

Computation Efficient Structure-Aware PSI From Incremental Function Secret Sharing

Gayathri Garimella
Brown University
gayathri_garimella@brown.edu

Benjamin Goff
Brown University
benjamin_goff@brown.edu

Peihan Miao
Brown University
peihan_miao@brown.edu

Abstract

Structure-Aware Private Set Intersection (**sa-PSI**), recently introduced by Garimella et al. (Crypto'22), is a PSI variant where Alice's input set S_A has a publicly known structure (for example, interval, ball or union of balls) and Bob's input S_B is an unstructured set of elements. Prior work achieves **sa-PSI** where the communication cost only scales with the description size of S_A instead of the set cardinality. However, the computation cost remains linear in the cardinality of S_A , which could be prohibitively large.

In this work, we present a new semi-honest **sa-PSI** framework where both computation and communication costs *only scale with the description size* of S_A . Our main building block is a new primitive that we introduce called Incremental Boolean Function Secret Sharing (**ibFSS**), which is a generalization of FSS that additionally allows for evaluation on input prefixes. We formalize definitions and construct a weak **ibFSS** for a d -dimensional ball with ℓ_∞ norm, which may be of independent interest. Independently, we improve spatial hashing techniques (from prior work) when S_A has structure union of d -dimensional balls in $(\{0, 1\}^u)^d$, each of diameter δ , from $\mathcal{O}(u \cdot d \cdot (\log \delta)^d)$ to $\mathcal{O}(\log \delta \cdot d)$ in terms of both computation and communication. Finally, we resolve the following open questions from prior work with communication and computation scaling with the description size of the structured set.

- Our PSI framework can handle a union of overlapping structures, while prior work strictly requires a disjoint union.
- We have a new construction that enables Bob with unstructured input S_B to learn the intersection.
- We extend to a richer class of functionalities like structure-aware PSI Cardinality and PSI-Sum of associated values.

Contents

1	Introduction	3
1.1	Our Contributions	3
1.2	Related Work	5
2	Technical Overview	6
3	Preliminaries	9
4	Incremental Boolean Function Secret Sharing	10
5	sa-PSI with Alice Learning Output	12
5.1	Multiple (Overlapping) Balls with Multiple Points	12
5.2	Spatial Hashing	16
6	ibFSS for d-dimensional ℓ_∞-balls	19
6.1	ibDCF Definition	19
6.2	ibDCF Constructions	21
7	Instantiation and Evaluation	23
7.1	Single Ball with Single Point	24
7.2	Multiple (Overlapping) Balls with Multiple Points	25
7.3	Spatial Hashing	25
8	Extending Functionality	27
8.1	sa-PSI with Bob Learning Output	27
8.2	Structure-Aware PSI Cardinality/Sum	28
9	Acknowledgments	29
A	Cryptographic Tools	37
A.1	Cryptographic Schemes	37
A.2	Secure Two-Party Computation	37
B	Security Proofs	38
B.1	Proof of Theorem 6	38
B.2	Proof of Theorem 8	42
B.3	Proof of Theorem 9	47
B.4	Proof of Theorem 10	47
B.5	Proof of Theorem 11	47

1 Introduction

Private Set Intersection (PSI) enables two distrusting parties, each holding a private set of elements, to jointly compute the intersection of their sets without revealing any additional information. Over the past decade, PSI has found numerous real-world applications, such as secure password breach checkup [ALP⁺21, Chr, Edg, Key], mobile private contact discovery [KRS⁺19, Sig], online advertising measurement [IKN⁺20, MPR⁺20, Ads], privacy-preserving contact tracing in a global pandemic [TSS⁺20, BBV⁺20, App], and more. There has been enormous progress towards realizing PSI efficiently to be deployed in practice, at scale. Most of these protocols have both the computational and communication complexity proportional to the size of the two sets [HFH99, HEK12, KKRT16, RR17, PSTY19, PRTY19, CM20, RS21, GPR⁺21, CRR21, RS21, RR22, CILO22, BPSY23].

Recent work by Garimella et al. [GRS22] considers the scenario where one party’s input set has a publicly known structure. They introduce *Structure-Aware PSI* (sa-PSI), where Alice has an input set $S_A \in \mathcal{S}$ from a known family of structured sets \mathcal{S} (for example, a single-dimensional interval, high-dimensional ball, disjoint union of balls) and Bob’s input set S_B consists of an unstructured collection of elements. The size of S_A could potentially be much larger than S_B , namely $|S_A| \gg |S_B|$, but the structure leads to a succinct *description size* of the set S_A . sa-PSI [GRS22, GRS23] allows Alice to learn the intersection $S_A \cap S_B$ with communication $\ll |S_A|$, that *scales with the description size* of S_A instead of its cardinality.

Structure-aware PSI finds many applications in practice. For instance, suppose Alice and Bob each have a set of points A and B respectively, and they want to identify which pairs of their points are *close by*, namely $(x, y) \in A \times B$ satisfying $D(x, y) \leq \delta$, where D is a public distance metric and δ is a public threshold. We can model this problem as an instance of sa-PSI, where Alice’s structured input S_A is a union of δ -balls (for each point $x \in A$, consider a ball centered at x with radius δ under the distance metric D), and Bob’s input is his unstructured input set B . This can be used for applications such as noisy/fuzzy biometric matching [FNP04, UCK⁺21] and privacy-preserving ride sharing [RADN22, GWSW22], among others.

Prior works [GRS22, GRS23] present a general framework for semi-honest and malicious secure sa-PSI, respectively, with specific construction for set family of union of d -dimensional balls in the ℓ_∞ norm, where the communication cost is *independent of the total volume of the balls* in the structured input. However, an inherent limitation is that Alice’s local *computation cost* remains linear in the size of her structured input $|S_A|$ (that is, total volume of the balls), which can be exponentially large. This leaves us with the following open problem that is explicitly mentioned in [GRS22]:

*Can we achieve sa-PSI where **both** computation and communication costs only scale with the description size of the structured set?*

1.1 Our Contributions

Structure-Aware PSI with efficient computation. In this work, we answer the above question in the affirmative. We present a new semi-honest sa-PSI protocol for union of ℓ_∞ -balls in a d -dimensional space, where *both* computation and communication costs only scale with the description size of the structured set. We achieve both communication and computation complexity of $\mathcal{O}(N_A \cdot u^d)$, in a d dimensional space $(\{0, 1\}^u)^d$ where N_A is the number of balls in Alice’s set, instead of scaling with $|S_A|$, which could potentially be as large as $2^{u \cdot d}$.

Incremental Boolean Function Secret Sharing. The main building block underneath our construction is a new notion we introduce, called *Incremental Boolean Function Secret Sharing* (ibFSS). Intuitively speaking, FSS [BGI15] splits a function f from a function family \mathcal{F} into two functions, f_0, f_1 , such that each function f_b hides the function f and that $f_0(x) + f_1(x) = f(x)$ for all input x . Our new notion ibFSS is an instance of Boolean FSS [GRS22], for a function that captures set membership of a structured set S_A , meaning that $f(x) = 0$ if $x \in S_A$ and $f(x) = 1$ otherwise. Additionally, ibFSS also enables evaluation on *prefixes* of all the points while preserving the FSS properties. We formalize the notion of ibFSS and give a construction for d -dimensional ℓ_∞ -balls. We then present a generic framework of constructing sa-PSI from ibFSS, and instantiate the protocol with d -dimensional ℓ_∞ -balls. As a result, Alice’s computation cost is reduced from $\mathcal{O}(|S_A|)$ to $\mathcal{O}(u^d)$ for a single ball (see Table 1 in Section 7.1). It is worth noting that our ibFSS construction only uses lightweight operations of pseudorandom generators (PRGs) and may be of independent technical interest.

