N})) \cdot 2^t m^4 \lambda$.

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = N^{o(1)} \cdot 2^{O(t)} \lambda$.

Corollary 2. *Assume a pseudorandom generator $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2(\lambda+1)}$. Suppose that $m \geq 2^t + t$. Let $N \in \text{poly}(\lambda)$. Then, there exists a computationally actively t -secure 1-round m -server protocol Π for INDEX_N such that*

- $\text{Comm}(\Pi) = O(\log N \cdot \lambda^2 \cdot t2^t m^4)$;
- $\text{c-Comp}(\Pi)$ is $O(\log N \cdot t2^t m^5)$ invocations of G ;
- $\text{s-Comp}(\Pi)$ is $O(N^2 \log N \cdot tm^4)$ invocations of G .

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = \log N \cdot 2^{O(t)} \cdot \text{poly}(\lambda)$.

By applying Theorem 3 to Proposition 3, we obtain an actively secure protocol for multivariate polynomials.

Corollary 3. *Let $N, D, M \in \text{poly}(\lambda)$. Let R be a ring such that for any $a \in \{1, 2, \dots, m-1\}$, an element $a \cdot 1_R$ has an inverse in R . Suppose that*

$$m > \left(\frac{D}{2} + 1 \right) t.$$

Then, there exists an actively t -secure 1-round m -server protocol Π for $\text{POLY}_{N,D,M}(R)$ such that

- $\text{Comm}(\Pi) = O(Nm^4 \lambda)$ ring elements;
- $\text{c-Comp}(\Pi) = O(Ntm^6 \lambda)$ ring operations;
- $\text{s-Comp}(\Pi) = O(NMDm^4 \lambda)$ ring operations.

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = N \cdot \text{poly}(m, \lambda)$.

By applying Theorem 3 to Proposition 4, we can reduce the required number of servers by a factor of d assuming homomorphic encryption for degree- d polynomials.

Corollary 4. *Let $d = O(1)$ and R be a ring such that for any $a \in \{1, 2, \dots, \max\{d, m-1\}\}$, an element $a \cdot 1_R$ has an inverse in R , where 1_R is the multiplicative identity of R . Assume a homomorphic encryption scheme HE for degree- d polynomials over R . Suppose that*

$$m > \left(\frac{D}{d+1} + 1 \right) t$$

Let $M, N \in \text{poly}(\lambda)$ and $D = O(1)$. Then there exists a computationally actively t -secure $O(m^2)$ -round m -server protocol Π for $\text{POLY}_{N,D,M}(R)$ such that

- $\text{Comm}(\Pi) = O(Nm^5\lambda \cdot \ell_{\text{ct}})$, where ℓ_{ct} is the description length of ciphertexts of HE;
- $\text{c-Comp}(\Pi) = O((Nt \cdot \tau_{\text{Enc}} + \tau_{\text{Dec}})m^6\lambda)$, where τ_{Dec} and τ_{Enc} are the running time of the decryption and encryption algorithms of HE, respectively;
- $\text{s-Comp}(\Pi) = O(M \cdot m^4\lambda\tau_{\text{Eval}})$, where τ_{Eval} is the running time per operation of the evaluation algorithm of HE.

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = N \cdot \text{poly}(m, \lambda)$.

By applying Theorem 3 to Proposition 5, we obtain an actively t -secure protocol for polynomials achieving the minimum number of servers $2t + 1$.

Corollary 5. *Suppose that $m = 2t + 1$. Assume that the (δ, q) -sLPN assumption holds for a constant $0 \leq \delta \leq 1$ and a sequence $q = (q(\lambda))_{\lambda \in \mathbb{N}}$ of prime powers that are computable in polynomial time in λ . Let $L, M, N \in \text{poly}(\lambda)$ and $D = O(\log \lambda / \log \log \lambda)$. Then, there exists a computationally actively t -secure $O(m^2)$ -round m -server protocol Π for $\text{POLY}_{N, D, M}^L(\mathbb{F}_q)$ such that*

- $\text{Comm}(\Pi) = \tilde{O}((M^{2/\delta}N + L)(\log q)m^5\lambda^2)$;
- $\text{c-Comp}(\Pi) = \tilde{O}((M^{2/\delta}N + L)(\log q)m^6\lambda^2)$;
- $\text{s-Comp}(\Pi) = \tilde{O}(M^{1/\delta+1}L(\log q)m^4\lambda^2)$.

