

# Private Collection Matching Protocols

Kasra EdalatNejad  
EPFL  
Lausanne, Switzerland  
kasra.edalat@epfl.ch

Wouter Lueks  
CISPA Helmholtz Center for Information Security  
Saarbrücken, Germany  
lueks@cispa.de

Mathilde Raynal  
EPFL  
Lausanne, Switzerland  
mathilde.raynal@epfl.ch

Carmela Troncoso  
EPFL  
Lausanne, Switzerland  
carmela.troncoso@epfl.ch

## ABSTRACT

We introduce *Private Collection Matching (PCM)* problems, in which a client aims to determine whether a collection of sets owned by a server matches their interests. Existing privacy-preserving cryptographic primitives cannot solve PCM problems efficiently without harming privacy. We propose a modular framework that enables designers to build privacy-preserving PCM systems that output one bit: whether a collection of server sets matches the client's set. The communication cost of our protocols scales linearly with the size of the client's set and is independent of the number of server elements. We demonstrate the potential of our framework by designing and implementing novel solutions for two real-world PCM problems: determining whether a dataset has chemical compounds of interest, and determining whether a document collection has relevant documents. Our evaluation shows that we offer a privacy gain with respect to existing works at a reasonable communication and computation cost.

## KEYWORDS

Private set intersection, private computation, homomorphic encryption, private aggregation

## 1 INTRODUCTION

In many scenarios, a server holds a collection of sets and clients wish to determine whether these server sets *match* their own set, while both client and server keep their privacy. We call these *Private Collection Matching (PCM)* problems. In this paper, we study for the first time the requirements of PCM problems.

We identify the privacy and efficiency requirements of PCM problems by analyzing three real-world use cases: determining whether a pharmaceutical database contains compounds that are chemically similar to the client's [42, 69, 73], determining whether an investigative journalist holds relevant documents [21] (or how many), and matching a user's profile to items or other users in mobile apps [72, 74, 75]. We find that PCM problems have three common characteristics: (1) Clients want to compare their *one* set with *all* sets at the server. (2) Clients do not need per-server set

results, only an aggregated output (e.g., whether any server set matches). (3) Clients and server want privacy: the server should learn nothing about the clients' set, and the clients should only learn the aggregated output. However, PCM problems differ in their definition of *when sets match* and *how to combine individual matching responses*. Now, we discuss these two aspects in more detail:

*Set matching.* Typically, set matching is defined as a function of the intersection of two sets. Hence, clients could detect a matching server set by using private set intersection (PSI) protocols [13–15, 38, 57, 57–60] to privately compute the intersection, then post-process the intersection to determine interest locally. PCM applications differ in their matching criteria and may decide interest using measures such as a cardinality threshold, containment, or set similarity. This local processing approach, unfortunately, *reduces privacy of the server's sets* by leaking information beyond the set's matching status to the client. Such leakage could, for instance, reveal secret chemical properties of compounds, or the content of journalists' sensitive documents.

*Many-set.* In PCM problems, the server holds a *collection of  $N$  sets*. This creates two challenges. First, running one matching (or PSI) interaction per server set is inefficient. Second, revealing individual set-matching statuses harms server privacy. While servers may be interested in selling data to or collaborating with clients, they want to ensure that clients cannot use the 'PCM solution' to extract information about sets. Clients, meanwhile, often only need an aggregated response summarizing the utility of a collection. Servers therefore enact application-dependent aggregation policies ensuring that clients can, e.g., determine only whether at least one set matches or learn only the the number of matching sets.

We construct a framework that leverages computation *in the encrypted domain* to solve PCM problems. In our framework, shown in Fig. 1, given an encrypted client set, the server uses a matching criteria  $f_M$  to compute per-server-set *binary* answers to "is this set of interest to the client?". Next, the server uses an aggregation policy  $f_A$  to combine per server-set responses into a collection-wide response. Finally, the client decrypts the aggregated result.

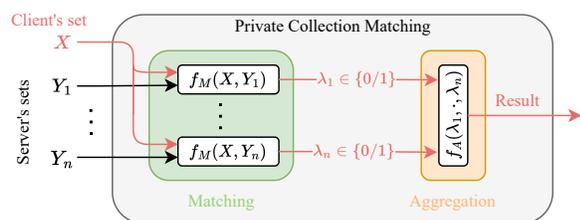
Our work makes the following contributions:

✓ We introduce Private Collection Matching (PCM) problems. We derive their requirements from three real-world problems.

✓ We design single-set protocols where the client learns a *one-bit output* – whether one server set is of interest to the client – and many-set protocols where the client learns a *collection-wide output*

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.  
*Proceedings on Privacy Enhancing Technologies 2023(3)*, 446–468  
© 2023 Copyright held by the owner/author(s).  
<https://doi.org/10.56553/popets-2023-0091>





**Figure 1: Structure of our PCM framework.** Red arrows show values encrypted under the client’s key.  $f_M$  designates a matching function: it outputs a binary value  $\lambda$  indicating whether two sets match.  $f_A$  is an aggregation function that combines  $n$  matching statuses into a collection-wide result.

that aggregates individual matching responses. The communication cost of our protocols scales linearly with the size of the client’s set and is independent of the number of server sets and their total size.

✓ We propose a modular design that separates flexible set matching criteria from many-set aggregation. Our modularity enables extending our design with new matching or aggregation policies and simplifies building privacy-preserving PCM solutions.

✓ We demonstrate our framework’s capability by solving chemical similarity and document search problems. We show that our framework offers improved privacy with competitive cost compared to custom-made solutions, and significantly improves the latency, client’s computation cost, and communication cost with respect to generic solutions that offer the same privacy guarantee.

## 2 PRIVATE COLLECTION MATCHING

In this section, we define the Private Collection Matching (PCM) problem. We derive its basic requirements from three real-world matching problems. We also explain why existing PSI solutions cannot satisfy the privacy requirements of PCM problems.

### 2.1 Case studies

We study three cases that can benefit from PCM.

*Chemical similarity.* Chemical research and development is a multi-billion dollar industry. When studying a new chemical compound, knowing the properties of similar compounds can speed up the research. In an effort to monetize research, companies sell datasets describing thousands to millions of compounds and their properties. Chemical R&D teams are willing to pay high prices for these datasets but *only if* they include compounds similar to their research target. Determining whether this is the case is tricky: buyers want to hide the compound they are currently investigating [69], and sellers want to hide information about the compounds in their dataset before the sale is finalized.

Chemical similarity of compounds is determined by comparing molecular fingerprints of compounds [8, 42, 48, 73, 79, 80]. Fingerprints are based on the substructure of compounds and are represented as fixed-size bit vectors – these vectors are between a few hundred and few thousand bits long. Measures such as Tversky [76] and Jaccard [33] determine the similarity of these fingerprints, and thus of the compounds.

Revealing pair-wise intersection cardinalities or even similarity scores between the fingerprints of a target compound and the seller’s compounds results in unacceptable leakage. A buyer can reconstruct an  $n$ -bit molecular fingerprint  $F$  by learning similarity values between  $F$  and  $n + 1$  known compounds [69]. To prevent inferences, the buyer should learn only the number of similar compounds in the seller’s dataset, or, better, only learn whether at least one similar compound exists.

*Peer-to-peer document search.* Privacy-preserving peer-to-peer search engines help users and organizations with strict privacy requirements to collaborate safely. We take the example of investigative journalists who, while unwilling to make their investigations or documents publicly available, want to find collaboration opportunities within their network [21].

To identify those opportunities, a journalist performs a search to learn whether a document owner has documents of interest. A search query consists of keywords relevant to the journalist’s investigation. The document owner compares the query to all documents in his collection. A document is deemed relevant if it contains *all* or *a sufficient number of* queried keywords. Journalists own a collection of a thousand documents (on average), and each document is represented by around a hundred keywords.

The sensitivity of journalists’ investigations demand that both the content of the documents and of the queries remain private [21]. Journalists only need to learn one bit of information – that at least one or a threshold number of documents in the owner’s collection is relevant – to determine whether they should contact the owner.

*Matching in mobile apps.* A common feature in mobile apps is enabling users to find records of interest in the app servers’ databases, e.g., restaurants [74], routes for running [72], or suitable dating partners [51, 75]. Users are typically interested in records that have at least a number of matching characteristics in common with their search criteria or that are a perfect match. Also, users need to be able to retrieve these records.

The user-provided criteria – typically range choices entered via radio buttons or drop-down menus – are compared to the attributes of records. An app database can have millions of records and records can have dozens to hundreds of attributes.

Both search criteria and records are sensitive. Knowing search criteria enables profiling of user interests. These are particularly sensitive for dating applications. Thus, search queries should be kept private. The secrecy of the records in the database is not only at the core of the business value of these apps but also required by law in cases where records contain personal data (e.g., dating apps).

### 2.2 PCM requirements

We extract requirements that PCM protocols should fulfill based on the commonalities between the use cases. These come in addition to basic PSI properties such as client privacy.

*RQ.1: Flexible set matching.* PCM protocols need to be able to determine matches between sets without revealing other information such as intersections or cardinalities to the client. In the use cases, clients do not need to know the intersection or its cardinality. They are interested only in whether there is a match. Matches in these examples are a function of the intersection between a client and a server set: a chemical compound is a match when the Tversky or

Jaccard similarity with the query exceeds a threshold; a document is a match when it contains some or all query keywords; and a record is a match when it includes a threshold of query attributes. PCM protocols must be able to detect matching sets and compute a *single* one-bit matching status per-set.

*RQ.2: Aggregate many-set responses.* PCM protocols need to have the capability to provide an aggregated response for a collection of sets without leaking information about individual sets. Our use cases highlight that in many applications, a client (buyer, journalist, user) may want to compare their input (compounds under investigation, keywords of interest, search criteria) with a *collection* of sets (compounds in a database, documents in a collection, records in a database). More importantly, we observe that clients wish to know how interesting the collection is as a whole. For example, a buyer is interested in a chemical dataset if it contains at least one similar compound and a querier journalist may contact a document owner if the owner has a number of relevant documents in their collection. Therefore, to satisfy clients’ needs yet protect the server’s privacy, PCM protocols should only reveal aggregated per-collection results.

*RQ.3: Extreme imbalance.* PCM problems have thin clients and imbalanced input sizes; thus, protocols must not require communication and computation linear to the server’s input size from clients. Drawing from our earlier scenarios, the total input size of the server may be as much as 6 orders of magnitude larger than the client’s input, as shown by the chemical similarity scenario. The server, holding many sets, can safely be assumed to be resourceful, while clients may be constrained in their capabilities, e.g., a client running the PCM protocol from their mobile phone. This can be in terms of computation, e.g., battery has to be preserved in mobile apps; or in terms of bandwidth, e.g., journalists that can be in locations with poor Internet access. Therefore, PCM protocols should not incur a large client-side cost.

### 2.3 Formal PCM definition

Let  $X$  be a client set with  $n_c$  elements  $\{x_1, \dots, x_{n_c}\}$  from input domain  $D$  and  $\mathcal{Y}$  be a collection of  $N$  server sets  $\{Y_1, \dots, Y_N\}$  where the  $i$ ’th server set  $Y_i = \{y_{i,1}, \dots, y_{i,n_s^i}\}$  has  $n_s^i$  elements, also from  $D$ , leading to a total server size of  $N_s = \sum_i n_s^i$ . We define two families of functions as follows:

Matching functions  $\lambda_i \leftarrow f_M(X, Y_i)$  take two sets  $X$  and  $Y_i$  as input and compute a binary matching status  $\lambda_i$  determining interest. This family represents our flexible matching criteria RQ.1. For example, in the document search scenario where a server set (document) is of interest when it contains all queried keywords, we define  $f_M(X, Y_i)$  as 1 if  $X \subseteq Y_i$  and 0 otherwise.

Aggregation functions  $A \leftarrow f_A(\lambda_1, \dots, \lambda_N)$  take  $N$  binary matching statuses ( $\lambda_i$ ) and aggregates them into a single response  $A$ . This family represents our aggregation requirement RQ.2. For example, if we want to count the number of relevant documents in a search, we define  $f_A(\lambda_1, \dots, \lambda_N)$  as  $\sum_j \lambda_j$ .

In Table 4 in Appendix A, we summarize the matching and aggregation functions that we implement.

**Definition 1** (PCM). PCM protocols are two-party computations between a client and a server with common inputs  $f_M$  and  $f_A$ , where the client learns an aggregated matching status and the

**Table 1: Overview of PSI approaches in the PCM setting.**

		Privacy		Efficiency
		RQ.1	RQ.2	RQ.3
Comparison	[22, 39, 53, 55, 59, 60]	×	×	×
OPRF	[13–15, 37, 66]	×	×	✓*
OPE	[11, 16, 24, 30]	×	×	✓
Generic SMC	[31, 34, 56]	✓	✓	×
Circuit-PSI	[10, 12, 36, 44, 57, 58, 64]	✓	✓	×
Flexible func.	[28, 32, 69, 82–84]	✓	×	×
<b>This paper</b>		✓	✓	✓

\*Efficient communication requires pre-processing.

server learns nothing. Formally:

$$(A = f_A(f_M(X, Y_1), \dots, f_M(X, Y_N)), \perp) \leftarrow PCM_{f_M, f_A}(X, \mathcal{Y})$$

We use this notation to define formal properties of PCM protocols.

**Definition 2** (Correctness). A PCM protocol is correct if the client output matches the result of  $A = f_A(f_M(X, Y_1), \dots, f_M(X, Y_N))$ .

**Definition 3** (Client privacy). A PCM protocol is client private if the server cannot learn any information about the client’s set beyond the size of the client’s set.

**Definition 4** (Server privacy). A PCM protocol is server private if the client cannot learn any information about the server elements beyond the number of server sets  $N$ , the maximum server set size  $\max n_s^i$ , and the client output  $A = f_A(f_M(X, Y_1), \dots, f_M(X, Y_N))$ .

### 3 RELATED WORK

While ad-hoc solutions for chemical (Shimizu et al. [69]) and document search (EdalatNejad et al. [21]) exist, most prior work focuses on building private set intersection (PSI) protocols, which are a special case of PCM. We introduce the PSI protocols most relevant to our work. We leave the detailed comparison with ad-hoc solutions to the evaluation (see Section 11).

We compare existing work on two critical aspects of PCM problems: privacy and efficiency. We summarize existing schemes and their suitability for the PCM scenario in Table 1.

For privacy, we assess whether existing approaches provide flexible matching (RQ.1) and aggregated many-set responses (RQ.2). We note that the majority of prior works do not consider, or support, many sets. When there is no natural extension to support many sets at once, we run a single-set interaction per server set, leading to an  $N \times$  increase in cost. This naive extension does not provide the privacy enhancement of many-set aggregation but enables us to reason about the efficiency of these approaches.

For efficiency, taking into account the extreme imbalance requirement (RQ.3), we focus on the client’s computation and communication cost and require schemes to have a client cost of  $\omega(N_s (\approx N n_s^i))$ .

**Traditional single-set PSI.** We first study protocols *solely focusing on single-set intersection or cardinality*. We study schemes that offer *enhanced functionality or privacy* below in the Custom PSI Protocols section.

In PSI, clients learn information about the intersection of two sets while (i) not learning anything about the server’s non-intersecting elements, and (ii) not leaking any information about their own set to the server. PSI protocols in the literature focus on providing two possible outputs: the *intersection* (e.g., finding common network intrusions [49], or discovering contacts [19]); and the *cardinality of the intersection* (e.g., privately counting the number of common friends between two social media users [50], or performing genomic tests [4]). Works in this area opt for a variety of trade-offs between the computational capability and the amount of bandwidth required to run the protocol [11, 34, 37, 54]. These works show that PSI can scale to large datasets [35, 57], and support light clients [37].

We classify PSI protocols according to the techniques they use to compute the PSI functionality:

**Comparison-based protocols.** The fastest class of PSI protocols uses bucketing to enable PSI protocols that run optimized comparison protocols between very small client and server buckets [22, 39, 53, 55, 59, 60]. These approaches use hashing to first map elements to small buckets. Then they compare the client and server elements in each bucket using comparison primitives that are efficient when buckets have very few elements, for example built using oblivious transfer (OT). These approaches reveal comparison results, and thus intersections or cardinalities, to the client; therefore, they do not satisfy our single-set privacy requirement (RQ.1). As these approaches reveal individual detailed set responses, *private aggregation* (RQ.2) is impossible. We discuss approaches that do not reveal the comparison result separately below as ‘Circuit-PSI’. Each client and server element must participate in at least one comparison; consequently, the communication cost is linear in the client size  $n_c$  and total server size  $N_s$ . Therefore, comparison-based approaches do not satisfy our efficiency requirement (RQ.3).

In Appendix D.3, we confirm our efficiency assessment by evaluating the cost of SpOT-light [53], one of the fastest comparison-based PSI protocols, in the PCM setting. We show significant improvement in latency (10–65x), client’s computation (1800–24,800x), and transfer cost (1.7–27x).

**OPRFs.** Another technique to construct PSI protocols is to evaluate a pseudo-random function (PRF) over client and server set elements and compare these. To protect privacy, the client obliviously evaluates the PRF over its elements together with the server, and the server sends the PRF evaluation of its elements to the client. Oblivious Pseudo-random Functions (OPRFs) can be constructed from asymmetric primitives such as RSA [14, 15, 37], Diffie-Hellman [37, 66], discrete logarithms [13], or by applying SMC on symmetric PRFs [34, 37].<sup>1</sup> What distinguishes these OPRF-based approaches from the previous category is that instead of comparing many *pairs* in small buckets, OPRF-based approaches compute a deterministic tag for each client and server element that can be compared locally. While the communication cost is linear in the size of both client and server sets, preprocessing and reusing tags for server elements can make the transfer cost independent of the server size [37]. Similar to comparison-based protocols, OPRF

approaches cannot compute flexible set matches without leaking intermediate data nor do they support aggregation.

**Oblivious polynomial evaluation.** Another approach is to use (partial) homomorphic encryption [65] to determine set intersection using oblivious polynomial evaluation (OPE) [11, 16, 24, 30]. The client encrypts their elements and sends them to the server. The server constructs a polynomial with its set elements as roots, evaluates the polynomial on encrypted client elements, randomizes the result, and sends them back to the client. The client decrypts the results, a 0 indicates a matching element. We use a similar approach in our schemes. Existing OPE-based approaches do not support flexible set matching or aggregation, but achieve cost independent of the server input size ( $N_s$ ) for the client.

**Generic SMC.** Some PSI protocols use generic SMC tools [6, 29, 81] to construct full circuits such as sort-compare-shuffle [31] to compute the intersection [34, 56]. They can be extended to support flexible set matching or many-set aggregation. However, circuits have communication linear in the size of their inputs (wires) which guarantees a transfer cost of  $O(n_c + N_s)$ . This is a fundamental limit. Thus, circuits cannot satisfy our efficiency requirement (RQ.3). Because circuits can satisfy our privacy requirements, we develop a generic alternative to our framework using an SMC compiler in Section 11.2 and show that our system improves latency (2–96x), client’s computation (75–2250x), and transfer cost (93–2800x). Besides for generic protocols, circuits are used to (1) extend OT-based protocols (discussed as ‘Circuit-PSI’ below) or (2) obliviously evaluate PRFs [34, 37] such as AES or LowMC [2] (discussed as ‘OPRF’).