New, efficient construction for union of (overlapping) structured sets. We develop a new approach for PSI with union of structures as input, that improves *both computation and communication costs* from exponential to linear in the dimension d , namely from $\mathcal{O}(u^d)$ to $\mathcal{O}(u \cdot d)$ (see Table 2 in Section 7.2). The construction in prior works [GRS22, GRS23] crucially requires that Alice’s structure is a union of *disjoint* balls, because it is unclear how to construct an efficient FSS scheme for a union of *overlapping* structures. However, such a requirement may not be feasible in certain applications such as privacy-preserving ride sharing. We answer the following question in the affirmative:

*Can we achieve sa-PSI for a union of **overlapping** balls?*

New spatial hashing techniques. Garimella et al. [GRS22] introduced a *spatial hashing* technique to improve the performance of structure-aware PSI for a union of N_A disjoint d -dimensional ℓ_∞ -balls of the same diameter δ . This technique was further improved in [GRS23]. In this work, we introduce a new perspective on spatial hashing, which improves both computation and communication costs from $\mathcal{O}(u \cdot d \cdot (\log \delta)^d)$ to $\mathcal{O}(\log \delta \cdot d)$ (see Table 3 in Section 7.3). Furthermore, we relax the restrictions on the balls, supporting overlapping balls with different diameters. The new spatial hashing techniques is compatible with the previous framework [GRS22, GRS23] and improves the communication cost of all prior constructions.

Extended functionalities. The previous construction [GRS22, GRS23] is critically restricted to *one-sided, plain* PSI, where Alice, the party holding the structured set, learns the entire set intersection. It is unclear how to extend it to more advanced functionalities. For instance, certain real-world applications may require *only Bob* to learn the output. A naive approach that does not work is to have Alice send the FSS evaluations of all her elements in S_A to Bob, which is highly inefficient. Even worse, these evaluations can reveal critical information about Alice’s structured set; in fact, they reveal Alice’s entire set to Bob! As another extension, we may want to reveal only aggregate information about the intersection, while hiding the actual contents of the intersection. One notable functionality, *PSI-Cardinality*, allows two parties to only learn the cardinality of their intersection $|S_A \cap S_B|$.

*Can we achieve sa-PSI that enables **only Bob** to learn the output?*
*Can we achieve sa-PSI that enables Alice (or Bob) to learn **PSI-Cardinality**?*

In this work, we present new *sa-PSI* protocols for these extended functionalities, allowing them to learn *PSI-Cardinality* as well as *PSI-Sum* [IKN⁺20], in which Bob has an integer value associated with every element in his set and they jointly learn the summation of all the associated values for elements in the intersection. We further allow *only Bob* to learn the output in all these functionalities, including PSI, PSI-Cardinality, and PSI-Sum. All our protocols have communication and computation costs that only scale with the description size of the structured set. Our constructions only use lightweight symmetric-key cryptographic operations (such as pseudorandom generators and hash functions) and oblivious transfer (OT) [Rab05], which can be instantiated efficiently with OT extension [IKNP03].

1.2 Related Work

Conventional PSI and related notions. In recent years, many paradigms for PSI protocols have evolved, including circuit-based [HEK12,PSSZ15,PSWW18,PSTY19], Diffie-Hellman [HFH99,DT10,JL10,IKN⁺20], oblivious polynomial evaluation [KS05,DMRY11], RSA [DT10,ADT11], fully homomorphic encryption [CLR17,CHLR18,CMdG⁺21], bloom filters [DCW13,RR17], oblivious transfer [PSZ14,KKRT16,PRTY19,CM20], and vector oblivious linear evaluation [RS21,GPR⁺21,CRR21,RR22] approaches, achieving both semi-honest and malicious security [RR17,OOS17,CHLR18,PRTY20,RS21,CILO22,BPSY23].

As discussed earlier, certain applications require PSI with more refined functionalities, such as allowing only a designated party to learn the output or enabling restricted computation on the elements in the intersection [PSTY19,KK20,HMS21,GMR⁺21,RS21,HMS21,CGS22]. For instance, PSI-Cardinality and PSI-Sum model many applications like aggregated ads measurement [IKN⁺20,MPR⁺20] and privacy-preserving contact tracing [TSS⁺20,BBV⁺20].

Structure-aware PSI. Recent work of Garimella et al. [GRS22] introduced the notion of structure-aware PSI, where Alice holds a set with a publicly known structure and Bob holds an unstructured set of points. In such a setting where one party’s set could be substantially larger than the other, namely *unbalanced PSI*, it is possible to construct protocols with communication sub-linear in the larger set using RSA accumulators [DT10,ADT11], leveled fully homomorphic encryption [CLR17,CHLR18,CMdG⁺21], or laconic PSI [ABD⁺21,ALOS22,DKL⁺23,GHMM24]. However, they require expensive public-key cryptographic operations, and the computation cost remains linear in the larger set.

In contrast, the work of Garimella et al. [GRS22] constructs a semi-honest structure-aware PSI using lightweight cryptographic tools including PRG, hash functions, and OT, taking advantage of the efficient OT extension [IKNP03]. A follow-up work [GRS23] extends the protocol to achieve malicious security. Both works achieve communication cost that only scales with the description size instead of cardinality of Alice’s structured set. However, Alice’s computation cost still scales with the cardinality of her set. Another recent work [vBP24] achieves both communication and computation costs that scale linearly with the diameter δ and exponentially in the dimension d , in two interaction messages. They rely on the Decisional Diffie-Hellman (DDH) assumption and public-key operations, especially leveraging the homomorphism of ElGamal encryption [ElG85].

The following works [AKB07,WG14,WXL⁺18,GS19,ZCL21] study fuzzy matching or threshold PSI, that reveals the intersection, only when it is above a certain threshold. Chakraborti et al. [CFR21] and [UCK⁺21] construct fuzzy PSI protocols (which they call *distance-aware PSI*) for Hamming distance metric, based on homomorphic encryption.

Function Secret Sharing. Function secret sharing (FSS) was first introduced by Boyle et al. [BGI15] who proposed efficient FSS constructions for point functions, comparison functions, and a few other interesting classes. These original FSS constructions were further optimized and extended in a sequence of works for point functions, multi-point functions, comparison functions and d -dimensional intervals [BGI16, BCG⁺21, BGIK22]. Boneh et al. [BBC⁺21] extended the construction of Distributed Point Function (DPF), namely FSS for point functions, to Incremental DCF, which allows for evaluation on any prefix of a string. In this work, we further extend the notion to Incremental Distributed Comparison Function (DCF) and give an efficient construction of it from PRG, which may be of independent interest. *Boolean FSS (bFSS)* was introduced by Garimella et al. [GRS22] to capture structured sets, and they present bFSS constructions for union of disjoint structured sets. We formalize an extended notion of Incremental Boolean FSS (ibFSS) and present an ibFSS construction for d -dimensional ℓ_∞ balls, from which we construct our sa-PSI protocol.

2 Technical Overview

Prior framework for sa-PSI. We first briefly revisit the framework in prior works [GRS22, GRS23]. The main building block is a weak Boolean FSS scheme t -bFSS (with evaluation output bit length t) for Alice’s set family (e.g., a single-dimensional interval, d -dimensional ball, or union of balls). The scheme succinctly secret shares a set S_A and produces a pair of keys $(k_0, k_1) \leftarrow \text{Share}(1^\kappa, S_A)$ with the property that each key share k_b hides the structure S_A and that $\text{Eval}(k_0, \vec{y}) \oplus \text{Eval}(k_1, \vec{y}) = 0^t$ when $\vec{y} \in S_A$ and $\text{Eval}(k_0, \vec{y}) \oplus \text{Eval}(k_1, \vec{y}) \neq 0^t$ otherwise. The authors leverage this tool for set-membership testing and realize a secure PSI scheme as follows.