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = (M^{2/\delta}N + L) \log q \cdot \text{poly}(m, \lambda)$.

Finally, by applying Theorem 3 to Proposition 6, we obtain an actively secure protocol for constant-depth circuits.

Corollary 6. *Let $N, M \in \text{poly}(\lambda)$ and $D = O(1)$. Suppose that*

$$m \geq \left(\frac{(\log M + 3)^{D-1}}{2} + 1 \right) t.$$

Then, there exists an actively t -secure $O(m^2)$ -round m -server protocol Π for $\text{CIRC}_{N, D, M}$ such that

- $\text{Comm}(\Pi) = O((\log M)^{D-1}N(\log N)\lambda^2m^5)$;
- $\text{c-Comp}(\Pi) = O((\log M)^{D-1}N(\log N)\lambda tm^6 + (\log M)^2\lambda^2m^6)$;
- $\text{s-Comp}(\Pi) = O(M(\log M)Nm^4\lambda^2)$.

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = N \cdot \text{poly}(m, \lambda)$.

5 Non-interactive Actively Secure Protocols

In this section, we show our compiler from one-round passively t -secure k -server protocols to one-round actively t -secure m -server protocols such that $m = O(k \log k) + 2t$. To this end, we introduce a novel combinatorial object of *locally surjective map families*, which is a variant of perfect hash families with a stronger property. We show a probabilistic construction of such families in Sect. 5.1 and then show a generic compiler from passively secure to actively secure protocols in Sect. 5.2.

5.1 Locally Surjective Map Family

We show the formal definition of locally surjective map families.

Definition 4. Let $m, h, k \in \mathbb{N}$ and \mathcal{L} be a family of maps from $[m]$ to $[k]$. We call \mathcal{L} an (m, h, k) -locally surjective map family if $|A_H| > |\mathcal{L}|/2$ for any $H \in \binom{[m]}{h}$, where $A_H = \{f \in \mathcal{L} : f(H) = [k]\}$.

A locally surjective map family satisfies a stronger property than a nearly perfect hash family \mathcal{L}' introduced in [11], which assumes that for any $H \in \binom{[m]}{h}$, there exists at least one map $f \in \mathcal{L}'$ such that $f(H) = [k]$.

We show a probabilistic construction of an (m, h, k) -locally surjective map family of size $O(m)$ for $k = O(h/\log h)$. The formal proof is deferred to the full version.

Proposition 8. Let $m, h, k \in \mathbb{N}$ be such that $h \geq 15$, $m \geq 15$ and $k \leq h/(\gamma \ln h)$, where $\gamma := 1 + (\ln 3 - \ln \ln 15)/(\ln 15) < 1.04$. Then, there exists an (m, h, k) -locally surjective map family \mathcal{L} such that $w := |\mathcal{L}| = 14m$.

5.2 Compiler from Passively Secure to Actively Secure Protocols

Based on locally surjective map families, we show our construction of one-round actively secure protocols from any one-round passively secure protocol. We give a sketch here and defer the formal proof to the full version.

Theorem 4. Suppose that there exists a 1-round (resp. computationally) passively t -secure k -server protocol $\Pi = (\text{Query}, \text{Answer}, \text{Output})$ for $\mathcal{F} = (F_\lambda : P_\lambda \times X_\lambda \rightarrow Y_\lambda)_{\lambda \in \mathbb{N}}$. If there exists an $(m, m - 2t, k)$ -locally surjective map family \mathcal{L} of size $w = \text{poly}(\lambda)$, there exists a 1-round (resp. computationally) actively t -secure m -server protocol $\Pi' = (\text{Query}', \text{Answer}', \text{Output}')$ for \mathcal{F} such that

- $\text{Comm}(\Pi') = O(twm \cdot \text{Comm}(\Pi))$;
- $\text{c-Comp}(\Pi') = O(twm \cdot \text{c-Comp}(\Pi))$;
- $\text{s-Comp}(\Pi') = O(tw \cdot \text{s-Comp}(\Pi))$.