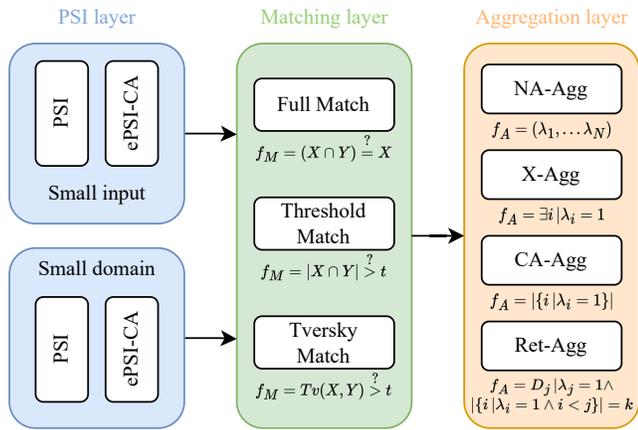
**Custom PSI protocols.** Some PSI protocols go beyond cardinality and compute more complex functions over the intersection.

**Circuit-PSI.** A new line of work extends comparison-based PSI protocols to support arbitrary extensions with generic SMCs [10, 12, 36, 44, 57, 58, 64]. These works compute the intersection of two sets but instead of revealing the plain result to the client, they secret share the intersection between the two parties. This secret shared output enables parties to privately compute arbitrary functions on top of the intersection. Unfortunately, these work focus on scenarios with equal client and server sizes as their cost is linear in the input of both parties  $O(n_c + N_s)$ . This linear communication cost is a fundamental limit. As these approaches can satisfy our privacy requirement, we evaluate Chandran et al. [10], a state-of-the-art Circuit-PSI paper, in the PCM setting in Section 11.2. We show that our framework significantly improves latency (580x), client’s computation (70,000x), and transfer cost (2360x).

**Flexible functionality.** Several privacy-preserving custom protocols provide functionality beyond computing intersection or cardinality. For instance, computing the sum or statistical functions over associated data [32, 82], evaluating a threshold on intersection size [28, 83, 84], or computing Tversky similarity [69]. These approaches improve privacy by supporting flexible set matching but do not extend well to many-sets scenarios and aggregation. They are optimized for a specific setting and do not achieve cost independent of the server input size.

**Orthogonal works.** We briefly mention two groups of related work that, while of interest, are orthogonal to PCM problems and solve different challenges. *Encrypted databases and ORAM* [45, 61,

<sup>1</sup>We only consider OPRF approaches where the parties can *choose the PRF key*. Primitives where a random key is generated during the execution [39, 53] can only be used for comparing-based approaches.



**Figure 2: An overview of our layers and their composition. Refer to Table 4 for a summary of protocol definitions.**

62, 71] let clients query outsourced data, or, in more recent work [9], subsets of outsourced data via an access control policy. Typically, these approaches do not limit what the client can learn about the entries they have access to. This is in contrast with the PCM setting where the client computes a function over the server’s sets, and clients should not learn more than the function’s outcome. Therefore, traditional applications of encrypted databases and ORAM do not seem to directly enable the construction of PCM solutions. However, we do not rule out the possibility that they can be extended or used in creative ways to solve PCM problems.

In *multi-party PSI* [5, 40, 78],  $p$  parties each with a set  $S_i$  compute one intersection  $I = \bigcap S_j$ . PCM problems, instead, are a two-party protocol where the server holds  $N$  sets.

#### 4 A FRAMEWORK FOR PCM SCHEMES

We showed that existing work – except for ad-hoc solutions – cannot solve PCM problems without losing either privacy or efficiency. We now introduce a modular framework that enables the design of PCM solutions with minimal effort and strong privacy while providing performance close to ad-hoc solutions.

The framework has three layers, shown in Fig. 2:

*PSI layer* protocols operate directly on a client’s and a server’s set. These protocols compute single-set PSI functionalities such as intersection or cardinality. Our implementation focuses on two scenarios: small input domain size and small constant-size client sets.

*Matching layer* protocols use PSI layer protocols to compute a binary answer determining whether each of the server sets matches the interest of the client according to a pre-defined matching function  $f_M$  (RQ.1). Computation in this layer is the same regardless of the scenario chosen in the PSI layer.

*Aggregation layer* protocols aggregate  $N$  single-set responses into one collection-wide response according to a pre-defined aggregation function  $f_A$  (RQ.2). This layer achieves constant size responses and ensures efficient client communication (RQ.3).

*Modularity.* While we provide a large number of protocols for the PSI, Matching, and Aggregations layers (see Fig. 2), an advantage of our framework is extensibility. Whenever an application requires new matching or aggregation criteria, designers can add (or adapt) a single functionality while taking advantage of the existing optimized layers. As an example, we extend our framework to support the single-set PSI-SUM functionality in Appendix F.

The layers can also be used as standalone protocols. We include blocks such as ‘Naive’ aggregation such that even if an application does not require matching or aggregation the designer can use the framework to enjoy its capability to tackle the many-set scenario.

*Security goals.* The PCM framework should protect the privacy of clients and servers against semi-honest adversaries. Furthermore, it is desirable that the framework provides server and client privacy against malicious adversaries, however, as we discuss in Section 9 our framework only guarantees client privacy. Servers are free to choose their input, allowing them to degrade the quality of the protocol’s result without any misbehavior. We accept this inherent weakness of PCM protocols and make the deliberate decision to aim for correctness only in the semi-honest setting, and *not* when the server is malicious.

### 5 TECHNICAL BACKGROUND

We introduce our notation and define the syntax of the fully homomorphic encryption scheme we use.

*Notation.* Let  $\ell$  be a security parameter. We write  $x \leftarrow X$  to denote that  $x$  is drawn uniformly at random from the set  $X$ . Let  $q$  be a positive integer, then  $\mathbb{Z}_q$  denotes the set of integers  $[0, \dots, q)$ , and  $\mathbb{Z}_q^*$  represent the elements of  $\mathbb{Z}_q$  that are co-prime with  $q$ . We write  $[n]$  to denote the set  $\{1, \dots, n\}$ , and use  $\langle a_i \rangle_m$  to present the list  $[a_1, \dots, a_m]$ . We drop the subscript  $m$  when the list length is clear from the context. We write  $\llbracket x \rrbracket$  to denote the encryption of  $x$ . We write  $\mathbb{1}[E]$  to denote the indicator function that returns ‘1’ if the event  $E$  is true, and ‘0’ otherwise. Table 3 in Appendix A summarizes our notation.

#### 5.1 Homomorphic encryption

Homomorphic encryption (HE) schemes enable arithmetic operations on encrypted values without decryption. We use HE schemes that operate over the ring  $\mathbb{Z}_q$  with prime  $q$ , such as BFV [23].

*Syntax.* HE is defined by the following procedures:

- $\text{params} \leftarrow \text{HE.ParamGen}(q)$ . Generates HE parameters with the plaintext domain  $\mathbb{Z}_q$ .
- $pk, sk \leftarrow \text{HE.KeyGen}(\text{params})$ . Takes the parameters  $\text{params}$  and generates a fresh pair of keys  $(pk, sk)$ . For brevity, we do not explicitly mention evaluation keys  $evk$  and consider them to be incorporated in the public key.
- $\llbracket x \rrbracket \leftarrow \text{HE.Enc}(pk, x)$ . Takes the public key  $pk$  and a message  $x \in \mathbb{Z}_q$  and returns the ciphertext  $\llbracket x \rrbracket$ .
- $x \leftarrow \text{HE.Dec}(sk, \llbracket x \rrbracket)$ . Takes the secret key  $sk$  and a ciphertext  $\llbracket x \rrbracket$  and returns the decrypted message  $x$ .

The correctness property of homomorphic encryption ensures that  $\text{HE.Dec}(sk, \text{HE.Enc}(pk, x)) \equiv x \pmod{q}$ .

---

**Algorithm 1** Check whether  $\llbracket x \rrbracket$  is zero.

---

```

function HE.IsZero( $pk, \llbracket x \rrbracket$ )
   $\llbracket b \rrbracket = 1 - \llbracket x \rrbracket^{(q-1)}$   $\triangleright b \leftarrow (x = 0)$ 
  return  $\llbracket b \rrbracket$ 

```

---

*Homomorphic operations.* HE schemes support homomorphic addition (denoted by  $+$ ) and subtraction (denoted by  $-$ ) of ciphertexts:  $\text{HE.Dec}(\llbracket a \rrbracket + \llbracket b \rrbracket) = a + b \bmod q$  and  $\text{HE.Dec}(\llbracket a \rrbracket - \llbracket b \rrbracket) = a - b \bmod q$ . HE schemes also support multiplication (denoted by  $\cdot$ ) of ciphertexts:  $\text{HE.Dec}(\llbracket a \rrbracket \cdot \llbracket b \rrbracket) = a \cdot b \bmod q$ .

Besides operating on two ciphertexts, it is possible to perform addition and multiplication with plaintext scalars. In many schemes, such scalar-ciphertext operations are more efficient than first encrypting the scalar and then performing a standard ciphertext-ciphertext operation. We abuse notation and write  $a\llbracket x \rrbracket + b$  to represent  $(\llbracket a \rrbracket \cdot \llbracket x \rrbracket) + \llbracket b \rrbracket = \llbracket ax + b \rrbracket$ .

*Multiplicative depth.* Our framework is designed with fully homomorphic encryption (FHE) in mind and assumes unbounded multiplication depth. For practical purposes, we port the majority, but not all, of our protocols to support execution with somewhat homomorphic encryption (SWHE) and optimize operations in Section 10.

## 5.2 Core functions

The complex functionality of PCM protocols can be reduced to a sequence of zero detection and inclusion test procedures. These two functions allow us to describe our protocol at a higher abstraction level. Moreover, any improvement to these basic functions automatically enhances our framework.

*Zero detection.* The function  $\llbracket b \rrbracket \leftarrow \text{HE.IsZero}(pk, \llbracket x \rrbracket)$  computes whether the ciphertext  $\llbracket x \rrbracket$  is an encryption of zero. The binary output  $b \in \{0, 1\}$  is defined as  $b = 1$  if  $x \equiv 0 \pmod{q}$  otherwise  $b = 0$ . We use Fermat's Little Theorem for zero detection [7]. We rely on the prime ring structure of  $\mathbb{Z}_q$  as any non-zero variable  $x \in \mathbb{Z}_q^*$  to the power  $q - 1$  is congruent to one modulo the prime  $q$ . We can perform this exponentiation with  $\lg(q)$  multiplications. See Algorithm 1 for the implementation.

The high multiplicative depth of  $\text{HE.IsZero}$  makes it impractical for use with most SWHE schemes. We hope that the research and advances in HE comparison enables efficient instantiations of this function and unlock our framework's full capabilities. When evaluating our framework in Section 11, we use ad-hoc techniques to avoid the need for this function.

*Inclusion test.* The function  $\llbracket I \rrbracket \leftarrow \text{HE.IsIn}(pk, \llbracket x \rrbracket, Y)$  checks if  $x$  is included in the set  $Y$  of cardinality  $n$ . We consider two variants. In the first,  $Y$  is a set of ciphertexts  $\llbracket y_i \rrbracket$ , in the second,  $Y$  is a set of plaintexts  $y_i$ . In both cases, the output  $I$  equals 0 if and only if a  $y_i$  exists such that  $x \equiv y_i \pmod{q}$ , otherwise  $I$  will be a uniformly random element in  $\mathbb{Z}_q^*$ . See Algorithm 2 for the implementation.

The function  $\text{HE.IsIn}$  relies on oblivious polynomial evaluation (OPE) [24, 30]. We create an (implicit) polynomial  $P$  with roots  $y_i$ , and evaluate  $\llbracket I \rrbracket = \llbracket r \cdot P(x) \rrbracket$ . If  $x$  is in the set, there exists a variable  $y_i$  where  $x \equiv y_i$ , thus  $I$  is zero. Otherwise,  $I$  is the product of  $n$  non-zero factors modulo  $q$ . Since  $q$  is prime, the product of non-zero values is non-zero. The random value  $r$  ensures uniformity in this case. The multiplicative depth of  $\text{HE.IsIn}$  scales with the size of

---

**Algorithm 2** Check inclusion of an encrypted variable  $x$  in a plain  $Y = \{y_1, \dots, y_n\}$  or an encrypted  $Y = \{\llbracket y_i \rrbracket, \dots, \llbracket y_n \rrbracket\}$  set.

---

```

function HE.IsIn( $pk, \llbracket x \rrbracket, Y = \{\llbracket y_1 \rrbracket, \dots, \llbracket y_n \rrbracket\}$ )
   $\llbracket I \rrbracket \leftarrow \prod_{i \in [n]} (\llbracket x \rrbracket - \llbracket y_i \rrbracket)$ 
   $r \leftarrow \mathbb{Z}_q^*$ 
  return  $r \cdot \llbracket I \rrbracket$ 

function HE.IsIn( $pk, \llbracket x \rrbracket, Y = \{y_1, \dots, y_n\}$ )
   $[a_0, \dots, a_n] \leftarrow \text{ToCoeffs}(Y)$   $\triangleright$  Such that  $\prod_i (x - y_i) = \sum_i a_i x^i$ 
   $\llbracket I \rrbracket \leftarrow \sum_{i \in [0..n]} a_i \cdot \llbracket x \rrbracket^i$ 
   $r \leftarrow \mathbb{Z}_q^*$ 
  return  $r \cdot \llbracket I \rrbracket$ 

```

---

$Y$ . We use the second form, where  $Y$  is a set of plaintexts, to lower the multiplicative depth when  $\llbracket x^i \rrbracket$  are known, see Section 10.3.

## 6 PSI LAYER

The PSI layer of our framework implements basic PSI functionalities: computing intersection or intersection cardinality. These protocols can be used in isolation, but in our framework they serve to form the input to the matching layer (see Fig. 2). We build PSI protocols for two scenarios: (1) scenarios where the client set has a small constant size (e.g., document search queries which typically have less than 10 keywords); and (2) scenarios where set elements come from a small input domain (e.g., gender and age in a dating profile). We build basic protocols assuming semi-honest clients, but allow for an extension – see Section 6.3 – that ensures queries represent a valid input set even if clients deviate from the protocol.

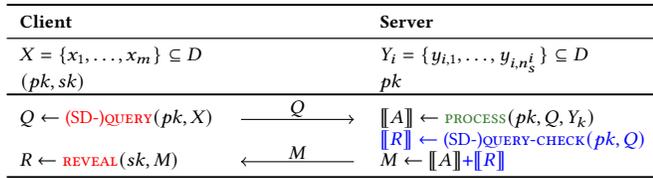
We structure our single-set protocols following Fig. 3. The client generates a HE key pair  $(pk, sk) \leftarrow \text{HE.KeyGen}(\text{params})$  and sends the public key  $pk$  to the server ahead of the protocol. Clients perform `QUERY` and send an encrypted representation of their set to the server. The server runs a protocol-specific processing function `PROCESS` to obtain the result  $M$ . The protocol is either used as the first layer and passes  $M$  into the second layer, or is stand-alone and returns  $M$  to the client. The server optionally runs `QUERY-CHECK` to randomize the result when a client submits a malformed query ( $\llbracket R \rrbracket$  is zero for correctly formed queries). Finally, the client runs the protocol-specific function `REVEAL` to compute the result. We denote algorithms run by the **client** in **red** and by the **server** in **green**. The *optional server-side* checks that ensure well-formedness of queries are denoted in **blue**.

### 6.1 Small constant-size client set

We start with scenarios where client sets are small and constant-size, typical for representing a search criteria. Clients use `QUERY` to encrypt their set elements  $x_i \in X$  as a query  $Q$  and send it to the server. Algorithm 3 instantiates small input functions.

*PSI.* The PSI protocol computes  $\text{PSI}(X, Y_k) = X \cap Y_k$ . The server uses the inclusion test  $\text{HE.IsIn}$  (see Section 5.1) to compute an inclusion status  $\llbracket s_i \rrbracket$  for each client element  $x_i$  (see `PSI-PROCESS`). An element  $x_i$  is in the intersection if and only if the corresponding inclusion status  $\llbracket s_i \rrbracket$  is zero (recall that the inclusion test produces zero for values *in* the set). When used as a stand-alone protocol, the server returns the list of encrypted inclusion values  $M$ , which the client then decrypts (see `PSI-REVEAL`).

*Cardinality.* The PSI cardinality protocols compute  $\text{PSI-CA}(X, Y_k) = |X \cap Y_k|$ . There exist two variants: the standard `PSI-CA` variant



**Figure 3: Single-set protocol structure.** SD refers to small domain variants. The optional blue parts ensure queries are well-formed.

**Algorithm 3** Single set procedures with *small input size*.

```

function QUERY(pk, X)
   $\llbracket x_i \rrbracket \leftarrow \text{HE.Enc}(pk, x_i)$ 
  return  $Q = \langle \llbracket x_i \rrbracket \rangle$ 

function PSI-PROCESS(pk, Q =  $\langle \llbracket x_i \rrbracket \rangle$ ,  $Y_k$ )
   $\llbracket s_i \rrbracket \leftarrow \text{HE.IsIn}(pk, \llbracket x_i \rrbracket, Y_k)$ 
  return  $M \leftarrow \langle \llbracket s_i \rrbracket \rangle$ 

function PSI-REVEAL(pk, M =  $\langle \llbracket s_i \rrbracket \rangle$ )
  return  $\{x_i \mid \text{HE.Dec}(sk, \llbracket s_i \rrbracket) = 0\}$ 

function ePSI-CA-PROCESS(pk, Q =  $\langle \llbracket x_i \rrbracket \rangle$ ,  $Y_k$ )
   $\langle \llbracket s_i \rrbracket \rangle \leftarrow \text{PSI-PROCESS}(pk, \langle \llbracket x_i \rrbracket \rangle, Y_k)$ 
   $\llbracket ca \rrbracket \leftarrow \sum_{i \in [m]} \text{HE.IsZero}(\llbracket s_i \rrbracket)$ 
  return  $M \leftarrow \llbracket ca \rrbracket$ 

function ePSI-CA-REVEAL(pk, M =  $\llbracket ca \rrbracket$ )
  return  $\text{HE.Dec}(sk, \llbracket ca \rrbracket)$ 

function QUERY-CHECK(pk, Q =  $\langle \llbracket x_i \rrbracket \rangle$ )
   $\llbracket T \rrbracket \leftarrow \prod_{i \in [Q], j \in [i-1]} (\llbracket x_i \rrbracket - \llbracket x_j \rrbracket)$ 
   $\llbracket R \rrbracket \leftarrow r \cdot \text{HE.IsZero}(pk, \llbracket T \rrbracket)$ 
  return  $\llbracket R \rrbracket$ 

```

in which the *client* learns the cardinality  $|X \cap Y_k|$  [13, 18], and the ePSI-CA variant in which the *server* learns an *encrypted* cardinality [84]. We focus on the latter to enable further computation on the intersection cardinality in the next layers.

Our ePSI-CA protocol (see ePSI-CA-PROCESS) first computes the inclusion statuses  $\llbracket s_i \rrbracket$  using PSI-PROCESS and then uses HE.IsZero to compute – in the ciphertext domain – the cardinality, i.e., the number of elements  $\llbracket s_i \rrbracket$  that are zero. When used as a stand-alone protocol, the server returns  $\llbracket ca \rrbracket$  to the client which decrypts it to obtain the answer (see ePSI-CA-REVEAL).

When the cardinality protocol is used as a stand-alone protocol without next layers, it is possible to mimic earlier work [13] and construct a cardinality protocol from the above-mentioned naive PSI protocol by shuffling server responses  $M$  before returning them.

*Efficiency.* While literature often dismisses OPE-based schemes due to their ‘quadratic’ total computation cost  $\mathcal{O}(|X| \cdot |Y_k|)$ , this approach excels in PCM scenarios with small client input. Our protocols achieve client computation and communication costs of  $\mathcal{O}(n_c)$ , which is independent of the server’s input size. While the ‘extra’ burden for the server is linear in the size of the client set which is a small constant.