Alice generates κ *independent* sharings of her set and obtains (short) keys $(k_0^{(i)}, k_1^{(i)}) \leftarrow \text{Share}(1^\kappa, S_A)$ for each instance $i \in [\kappa]$, where κ is the security parameter. Bob randomly samples a string $s \leftarrow \{0, 1\}^\kappa$ to indicate his choice bits for each sharing. Alice and Bob perform an oblivious transfer (OT) protocol [Rab05], where Bob, as the receiver, learns one of the two shares $k_{s[i]}^{(i)}$ based on his choice bit $s[i]$, while Alice remains oblivious to Bob’s choice bits. Bob defines a special function that can be computed as follow:

$$F(\vec{y}) = \text{H} \left(\text{Eval}(k_{s[1]}^{(1)}, \vec{y}) \parallel \dots \parallel \text{Eval}(k_{s[\kappa]}^{(\kappa)}, \vec{y}) \right).$$

Bob sends $\{F(\vec{y}) \mid \vec{y} \in S_B\}$ to Alice. Alice can *only* compute this special function if $\vec{y} \in S_A$, otherwise the output looks pseudorandom. Why? When $\vec{y} \in S_A$, both key shares evaluate to the same value, namely $\text{Eval}(k_0^{(i)}, \vec{y}) = \text{Eval}(k_1^{(i)}, \vec{y})$ by the t -bFSS property, hence the function evaluation is independent of Bob’s choice bits. She simply computes the function as follows:

$$F(\vec{y}) = \text{H} \left(\text{Eval}(k_0^{(1)}, \vec{y}) \parallel \dots \parallel \text{Eval}(k_0^{(\kappa)}, \vec{y}) \right).$$

When $\vec{y} \notin S_A$, it holds that $\text{Eval}(k_0^{(i)}, \vec{y}) \neq \text{Eval}(k_1^{(i)}, \vec{y})$, so Alice must simultaneously guess all of Bob’s choice bits (which is negligibly likely), to compute the function. The output $F(\vec{y})$ is pseudorandom to Alice assuming H is a correlation robust hash function. Therefore, Alice only recognizes evaluations on values in the intersection.

Cost Analysis. If there exists a t -bFSS scheme with succinct keys (of size $\sigma \ll |S_A|$) for Alice’s set family, then the sa-PSI protocol is communication-efficient with cost $\mathcal{O}((\sigma + |S_B|) \cdot \kappa)$. However,

an inherent drawback of this approach is that Alice must compute the following set of size $|S_A|$

$$\{\mathsf{H}(\mathsf{Eval}(k_0^{(1)}, \vec{x}) \parallel \dots \parallel \mathsf{Eval}(k_0^{(\kappa)}, \vec{x})) \mid \vec{x} \in S_A\}$$

in order to locally compare with Bob’s message of hash evaluations, and identify matches to learn the intersection. Alice’s computation cost scales with the cardinality of her structured input $|S_A|$, which can be prohibitively large for many applications.

Our new framework for sa-PSI. Our first contribution is a new framework for sa-PSI where *both computation and communication costs scale with the (short) description size* of Alice’s set, from a new tool called weak Incremental Boolean Function Secret Sharing **ibFSS**. At a high-level, the idea is that Bob will craft and send additional “hints” along with hash evaluations on his inputs. The hints will help Alice *efficiently identify and search for values* that are in the intersection, without having to expand and compute on every item in her full set.

Firstly, we observe that a structured set from a set family $S_A \in \mathcal{S}$ can be succinctly represented by a distinct and bounded set of w *critical prefixes*, where w is a direct function of \mathcal{S} . Each of these critical prefixes is a placeholder for all inputs in S_A that share the same prefix. Stated differently, every input in S_A has one of the w critical prefixes as its prefix. For instance, if Alice’s structured set is a one-sided interval $(0, \alpha]$, where the partition point $\alpha \in \{0, 1\}^u$, then at most u critical prefixes will span her interval (since α is a string of length u). Alternatively, if the structured set is a d -dimensional ℓ_∞ -ball, which is a cross product of $2 \cdot d$ one-dimensional intervals, then at most $u^{2 \cdot d}$ critical prefixes will span such an input. In our new sa-PSI framework, we exploit the fact that the set of critical prefixes is much smaller than the total set size to reduce computation. For this, we will require an additional property on the weak Boolean FSS (from prior work [GRS22]) that allows evaluation not only on the inputs strings but also on all its prefixes. The guarantee will be $\mathsf{Eval}(k_0, \vec{y}) \oplus \mathsf{Eval}(k_1, \vec{y}) = 0^t$ for all \vec{y} that is a critical prefix or contains a critical prefix (this implicitly includes every input in Alice’s set S_A). For all other \vec{y} , it holds that $\mathsf{Eval}(k_0, \vec{y}) \oplus \mathsf{Eval}(k_1, \vec{y}) \neq 0^t$. We formalize these requirements and define **weak Incremental Boolean Function Secret Sharing** **ibFSS** in in Section 4, and present a construction for d -dimensional ℓ_∞ -balls in Section 6.

So, how does our sa-PSI protocol work and what are these “hints”? Let’s use a simple example. Alice’s input S_A is a single d -dimensional ball (or any other structure in d -dimensional space), while Bob’s input S_B contains a single point $\vec{y} = (y_1, \dots, y_d) \in [2^u]^d$. Following prior work, Alice secret shares her input κ times using **ibFSS**. Bob then samples a uniform string $s \leftarrow \{0, 1\}^\kappa$ and via OT learns one of the two keys for each sharing. Now, Bob defines a special function using his shares and evaluates on his input \vec{y} :

$$\mathsf{F}(\vec{y}) = \mathsf{H}\left(\mathsf{Eval}(k_{s[i]}^{(1)}, \vec{y}) \parallel \dots \parallel \mathsf{Eval}(k_{s[i]}^{(\kappa)}, \vec{y})\right).$$

Additionally, he evaluates on all prefixes \vec{y}' of his input as “hints” and sends to Alice:

$$\left\{ \mathsf{F}(\vec{y}) \mid \begin{array}{l} \ell_1 \in [u], \dots, \ell_d \in [u], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right\}.$$

Alice identifies a set of critical prefixes for her input ball. The claim is that if Bob’s input $\vec{y} \in S_A$, then there is exactly one critical prefix $\vec{p} = (p_1, \dots, p_d)$ that is a prefix match of \vec{y} . As guaranteed by the **ibFSS** scheme, Alice can compute Bob’s function on this prefix \vec{p} as follows:

$$\mathsf{F}(\vec{p}) = \mathsf{H}\left(\mathsf{Eval}(k_0^{(1)}, \vec{p}) \parallel \dots \parallel \mathsf{Eval}(k_0^{(\kappa)}, \vec{p})\right).$$

Alice can recognize $F(\vec{p})$ from Bob’s set of hints. Furthermore, it’s not hard to see that exactly one of the two prefixes $(p_1\|0, p_2, \dots, p_d), (p_1\|1, p_2, \dots, p_d)$ is a prefix of \vec{y} . Alice can correctly compute the function on the matching prefix and is guaranteed to find this value in Bob’s set of hints (since he sends evaluations on every prefix of his input). Similarly, Alice can iteratively figure out the rest of the bits of \vec{y} using binary search, to learn the intersection. Finally, if Bob’s input $\vec{y} \notin S_A$, then Alice cannot compute Bob’s function correctly on any of her prefixes, and as a result, learns nothing about his private input. This is the crux of our new approach and how we achieve savings in computation at the cost of slightly increased communication from Bob.

Multiple (overlapping) balls with multiple points. Next, we consider the case when Alice’s input is a union of multiple balls (or any structure), namely $S_A = \vec{S}_A^{(1)} \cup \dots \cup \vec{S}_A^{(N_A)}$, and Bob’s input contains multiple points $|S_B| > 1$. Prior work [GRS22, GRS23] introduced an elegant sum transformation to obtain a 1-bFSS scheme for a *disjoint* union of structures. However, this technique is only compatible with a restricted class of *strong* 1-bFSS schemes that output a single bit, whose constructions are significantly more expensive. In particular, both the key share sizes and evaluation costs scale exponentially in the dimension d .