Proof (sketch). Let $\mathcal{L} = \{f_1, \dots, f_w\}$ be an (m, h, k) -locally surjective map family, where $h = m - 2t$. For $u \in [w]$ and $j \in [k]$, define $G_{u,j} = f_u^{-1}(j) = \{i \in [m] : f_u(i) = j\}$. Consider the following protocol Π' : For all $u \in [w]$ and $\ell \in [t + 1]$ (in parallel),

1. The client C computes k queries $(\text{que}_1^{(u,\ell)}, \dots, \text{que}_k^{(u,\ell)})$ of Π .
2. C sends $\text{que}_{f_u(i)}^{(u,\ell)}$ to each server S_i .
3. Each S_i returns an answer $\text{ans}_i^{(u,\ell)}$ as the $f_u(i)$ -th server would answer to $\text{que}_{f_u(i)}^{(u,\ell)}$ in Π .

To obtain an output, C sets $S \leftarrow [m]$ and $L \leftarrow 1$, and does the following:

1. Check whether for all $u \in [w]$ and $j \in [k]$, the answers $\text{ans}_i^{(u,L)}$ returned by servers S_i in $G_{u,j}$ are identical with each other.
2. If so, let $\alpha_{u,j}$ be the unique answer by servers in $G_{u,j}$ and run the output algorithm of Π on $(\alpha_{u,1}, \dots, \alpha_{u,k})$ to obtain y_u . Then, output the majority of y_1, \dots, y_w .
3. Otherwise, find a pair (i_1, i_2) of servers who are mapped to the same group $G_{u,j}$ but returned different answers. That is, $f_u(i_1) = f_u(i_2)$ and $\text{ans}_{i_1}^{(u,L)} \neq \text{ans}_{i_2}^{(u,L)}$ for some $u \in [w]$. Note that at least one of them are malicious. Then, update $S \leftarrow S \setminus \{i_1, i_2\}$ and $L \leftarrow L + 1$, and go back to Step 1.

Privacy. An adversary corrupting a set T of at most t servers can only learn queries received by a set $f_u(T)$ of servers in Π . Since $|f_u(T)| \leq |T| \leq t$, the privacy of Π' follows from that of Π .

Byzantine-Robustness. An adversary succeeds in letting the client accept an incorrect result only if at least $w/2$ out of y_1, \dots, y_w are incorrect in some iteration (say, L) in the output phase of C. This implies that for at least $w/2$ u 's, there exists a remaining corrupted server $i \in T \cap S$ who submits an incorrect answer $\widetilde{\text{ans}}_i^{(u,L)} \neq \text{ans}_i^{(u,L)}$. On the other hand, since at most one honest server is eliminated from S in each iteration, it holds that $|H \cap S| \geq (m - t) - t = m - 2t$, where H is the set of all honest servers. Therefore, the property of locally surjective map families ensures that $f_u(H \cap S) = [k]$ holds for at least one of the above $w/2$ u 's. In other words, there exists a remaining honest server $i' \in H \cap S$ such that $f_u(i') = f_u(i)$, and the answer $\widetilde{\text{ans}}_i^{(u,L)}$ is compared with the correct answer $\text{ans}_{i'}^{(u,L)}$ from the honest server i' . Thus, the client can detect the malicious behavior of the corrupted server i . Therefore, the client can successfully eliminate at least one malicious server in each iteration and obtain the correct result after at most t iterations. \square

To obtain a concrete compiler from Theorem 4, we plug in the (m, h, k) -locally surjective map family in Proposition 8 with $h = m - 2t$.