## 6.2 Small input domain

When the set’s input domain  $D$  is small, sets can be efficiently represented and manipulated as bit-vectors [5, 67, 69]. Parties agree on a fixed ordering  $d_1, \dots, d_{|D|}$  of the elements in  $D$ . Then, clients use SD-QUERY to compute a vector of encrypted inclusion statuses

**Algorithm 4** Single set procedures with *small input domain*.

```

function SD-QUERY(pk, X)
   $z_i \leftarrow (d_i \in X)$ 
   $\llbracket z_i \rrbracket \leftarrow \text{HE.Enc}(pk, z_i)$ 
  return  $Q = \langle \llbracket z_i \rrbracket \rangle$ 

function PSI-SD-PROCESS(pk, Q =  $\langle \llbracket z_i \rrbracket \rangle$ ,  $Y_k$ )
   $v_i \leftarrow (d_i \in Y_k)$ 
   $\llbracket s_i \rrbracket \leftarrow \llbracket z_i \rrbracket \cdot v_i$ 
  return  $M \leftarrow \langle \llbracket s_i \rrbracket \rangle$ 

function ePSI-CA-SD-PROCESS(pk, Q =  $\langle \llbracket z_i \rrbracket \rangle$ ,  $Y_k$ )
   $\langle \llbracket s_i \rrbracket \rangle \leftarrow \text{PSI-SD-PROCESS}(pk, \langle \llbracket z_i \rrbracket \rangle, Y_k)$ 
   $\llbracket ca \rrbracket \leftarrow \sum_{d_i \in D} \llbracket s_i \rrbracket$ 
  return  $M \leftarrow \llbracket ca \rrbracket$ 

function SD-QUERY-CHECK(pk, Q =  $\langle \llbracket z_i \rrbracket \rangle$ )
   $\llbracket t_i \rrbracket \leftarrow \text{HE.IsIn}(pk, \llbracket z_d \rrbracket, \{0, 1\})$ 
   $\llbracket R \rrbracket \leftarrow \sum_{d \in D} \llbracket t_i \rrbracket$ 
  return  $\llbracket R \rrbracket$ 

```

$\langle \llbracket z_i \rrbracket \rangle$ , where  $z_i = 1$  iff  $d_i \in X$  and 0 otherwise, for all  $d_i$ s in  $D$ . We instantiate small domain procedures, except for REVEAL processes that are not impacted by the domain size, in Algorithm 4.

The function PSI-SD-PROCESS creates a bit vector of the intersection. The status  $\llbracket s_i \rrbracket$  is an encryption of 1 if  $d_i$  is present in both sets and 0 otherwise. To do so, the server multiplies the indicator  $\llbracket z_i \rrbracket$  with another binary indicator  $v_i$  determining whether the element  $d_i$  is present in the server set  $Y_k$ . The function ePSI-CA-SD-PROCESS computes the sum of the inclusion statuses  $\llbracket s_i \rrbracket$  for all domain values  $d_i \in D$ . The functions QUERY, QUERY-CHECK, and base layer PROCESS must have the same domain size. The rest of the functions and layers are not impacted by the choice of domain.

*Efficiency.* While the idea of representing sets as bit-vectors is not new [5, 67, 69], existing works dismiss FHE protocols as too costly and focus on additively homomorphic solutions. We use the inherent parallelism of schemes such as BFV [23], that we discuss in Section 10, to achieve lower computation and communication costs, especially in the many-set scenario. In Appendix D.1, we show that PSI-CA-SD has a competitive performance to the existing schemes such as Ruan et al. [67].

## 6.3 Ensuring well-formed queries

In this section we consider a specific class of misbehaving clients: those that encrypt a malformed input and submit it as a query. Since the client’s query is encrypted, the server cannot directly verify well-formedness of the query. In this section we show how we can ensure that if a client’s query is malformed, i.e., it does not correspond to a *valid set*  $X$ , the client’s output is random.

Zero-knowledge proofs are not practical in the FHE setting, so we rely on an HE technique to ensure queries are well-formed. The server uses (SD-)QUERY-CHECK to compute a randomizer term  $\llbracket R \rrbracket$  that is random in  $\mathbb{Z}_q$  if the client misbehaves and 0 otherwise. By adding  $\llbracket R \rrbracket$  to the result  $\llbracket A \rrbracket$ , misbehaving clients learn nothing about the real result. With abuse of notation, the server adds a vector of fresh randomizers when the result is a vector such as the output of PSI-SD-PROCESS. The server can amortize the cost of computing  $N$  randomizers  $\llbracket R_i \rrbracket$  for  $N$  variables  $i$ : The server first computes  $\llbracket R \rrbracket$  as before, then picks a fresh randomness  $\delta_i \leftarrow \mathbb{Z}_q^*$  and sets  $\llbracket R_i \rrbracket \leftarrow \delta_i \cdot \llbracket R \rrbracket$ .

*Small domain.* Misbehaving clients can submit non-binary ciphertexts  $\llbracket z_i \rrbracket$  to learn more than the cardinality. We compute a term  $\llbracket R \rrbracket$  which is zero when all  $z_i$ s are binary, and is random otherwise. The term  $\llbracket t_i \rrbracket \leftarrow \text{HE.IsIn}(pk, \llbracket z_i \rrbracket, \{0, 1\})$  evaluates to 0 if  $z_i \in \{0, 1\}$  and to a uniformly random element in  $\mathbb{Z}_q^*$  otherwise. Therefore, the distribution of  $\llbracket R \rrbracket \leftarrow \sum_{d \in D} \llbracket t_i \rrbracket$  will be close to uniformly random in  $\mathbb{Z}_q$  as long as at least one non-binary  $z_i$  exists in the client's query. See Appendix B for the exact distribution.

*Small input.* The client sends a list of encrypted values  $\llbracket x_i \rrbracket$  to the server. This list represents a set as long as all elements are distinct, so the server needs to ensure that no two client elements are equal. First, the server computes  $\llbracket T \rrbracket$ , the product of pairwise differences of the client elements. Since these multiplications are performed in a prime group  $\mathbb{Z}_q^*$ , the product will be zero if and only if two equal elements exist. Second, the server uses a uniformly random element  $r \leftarrow \mathbb{Z}_q$  in combination with the zero detection on  $\llbracket T \rrbracket$  to compute the additive randomizer  $\llbracket R \rrbracket$ .

Since the zero detection function is impractical, we provide a practical alternative protection method which deviates from our structure. The client can deterministically compute  $T$ , allowing us to protect the result in a multiplicative way by returning  $M \leftarrow \llbracket A \rrbracket \cdot \llbracket T \rrbracket$ . As long as  $T$  is not zero, which signifies a malicious query, the client can reverse  $T$  and recover  $A \leftarrow M \cdot T^{-1}$ .

Without this protection, misbehaving clients are still limited to submitting a list of scalar values, with a limited size, which is equivalent to a *multi-set*. Depending on the flexible matching function, allowing multi-set queries may or may not impact the security. On one hand, if the server reveals PSI-CA, allowing multi-sets may lead to the extraction of the intersection from the cardinality. On the other hand, if the server is computing F-Match, which checks  $X \subseteq Y_k$ , there is no difference between querying a set or a multi-set.

At first glance, QUERY-CHECK's quadratic computation cost  $\mathcal{O}(n_c^2)$  seems expensive. However, we target imbalanced scenarios where  $n_c \ll N_s$ . In Section 10.1, we discuss how this optional protection cost will be overshadowed by the cost of our PSI layer  $\mathcal{O}(n_c \cdot N_s)$ ; thus, not having any impact on our final cost.

*Next layers.* We note that the checks on well-formedness in the PSI layer extend to the matching and aggregation layers by instead applying the randomizer  $\llbracket R \rrbracket$  computed on the client's query to the results that the matching and aggregation layers compute.

## 7 MATCHING LAYER

Given the output of the PSI layer, the matching layer determines whether the server set  $Y_k$  is of interest to the client. The matching layer outputs a matching status  $\llbracket \gamma_k \rrbracket$ . Similar to the inclusion test, HE.IsIn in Section 5.1,  $\gamma_k$  is zero for sets of interest and a random value in  $\mathbb{Z}_q^*$  otherwise. The value  $\llbracket \gamma_k \rrbracket$  can be revealed as a binary output  $\lambda_k$  or passed-on to the next layer.

These matching operations can use either the small input or the small domain PSI layer protocols. To instantiate a matching protocol, the client and the server proceed as in Fig. 3, but plug in the desired PROCESS variant based on  $f_M$ . As MATCH-PROCESS functions have identical outputs, they share the same MATCH-REVEAL method.

We provide three matching functions  $f_M$ : full matching (F-Match), which determines if the client's query set is fully contained in the server set; threshold matching (Th-Match), which determines if

---

### Algorithm 5 Processing matching variants.

---

```

function F-MATCH-PROCESS( $pk, Q, Y_k$ ) ▷ Full match
   $\langle \llbracket s_i \rrbracket \rangle \leftarrow \text{PSI-PROCESS}(pk, Q, Y_k)$ 
   $\llbracket \gamma_k \rrbracket \leftarrow \sum_{i \in [n]} \llbracket s_i \rrbracket$ 
  return  $\llbracket \gamma_k \rrbracket$ 

function TH-MATCH-PROCESS( $pk, Q, Y_k, \Delta = t_{\min}$ ) ▷ Threshold
   $\llbracket ca \rrbracket \leftarrow \text{EPSI-CA-PROCESS}(pk, Q, Y_k)$ 
   $T \leftarrow \{t \mid t \in \mathbb{Z}_q, t_{\min} \leq t \leq \min(|Q|, |Y_k|)\}$  ▷ Threshold to set
   $\llbracket \gamma_k \rrbracket \leftarrow \text{HE.IsIn}(pk, \llbracket ca \rrbracket, T)$ 
  return  $\llbracket \gamma_k \rrbracket$ 

function TV-MATCH-PROCESS( $pk, Q, Y_k, \Delta = (t, \alpha, \beta)$ ) ▷ Tversky
   $\llbracket ca \rrbracket \leftarrow \text{EPSI-CA-PROCESS}(pk, Q, Y_k)$ 
   $(a, b, c) \leftarrow \text{Tversky-param-process}(\alpha, \beta, t)$ 
   $T = \{t \mid t \in \mathbb{Z}_q, 0 \leq t \leq (a - b - c)|Y_k|\}$ 
   $\llbracket \text{sizeX} \rrbracket \leftarrow \begin{cases} \llbracket |Q| \rrbracket & \text{For small input size variants} \\ \sum_i \llbracket z_i \rrbracket & \text{For small domain var. with } Q = \langle \llbracket z_i \rrbracket \rangle \end{cases}$ 
   $\llbracket T_v \rrbracket \leftarrow a \cdot \llbracket ca \rrbracket - b \llbracket \text{sizeX} \rrbracket - c|Y_k|$ 
   $\llbracket \gamma_k \rrbracket \leftarrow \text{HE.IsIn}(pk, \llbracket T_v \rrbracket, T)$ 
  return  $\llbracket \gamma_k \rrbracket$ 

function MATCH-REVEAL( $sk, M = \llbracket \gamma_k \rrbracket$ ) ▷ Reveal matching output
   $\gamma_k \leftarrow \text{HE.Dec}(sk, \llbracket \gamma_k \rrbracket)$ 
   $\lambda_k \leftarrow \mathbb{1}[\lambda_k = 0]$ 
  return  $\lambda_k$ 

```

---

the size of intersection exceeds a threshold; and Tversky matching (Tv-Match), which determines if the Tversky similarity between the client's and the server's set exceeds a threshold. The associated PROCESS variants are described in Algorithm 5.

*Full matching.* The F-Match variant determines if all the client elements are inside the server's set, i.e.,  $f_M(X, Y_k) = \mathbb{1}[X \subseteq Y_k]$ . The server first computes the inclusion statuses  $\llbracket s_i \rrbracket$  by calling PSI-PROCESS (see F-MATCH-PROCESS). Recall  $\llbracket s_i \rrbracket$  is zero when  $x_i \in Y_k$ . Therefore, when  $X \subseteq Y_k$ , the sum  $\llbracket \gamma_k \rrbracket$  of all  $\llbracket s_i \rrbracket$  is zero. When an element  $x_i$  is not in the server set, its inclusion status  $s_i$  is uniformly random in  $\mathbb{Z}_q^*$ , and therefore the sum  $\lambda_k$  is also random.

The F-Match protocol has a small false-positive probability when more than one  $x_i$  exists such that  $x_i \notin Y_k$ . Adding multiple random  $\llbracket s_i \rrbracket \in \mathbb{Z}_q^*$  PSI responses can, incorrectly, lead to a zero sum  $\llbracket \gamma_k \rrbracket$ . In Appendix B, we bound the probability of a false-positive to  $1/(q-1)$ . Moreover, we bound the difference between the distribution of the sum  $R$  and uniformly random over  $\mathbb{Z}_q$  to  $1/(q-1)^2$  at all points when more than one  $x_i$  is missing. The false-positive probability is zero when only one  $x_i$  is missing.

Note that the F-Match protocol computes containment and not equality. Thus, hashing and comparing the client set with server's set does not work as the server would need to hash every combination of  $|X|$  server set elements resulting in an exponential cost.

*Threshold matching.* The Th-Match variant determines if the two sets have at least  $t_{\min}$  elements in common, i.e.,  $f_M(X, Y_k; t_{\min}) = \mathbb{1}[|X \cap Y_k| \geq t_{\min}]$ . The server first computes the encrypted cardinality  $\llbracket ca \rrbracket$  using EPSI-CA-PROCESS, then evaluates the inequality  $ca \geq t_{\min}$  (see TH-MATCH-PROCESS).

Directly computing this one-sided inequality over encrypted values is costly. However, we know that  $|X \cap Y_k| \leq \min(|Q|, |Y_k|) = t_{\max}$ , where we bound the client set size by the size of the query. The server evaluates the inequality  $t_{\min} \leq ca \leq t_{\max}$  by performing the inclusion test  $ca \in \{t \mid t_{\min} \leq t \leq t_{\max}\}$  using HE.IsIn.

*Tversky similarity.* The Tv-Match variant determines if the Tversky similarity of the two sets exceeds a threshold  $t$ . Formally, the

protocol computes  $f_M(X, Y_k; \alpha, \beta, t) = \mathbb{1}[\text{Tv}_{\alpha, \beta}(X, Y_k) \geq t]$  where

$$\text{Tv}_{\alpha, \beta}(X, Y_k) = \frac{|X \cap Y_k|}{|X \cap Y_k| + \alpha|X - Y_k| + \beta|Y_k - X|}$$

is the Tversky similarity with parameters  $\alpha$  and  $\beta$ . Computing the Tversky similarity in this form is difficult as it requires floating-point operations. We assume  $t, \alpha$ , and  $\beta$  are rational, and known to both the client and the server. We follow the approach of Shimizu et al. [69] and transform the inequality  $\text{Tv}_{\alpha, \beta}(X, Y) \geq t$  to

$$\begin{aligned} (t^{-1} - 1 + \alpha + \beta)|X \cap Y| - \alpha|X| - \beta|Y| &\geq 0 \\ \Rightarrow (a, b, c) \in \mathbb{Z}_q^3, \quad a|X \cap Y| - b|X| - c|Y| &\geq 0 \end{aligned} \quad (1)$$

for appropriate integer values of  $a, b, c$ . The server either knows  $|X|$  ( $|X| = |Q|$ ) or can compute it ( $\llbracket |X| \rrbracket = \sum_i \llbracket z_i \rrbracket$  for small domain protocols). The server also knows  $|Y_k|$  and can compute  $\llbracket ca \rrbracket = |X \cap Y_k|$  using ePSI-CA-PROCESS. Evaluating the inequality requires two steps (see TV-MATCH-PROCESS):

*Step 1.* Transform coefficients  $(t^{-1} - 1 + \alpha + \beta)$ ,  $\alpha$  and  $\beta$  to equivalent integer coefficients  $a, b, c$ . We describe this in detail in Appendix A. *Step 2.* Evaluate the Tversky similarity inequality (1). We convert this inequality to a two-sided equation. We know  $|X| \geq |X \cap Y_k|$  and  $|Y_k| \geq |X \cap Y_k|$ , thus

$$a|X \cap Y_k| - b|X| - c|Y_k| \leq (a - b - c)|X \cap Y_k| \leq (a - b - c)|Y_k|.$$

Therefore, two sets  $X$  and  $Y_k$  satisfy  $\text{Tv}_{\alpha, \beta}(X, Y_k) \geq t$  iff

$$0 \leq a\llbracket ca \rrbracket - b|X| - c|Y_k| \leq (a - b - c)|Y_k|.$$

The server uses an inclusion test to evaluate this inequality.

## 8 AGGREGATION LAYER

The aggregation layer combines the outputs of the matching layer, over many sets, into a single collection-wide result. We provide four aggregation functions  $f_A$ : naive aggregation (NA-Agg), which returns outputs as is; existential search (X-Agg), which returns whether at least one server set matched; cardinality search (CA-Agg) which returns the number of matching server sets; and retrieval (Ret-Agg) which returns the index of the  $\kappa$ 'th matching server set.

Figure 4 shows the structure of the aggregation protocols. Upon receiving the query  $Q$ , the server runs the desired matching protocol MATCH-PROCESS (e.g., one from Algorithm 5) on each of its sets to compute the matching output  $\llbracket Y_j \rrbracket$ . The response  $\llbracket Y_j \rrbracket$  is zero if the set  $Y_j$  is interesting for the client and random otherwise. The server next runs an aggregation function AGG-PROCESS that takes  $N$  matching responses  $\llbracket Y_j \rrbracket$  as input and computes the final result  $A$ . We show how to instantiate AGG-PROCESS and REVEAL in Algorithm 6. If using naive aggregation (NA-Agg), the client runs NA-REVEAL. Otherwise, the client runs REVEAL to compute the result. Finally, the server can run (SD-)QUERY-CHECK on the query and apply it to the final result  $A$  for protection.

*Naive aggregation (NA-Agg).* The naive variant runs the matching protocol on all  $N$  server sets and returns  $N$  results to the client, i.e.,  $f_A(\lambda_1, \dots, \lambda_N) = \lambda_1, \dots, \lambda_N$ . This enables our framework to support many-sets when there is no need for aggregation and reduces the client's cost by computing and sending the query  $Q$  only once.

*Existential search (X-Agg).* The existential search variant determines if at least one server set  $Y_j$  is of interest to the client. Formally,

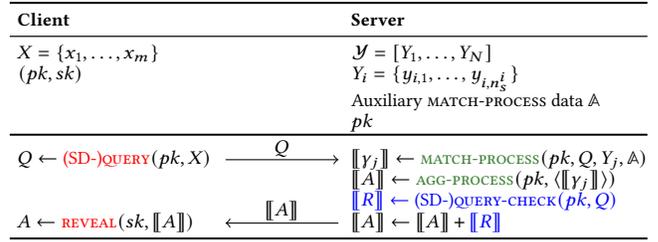


Figure 4: Aggregation protocol structure.