In this work, we circumvent this limitation to realize sa-PSI for union of structures (even possibly *overlapping* structures, under certain conditions) from the more efficient *weak t*-bFSS schemes with both key share sizes and evaluation costs linear in d . Our improvement comes from the following observation. If Alice learns from the PSI ideal functionality that an element \vec{y} is in the intersection, then it implicitly reveals the underlying structured set(s) that contain \vec{y} , that is, every $i \in [N_A]$ for which $y \in \vec{S}_A^{(i)}$. Hence, we propose a new protocol where Alice compares each of her structures $\vec{S}_A^{(i)}, i \in [N_A]$, individually, with Bob’s input \vec{y} (effectively, N_A instances of PSI with a single structure against Bob’s input \vec{y}). Since we evaluate each structure separately, the underlying *t*-bFSS scheme can be instantiated more efficiently (no longer restricted by the sum transformation) and it’s compatible even when the structures are overlapping.

To handle multiple points in Bob’s set, he will send prefix evaluations for each of his inputs to Alice. It could happen that some prefixes overlap across different inputs, which may leak extra information to Alice. A simple fix is to have Bob send the union of his evaluations across all inputs and pad with dummy values, which we will show is sufficient. We present the full construction in [Section 5.1](#).

Spatial hashing. One drawback of our sa-PSI protocol discussed so far is that Bob’s communication cost scales with the number of structures in Alice’s set N_A . Garimella et al. [GRS22] introduced a technique called *spatial hashing*, that removes this dependence for a *disjoint union* of d -dimensional ℓ_∞ -balls of the same, fixed diameter δ . This technique was further improved in [GRS23]. Their high-level idea is to divide the input space into contiguous grid cells, construct FSS keys for all the active grid cells that contain or intersect with Alice’s input balls, and then pack these FSS keys into an oblivious key value store (OKVS) data structure [GPR⁺21]. In our work, we present a new, improved spatial hashing technique, that is composable with our new (and prior) framework instantiated with *weak* FSS, and supports overlapping structures of varying size. We discuss more in [Section 5.2](#).

Extended functionalities. We solve another open question from prior work – can we construct sa-PSI where Bob (with unstructured set) learns the intersection? We leverage the fact that, Alice can identify a (bounded) set of $w \ll |S_A|$ critical prefixes to represent her set. For a given input $\vec{y} \in S_B$, Bob only needs to compare all its prefixes with Alice’s critical prefixes. If Bob learns that

there is exactly one match, then he includes \bar{y} in the intersection; otherwise not. A straightforward method is to run a standard PSI-Cardinality between both sets of prefixes, for each of Bob’s input. We present the construction in [Section 8.1](#).

Finally, we present the first construction for Structure-Aware PSI-Cardinality. A naive approach is that both parties run the best known semi-honest PSI-Cardinality protocol on their inputs. However, since the communication and computation will scale with Alice’s set size $|S_A|$, the protocol is not efficient. Instead, since Alice can identify a set of $w \ll |S_A|$ critical prefixes, if Bob’s input intersects/matches with any of the critical prefixes, then it should be counted for intersection cardinality, otherwise not. This idea can be further generalized to compute PSI-Sum. We discuss the full details in [Section 8.2](#).

3 Preliminaries

Notation. Let κ and λ denote the computational and statistical security parameter, respectively. For an integer $m \in \mathbb{N}$, let $[m] = \{1, 2, \dots, m\}$. PPT stands for probabilistic polynomial time. By \cong_κ and \cong_λ we mean two distributions are computationally indistinguishable and statistically close, respectively. Given a string $\alpha \in \{0, 1\}^*$, we use $|\alpha|$ to denote its length, and $\alpha_{[1:i]}$ to represent the prefix of length i of α , where $i \in [|\alpha|]$. For two strings $\alpha, \beta \in \{0, 1\}^*$, we use $\alpha\|\beta$ to denote string concatenation.

Ball of diameter δ in ℓ_∞ metric space. Let $\mathbf{x} = (x_1, \dots, x_d)$ be a point in a d -dimensional space. The ℓ_∞ norm of \mathbf{x} is defined as $\|\mathbf{x}\|_\infty \stackrel{\text{def}}{=} \max_i |x_i|$. A **ball** consists of the set of points within some distance of a center point. The ball of diameter δ (or radius $\frac{\delta}{2}$) centered at x is defined as $\mathcal{B}(x, \frac{\delta}{2}) \stackrel{\text{def}}{=} \{y \mid d(x, y) \leq \frac{\delta}{2}\}$. Under the ℓ_∞ norm, a ball $\mathcal{B}(x, \frac{\delta}{2})$ is a polytope with $2d$ faces, i.e., the intersection of $2d$ half-spaces. Namely, the ball $\mathcal{B}(\mathbf{0}, \frac{\delta}{2})$ centered at the origin is the intersection of all half-spaces of the form $\pm x_i \leq \frac{\delta}{2}$ for every i and every choice of sign.

Oblivious Key Value Stores. An oblivious key value stores (OKVS) (introduced in [\[GPR⁺21\]](#)) is a data structure that encodes a set of key value mappings, which effectively hides the underlying key set when the associated values are pseudorandom.

Definition 1 (Oblivious Key Value Stores (OKVS) [\[GPR⁺21\]](#)). *An OKVS consists of algorithms Encode and Decode, with an associated key space \mathcal{K} and value space \mathcal{V} .*

- *Encode takes as input a set of $A \in \mathcal{K} \times \mathcal{V}$ of key-value pairs and outputs an object \mathcal{KV} (or, with statistically small probability, an error indicator \perp).*
- *Decode takes as input an object \mathcal{KV} , any key $k \in \mathcal{K}$, and outputs a value v .*

An OKVS must satisfy the following properties:

- **Correctness:** *For all $A \subseteq \mathcal{K} \times \mathcal{V}$ with distinct keys, and all $(k, v) \in A$:*

$$\Pr[\text{Decode}(\text{Encode}(A), k) = v] = 1$$

Note that, we can invoke $\text{Decode}(\mathcal{KV}, k)$ on any key k on a given object \mathcal{KV} .

- **Obliviousness:** *For all distinct $\{k_1^0, \dots, k_n^0\}$ and distinct $\{k_1^1, \dots, k_n^1\}$, the output of $\mathcal{R}(k_1^0, \dots, k_n^0)$*

is computationally indistinguishable from that of $\mathcal{R}(k_1^1, \dots, k_n^1)$, where:

```

 $\mathcal{R}(k_1, \dots, k_n):$ 
for  $i \in [n]: v_i \leftarrow \mathcal{V}$ 
return  $\text{Encode}(\{(k_1, v_1), \dots, (k_n, v_n)\})$ 

```

The obliviousness property guarantees that if the values are pseudorandom, then any two outputs of experiment \mathcal{R} , encoding different sets of keys are computationally indistinguishable. Additionally, our construction requires an additional independence property, that is satisfied by all the constructions presented in [GPR⁺21].

Definition 2. An OKVS satisfies the *independence property* if for all $A \subseteq \mathcal{K} \times \mathcal{V}$ with distinct \mathcal{K} -values, and any k^* not appearing in the first component of any pair in A , the output of $\text{Decode}(\text{Encode}(A), k^*)$ is indistinguishable from random, over the randomness in Encode .

Other cryptographic tools. We give definitions for other cryptographic tools including pseudorandom generators, Hamming correlation robust hash functions and secure two-party computation (2PC) and specific 2PC functionalities, including oblivious transfer (OT), PSI cardinality, and PSI with associated sum in the full version of the paper.

4 Incremental Boolean Function Secret Sharing

In this section, we define a weak multi-dimensional incremental Boolean Function Secret Sharing (ibFSS) scheme. Looking ahead, we will use this scheme as a building block to construct our structure-aware PSI protocol in Section 5. We then give a construction of ibFSS for d -dimensional ℓ_∞ -balls in Section 6 and discuss our PSI protocol instantiated with it in Section 7.