Theorem 5. *Suppose that there exists a 1-round (resp. computationally) passively t -secure k -server protocol Π for \mathcal{F} . If*

$$m \geq 2t + 15 \text{ and } \frac{m - 2t}{\gamma \ln(m - 2t)} \geq k,$$

where $1 < \gamma < 1.04$ is the constant in Proposition 8, then there exists a 1-round (resp. computationally) actively t -secure m -server protocol Π' for \mathcal{F} such that

- $\text{Comm}(\Pi') = O(tm^2 \cdot \text{Comm}(\Pi))$;
- $\text{c-Comp}(\Pi') = O(tm^2 \cdot \text{c-Comp}(\Pi))$;
- $\text{s-Comp}(\Pi') = O(tm \cdot \text{s-Comp}(\Pi))$.

Remark 3. The computational complexity of the construction in Theorem 5 does not take into account that of finding a locally surjective map family \mathcal{L} . We note that the choice of \mathcal{L} does not affect the security of a protocol. Hence we can construct it before the protocol starts and the family is reusable any number of times.

5.3 Instantiations

By applying our compiler in Theorem 5 to the protocols in Propositions 1 and 2, we obtain the following corollaries. The formal proof appears in the full version.

Corollary 7. *Suppose that $m \geq \max\{2t3^t + 2t, 2t + 15\}$. Let $N \in \text{poly}(\lambda)$. Then, there exists a computationally actively t -secure 1-round m -server protocol Π for INDEX_N such that*

- $\text{Comm}(\Pi) = \exp(O(\sqrt{\log N \log \log N})) \cdot t^2 3^t m^2$;
- $\text{c-Comp}(\Pi) = \exp(O(\sqrt{\log N \log \log N})) \cdot t^2 3^t m^2$;
- $\text{s-Comp}(\Pi) = N^2 \cdot \exp(O(\sqrt{\log N \log \log N})) \cdot t 2^t m$.

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = N^{o(1)} \cdot 2^{O(t)}$.

Corollary 8. *Assume a pseudorandom function $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2(\lambda+1)}$. Suppose that $m \geq \max\{t2^{t+1} + 2t, 2t + 15\}$. Let $N \in \text{poly}(\lambda)$. Then, there exists an actively t -secure 1-round m -server protocol Π for INDEX_N such that*

- $\text{Comm}(\Pi) = O(\log N \cdot \lambda \cdot t^2 2^t m^2)$;
- $\text{c-Comp}(\Pi)$ is $O(\log N \cdot t^2 2^t m^2)$ invocations of G ;
- $\text{s-Comp}(\Pi)$ is $O(N^2 \log N \cdot t 2^t m)$ invocations of G .

In particular, $\max\{\text{Comm}(\Pi), \text{c-Comp}(\Pi)\} = \log N \cdot 2^{O(t)} \cdot \text{poly}(\lambda)$.

Note that it is possible to apply the compiler in Theorem 5 to the passively secure k -server protocols in Propositions 3, 4, 5, and 6. Since $k > t$, the number of servers of the resulting protocols is $\Omega(k \log k) + 2t = \Omega(t \log t)$. On the other hand, these protocols can also be made actively secure by using the standard error correction algorithm [41] or the technique of [38], and one can then obtain actively secure protocols that has a smaller number of servers $O(t)$. We thus do not show instantiations based on these protocols.

Acknowledgements. This research was partially supported by JSPS KAKENHI Grant Numbers JP20J20797 and JP19H01109, Japan, JST CREST Grant Numbers JPMJCR2113 and JPMJCR22M1, Japan, and JST AIP Acceleration Research JPMJCR22U5, Japan.

References

1. Ananth, P., Chandran, N., Goyal, V., Kanukurthi, B., Ostrovsky, R.: Achieving privacy in verifiable computation with multiple servers – without FHE and without pre-processing. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 149–166. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54631-0_9
2. Applebaum, B., Ishai, Y., Kushilevitz, E.: From secrecy to soundness: efficient verification via secure computation. In: Abramsky, S., Gavioille, C., Kirchner, C., Meyer auf der Heide, F., Spirakis, P.G. (eds.) ICALP 2010. LNCS, vol. 6198, pp. 152–163. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14165-2_14