Algorithm 6 Processing  $f_M$ .

```

function NA-AGG-PROCESS( $pk, \langle \llbracket Y_j \rrbracket \rangle$ )
    return  $\langle \llbracket Y_j \rrbracket \rangle$ 
function NA-REVEAL( $sk, M = \langle \llbracket Y_j \rrbracket \rangle$ ) ▷ Naive output
     $\text{match}_j \leftarrow \text{MATCH-REVEAL}(sk, \llbracket Y_j \rrbracket)$ 
    return  $\langle \text{match}_j \rangle$ 
function X-AGG-PROCESS( $pk, \langle \llbracket Y_j \rrbracket \rangle$ ) ▷ At least one match
     $\llbracket A \rrbracket \leftarrow \prod_{j \in [N]} \llbracket Y_j \rrbracket$ 
    return  $\llbracket A \rrbracket$ 
function CA-AGG-PROCESS( $pk, \langle \llbracket Y_j \rrbracket \rangle$ ) ▷ Number of matches
     $\llbracket b_j \rrbracket \leftarrow \text{HE.IsZero}(pk, \llbracket Y_j \rrbracket)$ 
     $\llbracket A \rrbracket \leftarrow \sum_{j \in [N]} \llbracket b_j \rrbracket$ 
    return  $\llbracket A \rrbracket$ 
function RET-AGG-PROCESS( $pk, \langle \llbracket Y_j \rrbracket \rangle, \mathcal{D}, \kappa$ ) ▷ Retrieve data for the  $\kappa$ 'th match
     $\llbracket b_j \rrbracket \leftarrow \text{HE.IsZero}(pk, \llbracket Y_j \rrbracket)$ 
     $\llbracket \text{ctr}_0 \rrbracket \leftarrow \llbracket 0 \rrbracket$ 
     $\llbracket \text{ctr}_j \rrbracket \leftarrow \llbracket \text{ctr}_{j-1} \rrbracket + \llbracket b_j \rrbracket$  ▷ #matches before  $j$ 'th set
     $\llbracket I_j \rrbracket \leftarrow \text{HE.IsZero}(pk, \llbracket \text{ctr}_j \rrbracket \cdot \llbracket b_j \rrbracket - \kappa)$  ▷ Create index
     $\llbracket A \rrbracket \leftarrow \sum_{j \in [N]} \mathcal{D}_j \cdot \llbracket I_j \rrbracket$ 
    return  $\llbracket A \rrbracket$ 
function REVEAL( $sk, M = \llbracket A \rrbracket$ )
    return  $\text{HE.Dec}(sk, \llbracket A \rrbracket)$ 
    
```

the aggregation computes  $f_A(\lambda_1, \dots, \lambda_N) = \mathbb{1}[\exists i | \lambda_i = 1]$ . Recall that interesting sets produce zero matching responses  $\llbracket Y_j \rrbracket$ , so a collection will have an interesting set if and only if the product of matching responses is zero (see X-AGG-PROCESS). As responses  $\gamma_j$  are elements of the prime field  $\mathbb{Z}_q$ , their product will never be zero without having a zero response; thus, there will be no false-positives.

*Cardinality search (CA-Agg).* The cardinality search variant counts the number of interesting server sets  $Y_j$ , i.e.,  $f_A(\lambda_1, \dots, \lambda_N) = |\{i | \lambda_i = 1\}|$ . This aggregation (see CA-AGG-PROCESS) follows the same process as ePSI-CA and uses HE.IsZero to turn the matching responses  $\llbracket Y_j \rrbracket$  into binary values  $b_j$  and computes their sum.

Similar to the single set PSI-CA, we can use shuffling to convert the naive aggregation into cardinality search with minimal computational overhead. This gain comes at the cost of increased communication as the protocol needs to send the  $N$  shuffled set responses to the client instead of a single encrypted cardinality.

*Retrieval (Ret-Agg).* The retrieval variant returns associated data  $\mathcal{D}_j$  of the  $\kappa$ th matching server set  $Y_j$ . Formally,  $f_A(\lambda_1, \dots, \lambda_N) = \mathcal{D}_j | (\lambda_j = 1) \wedge (|\{i | \lambda_i = 1 \wedge i < j\}| = \kappa)$ . Clients use this variant when they are not concerned about whether a matching set exists, but rather about information related to this matching set – such as an index ( $\mathcal{D}_j = j$ ) for retrieving records. A good example is the matching scenario where apps want to retrieve a lot of data about

the matching records. Apps would first run the Ret-Agg protocol to retrieve the index of the matching record and then follow with a PIR request to retrieve the matching set's associated data.

The Ret-Agg protocol takes an input parameter  $\kappa$  denoting that the client wants to retrieve the associated data of the  $\kappa$ th matching set. The server builds an encrypted index of interesting sets in three steps (see RET-AGG-PROCESS): (1) The server uses HE.IsZero to compute  $\llbracket b_j \rrbracket$  indicating if a set is interesting. (2) The server computes a counter  $\llbracket \text{ctr}_j \rrbracket$  to track how many interesting sets exist in the first  $j$  sets. (3) The server combines  $\llbracket b_j \rrbracket$  and  $\llbracket \text{ctr}_j \rrbracket$  to compute an index  $\llbracket I_j \rrbracket$  where  $I_j$  is 1 if  $Y_j$  is the  $\kappa$ th interesting set, and zero otherwise. Adding weighted  $\llbracket I_j \rrbracket$  values produces the result.

## 9 SECURITY AND PRIVACY

Section 2.3 defines correctness (Definition 2), client privacy (Definition 3), and server privacy (Definition 4) of PCM protocols. We now discuss the correctness and privacy properties of our framework in the semi-honest and malicious settings. Our theorems rely on properties of the HE scheme, which are defined formally in Appendix C.1.

**Semi-honest setting.** When both parties are semi-honest, our protocols achieve all three security properties.

**THEOREM 1.** *Our protocols are correct, client private, and server private against semi-honest adversaries as long as the HE scheme is IND-CPA secure and circuit private.*

We simulate our protocols in a real-world/ideal-world setting to prove Theorem 1 in Appendix C.2.

**Malicious setting.** When the server is malicious, server privacy does not apply. Moreover, our framework does *not* provide any correctness guarantee to clients. Malicious servers can produce arbitrary, fixed (always match/no match) or input-dependent, responses to the clients without detection. However, servers cannot learn information about the clients' private input even if they misbehave.

**THEOREM 2.** *Our protocols provide client privacy against malicious servers, as long as the HE scheme is IND-CPA.*

In Appendix C.3, we prove Theorem 2 by reducing client privacy to the IND-CPA security of the underlying HE scheme. The client's sole interaction is sending encrypted queries. The server cannot extract information from these encrypted queries without knowing the secret key or breaking the security of the HE scheme.

When the client is malicious, client privacy does not apply. The correctness property, which requires that the malicious client's output is correct, does not apply either. In Appendix C.4, we reason about server privacy when the client is malicious in a non-standard model.

**System-wide security.** We proved the security of our framework in isolation, where there is no honest interaction outside our framework. In practice, when integrating our framework into a larger system, e.g., allowing journalists to contact document owners after a document search, user interactions may leak information such as the outcome of the search. Therefore, whenever protocol designers decide to integrate our framework into a larger system, they *must perform a system-wide security analysis*.

**Table 2: Asymptotic cost of our protocols. See Table 3 for a summary of our parameter definition. SD stands for small domain protocols. The column Calls denotes which protocols from earlier layers are called.**

	Calls	Add	Mult.	Exp.
PSI	-	$n_c n_s^i$	$n_c n_s^i$	
PSI-SD	-		$ D $	
ePSI-CA	1×PSI	$n_c$		
ePSI-CA-SD	1×PSI-SD	$ D $		
F-Match	1×PSI	$n_c$		
Th-Match	1×ePSI-CA	$ T $	$ T $	
Tv-Match	1×ePSI-CA	$ T $	$ T $	
NA-Agg	$N \times \text{Match}$			
X-Agg	$N \times \text{Match}$		$N$	
CA-Agg	$N \times \text{Match}$	$N$		$N$
Ret-Agg	$N \times \text{Match}$	$2N$	$2N$	$2N$
Query check	-		$n_c^2$	
SD-Query check	-	$ D $	$ D $	

## 10 FROM THEORY TO PRACTICE

We first analyze the asymptotic cost of our schemes, and explain the optimizations we implement to make our schemes practical.

### 10.1 Asymptotic cost

Our framework is modular and supports arbitrary protocol combinations. Hence, we report the cost of layers separately. Table 2 summarizes the asymptotic costs of modules in our framework. We report the number of homomorphic additions, multiplications, and exponentiations (counting scalar-ciphertext operations as ciphertext-ciphertext operations; but excluding operations in earlier layers). The PSI and matching layers' costs are reported for a single set (recall  $n_s^i$  is the size of the server's  $i$ 'th set). As the well-formedness check is optional, we report its one-time cost separately.

As an example, we compute the cost of using existential search with full-match interest criteria with no input domain restriction. The client runs a single X-Agg protocol, which per server set calls one F-Match and one small input PSI protocol. This leads to  $N$  multiplications in the aggregation layer,  $\sum_i^N (n_c) = N n_c$  additions in the matching layer, and  $\sum_i^N n_c n_s^i = n_c N_s$  multiplications in the PSI layer. Additionally, the server needs to perform the well-formedness check once, leading to  $n_c^2$  extra multiplications. The total cost will be  $N + N n_c + n_c N_s + n_c^2$ , but knowing that the total number of server elements ( $N_s$ ) is larger than the set number ( $N$ ) or the client size ( $n_c$ ), we simplify the cost to  $O(n_c N_s)$ .

**Communication.** The client sends a query  $Q$  in the PSI layer and receives a response  $A$  from the aggregation layer. The query's size only depends on the PSI layer variant and is independent from the server's input. Small input queries are an encrypted list of client elements containing  $n_c$  encrypted scalars while small domain queries are encrypted bit-vectors containing  $|D|$  scalar elements. The response size depends on the aggregation method. Naive aggregation has to produce  $N$  individual responses for  $N$  sets, while other aggregation methods produce a single scalar value.

This results in a total cost of  $O(|Q|)$  for non-naive and  $O(|Q| + N)$  for naive protocols which is the minimum achievable cost.

*Computation.* The PSI layer dominates the cost of our framework. Each PCM run calls  $N$  instances of PSI leading to a cost of  $O(n_c N_s)$  for small client input and  $O(N \cdot |D|)$  for small input domain variants. Our cost is asymptotically higher than alternatives [13, 37, 54] and is often dismissed as “quadratic” and too expensive. In contrast, approaches [3, 12, 14, 15, 57] with “linear” cost  $O(k(n_c + N_s))$ , where  $10 < k < 100$  determines the protocol’s false positive rate, are considered acceptable. The extreme imbalance requirement (RQ.3) leads to scenarios where  $n_c \ll N_s$  or even  $n_c < \log(N_s)$  and impacts how asymptotic costs should be interpreted. We show a how “quadratic” cost  $O(n_c N_s)$  can outperform the “linear” cost  $O(k(n_c + N_s))$  in Section 11.

## 10.2 Implementation

Fully homomorphic encryption schemes assume unbounded multiplication depth and rely on bootstrapping, which is prohibitively expensive. Thus, we use the somewhat homomorphic BFV cryptosystem [23] with a fixed multiplicative depth.

Let  $N_{deg}$  be the degree of the RLWE polynomial,  $m_{pt}$  be the plaintext modulus, and  $m_{ct}$  be the ciphertext modulus. The polynomial degree  $N_{deg}$  and ciphertext modulus  $m_{ct}$  determine the multiplicative depth of the scheme. The plaintext modulus determines the input domain. We define two sets of parameters:  $P_{8k}(N_{deg} = 8192, m_{pt} = 4,079,617)$  and  $P_{32k}(N_{deg} = 32,768, m_{pt} = 786,433)$ , and follow the Homomorphic Encryption Security Standard guidelines [1] to choose  $m_{cts}$  that provide 128 bits of security. We use relinearization keys to support multiplication, and rotation keys to support some of our optimizations (see next section). Generating and communicating keys is expensive. Therefore, we assume that clients generate these keys once at setup and use them for all subsequent protocols. In Appendix A, we provide more details on our parameters in Table 5 and a microbenchmark of basic BFV operations and key sizes in Table 6.

We implement a subset of our protocols using the Go language. Our code is open-source and 1,620 lines long.<sup>2</sup> We use the Lattigo library [41, 47] for BFV operations. Unfortunately, Lattigo does not support circuit privacy. We discuss the implications of this lack of support and possible countermeasures in Appendix C.1. We run experiments on a machine with an Intel i7-9700 @ 3.00 GHz processor and 16 GiB of RAM. Reported numbers are single-core costs. As the costly operations are inherently parallel, we expect our scheme to scale linearly with the number of cores.

## 10.3 Optimizations

We explain how we optimize our implementation and limit the multiplicative depth of (some of) our algorithms to improve efficiency.

*NTT batching.* We use BFV in combination with the number-theoretic transformation (NTT) so that a BFV ciphertext encodes a vector of  $N_{deg}$  elements [70]; BFV additions and multiplications act as element-wise vector operations. This batching enables single instruction multiple data (SIMD) operations on encrypted scalars. Performing operations between scalars in the same ciphertext (such

as computing the sum or product of elements) requires modifying their position through rotations. Rotations require rotation keys.

Batching renders HE.IsZero infeasible. The exponentiation with  $m_{pt} - 1$  consumes a multiplicative depth of  $\lg(m_{pt})$  in HE.IsZero, so  $m_{pt}$  must be small. To use batching, however, the plaintext modulus  $m_{pt}$  must be prime and large enough that  $2N_{deg} \mid m_{pt} - 1$ . Batching does not support small  $m_{pts}$  and consequently HE.IsZero; we prioritize efficiency and only evaluate variants that do not require zero detection. The parameters we use support batching.

*Replication.* The client’s query is small with respect to the capacity of batched ciphertexts, i.e.,  $|Q| \ll N_{deg}$ . We use two forms of replications to make full use of SIMD operations: powers and duplicates. When using a small input PSI variant, the client encodes powers of each element  $x$  (e.g.,  $[x, x^2, \dots, x^t]$ ) into the query ciphertext to reduce the multiplicative depth of HE.IsIn, see the second variant in Algorithm 2. Next, the client encodes  $k$  duplicates of the full query (including powers, when in use) regardless of the PSI variant.

Replication is straightforward when the client is semi-honest, but impacts security when the clients can misbehave. We follow a similar process to QUERY-CHECK to enforce correct replication. The server computes a second randomizer  $\llbracket R \rrbracket$  such that  $R$  will be zero when (1) all duplicates are equal and (2) for all consecutive power elements  $x_i$  and  $x_{i+1}$ , the equality  $x_{i+1} = x_1 \cdot x_i$  holds. The server needs to compute this check only once per query. We implemented this check and included its cost in all figures.

## 11 PCM IN PRACTICE

To demonstrate our framework’s capability, we solve the chemical similarity and document search problems. We discuss matching in mobile apps in Appendix E. We focus on end-to-end PCM solutions and do not evaluate single-set protocols or scenarios.

### 11.1 Chemical similarity

Recall from Section 2.1 that chemical similarity of compounds is determined by computing and comparing molecular fingerprints [8, 42, 48, 73, 79, 80]. We use the Tversky similarity matching algorithm, Tv-Match, to compute whether a compound in the seller’s database is similar to the query compound. As fingerprints are short, we instantiate Tv-Match with the small-domain ePSI-CA-SD and represent molecular fingerprint as bit-vectors.

We follow a popular configuration [69] where compounds are represented by 166-bit MACCS fingerprints [20] and Tversky parameters are set to  $\alpha = 1, \beta = 1, t = 80\%$ . Processing these raw parameters (see Algorithm 5) leads to the inequality

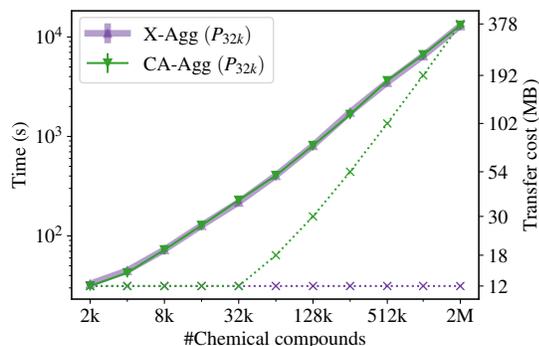
$$a, b, c = 9, 4, 4 \leftarrow \text{Tversky-param-process}(1, 1, 0.8)$$

$$0 \leq 9\llbracket ca \rrbracket - 4|X| - 4|Y| \leq |Y|.$$

We evaluate two aggregation policies. We apply X-Agg aggregation to determine whether at least one compound in the database matches and CA-Agg to count the number of matching compounds. These variants have high multiplicative depth, so we modify them to enable efficient deployment without relying on bootstrapping.

*Existential search.* The X-Agg protocol applied to a collection of  $N$  compounds requires a multiplicative depth of  $\lg(N)$  to compute  $\prod_{j \in [N]} \llbracket Y_j \rrbracket$ . This depth is too high for our parameters. Instead,

<sup>2</sup>Code available online: <https://github.com/spring-epfl/private-collection-matching>



**Figure 5: Computation time (solid lines) and transfer cost (dotted lines) for computing aggregated chemical similarity.**

we relax the output requirements and reduce the output as much as possible: For a fixed depth  $l$ , we return  $\lceil N/2^l \rceil$  encrypted scalars  $\llbracket \lambda_k \rrbracket = \prod_{k2^l \leq j < (k+1)2^l} \llbracket y_j \rrbracket$  to the client. This relaxation reduces the privacy of X-Agg – it is less than full X-Agg but better than CA-Agg – at the gain of efficiency. If this reduced privacy is unacceptable, the client just needs to choose larger HE parameters.

*Cardinality search.* We use the shuffling variant of the CA-Agg protocol due to its lower multiplicative depth. When the server shuffles the  $N$  matching responses, the client only learns the number of interesting compounds. Shuffling batched encrypted values is hard, so the server shuffles the compounds (server sets) as plaintext before processing the query to the same effect.

Our modifications to both aggregators increase the transfer cost from a single aggregated result to linear in the number of server sets. Yet, the X-Agg protocol sends 1 scalar value per  $2^l$  compounds.

*Evaluation.* We evaluate our similarity search on the ChEMBL database [26, 46], which contains more than 2 million compounds and is one of the largest public chemical databases. The database contains compounds in the SMILES format. We use the RDKit library [63] to compute MACCS fingerprints from this format.

Figure 5 shows the cost of running our protocols with the BFV parameter set  $P_{32k}$ . We ran 5 times the experiments with databases containing up to 256k compounds, and 3 times the larger experiments. We report the average cost. Standard errors of the mean are small, so error bars are not visible.

The server batches 128 compounds per ciphertext. Performing the PSI layer protocol ePSI-CA only requires 1 multiplication per ciphertext, while computing the binary Tversky score requires a depth of 7. For the X-Agg protocol, we set  $l = 6$  and aggregate up to 64 matching results into each scalar  $\llbracket y_k \rrbracket$ . As each ciphertext holds 32k scalars, the aggregation of up to  $2M$  server sets requires 1 ciphertext. CA-Agg transfers 1 ciphertext per 32k compounds. The cost of the extra X-Agg multiplications is insignificant compared to computing similarity, leading to similar computation costs for both protocols.

The client computation of searching among 2 million compounds is less than a second and the transfer cost is 12MB for X-Agg and 378MB for CA-Agg. These protocols can be run by a thin client. The server, however, requires 3.5 hours of single-core computation.

*Comparison to ad-hoc solution.* Shimizu et al. [69] build a custom chemical search that computes the number of matching

compounds. They offer the same functionality/privacy as our CA-Agg protocol, but reveal more information than our X-Agg variant. Shimizu et al. protect privacy against malicious adversaries, but malicious servers can violate correctness. They use the secp192k1 curve which provides less than 100-bit of security, while we offer full 128-bit security. Moreover, their use of differential privacy requires distributional assumptions. Shimizu et al. report the cost of searching 1.3 million compounds as: 167 seconds of server computation, 172 seconds of client computation, and 265MB of data transfer. Our X-Agg and CA-Agg solutions achieve higher security, lower client computation, and lower ratio of bandwidth consumed per compound, at the cost of higher server computation.