A d -dimensional ibFSS scheme is defined for a family of d -dimensional sets $\mathcal{S} \subseteq \underbrace{2^{\mathcal{U}} \times \dots \times 2^{\mathcal{U}}}_d$

where $\mathcal{U} = \{0, 1\}^u$.¹ As a concrete example, one can think of \mathcal{S} as the family of all d -dimensional ℓ_∞ -balls. For each d -dimensional set in the family \mathcal{S} , namely $\vec{S} = (S_1 \times \dots \times S_d) \in \mathcal{S}$ (in the example, \vec{S} is a d -dimensional ℓ_∞ -ball), the set \vec{S} implicitly defines a Boolean function f on input space $(\{0, 1\}^u)^d$ as follows. For all $\vec{x} = (x_1, \dots, x_d)$ where each $x_i \in \{0, 1\}^u$, we define $f(\vec{x}) = 0$ if $\vec{x} \in \vec{S}$, and $f(\vec{x}) = 1$ otherwise. In *incremental* Boolean FSS, we additionally define the function f evaluated on all prefixes of \vec{x} . Specifically, for all $\vec{x} = (x_1, \dots, x_d)$ where each $x_i \in \bigcup_{\ell=1}^u \{0, 1\}^\ell$ can be of arbitrary length (at most u bits), we define $f(\vec{x}) = 0$ if the prefix \vec{x} spans a set that is entirely contained in \vec{S} (in our formal definition below, $\text{PreS}_{\vec{x}} \subseteq \vec{S}$), and $f(\vec{x}) = 1$ otherwise. An ibFSS scheme splits the function f into two functions f_0, f_1 (defined by keys k_0, k_1) such that each function f_b hides f , and that $f_0(\vec{x}) \oplus f_1(\vec{x}) = f(\vec{x})$ for all prefixes \vec{x} .

What we need for structure-aware PSI is a *weak* ibFSS, which relaxes the above definition in two ways. First, the function f outputs a t -bit string instead of a single bit. Second, for each prefix \vec{x} , it holds that $f_0(\vec{x}) \oplus f_1(\vec{x}) = 0^t$ if $\text{PreS}_{\vec{x}} \subseteq \vec{S}$ and $f_0(\vec{x}) \oplus f_1(\vec{x}) \neq 0^t$ otherwise. This weak definition follows from prior work [GRS22, GRS23]. In fact, they further relaxed the definition to allow for false positives, which we don't need to consider in this work. Throughout the rest of the paper, we refer to ibFSS as the *weak* variant unless specified otherwise.

¹In the ibFSS definition, we consider all dimensions to have the same input domain for simplicity. The same definition and our constructions also work for dimensions with different input domains.

Definition 3. *[(u, d, t)-ibFSS: Syntax]* A (two-party) d -dimensional weak incremental Boolean function secret sharing scheme (u, d, t)-ibFSS for a family of sets $\mathcal{S} \subseteq \underbrace{2^{\mathcal{U}} \times \dots \times 2^{\mathcal{U}}}_d$, where $\mathcal{U} = \{0, 1\}^u$,

consists of a pair of algorithms (Share, Eval) with the following syntax:

- $(k_0, k_1) \leftarrow \text{Share}(1^\kappa, \hat{S})$: The randomized share function takes as input the security parameter κ and (the description of) a d -dimensional set $\vec{S} = (S_1 \times \dots \times S_d) \in \mathcal{S}$, and outputs two key shares.
- $y_{idx} \leftarrow \text{Eval}(idx, k_{idx}, \vec{x})$: The deterministic evaluation function takes as input a party index $idx \in \{0, 1\}$, the corresponding key share k_{idx} , and input $\vec{x} = (x_1, \dots, x_d)$ where each $x_i \in \bigcup_{\ell=1}^u \{0, 1\}^\ell$. It outputs a string $y_{idx} \in \{0, 1\}^t$.

Notation. For a string $x \in \bigcup_{\ell=1}^u \{0, 1\}^\ell$ and $\mathcal{U} = \{0, 1\}^u$, we define a prefix set $\text{PreS}_x := \{s \mid s_{[1:|x|]} = x, s \in \mathcal{U}\}$ as the set of all the strings in \mathcal{U} with a prefix x . For a vector $\vec{x} = (x_1, \dots, x_d)$, where each $x_i \in \bigcup_{\ell=1}^u \{0, 1\}^\ell$, we define $\text{PreS}_{\vec{x}} := \{\vec{s} = (s_1, \dots, s_d) \mid (s_i)_{[1:|x_i|]} = x_i \text{ for all } i \in [d], s_i \in \mathcal{U}\}$. Note that $\text{PreS}_{\vec{x}} = \text{PreS}_{x_1} \times \dots \times \text{PreS}_{x_d}$.

Definition 4. *[(u, d, t)-ibFSS: Security]* A (two-party) (u, d, t)-ibFSS scheme (Share, Eval) for \mathcal{S} is secure if it satisfies the following conditions:

- **Correctness for yes-instances:** For every d -dimensional set $\vec{S} = (S_1 \times \dots \times S_d) \in \mathcal{S}$, every input $\vec{x} = (x_1, \dots, x_d)$ for which $\text{PreS}_{x_1} \subseteq S_1, \dots, \text{PreS}_{x_d} \subseteq S_d$ (or equivalently, $\text{PreS}_{\vec{x}} \subseteq \vec{S}$), and security parameter κ ,

$$\Pr \left(y_0 \oplus y_1 = 0^t \mid \begin{array}{l} (k_0, k_1) \leftarrow \text{Share}(1^\kappa, \vec{S}) \\ y_0 \leftarrow \text{Eval}(0, k_0, \vec{x}) \\ y_1 \leftarrow \text{Eval}(1, k_1, \vec{x}) \end{array} \right) = 1.$$

- **Correctness for no-instances:** For every d -dimensional set $\vec{S} = (S_1 \times \dots \times S_d) \in \mathcal{S}$, every input $\vec{x} = (x_1, \dots, x_d)$ for which $\text{PreS}_{x_1} \not\subseteq S_1$ or \dots or $\text{PreS}_{x_d} \not\subseteq S_d$ (or equivalently, $\text{PreS}_{\vec{x}} \not\subseteq \vec{S}$), and security parameter κ ,

$$\Pr \left(y_0 \oplus y_1 \neq 0^t \mid \begin{array}{l} (k_0, k_1) \leftarrow \text{Share}(1^\kappa, \vec{S}) \\ y_0 \leftarrow \text{Eval}(0, k_0, \vec{x}) \\ y_1 \leftarrow \text{Eval}(1, k_1, \vec{x}) \end{array} \right) = 1.$$

- **Privacy:** There exists a PPT simulator Sim such that for all $idx \in \{0, 1\}$ and all $\vec{S} \in \mathcal{S}$, the following distributions are computationally indistinguishable in the security parameter κ :

$$\boxed{\begin{array}{l} (k_0, k_1) \leftarrow \text{Share}(1^\kappa, \vec{S}) \\ \text{return } k_{idx} \end{array}} \cong_{\kappa} \text{Sim}(1^\kappa, idx).$$

We say the ibFSS scheme has pseudorandom keys property, if the output of the simulator is a random string of fixed length.

Decomposing a set into prefix sets. In our structure-aware PSI protocol, we will decompose each structured set into a disjoint union of prefix sets. Specifically, for any single-dimensional set $S \in 2^{\mathcal{U}}$, there is a unique way to decompose it into a disjoint union of prefix sets, $S = \dot{\bigcup}_{j \in [w]} \text{PreS}_{x^j}$

where $x^j \in \bigcup_{\ell=1}^u \{0, 1\}^\ell$. Similarly, for any d -dimensional set $\vec{S} = (S_1 \times \cdots \times S_d) \in \underbrace{2^\mathcal{U} \times \cdots \times 2^\mathcal{U}}_d$, there is a unique way to decompose it into a disjoint union of prefix sets, $\vec{S} = \dot{\bigcup}_{j \in [w]} \text{PreS}_{x^j}^{\vec{S}}$ where $\vec{x}^j = (x_1^j, \dots, x_d^j)$ and $x_i^j \in \bigcup_{\ell=1}^u \{0, 1\}^\ell$

For instance, consider a single-dimensional interval $S = \{x \mid x < \alpha, x \in \{0, 1\}^u\}$, where $\alpha = \overline{\alpha^{(1)}\alpha^{(2)} \dots \alpha^{(u)}} \in \{0, 1\}^u$ (bit representation of α). The set S can be decomposed into a disjoint union of at most u prefix sets, namely $S = \bigcup_{j \in [u]: \alpha^{(j)}=1} \text{PreS}_{\overline{\alpha^{(1)} \dots \alpha^{(j-1)} 0}}$. Similarly, a d -dimensional interval can be decomposed into a disjoint union of at most u^d prefix sets.