3. Augot, D., Levy-dit-Vehel, F., Shikfa, A.: A storage-efficient and robust private information retrieval scheme allowing few servers. In: Gritzalis, D., Kiayias, A., Askoxylakis, I. (eds.) CANS 2014. LNCS, vol. 8813, pp. 222–239. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-12280-9_15
4. Banawan, K., Ulukus, S.: The capacity of private information retrieval from Byzantine and colluding databases. *IEEE Trans. Inf. Theory* **65**(2), 1206–1219 (2019)
5. Barkol, O., Ishai, Y.: Secure computation of constant-depth circuits with applications to database search problems. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 395–411. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_24
6. Barkol, O., Ishai, Y., Weinreb, E.: On d -multiplicative secret sharing. *J. Cryptol.* **23**(4), 580–593 (2010)
7. Barkol, O., Ishai, Y., Weinreb, E.: On locally decodable codes, self-correctable codes, and t -private PIR. *Algorithmica* **58**(4), 831–859 (2010)
8. Beimel, A., Ishai, Y., Kushilevitz, E., Raymond, J.F.: Breaking the $o(n/\sup 1/(2k-1))$ barrier for information-theoretic private information retrieval. In: The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings, pp. 261–270 (2002)
9. Beimel, A., Ishai, Y.: Information-theoretic private information retrieval: a unified construction. In: Automata, Languages and Programming, pp. 912–926 (2001)
10. Beimel, A., Ishai, Y., Kushilevitz, E., Orlov, I.: Share conversion and private information retrieval. In: 2012 IEEE 27th Conference on Computational Complexity, pp. 258–268 (2012)
11. Beimel, A., Stahl, Y.: Robust information-theoretic private information retrieval. *J. Cryptol.* **20**(3), 295–321 (2007)
12. Boyle, E., Gilboa, N., Ishai, Y.: Function secret sharing. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 337–367. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_12
13. Boyle, E., Gilboa, N., Ishai, Y.: Breaking the circuit size barrier for secure computation under DDH. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 509–539. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_19
14. Boyle, E., Gilboa, N., Ishai, Y.: Function secret sharing: improvements and extensions. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1292–1303, CCS 2016 (2016)
15. Boyle, E., Gilboa, N., Ishai, Y., Nof, A.: Sublinear GMW-style compiler for MPC with preprocessing. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12826, pp. 457–485. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84245-1_16
16. Boyle, E., Kohl, L., Scholl, P.: Homomorphic secret sharing from lattices without FHE. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part II. LNCS, vol. 11477, pp. 3–33. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_1
17. Bunn, P., Kushilevitz, E., Ostrovsky, R.: CNF-FSS and its applications. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) Public-Key Cryptography – PKC 2022, vol. 13177, pp. 283–314. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-97121-2_11
18. Canetti, R., Riva, B., Rothblum, G.N.: Refereed delegation of computation. *Inf. Comput.* **226**, 16–36 (2013)