## 11.2 Peer-to-Peer document search

To implement peer-to-peer document search, we represent queries and documents by the sets of keywords they contain. A single document, represented by the keyword set  $S$ , is of interest to the querier if it contains all query keywords  $Q$ , i.e.,  $Q \subseteq S$ . This functionality can be implemented with the full matching (F-Match) variant.

The client and the server must agree on how keywords are represented in  $\mathbb{Z}_q$ . We use hash functions to do the conversion. As the search queries typically contain few keywords and the domain for searchable keywords is too large for our small domain variant, we construct F-Match with PSI with a small client input. There are two sources of false-positive in our setup:

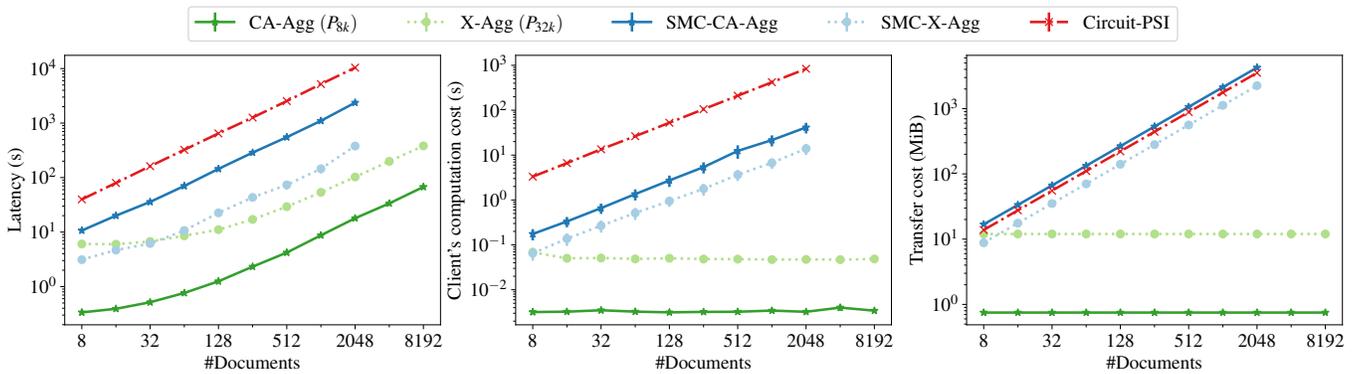
*False positive of mapping words.* The parameter  $m_{pt}$  (recall  $m_{pt} = q$ ) determines the input domain. Since  $m_{pt}$  is small, multiple words could be mapped to the same  $\mathbb{Z}_q$  element. This can lead to F-Match claiming there is a match, even though one of the colliding keywords is absent in the server’s set. Since  $m_{pt}$  impacts the multiplicative depth of our HE schemes, choosing a large value is impractical.

Instead of directly increasing the size of  $m_{pt}$  to reduce the false positive rate, the client hashes the  $n_c$  query keywords with  $t$  hash functions and encodes them into  $t n_c$  scalar values, which reduces the false-positive rate to  $1/q^t$ . When running PSI-PROCESS, the server knows the corresponding hash function for each scalar value and hashes them accordingly. Afterward, the server runs the F-Match protocol on all  $t n_c$  PSI outputs together; a document matches if and only if all hashed keywords are present.

Using multiple hashes to reduce false positives does not impact privacy, as it is straightforward to simulate a query with  $t$  hashes, with  $t$  F-Match queries. This modification increases the computation cost and the number of scalar values in the query by a factor of  $t$ . However, there is no concrete change in the communication cost as the client can still pack  $t n_c$  scalars into one ciphertext.

*False positive of F-Match.* The F-Match protocol itself has a false positive rate of  $\sim 1/q$  (see Section 7) caused by internal randomness. An easy way to reduce this FP rate is to run  $r$  instances of F-Match with different randomness and reveal the  $r$  responses. This repetition reduces the FP rate of a single matching response to  $\sim 1/q^r$ , while increasing the server’s computation cost by a factor  $r$ .

*Aggregation.* We explore two aggregation policies: existential (X-Agg) to determine whether at least one document in the collection matches; and cardinality (CA-Agg) to count matching documents. Since the multiplicative depth of F-Match is low, we can fully reduce



**Figure 6: The end-to-end search latency (left), client’s computation cost (middle), and communication cost (right) of document search. We limit execution to a single CPU core and enforce a bandwidth of 100 Mbps and an RTT of 100 ms on the connection.**

the X-Agg output to one encrypted scalar. For the CA-Agg variant, we still use the shuffling variant for lower multiplicative depth.

*Evaluation.* We use the parameters  $P_{8k}$  for CA-Agg and  $P_{32k}$  for X-Agg protocols. We base our evaluation on the requirements set out in EdalatNejad et al. [21]. We limit the number of keywords in each query to 8 and generate random documents of 128 keywords. We represent each keyword with two hash functions leading to false-positive rates of  $2^{-44}$  for CA-Agg and  $2^{-39}$  for X-Agg due to the mapping to  $\mathbb{Z}_q$ . We need to account for false-positive errors as we run a single F-Match per document. This error increases with the number of documents and reaches its peak at 8k documents where the probability of overestimating the cardinality is 0.2% (CA-Agg) and existence is 1% (X-Agg). Our protocols do not have false negatives. We skip QUERY-CHECK from Section 6.3 since F-Match is not impacted by repeated keywords in the query, but still perform checks from Section 10.3 to enforce honest batching. Our use of power replication leads to a multiplicative depth of 1 in the PSI layer while F-Match only uses addition. We report the cost of our document search in Fig. 6 and discuss our performance in Section 11.3.

*Comparison to ad-hoc solution.* EdalatNejad et al. [21] build a document search engine that performs PSI-CA in a many-set setting and post-processes the output, in the client, to detect relevant documents. EdalatNejad et al. protect privacy against malicious adversaries, but malicious servers can violate correctness. To enhance performance, they sacrifice privacy and leak more information than individual intersection cardinalities, yet less than vanilla PSI. The protocol of EdalatNejad et al. has a latency of less than 1 second to search 1k documents while our framework requires 8.7s for CA-Agg and 54s for X-Agg. Our CA-Agg search does not reveal information about individual documents and X-Agg *only reveals a single bit about the collection*. As expected, this privacy gain comes at a performance cost.

### 11.3 Comparison with generic solutions

After comparing against ad-hoc solutions, we compare our *document search* against generic SMC and circuit PSI that can offer the same privacy and functionality (see Table 1). See Appendix D.2 for

implementation details. In Appendix D.3, we compare our document search to one of the fastest OT-based solutions, which *cannot* satisfy our privacy requirements.

*Generic SMC.* We use a semi-honest SMC compiler, EMP toolkit [77], to build two custom search engines. Our circuits take 1 client and  $N$  server sets as input. They compute PSI and F-Match using boolean logic then aggregate the  $N$  matching results following either X-Agg or CA-Agg policy. We call these two approaches ‘SMC-X-Agg’ and ‘SMC-CA-Agg’.

*Circuit-PSI.* We compare against Chandran et al. [10] which is a state-of-the-art single-set circuit-PSI protocol. To support many-set, we run one PSI per document. Chandran et al. support extending the intersection computation with arbitrary circuits allowing us to perform F-Match and then X-Agg or CA-Agg aggregation. However, we consider the cost of computing the intersections as a lower bound for the cost of search and do not extend the circuit.

Figure 6 reports the latency, client’s computation, and transfer cost of document search. We repeat each experiment 4 times and report the average cost and the standard error of the mean. Unfortunately, the SMC and circuit-PSI solutions crash on runs with 4k documents and more, hence we only show entries up to 2k documents. The stable trend of our measurements let us think that one could easily extrapolate results for those settings from the measurements that we report. To estimate the latency of our framework, we run it in a single process and add the expected network time (as we have one round, we consider 1 round trip time plus transfer time).

Now we show that our framework supports thin clients by significantly reducing the clients’ computation and communication costs. Moreover, our framework provides better latency and in Appendix D.2 we show that we have a competitive server cost.

*Latency.* Both SMC and Circuit-PSI require many rounds of communication and are thus heavily impacted by the network’s round-trip time (RTT). We assume a transatlantic RTT of 100 ms. In this setting, our framework consistently outperforms competitors. When performing existential search (X-Agg), our framework cuts the end-to-end latency in half while our improvement factor increases to 96x when doing a cardinality search (CA-Agg) on 1k documents.

*Client's computation.* Circuit-based approaches have a balanced computation load between clients and servers while we outsource almost all the computations to the server. The client's computation cost of our framework is independent of the number of the server sets and is only 50 ms for the X-Agg and 5 ms for the CA-Agg search. While slower, the SMC approach is computationally efficient and only requires 3.7 s for the existential (X-Agg) and 11.3 s for the cardinality (CA-Agg) search for a 1k document collection. SMC's cost is still acceptable for thin clients, however using our framework leads to a saving factor of 75–2250x on client computation. In contrast, Chandran et al. require a prohibitive cost of 352 seconds, which is 70,000x higher than ours.

*Communication.* The X-Agg protocol has a constant communication cost of 12 MB while the cost of CA-Agg grows linearly with the number of documents. However, this cost is fixed to 768KB for CA-Agg in our evaluation since we can pack up to 8k results in a  $P_{8k}$  ciphertext. Both SMC and Circuit-PSI have costs linear in the inputs size. The SMC search requires 1.1 GiB for X-Agg and 2.1 GiB for CA-Agg queries when searching 1k documents while Chandran et al. require 1.7 GiB; thus, both approaches are prohibitively expensive. Ultimately, our framework reduces communication by a factor of 93–2800x.

## 12 TAKEAWAYS AND FUTURE WORK

In this work, we introduce and formalize private collection matching problems. Using homomorphic encryption, we build a modular framework for solving them. Our framework and its layers-based design simplify the construction of PCM schemes that achieve complex goals while limiting the leakage to what is required by the functionality, nothing more. Relying on homomorphic encryption is extremely advantageous in theory. Our work shows, however, that using it in practice is challenging. We have overcome these challenges by using optimizations from the literature, at the cost of reduced flexibility of our framework. Ultimately, our framework is competitive with ad-hoc solutions and outperforms generic approaches in all three latency, communication, and computation costs, sometimes by several orders of magnitude. As example, our framework requires 12MB of communication and less than a second of client computation to determine matching of a chemical compound against a database of 2 million compounds; or respectively 768KB and 50ms to determine whether the owner of a thousand documents has content of interest to a querier journalist. We believe further work on homomorphic encryption schemes – and ciphertext-based comparison methods in particular – will allow our framework to operate in a wider range of settings. We hope that our work fosters the evolution of homomorphic encryption in further areas so that the community can build a wider range of privacy-preserving applications.

## ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive feedback and the shepherd for helping us improve the paper. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

## REFERENCES

- [1] Martin Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin Lauter, Satya Lokam, Daniele Micciancio, Dustin Moody, Travis Morrison, Amit Sahai, and Vinod Vaikuntanathan. 2018. *Homomorphic Encryption Security Standard*. Technical Report. HomomorphicEncryption.org.
- [2] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. 2015. Ciphers for MPC and FHE. In *EUROCRYPT*.
- [3] Giuseppe Ateniese, Emiliano De Cristofaro, and Gene Tsudik. 2011. (If) Size Matters: Size-Hiding Private Set Intersection. In *PKC*.
- [4] Pierre Baldi, Roberta Baronio, Emiliano De Cristofaro, Paolo Gasti, and Gene Tsudik. 2011. Countering GATTACA: efficient and secure testing of fully-sequenced human genomes. In *CCS*.
- [5] Asli Bay, Zeki Erkin, Mina Alishahi, and Jelle Vos. 2021. Multi-Party Private Set Intersection Protocols for Practical Applications. In *SECRYPT*.
- [6] Donald Beaver, Silvio Micali, and Phillip Rogaway. 1990. The Round Complexity of Secure Protocols (Extended Abstract). In *STOC*.
- [7] Martin Burkhart, Mario Strasser, Dilip Many, and Xenofontas A. Dimitropoulos. 2010. SEPIA: Privacy-Preserving Aggregation of Multi-Domain Network Events and Statistics. In *USENIX*.
- [8] Adrià Cereto-Massagué, María José Ojeda, Cristina Valls, Miquel Mulero, Santiago Garcia-Vallvé, and Gerard Pujadas. 2015. Molecular fingerprint similarity search in virtual screening. *Methods* (2015).
- [9] Anrin Chakraborti and Radu Sion. 2019. ConcurORAM: High-Throughput Stateless Parallel Multi-Client ORAM. In *NDSS*.
- [10] Nishanth Chandran, Divya Gupta, and Akash Shah. 2022. Circuit-PSI With Linear Complexity via Relaxed Batch OPRF. *POPETs* (2022).
- [11] Hao Chen, Kim Laine, and Peter Rindal. 2017. Fast Private Set Intersection from Homomorphic Encryption. In *CCS*.
- [12] Michele Ciampi and Claudio Orlandi. 2018. Combining Private Set-Intersection with Secure Two-Party Computation. In *SCN*.
- [13] Emiliano De Cristofaro, Paolo Gasti, and Gene Tsudik. 2012. Fast and Private Computation of Cardinality of Set Intersection and Union. In *CANS*.
- [14] Emiliano De Cristofaro, Jihye Kim, and Gene Tsudik. 2010. Linear-Complexity Private Set Intersection Protocols Secure in Malicious Model. In *ASIACRYPT*.
- [15] Emiliano De Cristofaro and Gene Tsudik. 2009. Practical Private Set Intersection Protocols with Linear Computational and Bandwidth Complexity. *IACR Cryptol. ePrint Arch.* (2009).
- [16] Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Moti Yung. 2009. Efficient robust private set intersection. In *ACNS*.
- [17] Leo de Castro, Chiraag Juvekar, and Vinod Vaikuntanathan. 2021. Fast Vector Oblivious Linear Evaluation from Ring Learning with Errors. In *WAHC@CCS*.
- [18] Sumit Kumar Debnath and Ratna Dutta. 2015. Secure and Efficient Private Set Intersection Cardinality Using Bloom Filter. In *ISC*.
- [19] Daniel Demmler, Peter Rindal, Mike Rosulek, and Ni Trieu. 2018. PIR-PSI: Scaling Private Contact Discovery. *POPETs* (2018).
- [20] Durant, Joseph L and Leland, Burton A and Henry, Douglas R and Nourse, James G. 2002. Reoptimization of MDL keys for use in drug discovery. *Journal of chemical information and computer sciences* (2002).
- [21] Kasra Edalatnejad, Wouter Lueks, Julien Pierre Martin, Soline Ledéser, Anne L'Hôte, Bruno Thomas, Laurent Girod, and Carmela Troncoso. 2020. *DashboardNetwork: A Decentralized Privacy-Preserving Search Engine for Investigative Journalists*. In *USENIX*.
- [22] Brett Hemenway Falk, Daniel Noble, and Rafail Ostrovsky. 2019. Private Set Intersection with Linear Communication from General Assumptions. In *WPES@CCS*.
- [23] Junfeng Fan and Frederik Vercauteren. 2012. Somewhat Practical Fully Homomorphic Encryption. *IACR Cryptol. ePrint Arch.* (2012).
- [24] Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. 2004. Efficient Private Matching and Set Intersection. In *EUROCRYPT*.
- [25] Taher El Gamal. 1985. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inf. Theory* (1985).
- [26] Gaulton, Anna and Bellis, Louisa J and Bento, A Patricia and Chambers, Jon and Davies, Mark and Hersey, Anne and Light, Yvonne and McGlinchey, Shaun and Michalovich, David and Al-Lazikani, Bissan and others. 2012. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic acids research* (2012).
- [27] Craig Gentry. 2009. *A fully homomorphic encryption scheme*. Ph. D. Dissertation. Stanford University, USA.
- [28] Satrajit Ghosh and Mark Simkin. 2019. The Communication Complexity of Threshold Private Set Intersection. In *CRYPTO*.
- [29] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *STOC*.
- [30] Carmit Hazay. 2015. Oblivious Polynomial Evaluation and Secure Set-Intersection from Algebraic PRFs. In *TCC*.
- [31] Yan Huang, David Evans, and Jonathan Katz. 2012. Private Set Intersection: Are Garbled Circuits Better than Custom Protocols?. In *NDSS*.
- [32] Mihaela Ion, Ben Kreuter, Erhan Nergiz, Sarvar Patel, Shobhit Saxena, Karn Seth, David Shanahan, and Moti Yung. 2017. Private Intersection-Sum Protocol with