5 sa-PSI with Alice Learning Output

In this section, we present our new sa-PSI protocols where Alice holds a structured set and Bob holds a set of unstructured points. We present the protocol where Alice’s input is a union of (overlapping) structures and Bob’s input contains multiple points in [Section 5.1](#) and our new spatial hashing techniques in [Section 5.2](#). In both protocols, only Alice learns the output.

5.1 Multiple (Overlapping) Balls with Multiple Points

We present in [Figure 2](#) a general framework for two-party, semi-honest secure sa-PSI protocol, where Alice’s input is a union of (overlapping) structures in a d -dimensional, well-defined metric space. We summarize the key features in our framework below.

- **Incremental Boolean FSS** is a new tool, that allows Alice and Bob to evaluate on the prefixes of their inputs. Bob evaluates his key share on *every prefix* of his input. Alice identifies a set of *critical prefixes*, such that, a match with any critical prefix indicates that the element is in the intersection.
- **Intersection Search** is a method for Alice to recursively search for Bob’s input in her set. For every critical prefix in Alice’s input, she checks if there is a matching FSS evaluation in Bob’s set Y . She then appends bit 0 to the prefix and checks if the FSS evaluation belongs to set Y . If not, she is guaranteed to find prefix append with bit 1 in set Y . This process continues until she figures out the matching input.
- **OR Observation** is a new method for sa-PSI when Alice’s input is a union of N_A structures. In our protocol, we break away from the ibFSS abstraction for the union of structures. Instead, Alice computes an independent ibFSS secret sharing *for each of her structures* and compares with *Bob’s entire input* set S_B . An immediate advantage is that the sets can be overlapping since they are handled independently. Furthermore, prior known FSS constructions for union of structures are not as efficient as FSS for individual structures, leading to an advantage for our approach (for the set families considered in [[GRS22](#), [GRS23](#)]). We discuss the comparison in more detail in [Section 7](#).

Theorem 5. *Given a two-party (u, d, t) -ibFSS scheme for a family of sets $\mathcal{S} \subseteq \underbrace{2^\mathcal{U} \times \cdots \times 2^\mathcal{U}}_d$ where*

$\mathcal{U} = \{0, 1\}^u$ and every set in \mathcal{S} is a disjoint union of at most w prefix sets, and a Hamming-correlation robust hash $H : \{0, 1\}^ \rightarrow \{0, 1\}^{\lambda + \log |N_A| + \log |N_B| + \log w + d \log u}$, the protocol in [Figure 2](#) realizes $\mathcal{F}_{\text{sa-PSI}}$ ([Figure 1](#)) in the \mathcal{F}_{OT} -hybrid model in the presence of semi-honest adversaries.*

Parameters: A family of sets $\mathcal{S} \subseteq \underbrace{2^{\mathcal{U}} \times \dots \times 2^{\mathcal{U}}}_d$ where $\mathcal{U} = \{0, 1\}^u$. Number of Alice's structured sets N_A and size of Bob's set N_B .

Functionality:

1. Receive input $S_A = \bigcup_{i \in N_A} \vec{S}_A^{(i)}$, where $\vec{S}_A^{(i)} \in \mathcal{S}$ (or a concise representation of S_A) from Alice.
2. Receive input $S_B \subseteq \mathcal{U}^d$ of size N_B from Bob.
3. **[output]** Send $S_A \cap S_B$ to Alice.

Figure 1: Ideal functionality $\mathcal{F}_{\text{sa-PSI}}$ for structure-aware PSI, where Alice learns the output.

Algorithm 1 Recursive Intersection Search

```

1: global variable: intersection set  $I$ 
2: procedure INTSEARCH( $\{k^{(\eta)}\}_{\eta \in [\kappa]}$ ,  $u$ , index,  $\vec{x}$ ,  $Y$ )
3:   Parse  $\vec{x} = (x_1, \dots, x_d)$ 
4:   Compute  $h := \mathbf{H}(\vec{x} \parallel \text{index}; \text{Eval}(k^{(1)}, \vec{x}) \parallel \dots \parallel \text{Eval}(k^{(\kappa)}, \vec{x}))$ 
5:   if  $h \notin Y$  then
6:     return
7:   else if  $|x_i| = u$  for all  $i \in [d]$  then
8:      $I := I \cup \{\vec{x}\}$ 
9:   return
10:  Let  $i \in [d]$  be the smallest index such that  $|x_i| \leq u - 1$ 
11:  for  $b = 0$  to  $1$  do
12:     $\vec{x}' := (x_1, \dots, x_{i-1}, x_i \parallel b, x_{i+1}, \dots, x_d)$ 
13:    INTSEARCH( $\{k^{(\eta)}\}_{\eta \in [\kappa]}$ ,  $u$ , index,  $\vec{x}'$ ,  $Y$ )

```

Proof. **Bob is corrupt.** In the protocol, Bob learns only one of the key shares from OT for each of the κ independent ibFSS sharings. We can simulate Bob's view by invoking the ibFSS simulator (Definition 4) that takes as input the security parameter κ and (public information) a description of Alice's set family and nothing else about her private input.

Alice is corrupt. First, we define the following useful functions, where each function's output $\{\{0, 1\}^t\}^\kappa$ is a vector of length κ , where each component (string length t) is the result of ibFSS evaluation:

$$\begin{aligned}
E_*(\vec{y}) &= \left(\text{Eval}(k_*^{(1)}, \vec{y}) \parallel \dots \parallel \text{Eval}(k_*^{(\kappa)}, \vec{y}) \right) \\
E_0(\vec{y}) &= \left(\text{Eval}(k_0^{(1)}, \vec{y}) \parallel \dots \parallel \text{Eval}(k_0^{(\kappa)}, \vec{y}) \right) \\
\Delta(\vec{y}) &= \left(\text{Eval}(k_0^{(1)}, \vec{y}) \oplus \text{Eval}(k_1^{(1)}, \vec{y}) \parallel \dots \parallel \text{Eval}(k_0^{(\kappa)}, \vec{y}) \oplus \text{Eval}(k_1^{(\kappa)}, \vec{y}) \right)
\end{aligned}$$

The goal is to simulate (through a sequence of hybrids) Alice's view from the honest protocol execution independent of Bob's input.

Hybrid 0. In the protocol execution, Alice's view consists of N_A sets of values from Bob, which can effectively be viewed as a single of set of values Y . We use the fact that

$$\text{Eval}(k_*^{(\eta)}, \vec{y}') = \text{Eval}(k_{s_i}^{(\eta)}, \vec{y}') = \text{Eval}(k_0^{(\eta)}, \vec{y}') \oplus s \odot (\text{Eval}(k_0^{(\eta)}, \vec{y}') \oplus \text{Eval}(k_1^{(\eta)}, \vec{y}'))$$

Parameters:

- computational security parameter κ and statistical security parameter λ
- family of sets \mathcal{S} with corresponding (u, d, t) -ibFSS scheme (Share, Eval)
- oblivious transfer functionality \mathcal{F}_{OT}
- hamming-correlation robust hash H with output length $\lambda + \log |N_A| + \log |N_B| + \log w + d \log u$

Inputs:

- Alice has N_A structured sets $S_A = \bigcup_{i \in N_A} \vec{S}_A^{(i)}$, where $\vec{S}_A^{(i)} \in \mathcal{S}$.
- Bob has an unstructured set $S_B \subseteq \mathcal{U}^d$ of size N_B .