19. de Castro, L., Lee, K.: VeriSimplePIR: verifiability in simplePIR at no online cost for honest servers. In: 33rd USENIX Security Symposium (USENIX Security 2024) (2024, to appear). <https://www.usenix.org/conference/usenixsecurity24/presentation/de-castro>
20. Choi, S.G., Katz, J., Kumaresan, R., Cid, C.: Multi-client non-interactive verifiable computation. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 499–518. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36594-2_28
21. Chor, B., Goldreich, O., Kushilevitz, E., Sudan, M.: Private information retrieval. *J. ACM* **45**(6), 965–982 (1998)
22. Colombo, S., Nikitin, K., Corrigan-Gibbs, H., Wu, D.J., Ford, B.: Authenticated private information retrieval. In: 32nd USENIX Security Symposium (USENIX Security 2023), pp. 3835–3851 (2023)
23. Damgård, I., Orlandi, C., Simkin, M.: Yet another compiler for active security or: efficient MPC over arbitrary rings. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10992, pp. 799–829. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96881-0_27
24. Dao, Q., Ishai, Y., Jain, A., Lin, H.: Multi-party homomorphic secret sharing and sublinear MPC from sparse LPN. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023. LNCS, vol. 14082, pp. 315–348. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38545-2_11
25. Devet, C., Goldberg, I., Heninger, N.: Optimally robust private information retrieval. In: 21st USENIX Security Symposium (USENIX Security 2012), pp. 269–283 (2012)
26. Di Crescenzo, G., Malkin, T., Ostrovsky, R.: Single database private information retrieval implies oblivious transfer. In: Preneel, B. (eds.) Advances in Cryptology – EUROCRYPT 2000. LNCS, vol. 1807, pp. 122–138. Springer, Cham (2000). https://doi.org/10.1007/3-540-45539-6_10
27. Dvir, Z., Gopi, S.: 2-server PIR with subpolynomial communication. *J. ACM* **63**(4), 1–15 (2016)
28. Efremenko, K.: 3-query locally decodable codes of subexponential length. *SIAM J. Comput.* **41**(6), 1694–1703 (2012)
29. Eriguchi, R., Kurosawa, K., Nuida, K.: On the optimal communication complexity of error-correcting multi-server PIR. In: Kiltz, E., Vaikuntanathan, V. (eds.) Theory of Cryptography, TCC 2022. LNCS, vol. 13749, pp. 60–88. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-22368-6_3
30. Gennaro, R., Gentry, C., Parno, B.: Non-interactive verifiable computing: outsourcing computation to untrusted workers. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 465–482. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14623-7_25
31. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, STOC 2009, pp. 169–178 (2009)
32. Gilboa, N., Ishai, Y.: Distributed point functions and their applications. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 640–658. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_35
33. Goldberg, I.: Improving the robustness of private information retrieval. In: 2007 IEEE Symposium on Security and Privacy (SP’07). pp. 131–148 (2007)
34. Hazay, C., Ishai, Y., Marcedone, A., Venkitasubramanian, M.: LevioSA: lightweight secure arithmetic computation. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pp. 327–344, CCS 2019 (2019)

35. Hazay, C., Venkatasubramanian, M., Weiss, M.: The price of active security in cryptographic protocols. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12106, pp. 184–215. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_7
36. Ishai, Y., Lai, R.W.F., Malavolta, G.: A geometric approach to homomorphic secret sharing. In: Garay, J.A. (ed.) PKC 2021. LNCS, vol. 12711, pp. 92–119. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-75248-4_4
37. Korte, B.H., Vygen, J.: Combinatorial Optimization, vol. 1. Springer, Cham (2011). <https://doi.org/10.1007/978-3-642-77489-8>
38. Kurosawa, K.: How to correct errors in multi-server PIR. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11922, pp. 564–574. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34621-8_20
39. Orlandi, C., Scholl, P., Yakoubov, S.: The Rise of Paillier: homomorphic secret sharing and public-key silent OT. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021. LNCS, vol. 12696, pp. 678–708. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77870-5_24
40. Roy, L., Singh, J.: Large message homomorphic secret sharing from DCR and applications. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12827, pp. 687–717. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84252-9_23
41. Rudra, A.: Lecture 27: Berlekamp-Welch algorithm. <https://cse.buffalo.edu/faculty/atri/courses/coding-theory/lectures/lect27.pdf>
42. van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 24–43. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_2
43. Woodruff, D., Yekhanin, S.: A geometric approach to information-theoretic private information retrieval. *SIAM J. Comput.* **37**(4), 1046–1056 (2007)
44. Yao, X., Liu, N., Kang, W.: The capacity of multi-round private information retrieval from Byzantine databases. In: 2019 IEEE International Symposium on Information Theory (ISIT), pp. 2124–2128 (2019)
45. Yoshida, M., Obana, S.: Verifiably multiplicative secret sharing. *IEEE Trans. Inf. Theory* **65**(5), 3233–3245 (2019)
46. Zhang, L.F., Wang, H.: Multi-server verifiable computation of low-degree polynomials. In: 2022 IEEE Symposium on Security and Privacy (SP), pp. 596–613 (2022)
47. Zhang, L.F., Wang, H., Wang, L.P.: Byzantine-robust private information retrieval with low communication and efficient decoding. In: Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security, pp. 1079–1085, ASIA CCS 2022 (2022)