- Applications to Attributing Aggregate Ad Conversions. *IACR Cryptol. ePrint Arch.* (2017).
- [33] Paul Jaccard. 1912. The distribution of the flora in the alpine zone. *New phytologist* (1912).
- [34] Daniel Kales, Christian Rechberger, Thomas Schneider, Matthias Senker, and Christian Weinert. 2019. Mobile Private Contact Discovery at Scale. In *USENIX*.
- [35] Seny Kamara, Payman Mohassel, Mariana Raykova, and Saeed Sadeghian. 2013. Scaling Private Set Intersection to Billion-Element Sets. In *FC*.
- [36] Ferhat Karakoç and Alptekin Küpçü. 2020. Linear Complexity Private Set Intersection for Secure Two-Party Protocols. In *CANS*.
- [37] Agnes Kiss, Jian Liu, Thomas Schneider, N. Asokan, and Benny Pinkas. 2017. Private Set Intersection for Unequal Set Sizes with Mobile Applications. *PoPETS* (2017).
- [38] Lea Kissner and Dawn Xiaodong Song. 2005. Privacy-Preserving Set Operations. In *CRYPTO*.
- [39] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. 2016. Efficient Batched Oblivious PRF with Applications to Private Set Intersection. In *CCS*.
- [40] Vladimir Kolesnikov, Naor Matania, Benny Pinkas, Mike Rosulek, and Ni Trieu. 2017. Practical Multi-party Private Set Intersection from Symmetric-Key Techniques. In *CCS*.
- [41] Lattigo v2.1.1 2020. Lattigo. <http://github.com/ldsec/lattigo>. EPFL-LDS.
- [42] Laufkötter, Oliver and Miyao, Tomoyuki and Bajorath, Jürgen. 2019. Large-Scale Comparison of Alternative Similarity Search Strategies with Varying Chemical Information Contents. *ACS Omega* (2019).
- [43] Yehuda Lindell. 2017. How to Simulate It - A Tutorial on the Simulation Proof Technique. In *Tutorials on the Foundations of Cryptography*.
- [44] Jack P. K. Ma and Sherman S. M. Chow. 2022. Secure-Computation-Friendly Private Set Intersection from Oblivious Compact Graph Evaluation. In *AsiaCCS*.
- [45] Matteo Maffei, Giulio Malavolta, Manuel Reinert, and Dominique Schröder. 2017. Maliciously Secure Multi-Client ORAM. In *ACNS*.
- [46] David Mendez, Anna Gaulton, A Patricia Bento, Jon Chambers, Marleen De Veij, Eloy Félix, Maria Paula Magariños, Juan F Mosquera, Prudence Mutowo, Michal Nowotka, et al. 2019. ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Research* (2019).
- [47] Christian Mouchet, Juan Ramón Troncoso-Pastoriza, and Jean-Pierre Hubaux. 2020. Multiparty Homomorphic Encryption: From Theory to Practice. *IACR Cryptol. ePrint Arch.* (2020).
- [48] Ingo Muegge and Prasenjit Mukherjee. 2016. An overview of molecular fingerprint similarity search in virtual screening. *Expert Opinion on Drug Discovery* (2016).
- [49] Shishir Nagaraja, Prateek Mittal, Chi-Yao Hong, Matthew Caesar, and Nikita Borisov. 2010. BotGrep: Finding P2P Bots with Structured Graph Analysis. In *USENIX*.
- [50] Marcin Nagy, Emiliano De Cristofaro, Alexandra Dmitrienko, N. Asokan, and Ahmad-Reza Sadeghi. 2013. Do I Know You? Efficient and Privacy-Preserving Common Friend-Finder Protocols and Applications. In *ACSAC*.
- [51] OkCupid. OkCupid: Free Online Dating. <https://www.okcupid.com/>.
- [52] Pascal Paillier. 1999. Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In *EUROCRYPT*.
- [53] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. 2019. SpOT-Light: Lightweight Private Set Intersection from Sparse OT Extension. In *CRYPTO*.
- [54] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. 2020. PSI from PaXoS: Fast, Malicious Private Set Intersection. In *EUROCRYPT*.
- [55] Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. 2015. Phasing: Private Set Intersection Using Permutation-based Hashing. In *USENIX*.
- [56] Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. 2009. Secure Two-Party Computation Is Practical. In *ASIACRYPT*.
- [57] Benny Pinkas, Thomas Schneider, Oleksandr Tkachenko, and Avishay Yanai. 2019. Efficient Circuit-Based PSI with Linear Communication. In *EUROCRYPT*.
- [58] Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. 2018. Efficient Circuit-Based PSI via Cuckoo Hashing. In *EUROCRYPT*.
- [59] Benny Pinkas, Thomas Schneider, and Michael Zohner. 2014. Faster Private Set Intersection Based on OT Extension. In *USENIX*.
- [60] Benny Pinkas, Thomas Schneider, and Michael Zohner. 2018. Scalable Private Set Intersection Based on OT Extension. *ACM Trans. Priv. Secur.* (2018).
- [61] Rishabh Poddar, Tobias Boelter, and Raluca Ada Popa. 2019. Arx: An Encrypted Database using Semantically Secure Encryption. *Proc. VLDB Endow.* (2019).
- [62] Raluca A. Popa, Catherine M. S. Redfield, Nikolai Zeldovich, and Hari Balakrishnan. 2011. CryptDB: protecting confidentiality with encrypted query processing. In *SOSP*.
- [63] RDKit. RDKit: Open-source cheminformatics. <http://www.rdkit.org>.
- [64] Peter Rindal and Philipp Schoppmann. 2021. VOLE-PSI: Fast OPRF and Circuit-PSI from Vector-OLE. In *EUROCRYPT*.
- [65] Ronald L Rivest, Len Adleman, and Michael L Dertouzos. 1978. On data banks and privacy homomorphisms. *Foundations of Secure Computation* (1978).
- [66] Mike Rosulek and Ni Trieu. 2021. Compact and Malicious Private Set Intersection for Small Sets. In *CCS*.
- [67] Ou Ruan, Zihao Wang, Jing Mi, and Mingwu Zhang. 2019. New Approach to Set Representation and Practical Private Set-Intersection Protocols. *IEEE Access* (2019).
- [68] SEAL 2022. Microsoft SEAL (release 4.0). <https://github.com/Microsoft/SEAL>. Microsoft Research, Redmond, WA..
- [69] Kana Shimizu, Koji Nuida, Hiromi Arai, Shigeo Mitsunari, Nuttapon Attrapadung, Michiaki Hamada, Koji Tsuda, Takatsugu Hirokawa, Jun Sakuma, Goichiro Hanaoka, et al. 2015. Privacy-preserving search for chemical compound databases. *BMC bioinformatics* (2015).
- [70] Nigel P. Smart and Frederik Vercauteren. 2014. Fully homomorphic SIMD operations. *Des. Codes Cryptogr.* (2014).
- [71] Emil Stefanov, Marten van Dijk, Elaine Shi, T.-H. Hubert Chan, Christopher W. Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. 2018. Path ORAM: An Extremely Simple Oblivious RAM Protocol. *J. ACM* (2018).
- [72] Strava Route Explorer. Strava Rolls Out Significant New Routes Feature. <https://www.dcrainmaker.com/2020/03/strava-rolls-out-significant-new-routes-feature.html>.
- [73] Dagmar Stumpfe and Jürgen Bajorath. 2011. Similarity searching. *Wiley Interdisciplinary Reviews: Computational Molecular Science* (2011).
- [74] The Fork. The Fork: Restaurant Booking System. <https://www.thefork.com/>.
- [75] Tinder. Tinder: An Online Dating Application. <https://tinder.com/>.
- [76] Amos Tversky. 1977. Features of similarity. *Psychological review* (1977).
- [77] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. 2016. EMP-toolkit: Efficient MultiParty computation toolkit. <https://github.com/emp-toolkit>.
- [78] Zhusheng Wang, Karim Banawan, and Sennur Ulukus. 2021. Multi-Party Private Set Intersection: An Information-Theoretic Approach. *IEEE J. Sel. Areas Inf. Theory* (2021).
- [79] Peter Willett, John M Barnard, and Geoffrey M Downs. 1998. Chemical similarity searching. *Journal of chemical information and computer sciences* (1998).
- [80] Ling Xue, Florence L Stahura, Jeffrey W Godden, and Jürgen Bajorath. 2001. Mini-fingerprints detect similar activity of receptor ligands previously recognized only by three-dimensional pharmacophore-based methods. *Journal of chemical information and computer sciences* (2001).
- [81] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets (Extended Abstract). In *FOCS*.
- [82] Jason H. M. Ying, Shuwei Cao, Geong Sen Poh, Jia Xu, and Hoon Wei Lim. 2022. PSI-Stats: Private Set Intersection Protocols Supporting Secure Statistical Functions. In *ACNS*.
- [83] Yongjun Zhao and Sherman S. M. Chow. 2017. Are you The One to Share? Secret Transfer with Access Structure. *PoPETS* (2017).
- [84] Yongjun Zhao and Sherman S. M. Chow. 2018. Can You Find The One for Me?. In *WPES@CCS*.

## A EXTRA MATERIAL

We provide extra materials in this section.

**Summary.** Table 3 summarizes our notation and asymptotic parameters and Table 4 summarizes the functionality of our protocols.

**Tversky similarity.** Algorithm 7 shows how to process rational Tversky parameters to enable computing similarity with modular arithmetic.

**BFV parameters.** We report full details of our BFV parameters including their supported multiplicative depth in Table 5. Next we provide a microbenchmark of basic operations and key sizes in Table 6. It is possible to reduce the size of the rotation keys in exchange for more costly rotation operations. We report key sizes that provide a balanced computation/communication trade-off.

## B SUM OF RANDOM $\mathbb{Z}_q^*$ ELEMENTS

In Section 7, we argued that the distribution of the sum  $s = \sum_{i=1}^k x_i$  of  $k$  random  $x_i \leftarrow_s \mathbb{Z}_q^*$  elements is close to uniform when  $q$  is prime. Now we prove that the probability of the sum being zero is bounded by  $1/(q-1)$  and the difference between the probability of the sum being zero vs a non-zero value  $a$  is at most  $1/(q-1)^2$  when  $k$  is larger than one.

**Table 3: Notation.**

$q, \mathbb{Z}_q, \mathbb{Z}_q^*$	A prime number, a prime ring, and a prime field.
$\ell$	The security parameter.
$a \leftarrow_s A$	Draw $a$ uniformly at random from the set $A$ .
$[n], \langle a_i \rangle$	Present the set $\{1, \dots, n\}$ and the list $[a_1, \dots, a_m]$ .
$\mathbb{1}[E]$	Function that returns '1' when $E$ is true and '0' otherwise.
HE	An IND-CPA circuit-private homomorphic scheme.
$pk, sk$	The client's public and private HE keys.
$\llbracket a \rrbracket$	An encryption of $a$ .
$N_{deg}$	The degree of the RLWE polynomial.
$m_{pt}, m_{ct}$	The plaintext and ciphertext modulus of the HE scheme.
$X, \mathcal{Y}, Y_i$	The client's set, server's collection, and server's $i$ 'th set.
$n_c, n_s^i$	The size of client set $ X $ and server's $i$ 'th set $ Y_i $ .
$N, N_s$	The number of server sets and their total size $N_s = \sum_i n_s^i$ .
$D$	The set input domain.
$T$	The matching threshold $[t_{min}, t_{max}]$ .
$\llbracket z_i \rrbracket$	The client's encrypted bit-vector $z_i \leftarrow (d_i \in X)$ .
$Q$	The client's query. Small input: $\langle \llbracket x_i \rrbracket \rangle$ , domain: $\langle \llbracket z_i \rrbracket \rangle$ .
$\llbracket s_i \rrbracket$	An encrypted status determining iff $x_i \in Y_k$ .
$\llbracket ca \rrbracket$	An encrypted cardinality of intersection $ca =  X \cap Y_k $ .
$\llbracket Y_k \rrbracket$	A matching response. $\lambda = 0$ iff the set $Y_k$ is interesting.
$\llbracket A \rrbracket$	An aggregated collection-wide response.
$\llbracket R \rrbracket$	A term to randomize the output of malicious users.

**Table 4: Summary of our protocols. We show the computed functionalities, their output range, and auxiliary input variables in the table.**

Protocol	Function	Range	Aux.
PSI	$X \cap Y$	$\{0, 1\}^{n_c}$	
PSI-CA	$ X \cap Y $	$\mathbb{Z}$	
Matching: $\lambda \leftarrow f_M(X, Y)$			
F-Match	$X \subseteq Y$	$\{0, 1\}$	
Th-Match	$ X \cap Y  \geq t$	$\{0, 1\}$	$t$
Tv-Match	$Tv_{\alpha, \beta}(X, Y) \geq t$	$\{0, 1\}$	$t, \alpha, \beta$
Aggregation: $A \leftarrow f_A(\lambda_1, \dots, \lambda_N)$			
NA-Agg	$\langle \lambda_i \rangle$	$\{0, 1\}^N$	
X-Agg	$\exists i   \lambda_i = 1$	$\{0, 1\}$	
CA-Agg	$ \{i   \lambda_i = 1\} $	$\mathbb{Z}$	
Ret-Agg	$D_j   \lambda_j = 1 \wedge  \{i   \lambda_i = 1 \wedge i \in [j]\}  = \kappa$	$\mathbb{Z}$	$D, \kappa$

**Table 5: BFV parameters with 128-bit security**

	$N_{deg}$	$m_{pt}$	$\lg(m_{ct})$	Mult. depth
$P_{8k}$	8192	4079617	218-bit	2
$P_{16k}$	16384	163841	438-bit	7
$P_{32k}$	32768	786433	880-bit	16

Let  $z^{[k]}$  be the probability that the of sum of  $k$  elements from  $\mathbb{Z}_q^*$  is zero and  $p_a^{[k]}$  be the probability that the of sum is a non-zero

**Algorithm 7** Process Tversky parameters  $t, \alpha$ , and  $\beta$  to compute integer coefficients  $(a, b, c)$ .

```

function TVERSKY-PARAM-PROCESS( $\alpha, \beta, t$ )
  ( $\alpha_1, \alpha_2$ )  $\leftarrow$  ToRational( $\alpha$ )
  ( $\beta_1, \beta_2$ )  $\leftarrow$  ToRational( $\beta$ )
  ( $t_1, t_2$ )  $\leftarrow$  ToRational( $t$ )
   $l \leftarrow$  LCM( $t_1, \alpha_2, \beta_2$ )
   $a \leftarrow l \cdot (t^{-1} - 1 + \alpha + \beta)$ 
   $b \leftarrow l \cdot \alpha$ 
   $c \leftarrow l \cdot \beta$ 
   $g \leftarrow$  GCD( $a, b, c$ )
  return ( $a/g, b/g, c/g$ )

```

$\triangleright \alpha = \alpha_1/\alpha_2$   
 $\triangleright \beta = \beta_1/\beta_2$   
 $\triangleright t = t_1/t_2$

**Table 6: Cost of basic BFV operations.**

	$P_{8k}$	$P_{16k}$	$P_{32k}$
Addition ( $\mu$ s)	29	115	530
Multiplication (ms)	7.3	36.3	182
Plaintext mult. (ms)	0.97	4.13	18.3
Rotation by 1 (ms)	2.16	10.8	57
Ciphertext (KB)	384	1536	6144
Public key (KB)	512	2048	7680
Relinearization key (MB)	3	12	60
Rotation key (MB)	22	96	510

value  $a$ . When  $k = 1$ , e.g., we sum one element, these probabilities are  $z^{[1]} = 0$  and  $p_a^{[1]} = 1/(q-1)$ . The  $k$  elements  $x_i$  are independent from each other so we choose the value of the last element and recursively compute the probability by using the distribution of  $k-1$  elements.

$$\begin{aligned}
 z^{[k]} &= \sum_{a=1}^{q-1} p_a^{[k-1]} / (q-1) = p_a^{[k-1]} \\
 p_a^{[k]} &= z^{[k-1]} / (q-1) + \sum_{i \in \mathbb{Z}_q^* - \{a\}} p_{a-i}^{[k-1]} / (q-1) \\
 &= z^{[k-1]} / (q-1) + (q-2) \cdot p_a^{[k-1]} / (q-1) \\
 &= p_a^{[k-1]} + (z^{[k-1]} - p_a^{[k-1]}) / (q-1)
 \end{aligned}$$

The probability gap of  $z^{[k]}$  and  $p_a^{[k]}$  gets narrower as the number of elements  $k$  increases as:

$$\begin{aligned}
 p_a^{[k]} - z^{[k]} &= p_a^{[k-1]} + (z^{[k-1]} - p_a^{[k-1]}) / (q-1) - p_a^{[k-1]} \\
 &= (z^{[k-1]} - p_a^{[k-1]}) / (q-1)
 \end{aligned}$$

The highest probability of a zero sum happens when summing two random elements and the probability is bound by  $z^{[2]} = p_a^{[1]} = 1/(q-1)$ . Similarly the highest probability difference happens when  $k = 2$  and is bound by  $z^{[2]} - p_a^{[2]} = (p_a^{[1]} - z^{[1]}) / (q-1) = 1/(q-1)^2$ . This means that the probability of false-positive in our approach is bounded by  $1/(q-1)$  while the probability of distinguishing the number of non-zero elements (when at least one non-zero element is present) is  $1/(q-1)^2$ .

## C PRIVACY PROOF

We prove the security and privacy of our framework. We throughout assume that honest users do not have any interaction outside our framework that can leak information.

**Roadmap.** We start by formally defining security properties of HE schemes such as IND-CPA security, circuit privacy, and strong input privacy in Appendix C.1. Next, we address our framework’s security in a semi-honest setting and use real-world/ideal-world simulation to prove Theorem 1 in Appendix C.2. We study malicious servers in Appendix C.3 and provide a tight reduction from our client privacy to the semantic security of our HE scheme to prove Theorem 2. Finally, we extend the notion of real-world/ideal-world simulation to a new paradigm called cipher-world, which provides better support for simulating HE protocols in a malicious setting. We use this new notion to prove the server privacy of our scheme against malicious clients in a non-standard manner in Appendix C.4.

### C.1 Security properties of HE schemes

We formally define the security properties of HE schemes in this section. We start with semantic security (IND-CPA). Next, we discuss noise in HE schemes and define circuit privacy, which ensures that the noise contained in ciphertexts does not leak information about the computation performed on them. Finally, we extend circuit privacy, which only applies in a semi-honest setting, to a malicious setting and introduce strong input privacy.

**Definition 5** (IND-CPA). An encryption scheme is indistinguishable against chosen plaintext attacks if no PPT adversary  $\mathcal{A}$  exists such that:

$$\Pr \left[ \begin{array}{l} p \leftarrow \text{HE.ParamGen}(1^\ell) \\ (pk, sk) \leftarrow \text{HE.KeyGen}(p) \\ (st, m_0, m_1) \leftarrow \mathcal{A}(pk) \\ b \leftarrow_s \{0, 1\} \\ c_b \leftarrow \text{HE.Enc}(pk, m_b) \\ b' \leftarrow \mathcal{A}(st, c_b) \end{array} : b = b' \right] > \frac{1}{2} + \epsilon$$

Each BFV ciphertext  $c$  contains noise. The amount of noise increases with each operation and can be measured. As a result, a ciphertext  $c$  contains more information than its decrypted value  $p \leftarrow \text{HE.Dec}(sk, c)$ . If the noise grows beyond the HE parameter’s noise budget, then the decryption fails  $\perp \leftarrow \text{HE.Dec}(sk, c)$ . Informally, circuit privacy states that the ciphertext  $c$  should not reveal any information about the computation performed on  $c$  beyond the decrypted result  $p$ . We follow the definition of Castro et al. [17]:

**Definition 6** (Circuit privacy). Let HE be a leveled homomorphic encryption scheme and let

$$\begin{aligned} \text{params} &\leftarrow \text{HE.ParamGen}(q) \\ (sk, pk) &\leftarrow \text{HE.KeyGen}(\text{params}) \\ c_i &\leftarrow \text{HE.Enc}(sk, m_i) \\ M &\leftarrow f(m_1, \dots, m_n, p_1, \dots, p_k) \end{aligned}$$

be an (honestly) generated key pair, ciphertexts, and output of the computation. The scheme HE is  $\epsilon$ -circuit private if a PPT simulator  $\mathcal{S}$  exists such that for all functions  $f$  of depth  $l \leq L$  all PPT

distinguisher algorithms  $\mathcal{D}$  are bounded by

$$\left| \Pr[\mathcal{D}(\text{HE.Eval}(pk, f, \langle c_i \rangle, \langle p_i \rangle), sk, pk, \langle c_i \rangle)) = 1] - \Pr[\mathcal{D}(\mathcal{S}(sk, pk, M), sk, pk, \langle c_i \rangle)) = 1] \right| \leq \epsilon.$$

The circuit privacy’s definition focuses on the semi-honest setting and requires honest generation of HE keys and ciphertexts. We extend this definition by (1) removing the honest generation requirement, which extends the property to the malicious setting, and (2) relaxing privacy by revealing the functionality  $f$  and only hiding the evaluator’s private data  $\langle p_i \rangle$ .

**Definition 7** (Strong input privacy). A leveled homomorphic encryption scheme HE is  $\epsilon$ -strong input private if a simulator  $\mathcal{S}_{sip}$  exists such that all PPT adversaries  $\mathcal{A}$  are bounded by:

$$\Pr \left[ \begin{array}{l} (st, pk, sk) \leftarrow \mathcal{A}(1^\ell) \\ (st, f, \langle c_i \rangle, \langle p_i \rangle) \leftarrow \mathcal{A}(st) \\ b \leftarrow_s \{0, 1\} \\ m_i \leftarrow \text{HE.Dec}(sk, c_i) \\ M \leftarrow f(\langle m_i \rangle, \langle p_i \rangle) \\ a_0 \leftarrow \text{HE.Eval}(pk, f, \langle c_i \rangle, \langle p_i \rangle) \\ a_1 \leftarrow \mathcal{S}_{sip}(sk, pk, f, \langle c_i \rangle, M) \\ b' \leftarrow \mathcal{A}(st, a_b) \end{array} : b = b' \right] \leq \frac{1}{2} + \epsilon$$

Similar to the circuit privacy definition,  $f$  is the computed functionality,  $\langle c_i \rangle$  is the client’s encrypted input, and  $\langle p_i \rangle$  is the server’s plaintext input while  $st$  is a state for the adversary for storing information between interactions.