Protocol:

1. Alice initializes the intersection as $I := \emptyset$.
2. Bob chooses a random string $s \leftarrow \{0, 1\}^\kappa$.
3. For each $i \in [N_A]$, the parties do the following (in parallel for each set):
 - (a) Alice generates κ independent ibFSS sharings of her structured set $S_A^{(i)}$. For each $\eta \in [\kappa]$, compute $(k_0^{(\eta)}, k_1^{(\eta)}) \leftarrow \text{Share}(1^\kappa, S_A^{(i)})$.
 - (b) The parties invoke κ (parallel) instances of oblivious transfer using \mathcal{F}_{OT} . In the η -th instance:
 - Alice is the sender with input $(k_0^{(\eta)}, k_1^{(\eta)})$;
 - Bob is the receiver with choice bit $s[\eta]$. He obtains output $k_*^{(\eta)} = k_{s[\eta]}^{(\eta)}$.
 - (c) Bob does the following:
 - i. Initialize an empty set $Y := \emptyset$.
 - ii. For each element $\vec{y} \in S_B$:
 - Write $\vec{y} = (y_1, \dots, y_d)$, where $y_j \in \mathcal{U}$ for all $j \in [d]$.
 - Compute

$$Y' = \left\{ H \left(\vec{y} \| i; \text{Eval}(k_*^{(1)}, \vec{y}) \| \dots \| \text{Eval}(k_*^{(\kappa)}, \vec{y}) \right) \mid \begin{array}{l} \ell_1 \in [u], \dots, \ell_d \in [u] \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right\}$$
 - Update $Y := Y \cup Y'$.
 - iii. Pad Y with dummy random strings such that $|Y| = N_B \cdot u^d$ and send it to Alice.
 - (d) Alice does the following:
 - i. Write $S_A^{(i)} = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x^j}$ as disjoint union of w' prefix sets (w is the upper bound for w').
 - ii. For each $j \in [w']$, run $\text{INTSEARCH} \left(\{k_0^{(\eta)}\}_{\eta \in [\kappa]}, u, i, x^j, Y \right)$, with I being a global variable.
4. **[output]** Alice outputs the intersection I .

Figure 2: Protocol realizing $\mathcal{F}_{\text{sa-PSI}}$ (Figure 1) using ibFSS in the \mathcal{F}_{OT} -hybrid model.

where \odot denotes component wise multiplication (of single bit with t bit string) to rewrite the

message Y as shown below:

$$\begin{aligned}
Y &= \left\{ \text{H} \left(\vec{y}' \| i; \text{Eval}(k_*^{(1)}, \vec{y}') \| \dots \| \text{Eval}(k_*^{(\kappa)}, \vec{y}') \right) \left| \begin{array}{l} i \in [N_A], \vec{y} = (y_1, \dots, y_d) \in S_B, \\ \ell_1 \in [u], \dots, \ell_d \in [u], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\} \\
&= \left\{ \text{H} \left(\vec{y}' \| i; E_*(\vec{y}') \right) \left| \begin{array}{l} i \in [N_A], \vec{y} = (y_1, \dots, y_d) \in S_B, \\ \ell_1 \in [u], \dots, \ell_d \in [u], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\} \\
&= \left\{ \text{H} \left(\vec{y}' \| i; E_0(\vec{y}') \oplus s \odot \Delta(\vec{y}') \right) \left| \begin{array}{l} i \in [N_A], \vec{y} = (y_1, \dots, y_d) \in S_B, \\ \ell_1 \in [u], \dots, \ell_d \in [u], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\} \\
&= Y_1 \dot{\cup} Y_2 \dot{\cup} Y_3 \dot{\cup} Y_4,
\end{aligned}$$

where we expand the terms as follows:

$$\begin{aligned}
Y_1 &= \left\{ \text{H} \left(\vec{y}' \| i; E_0(\vec{y}') \right) \left| \begin{array}{l} i \in [N_A], S_A^{(i)} = \dot{\cup}_{j \in [w']} \text{PreS}_{x_j^i}, \vec{x}^j = (x_1^j, \dots, x_d^j), \\ \vec{y} = (y_1, \dots, y_d) \in \text{PreS}_{x_j^i} \cap S_B, \\ \ell_1 \in [|x_1^j| : u], \dots, \ell_d \in [|x_d^j| : u], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\} \\
Y_2 &= \left\{ \text{H} \left(\vec{y}' \| i; E_0(\vec{y}') \oplus s \odot \Delta(\vec{y}') \right) \left| \begin{array}{l} i \in [N_A], S_A^{(i)} = \dot{\cup}_{j \in [w']} \text{PreS}_{x_j^i}, \vec{x}^j = (x_1^j, \dots, x_d^j), \\ \vec{y} = (y_1, \dots, y_d) \in \text{PreS}_{x_j^i} \cap S_B, \\ (\ell_1, \dots, \ell_d) \in [u]^d, \text{ where } \ell_1 \in [|x_1^j| - 1] \vee \dots \vee \ell_d \in [|x_d^j| - 1], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\} \\
Y_3 &= \left\{ \text{H} \left(\vec{y}' \| i; E_0(\vec{y}') \oplus s \odot \Delta(\vec{y}') \right) \left| \begin{array}{l} i \in [N_A], \vec{y} = (y_1, \dots, y_d) \in S_B \setminus \vec{S}_A^{(i)}, \\ \ell_1 \in [u], \dots, \ell_d \in [u], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\}
\end{aligned}$$

and Y_4 consists of dummy pseudorandom strings with length equal to the hash outputs, such that $|Y| = N_B \cdot u^d$. Recall that, Bob sends hash evaluations on all prefix combinations of each of his inputs for every input $S_A^{(i)} = \dot{\cup}_{j \in [w']} \text{PreS}_{x_j^i}$ in Alice's private input. For Bob's protocol message Y we get the following cases for every set $S_A^{(i)}$:

1. If Bob's input $\vec{y} \in \text{PreS}_{x_j^i}$ and input prefix $\vec{y}' \in \text{PreS}_{x_j^i}$: set Y_1 includes hash evaluation on prefix $\text{H}(\vec{y}' \| i; E_*(\vec{y}'))$. The ibFSS scheme correctness guarantees $\Delta(\vec{y}') = 0^t$ when evaluated on an input \vec{y}' in some prefix set $\vec{y}' \in \text{PreS}_{x_j^i}$ that comprises one of Alice's sets. Thus, we can simplify $\text{H}(\vec{y}' \| i; E_0(\vec{y}') \oplus s \odot \Delta(\vec{y}')) = \text{H}(\vec{y}' \| i; E_0(\vec{y}'))$.
2. If Bob's input $\vec{y} \in \text{PreS}_{x_j^i}$ and input prefix $\vec{y}' \notin \text{PreS}_{x_j^i}$: set Y_2 includes hash evaluation on prefix $\text{H}(\vec{y}' \| i; E_*(\vec{y}'))$. The ibFSS scheme correctness guarantees $\Delta(\vec{y}')$ consists of κ components, where every component (string of length t) is non-zero, when evaluated on an input \vec{y}' outside the prefix set.

3. If Bob's input $\vec{y} \notin S_A^{(i)}$ and input prefix \vec{y}' : set Y_3 includes hash evaluation on prefix $H(\vec{y}'\|i; E_*(\vec{y}'))$. Again, the ibFSS scheme correctness guarantees $\Delta(\vec{y}')$ consists of κ components, where every component (string of length t) is non-zero, when evaluated on an input \vec{y}' outside every prefix set of $S_A^{(i)}$.

Hybrid 1. In this hybrid, we replace all the hash outputs in Y_2 and Y_3 with uniform strings of length $\lambda + \log |N_A| + \log |N_B| + \log w + d \log u$ as shown below:

$$Y = Y_1 \dot{\cup} \{h_1, \dots, h_{|Y_2|}\} \dot{\cup} \{h'_1, \dots, h_{|Y_3|}\} \dot{\cup} Y_4$$

where each $h_i, h'_i \leftarrow \{0, 1\}^{\lambda + \log |N_A| + \log |N_B| + \log w + d \log u}$ are sampled uniformly. This hybrid is indistinguishable from previous hybrid by the hamming correlation robust property ([Definition 13](#)) which states that hash values of the form $H(\vec{y}'\|i; t_i \oplus s \odot \Delta(\vec{y}'))$ are jointly pseudorandom, if they are never queried on repeated $\vec{y}'\|i$ values and $\Delta(\vec{y}')$ has at least κ non-zero components. As previously discussed all the hash outputs in Y_2 and Y_3 are computed on unique prefix concatenate with structure index values and have non-zero $\Delta(\vec{y}')$ evaluation with κ non-zero components.