**Ciphertext indistinguishability.** When simulating our protocols, in the next section, the simulator generates ciphertexts. As part of our proof, we need to show that these ciphertexts are indistinguishable from the output of our protocols. Hence, we discuss when a distinguisher  $\mathcal{D}$  can distinguish two ciphertexts  $c$  and  $c'$  in three settings:

**Known public key.** We first consider the case where the distinguisher only knows the public key ( $pk$ ). This scenario directly follows from the IND-CPA property. As long as the HE scheme is IND-CPA secure and ciphertexts have the same size  $|c| = |c'|$ , then the distinguisher  $\mathcal{D}(pk, c, c')$  has a negligible chance.

**Semi-honest with known secret key.** Next, we consider the case where the distinguisher knows both the secret and public keys ( $sk, pk$ ) in the semi-honest setting, i.e., keys and ciphertexts are honestly generated. In this scenario, the distinguisher can decrypt ciphertexts so IND-CPA is not enough. Decrypting a ciphertext  $p \leftarrow \text{HE.Dec}(sk, c)$  results in a plaintext  $p$  and a measurable noise  $\epsilon$ . This transforms ciphertext indistinguishability to showing two statements: both decrypted plaintexts and ciphertext noises are indistinguishable. While comparing decrypted values  $p$  is straightforward, we rely on circuit privacy to ensure noises are indistinguishable. As long as the HE scheme is IND-CPA secure and circuit-private, and  $p \leftarrow \text{HE.Dec}(sk, c) \wedge p' \leftarrow \text{HE.Dec}(sk, c') \wedge p \stackrel{c}{\equiv} p'$ , then the distinguisher  $\mathcal{D}(sk, pk, c, c')$  has a negligible chance in a semi-honest setting.

**Malicious with known secret key.** The distinguisher knows both the secret and public keys ( $sk, pk$ ) in a malicious setting. We require our scheme to be IND-CPA secure and strongly input private in this

setting and use the simulator of strong input privacy ( $\mathcal{S}_{sip}$ ) to produce indistinguishable ciphertexts.

**Lack of circuit privacy.** The Lattigo library does not provide circuit privacy. This is not surprising since other popular HE libraries such as Microsoft Seal [68] do not provide circuit privacy either. While there are possible mitigations in the semi-honest setting such as noise-flooding [27] (alternatively called noise smudging) to achieve circuit privacy, there is no known mechanism for the malicious setting (i.e., no HE scheme achieves strong input privacy). Noise-flooding is only proven private in a semi-honest scenario where keys are generated honestly and the computation is guaranteed to start on freshly encrypted ciphertexts. Beyond flooding, there is a new line of work [17] that uses careful parameter selection in RNS or DCRT representation of ciphertexts to achieve lightweight circuit privacy. We hope this approach will be adopted by HE libraries.

Our implementation does not add extra defense mechanisms to prevent possible leakage from noise, due to the extra cost associated with these defenses. Despite this leakage, practical attacks using noise are limited. The complexity and depth of our functions make extracting information from this noise more challenging; especially since the size of the server's private data,  $N_s$ , is significantly larger than the capacity of the noise for storing information.

## C.2 Semi-honest security

We use real-world/ideal-world simulation to prove the security of our PCM protocols in the semi-honest setting (see Theorem 1). Our well-formedness checks, computed using (SD-)QUERY-CHECK, have no impact on honest execution as they only add  $\llbracket 0 \rrbracket$  to the result when the query is generated honestly. Therefore, we ignore these functions in this section. Our framework can be instantiated to support different functionalities, but they share a similar structure, which enables us to write a single proof that is customizable depending on the protocol variation. As seen in Definition 1, PCM is a two-party interaction that computes

$$(f_c(X, \mathcal{Y}), \perp) \leftarrow PCM_{f_M, f_A}(X, \mathcal{Y})$$

where  $f_c(X, \mathcal{Y}) = f_A(f_M(X, Y_1), \dots, f_M(X, Y_N))$  is a *deterministic* function selected from the Table 4. To reason about the properties the matching layer, we allow  $f_A$  to be equal to the identity function. To reason about the PSI layer, we allow  $f_c$  to have the natural PSI and PSI-CA definition.

Our semi-honest scenario has deterministic output functions. Therefore, we can use the simpler formulation of security in Lindell [43] which requires schemes to satisfy two properties to be secure. *Correctness*: the output of parties is correct. *Privacy*: the view of parties can be separately simulated as follows:

$$\left\{ \mathcal{S}_C(1^\ell, X, f_c(X, \mathcal{Y})) \right\}_{X, \mathcal{Y}, \ell} \stackrel{c}{\equiv} \{\text{View}_{\text{client}}(\ell, X, \mathcal{Y})\}_{X, \mathcal{Y}, \ell},$$

$$\left\{ \mathcal{S}_S(1^\ell, \mathcal{Y}, \perp) \right\}_{X, \mathcal{Y}, \ell} \stackrel{c}{\equiv} \{\text{View}_{\text{server}}(\ell, X, \mathcal{Y})\}_{X, \mathcal{Y}, \ell}$$

where  $\ell$  is the security parameter,  $X$  is the client input, and  $\mathcal{Y}$  is the server input. We omit  $\ell$  in the rest of this section. We assume that the client honestly generates the key pair  $(pk, sk) \leftarrow \text{HE.KeyGen}(\text{params})$  and sends the public key  $pk$  to the server before running the protocol.

**Correctness.** In a semi-honest scenario where both parties follow the protocol specification, showing that Algorithms 3 to 6 compute the functionality described in Table 4 is straightforward math. We have described the intuition behind these algorithms in Sections 6 to 8, so we do not repeat the argument here.

**Server privacy.** We simulate the view of clients to ensure server privacy. Let the view of the client be

$$\text{View}_{\text{client}}(X, \mathcal{Y}) = (X, \text{rnd}, Q, R, A = f_c(X, \mathcal{Y}))$$

where  $\text{rnd}$  is the internal random tape,  $Q$  is the encrypted query,  $R$  is the server's encrypted response, and  $A$  is the client's output.

All our protocols start with clients sending an encrypted query  $Q$  to the server and receiving an encrypted response  $R$ . Afterward, clients apply the appropriate reveal function on  $R$  to compute the output  $A \leftarrow \text{REVEAL}(sk, R)$ . Correctness ensures that the output  $A$  computed by the reveal algorithm is equal to the expected output  $f_c(X, \mathcal{Y})$  summarized in Table 4.

Now we build a simulator  $\mathcal{S}_C$  that given the input  $(X, \text{rnd}, A = f_c(X, \mathcal{Y}))$  simulates the clients view as follows:

- (1) The simulator uses  $Q' \leftarrow (\text{SD-})\text{QUERY}(pk, X)$  to compute a query depending on the domain size with  $\text{rnd}$  as internal randomness.
- (2) We categorize the client's output based on the range of  $f_c(X, \mathcal{Y})$  into 3 groups:

$$A \in \begin{cases} \{0, 1\} & \text{For: F-Match, Th-Match, Tv-Match, X-Agg} \\ \{0, 1\}^k & \text{For: PSI, NA-Agg} \\ \mathbb{Z} & \text{For: PSI-CA, CA-Agg, Ret-Agg} \end{cases}$$

Depending on the range category, the simulator computes the server response  $R'$  as follows using the circuit-privacy simulator to ensure equivalent noise levels:

$$R' \leftarrow \begin{cases} \text{HE.Enc}(pk, r \cdot A) & A \in \{0, 1\} \\ \langle \text{HE.Enc}(pk, r_j \cdot A[j]) \rangle & A \in \{0, 1\}^k \\ \text{HE.Enc}(pk, A) & A \in \mathbb{Z} \end{cases}$$

where  $r \leftarrow_s \mathbb{Z}_q^*$  and  $r_j \leftarrow_s \mathbb{Z}_q^*$  are random values.<sup>3</sup>

- (3) The simulator returns  $(X, \text{rnd}, Q', R', A)$ .

Now we show that the distribution returned from the simulator  $\mathcal{S}_C$  is indistinguishable from  $\text{View}_{\text{client}}$ . The three variables  $X, A, \text{rnd}$  are taken from the input and are guaranteed to have the same distribution between the simulation and the view. Therefore, we only need to show that the joint distribution for the query and the server response are indistinguishable – conditional on the common  $X, A$ , and  $\text{rnd}$  variables.

Both the query  $Q$  and response  $R$  are ciphertexts. As discussed in ‘ciphertext indistinguishability’ in Appendix C.1, we are in a semi-honest scenario where the distinguisher knows both the public and private keys  $(pk, sk)$ . Our HE scheme is both IND-CPA and circuit private so we only need to show that these ciphertext pairs,  $(Q, Q')$  and  $(R, R')$ , decrypt to the same value. Queries are computed following  $Q \leftarrow (\text{SD-})\text{QUERY}(pk, X)$ . The query function only encrypts the input set  $X$  as the query so both  $Q$  and  $Q'$  should decrypt to the same value. The server's response  $R$  depends on the client's output  $A$  and this relation is specified in our REVEAL functions. Our

<sup>3</sup>When evaluating F-Match, the simulator chooses random variables from  $\mathbb{Z}_q$  instead of  $\mathbb{Z}_q^*$  to ensure the same false positive probability as the real-world execution.

correctness property ensures that the following relation between  $R$  and  $A$  holds in the real-world:

$$R \leftarrow \begin{cases} \llbracket b \rrbracket & b = \begin{cases} 0 & A = 0 \\ \leftarrow_s \mathbb{Z}_q^* & A = 1 \end{cases} & A \in \{0, 1\} \\ \langle \llbracket b_j \rrbracket \rangle & b_j = \begin{cases} 0 & A[j] = 0 \\ \leftarrow_s \mathbb{Z}_q^* & A[j] = 1 \end{cases} & A \in \{0, 1\}^k \\ \llbracket A \rrbracket & A \in \mathbb{Z} \end{cases} .$$

It is straightforward to see that the underlying plaintext of  $R'$ , produced in step (2) by  $\mathcal{S}_C$ , follows the same distribution.  $\square$

**Client privacy.** We simulate the view of servers to ensure client privacy. Simulating the server’s view is considerably simpler than the client’s as the server only observes ciphertexts encrypted under the client’s key. In this simulation, the adversary is a semi-honest server and the distinguisher does not have access to the client’s secret key. Moreover, our HE scheme is IND-CPA, which simplifies ciphertext indistinguishability to ensuring  $|c| = |c'|$  (see Appendix C.1).

Let the server’s view be  $\text{View}_{\text{server}} = (\mathcal{Y}, \text{rnd}, Q, R)$ . We build a simulator  $\mathcal{S}_S$  that given the input  $(\mathcal{Y}, \text{rnd})$  proceeds as follows:

- (1) The simulator chooses two random variables  $r_1$  and  $r_2$  with the same size as the query and the server response then encrypts them.

$$Q' \leftarrow \text{HE.Enc}(pk, r_1), R' \leftarrow \text{HE.Enc}(pk, r_2)$$

- (2) The simulators return  $(\mathcal{Y}, \text{rnd}, Q', R')$ .

The variables  $\mathcal{Y}$  and  $\text{rnd}$  are taken from the input and have the same distribution between the simulation and the view. The two variables  $Q'$  and  $R'$  are encrypted under the client’s key. Without the knowledge of the secret key  $Q \stackrel{c}{\equiv} Q'$  and  $R \stackrel{c}{\equiv} R'$  hold independent of their content.  $\square$

### C.3 Malicious server

Now, we study the setting where the server is malicious. In this setting, the client is honest and server privacy is not applicable. Our framework provides *no correctness guarantee* in this setting. The malicious server can force a corrupted response that depends on the client input or fix the outcome independent of the query. Despite the lack of correctness guarantee, malicious servers cannot learn any information about the client’s private data. We restate client privacy (Definition 3) more formally, then prove Theorem 2 by giving a tight reduction from our scheme’s client privacy to the IND-CPA security of our HE encryption scheme. During this proof, we assume that the client’s input fits into a single batched ciphertext, which holds throughout our evaluation. Adjusting the proof to support input encrypted in  $k$  ciphertexts is straightforward and gives a  $k$ -fold advantage to the adversary.

**Definition** (Client privacy). A PCM protocol is client private if no PPT adversary  $\mathcal{A}$  exists such that:

$$\Pr \left[ \begin{array}{l} p \leftarrow \text{HE.ParamGen}(1^\ell) \\ (pk, sk) \leftarrow \text{HE.KeyGen}(p) \\ (st, X_0, X_1) \leftarrow \mathcal{A}(pk) \\ b \leftarrow_s \{0, 1\} \\ Q_b \leftarrow (\text{SD-})\text{QUERY}(pk, X_b) \\ b' \leftarrow \mathcal{A}(st, Q_b) \end{array} : \begin{array}{l} |X_0| = |X_1| \\ b = b' \end{array} \right] > \frac{1}{2} + \epsilon.$$

**PROOF.** Let  $\mathcal{A}$  be an adversary that can break the client privacy property with a non-negligible probability  $\epsilon$ . We build a new adversary  $\mathcal{A}'$  that can break the IND-CPA security (Definition 5) of our HE scheme with the same probability  $\epsilon$ .

- (1) The adversary  $\mathcal{A}'$  starts an IND-CPA challenge and receives a public key  $\mathcal{A}'(pk)$ .
- (2) The adversary  $\mathcal{A}'$  calls  $(st, X_0, X_1) \leftarrow \mathcal{A}(pk)$ .
- (3) Depending on the domain size,  $\mathcal{A}'$  converts sets  $X_a$  to  $m_a = \langle x_i \rangle$  or  $m_a = \langle z_i \rangle$  following the logic of (SD-)QUERY.
- (4) The adversary  $\mathcal{A}'$  continues the IND-CPA challenge with  $(st, m_0, m_1)$  and receives  $\mathcal{A}'(st, c_b)$ .
- (5) The adversary  $\mathcal{A}'$  passes the challenge  $b' \leftarrow \mathcal{A}(st, Q_b = c_b)$  to  $\mathcal{A}$  and returns  $b'$  as the output.

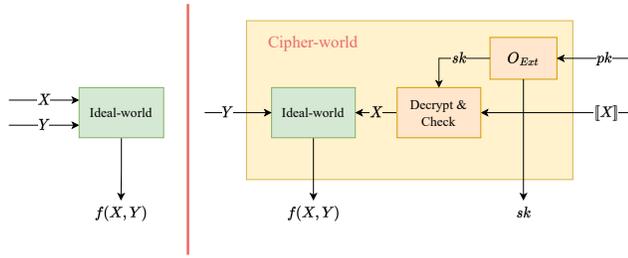
We show that the adversary  $\mathcal{A}$  cannot distinguish interaction with  $\mathcal{A}'$  from our protocol. The adversary  $\mathcal{A}$  has two interactions in the client privacy challenge. The first interaction is getting a fresh public key  $pk$  which is the same between the IND-CPA and client privacy challenges. In the second interaction,  $\mathcal{A}$  receives a query  $Q_b$  produced by (SD-)QUERY. We know that (SD-)QUERY consists of two steps: convert  $X$  into  $p = \langle x_i \rangle$  or  $p = \langle z_i \rangle$  which  $\mathcal{A}'$  performs in (3) and encrypting  $p$  with  $pk$  which is performed as part of IND-CPA challenge. This ensures that  $c_b$  is computed in the same way as  $Q_b$  and follows the same distribution. The adversary  $\mathcal{A}'$  runs one instance of client privacy with  $\mathcal{A}$  and succeeds the IND-CPA as long as the  $\mathcal{A}$  succeeds leading to the same  $\epsilon$  advantage.  $\square$

### C.4 Malicious clients

When addressing malicious clients, client privacy is irrelevant and correctness does not apply either as the server has no output. Therefore, we only need to address our protocols’ server privacy.

**THEOREM 3.** *Our protocols provide server privacy against malicious clients if the HE scheme is IND-CPA and strongly input private.*

Unfortunately, direct application of real-world/ideal-world simulation on our protocols in the malicious setting is not possible. The biggest challenges in simulation proofs in the malicious setting is that the adversary is not required to use its input and random tape during the execution. Therefore, the simulation needs to extract the effective input that the malicious adversary uses to determine the corresponding output. In our protocols, the server only receives an encrypted query from the client and there is no further interaction, such as having ROM calls in the client, to provide any extraction opportunity. Therefore, the semantic security of our HE scheme prevents extracting the effective inputs. To address this challenge, we adapt the ideal-world paradigm and introduce a new notion for simulating HE protocols called cipher-world. Cipher-world is inspired by trapdoor commitments and allows the trusted party to decrypt



**Figure 7: Structure of the cipher-world and its difference with the ideal-world.**

the query and compute the ideal functionality. We first prove that cipher-world provides the same server privacy guarantee as the ideal-world then use it to prove our framework server-private. The introduction of cipher-world as a new paradigm to prove security makes our proof highly non-standard.

### Cipher world

We extend the notion of real-world/ideal-world simulation for two-party HE-based schemes in which exactly one party, called the querier, holds a pair of HE keys while the other party, called the evaluator, performs computation in the encrypted domain. This extension is *one-sided* and only addresses the privacy of the evaluator (i.e., the server in our protocols). We call this extension cipher-world.

Cipher-world assumes that the trusted party defined in the ideal-world accepts encrypted inputs, is *computationally unbounded*, and can extract the secret key  $sk$  of the querier from its public key  $pk$  – breaking the security of the HE scheme in the process. The trusted party uses this secret key to extract the effective input  $X$  of malicious queriers. Unlike the ideal-world, there is no guarantee that the decrypted input of the cipher oracle is valid and follows the input restrictions. The cipher oracle first verifies input restrictions and outputs  $\perp$  if any check fails; otherwise, it computes and outputs  $f(X, Y)$ . Additionally, the cipher-world reveals the secret key  $sk$  to the simulator. This key is solely used for the purpose of simulating ciphertext noise. Figure 7 shows the structure of the cipher-world and how it compares to the standard ideal-world setting. We define both ideal-world and cipher-world oracles below to highlight their differences:

$$\begin{aligned} f(X, Y) &\leftarrow \text{IDEAL}(\ell, X, Y) \\ (sk, f(X, Y)/\perp) &\leftarrow \text{CIPHER}(\ell, (pk, c = \llbracket X \rrbracket), Y). \end{aligned}$$

Recall that here  $Y$  contains the evaluator’s private input.

**THEOREM 4.** *The cipher-world provides the same privacy guarantee for the evaluator’s private information as the ideal-world.*

**PROOF.** It is clear that queriers can learn more information from interacting with cipher oracles than ideal oracles. However, we show that this leakage does not impact evaluator’s privacy. First, we formally prove that this leakage  $\mathcal{L}$  is bounded to providing an oracle  $sk \leftarrow \mathcal{O}_{\text{Ext}}(pk)$  which extracts the secret key from HE public keys. Next, we prove that the leakage  $\mathcal{L}$  is independent of the evaluator’s private data  $Y$ .

We assume that a PPT adversary  $\mathcal{A}$  exists such that  $\mathcal{A}$  gains more advantage from interacting with a cipher oracle instead of an ideal oracle than  $\mathcal{L} = \{\mathcal{O}_{\text{Ext}}\}$ . We build a new adversary  $\mathcal{A}'$  that given an ideal oracle  $\text{IDEAL}(Y, \cdot)$  and an extraction oracle  $\mathcal{O}_{\text{Ext}}$  can simulate the view of  $\mathcal{A}$ .

- (1) Adversary  $\mathcal{A}'$  initiates a new interaction with  $\mathcal{A}$  and receives  $(pk, c = \llbracket X \rrbracket) \leftarrow \mathcal{A}(\cdot)$ .
- (2) Adversary  $\mathcal{A}'$  uses the extraction oracle within  $\mathcal{L}$  to extract the secret  $sk \leftarrow \mathcal{O}_{\text{Ext}}(pk)$ .
- (3) Adversary  $\mathcal{A}'$  decrypts  $\mathcal{A}$ ’s query  $X \leftarrow \text{HE.Dec}(sk, c)$ .
- (4) Adversary  $\mathcal{A}'$  verifies whether  $X$  passes input restrictions.
  - (4.a) If any check fails,  $\mathcal{A}'$  sets  $A \leftarrow \perp$ .
  - (4.b) Otherwise,  $\mathcal{A}'$  interacts with the ideal world oracle with the decrypted input and sets  $A \leftarrow \text{IDEAL}(X, Y)$ .
- (5) Adversary  $\mathcal{A}'$  finishes the execution  $\mathcal{A}(sk, A)$ .

We need to show that the adversary  $\mathcal{A}$  cannot distinguish  $(sk, A)$  produced in our simulation from the output of the cipher-world. Both the simulation and the cipher-world oracle are directly using the extraction oracle  $\mathcal{O}_{\text{Ext}}$  to produce the secret key  $sk$  which ensures secret key indistinguishability. We study two cases for  $\mathcal{A}$ : (1) The adversary  $\mathcal{A}$  does not follow the input restriction. In this case, the cipher oracle responds with the failure symbol  $\perp$ . The adversary  $\mathcal{A}'$  performs the same input verification process as the cipher oracle which leads to setting  $A \leftarrow \perp$  when one of step (4.a) checks fail. (2) The adversary  $\mathcal{A}$  follows the input restriction. In this scenario, both the cipher and ideal oracles compute the same output ensuring that  $A = f(X, Y)$ . This proves that the leakage of cipher-world can be bound to  $\mathcal{L} = \{\mathcal{O}_{\text{Ext}}\}$ .

Now we need to show that the leakage of the cipher-world is independent of the evaluator’s private data. As we bound the leakage to an extraction oracle  $\{\mathcal{O}_{\text{Ext}}\}$ , this independence is clear since extraction is not impacted by changing the evaluator’s private data  $Y$ . Note that in our protocol, exactly one party, the querier, generates HE keys, so assuming that HE key extraction is easy has no impact on the evaluator. Therefore, our cipher-world provides the same evaluator privacy guarantee as the ideal world.  $\square$

We showed that our cipher-world provides the same server (evaluator) privacy as the original ideal-world. Note that the cipher-world is one-sided and does not make any security claim about clients (queriers). To prove that our protocol provides server privacy, we have to show that the real-world view can be simulated given access to a cipher-world oracle.

$$\left\{ \text{CIPHER}_{S'_C}^{\text{Client}}(X, \mathcal{Y}) \right\} \stackrel{c}{\equiv} \left\{ \text{REAL}_{\mathcal{A}}^{\text{Client}}(X, \mathcal{Y}) \right\}$$

$$\updownarrow$$

$$S'_C(X, \text{rnd}, \mathcal{A}, \text{CIPHER}(Y, \cdot)) \stackrel{c}{\equiv} \text{View}_{\mathcal{A}}(X, \mathcal{Y}) = (X, \text{rnd}, Q, R, A)$$

We build a simulator  $S'_C$  that given a PPT real-world malicious client  $\mathcal{A}$  and a cipher-world oracle  $\text{CIPHER}(\mathcal{Y}, \cdot)$  fixed with the server’s input, simulates the real-world as follows:

- (1) Simulator  $S'_C$  initiates a new interaction with  $\mathcal{A}$  and receives  $(pk, Q = \llbracket X' \rrbracket) \leftarrow \mathcal{A}(\cdot)$ .
- (2) Simulator  $S'_C$  interacts with the cipher oracle and learns  $(sk, A') \leftarrow \text{CIPHER}((pk, Q))$ .

- (3) If the cipher-world detects a malicious query  $Q$  which does not respect input checks (i.e.,  $A' = \perp$ ), the simulator  $\mathcal{S}'_C$  chooses a random response  $t \leftarrow \mathbb{Z}_q$ , sets the output accordingly  $A' \leftarrow \text{REVEAL}(sk, \llbracket t \rrbracket)$ , and skips to the step 5.
- (4) Otherwise,  $\mathcal{S}'_C$  computes the response  $t$  as follows:

$$t \leftarrow \begin{cases} r \cdot A' & A \in \{0, 1\} \\ \langle r_j \cdot A'[j] \rangle & A \in \{0, 1\}^k \\ A' & A \in \mathbb{Z} \end{cases} .$$

Note that the range of  $A$  is determined by functionality  $f$  and is independent of the parties' input.

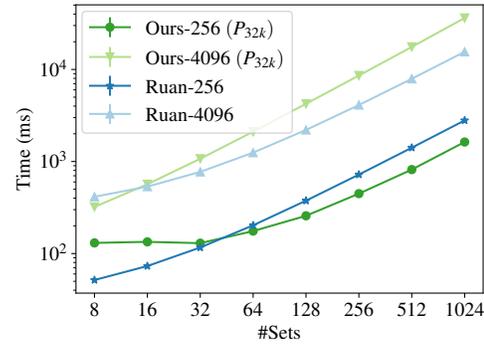
- (5) Simulator  $\mathcal{S}'_C$  performs  $R' \leftarrow \mathcal{S}_{sip}(sk, f, Q, t)$ .
- (6) Adversary  $\mathcal{A}$  only learns the response  $R$  if the decryption succeeds. Therefore, Simulator  $\mathcal{S}'_C$  checks if the decryption  $\text{HE.Dec}(sk, R')$  succeeds. Otherwise,  $\mathcal{S}'_C$  sets  $A' \leftarrow \perp$ .
- (7) The simulator returns  $(X, \text{rnd}, Q = \llbracket X' \rrbracket, R', A')$ .

Now, we show that the real view  $(X, \text{rnd}, Q = \llbracket X' \rrbracket, R, A)$  is indistinguishable from the simulated view  $(X, \text{rnd}, Q, R', A')$ . The simulator  $\mathcal{S}'_C$  does not use variables  $X$  and  $\text{rnd}$  as there is no guarantee that malicious clients will use their input or random tapes. As  $(X, \text{rnd}, Q)$  are directly taken from the input, we only need to show that  $(R, A) \stackrel{c}{\equiv} (R', A')$  conditional on the common variables. We split our analysis into three cases:

*Malicious queries which do not represent a set.* When the client is malicious there is no guarantee that the query represents a set. In Section 6.3, we designed (SD-)QUERY-CHECK functions and proved that they randomize the output of our protocols when the query does not represent a set as long as the HE abstraction holds. Our simulator  $\mathcal{S}'_C$  relies on the cipher oracle to detect when the query does not represent a valid set in step 3 and assigns a uniformly random value  $t$  to be encrypted as the response  $R'$ . To ensure that  $R'$  has an indistinguishable noise from the real response  $R$ , instead of directly using the encryption in step 5,  $\mathcal{S}'_C$  uses the simulator  $\mathcal{S}_{sip}$  from the strong input privacy property (Definition 7). Our simulator follows the same REVEAL process to compute the output  $A$  from the response as the real protocol. Since  $R \stackrel{c}{\equiv} R'$ , we will have  $A \stackrel{c}{\equiv} A'$  as long as the server response  $R'$  decrypts successfully.

*The malicious query decrypts to the set  $X'$ :* When the query represents a set, (SD-)QUERY-CHECK produces an encrypted zero and does not impact the output of the protocol (i.e., adding  $\llbracket 0 \rrbracket$  is neutral). Knowing that the query is an encryption of the set  $X'$ , our protocols ensure that  $A = f(X', \mathcal{Y})$  as long as the decryption of  $R$  succeeds. This guarantee follows from combining (1) our semi-honest correctness from Appendix C.2 and (2) knowing that the query decrypts to the same value as (SD-)QUERY( $pk, X'$ ). The simulator  $\mathcal{S}'_C$  follows a similar process to compute  $R$  from  $A$  as our semi-honest simulator  $\mathcal{S}_C$  with the difference of replacing HE.Enc with  $\mathcal{S}_{sip}$  (i.e., relying on strong input privacy instead of circuit privacy).

*Failed decryption.* Unlike the semi-honest setting where we know that the client's query is freshly encrypted, the server may receive queries with a high noise level. This may lead to producing server responses that fail the decryption. The use of  $\mathcal{S}_{sip}$  in step 5 ensures that the response of  $\mathcal{S}'_C$  has the same noise level as our real-world response. Therefore, the decryption of  $R'$  fails if and only if the decryption of  $R$  fails. When the decryption fails, the real-world



**Figure 8: Computation cost for performing small domain PSI-CA. Two systems provide different security levels: Ruan et al. support 80-bit security while ours provide 128-bit security.**

cannot compute the output ( $A = \perp$ ) while the simulator  $\mathcal{S}'_C$  sets  $A' \leftarrow \perp$  in step 6.

## D EXTRA BENCHMARKS

In this section, we provide extra details on our evaluation and add more benchmarks. First, we compare our small domain PSI layer to an existing small domain paper. Second, we provide more detail on how we design and evaluate generic solutions with the same privacy and functionality as our document search engine and provide extra performance plots. Third, we compare our document search engine to one of the fastest OT-based PSI protocols which does not satisfy our privacy requirements.

### D.1 Small-domain protocols

In this section, we evaluate the performance of our small domain PSI-CA protocol and compare it to existing work. We focus on Ruan et al. [67] in this section. We compare with Shimizu et al. [69] in Section 11. We do not consider Bay et al. [5] here, as it focuses on a multi-party scenario, which leads to higher costs.

The code for the protocol of Ruan et al. is not available but the paper provides a detailed cost analysis. We use the same scenario and the same CPU (Intel Core i7-7700) to allow us to directly compare performance without requiring us to rerun their protocol from scratch. Moreover, we extend and optimize their approach for a many-set scenario. Ruan et al. use bit-vectors encrypted with El-Gamal [25] or Paillier [52] encryption to perform PSI. We extend their approach to only compute the query once and apply it to many sets. Their detailed performance benchmark allows us to compute the cost of performing a many-set query with  $N$  sets. In Fig. 8, we show the computation cost with a fixed input domain size ( $|D| \in \{256, 4096\}$ ) and varied the number of sets. The cost does not include the key exchange. The performance of our scheme is comparable to Ruan et al. – which scheme is performing best depends on the specific scenario.

Despite having the same operating point for both approaches, they have very different security guarantees. Ruan et al. assume a semi-honest privacy model and use Paillier keys with 1024-bit RSA primes which only provide 80-bit security. Extracting more

information than cardinality from their protocol is trivial for malicious clients. On the other hand, we provide full *128-bit* security and protect against malformed queries by misbehaving clients (see Section 6.3).

## D.2 Circuit-based protocols

In this section, we provide more details on our evaluation of the generic solutions in the document search scenario of Section 11.2. We expand on their threat model and properties. Moreover, we report and compare the server computation cost of all approaches.

**Generic SMC.** We use a high-level SMC compiler, EMP toolkit [77], to design and evaluate circuits providing the same properties as our search engine. More specifically, we use the ‘EMP-sh2pc’ branch of the compiler that provides security in a two-party semi-honest setting and supports both boolean and arithmetic circuits.

We encode each input as a 32 bits binary value, which results in a higher false-positive rate due to encoding keywords than our framework, where we encode keywords using 39 or 44 bits. We use a private equality check offered by EMP for comparing encrypted set elements. To determine whether one keyword of the client set has a match in a given document, we perform equality tests against all keywords of a document and perform an OR over the comparison results. To perform the full matching, we do an AND over the matching status of all client keywords. This produces a 1-bit result for each document determining its relevance. The F-Match process does not create any extra false-positive in this approach.

Now we have to aggregate  $N$  1-bit document matching statuses according to our X-Agg and CA-Agg policies. The X-Agg variant is straightforward and we use an OR to check if any document is relevant. To count the number of relevant documents, we first convert binary matching statuses to integers encoding ‘0’ or ‘1’ to enable us to continue the computation with an arithmetic circuit. Then, we compute the sum of these integer statuses.

**Circuit-PSI.** We choose Chandran et al. [10] as a state-of-the-art circuit-PSI paper that is secure against semi-honest adversaries. Circuit-PSI protocols perform an intersection between the sets of two parties and secret share the output among them. This enables using circuits to privately compute arbitrary functions over the intersection. Despite the capability of Chandran et al. to be extended with circuits to compute F-Match matching and X-Agg or CA-Agg aggregation, we decide to not extend their circuit and use the time necessary for computing the intersection as a lower bound on the cost of searching. Since the PSI protocol of Chandran et al. is a single set protocol, we run  $N$  instances of Circuit-PSI sequentially to simulate searching  $N$  documents. We use the default parameters decided by Chandran et al., meaning that each item is encoded using 32 bits with an additional false positive rate of  $2^{-40}$  due to computation.

**Server’s computation cost.** We have already reported the end-to-end latency, communication cost, and client’s computation cost in the main body (Fig. 6). We report the server’s computation cost in Fig. 9. Starting from 16 documents, our framework has lower server computation than Chandran et al. [10]. On the other hand, the EMP solutions have better server efficiency than our scheme which is not surprising as our framework outsources the computation load from

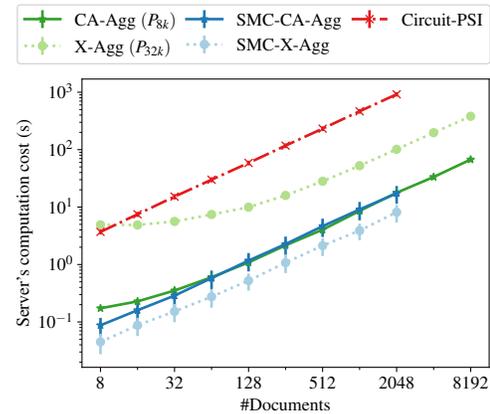


Figure 9: Server’s computation cost for document search.

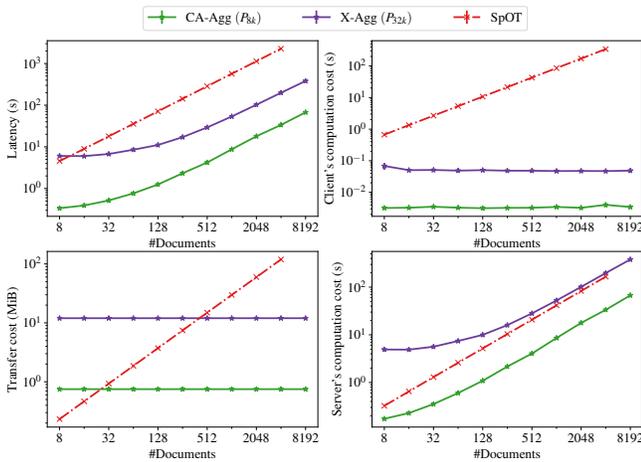
thin clients to the server. Despite our outsourcing, our CA-Agg search only increases the server cost by a factor of 2.5x and X-Agg by a factor of 30x when searching 1k documents.

## D.3 OT-based protocol

There are efficient PSI protocols that are based on the oblivious transfer in both the semi-honest (such as SpOT-light [53]) and the malicious setting (such as PaXoS [54]). As discussed in Section 3, this line of research focuses on computing one-to-one equality tests between the client and server which leaks information about each server set and cannot satisfy our privacy requirements. Despite providing a lower privacy guarantee, we compare our approach to the SpOT-light protocol as a baseline cost.

**Document search with SpOT-light.** We follow the document search setting from the Section 11.2 and evaluate the cost of using SpOT-light to search  $N$  documents. SpOT offers 128-bit security in a semi-honest setting and accepts 256-bit input elements which bypasses the false-positive rate of mapping keywords. However, SpOT is (1) a single-set protocol and (2) does not support privacy extensions such as computing relevance without leaking the intersection cardinality to the client (i.e., private set matching) or aggregating the search result of multiple documents similar to our X-Agg and CA-Agg variant. We handle the single-set limitation by running  $N$  sequential PSI interactions to search  $N$  documents, but we do not add any countermeasure for the lack of matching or aggregation functionality. We encountered concurrency issues (with async IO) when running SpOT. This issue gets amplified when repeating the protocol  $N$  times. Instead of directly running the code, we benchmarked the cost of each interaction through multiple runs with  $N = 8$ , then extrapolate the cost for all entries.

We report the end-to-end latency, computation, and communication costs of SpOT-light in Fig. 10. In the single-set setting and for a small number of documents, SpOT provides better performance. However, as soon as reaching 32 documents, our CA-Agg variant starts to provide better latency, computation, and communication than SpOT despite providing better privacy. When searching 1k documents, our framework improves latency by a factor of 10–65x, communication by a factor of 1.7–27, and client’s computation by a factor of 1800–24,800x depending on the search functionality.



**Figure 10: The end-to-end latency (upper-left), communication cost (lower-left), and client and server computation cost (right) of the document search.**

### E SOLVING MATCHING IN MOBILE APPS

We do not separately evaluate the matching in the mobile apps scenario as it is similar to the chemical similarity scenario. The set of attributes to be matched can be represented using a small domain as the number of attributes is limited and they have few possible values. Matching of individual records can be implemented using the threshold matching (Th-Match) protocol. This allows for approximate matches. The results can then be combined using naive aggregation to reveal the matching indices to the querier. Since Th-Match is simpler than the Tv-Match protocol and the threshold for matching is likely smaller than the chemical similarity case, we expect better performance for matching than chemical similarity.

### F PSI-SUM

A benefit of our framework’s modular design is extensibility. To show the ease of adding new functionality, we design a new protocol called PSI-SUM in this section which is getting more popular in the literature due to its use in private ad-monetization systems [32, 44, 82]. In the PSI-SUM protocol, the server assigns a weight to each of its elements and the client wants to compute the sum of weights of common elements, i.e.,  $\text{PSI-Sum}(X, (Y, W)) = \sum_{\{i|y_i \in X\}} w_i$ .

We add this new protocol in our PSI layer following the structure in Fig. 3. We define the PSI-SUM algorithms in Algorithm 8. The PSI-SUM-PROCESS computes a binary inclusion status  $\llbracket t_i \rrbracket$  for each server element  $y_i$ . The server then proceeds similarly to ePSI-CA to compute the encrypted weighted sum. The server can continue processing this value in the next layers, such as checking for a threshold value on the sum, or return  $\llbracket W \rrbracket$  to the client which decrypts it to obtain the answer (see PSI-SUM-REVEAL). Our extensions such as ensuring query well-formedness and many-set can directly apply to this protocol without any extra effort.

**Algorithm 8** Adding PSI-SUM capabilities.

```

function PSI-SUM-PROCESS( $pk, Q = \langle \llbracket x_i \rrbracket \rrbracket, Y, W = \langle w_i \rangle \rangle$ )
     $\llbracket t_i \rrbracket \leftarrow \text{HE.IsZero}(pk, \text{HE.IsIn}(pk, y_i, \langle \llbracket x_j \rrbracket \rangle))$ 
     $\llbracket W \rrbracket \leftarrow \sum_{i \in [n]} w_i \cdot \llbracket t_i \rrbracket$ 
    return  $\llbracket W \rrbracket$ 

function PSI-SUM-SD-PROCESS( $pk, Q = \langle \llbracket z_i \rrbracket \rrbracket, Y, W = \langle w_i \rangle \rangle$ )
     $\langle \llbracket t_i \rrbracket \rangle \leftarrow \text{PSI-SD-PROCESS}(pk, \langle \llbracket z_i \rrbracket \rangle, Y)$ 
     $\llbracket W \rrbracket \leftarrow \sum_{d_i \in D} w_i \cdot \llbracket t_i \rrbracket$ 
    return  $M \leftarrow \llbracket W \rrbracket$ 

function PSI-SUM-REVEAL( $pk, M = \llbracket W \rrbracket$ )
    return  $\text{HE.Dec}(sk, \llbracket W \rrbracket)$ 
    
```