Simulator. Hybrid 1 defines a valid simulation in the ideal world computed just using Alice's input S_A and output $S_A \cap S_B$. Note that, the output can be used to determine the cardinality of sets Y_2, Y_3 and Y_4 .

Correctness. Only Alice learns the output and our goal is to show that the protocol satisfies correctness. It suffices to consider Alice's simulated view since it is indistinguishable from her view in the honest execution of the protocol.

For every input $\vec{x} \in S_B \cap \text{PreS}_{\vec{x}^j}$ in the intersection (over all choices of $i \in [N_A], S_A^{(i)} = \dot{\cup}_{j \in [w]} \text{PreS}_{\vec{x}^j}, \vec{x}^j = (x_1^j, \dots, x_d^j)$), ibFSS scheme guarantees $\Delta(\vec{x}^j) = 0^t$ and the simulator includes the following hash values in the message to Alice:

$$\left\{ H(\vec{x}^j\|i; E_0(\vec{x}^j)) \left| \begin{array}{l} \ell_1 \in [|x_1^j| : u], \dots, \ell_d \in [|x_d^j| : u] \\ \vec{x}^j := ((x_1)_{[1:\ell_1]}, \dots, (x_d)_{[1:\ell_d]}) \end{array} \right. \right\}$$

By the collision resistant hash property, Alice can uniquely compute and recognize hash values from Bob's message and will correctly include every such element \vec{x} in the output.

What is the probability that Alice wrongly includes an element $\vec{x} \in S_A \setminus S_B$ in her output? This value is upper bounded by the probability that Alice finds the hash output of prefix of \vec{x} in Bob's message, that is, $H(\vec{x}^j\|i; E_0(\vec{x}^j)) \in Y$, where $\vec{x} \in \text{PreS}_{\vec{x}^j}$. By union bound, the probability that Alice finds a value matching a specific prefix $\text{PreS}_{\vec{x}^j}$ is $|Y'| \cdot 2^{-\lambda - \log |N_A| - \log |N_B| - \log w - d \log u} = 2^{-\lambda - \log |N_A| - \log w}$. Again, by the union bound over total number of prefixes, the probability that Alice includes a wrong element \vec{x} in her output is less than $w \cdot |N_A| \cdot 2^{-\lambda - \log |N_A| - \log w} = 2^{-\lambda}$, which is negligible. \square

5.2 Spatial Hashing

In the above construction, both parties compute ibFSS over the entire universe. Improving upon this construction framework, Garimella et al. [[GRS22](#)] introduced a spatial hashing technique to reduce the input domain for better efficiency. In this section, we present a new way of looking at spatial

hashing, which relaxes the restrictions on the structured sets and improves both computation and communication costs.

Overview of our spatial hashing technique. We start with the same setting as prior work [GRS22, GRS23], where Alice holds a union of d -dimensional ℓ_∞ -balls of diameter δ , namely every ball can be written as $[x_1, x_1 + \delta) \times \dots \times [x_d, x_d + \delta)$. We denote the ball as $\text{Ball}_{\vec{x}}$ for its “left-bottom corner” $\vec{x} = (x_1, \dots, x_d)$. We partition the entire universe $(\{0, 1\}^u)^d$ into contiguous grid cells of side length δ . For each vertex in the partitioned space, namely $\vec{o} = (o_1, o_2, \dots, o_d)$ where each o_i is a multiple of δ , we define a *mini-universe* with *origin* \vec{o} and diameter 2δ . That is, $\text{Univ}_{\vec{o}} := [o_1, o_1 + 2\delta) \times \dots \times [o_d, o_d + 2\delta)$. If all of Alice’s balls are disjoint and have the same diameter δ , then each $\text{Ball}_{\vec{x}}$ corresponds to a distinct grid cell, which is where \vec{x} is located at. Let \vec{o} be the “left-bottom corner” of that grid cell, then it is guaranteed that $\text{Ball}_{\vec{x}} \subseteq \text{Univ}_{\vec{o}}$. We can construct an **ibFSS** for $\text{Ball}_{\vec{x}}$ in the mini-universe $\text{Univ}_{\vec{o}}$ with diameter 2δ .

Recall that each ball corresponds to a unique origin \vec{o} , which we refer to as *active origins* with *active mini-universes*. We prepare an **ibFSS** key for each ball $\text{Ball}_{\vec{x}}$ in its corresponding mini-universe $\text{Univ}_{\vec{o}}$. Then we pack them into an oblivious-key value storage (OKVS) data structure [GPR⁺21], where the origin is treated as the OKVS key, and the corresponding **ibFSS** key is treated as its value.

To check if a point \vec{y} in Bob’s set is contained in any of Alice’s input balls, we only need to check all the mini-universes that contains \vec{y} , namely $\vec{y} \in \text{Univ}_{\vec{o}}$ for origins \vec{o} (there are 2^d of them). Bob can probe the OKVS data structure to obtain the **ibFSS** key for each relevant origin \vec{o} , and evaluate **ibFSS** on \vec{y} in the corresponding mini-universe $\text{Univ}_{\vec{o}}$. Thanks to our new OR observation, Bob can simply send **ibFSS** evaluations for all the structures within any relevant origin \vec{o} back to Alice, without having to perform a **sum** transformation.

Illustrative examples. We give a few examples in Figure 3 in 2-dimensional space to illustrate our new spacial hashing techniques. A structured set \vec{S} can be assigned to an origin \vec{o} if \vec{S} is completely contained in the mini-universe with origin \vec{o} and diameter 2δ , namely $\vec{S} \subseteq \text{Univ}_{\vec{o}}$. For instance, Figure 3a shows a mini-universe with origin \vec{z}_1 , and $\vec{S}_A^{(1)} \subseteq \text{Univ}_{\vec{z}_1}$. Therefore, the structured set $\vec{S}_A^{(3)}$ in Figure 3b can be assigned to any of the four origins $\vec{z}_2, \vec{z}_3, \vec{z}_6, \vec{z}_7$. In Figure 3c, we consider 3 structured sets in Alice’s set. $\vec{S}_A^{(1)}$ can be assigned to \vec{z}_1 ; $\vec{S}_A^{(2)}$ can be assigned to origins \vec{z}_2 or \vec{z}_6 ; $\vec{S}_A^{(3)}$ can be assigned to any of the four origins $\vec{z}_2, \vec{z}_3, \vec{z}_6, \vec{z}_7$. In this example, we assign $\vec{S}_A^{(1)}, \vec{S}_A^{(2)}, \vec{S}_A^{(3)}$ to $\vec{z}_1, \vec{z}_6, \vec{z}_3$, respectively (highlighted in the figure). Alice generates an **ibFSS** key share for each structured set and packs them into an OKVS, as shown in Figure 3e. When Bob evaluates at point \vec{y} (orange point in the figure), he needs to consider all the mini-universes that contain \vec{y} , namely $\text{Univ}_{\vec{z}_1}, \text{Univ}_{\vec{z}_2}, \text{Univ}_{\vec{z}_5}, \text{Univ}_{\vec{z}_6}$. In our new spatial hashing framework, Alice’s structured sets can overlap as long as every $\vec{S}_A^{(i)}$ can be assigned to a distinct origin $\vec{o}^{(i)}$ such that $\vec{S}_A^{(i)} \subseteq \text{Univ}_{\vec{o}^{(i)}}$. In another example in Figure 3d, we cannot find such a mapping.

What if there is no one-to-one mapping? When it is impossible to find a balls-to-origins mapping in spatial hashing (e.g., Figure 3d), one potential solution is to increase the number of structures that can be mapped to each universe, and then perform multi-ball-multi-point **sa-PSI** for each universe. Another potential solution is to have smaller grid cells and a larger universe size, so that Alice can find a unique mapping. Both solutions would incur higher communication and computation costs. The protocol can be adapted depending on the specific application.

Extensions and improvements over prior work. First, we improve both computation and

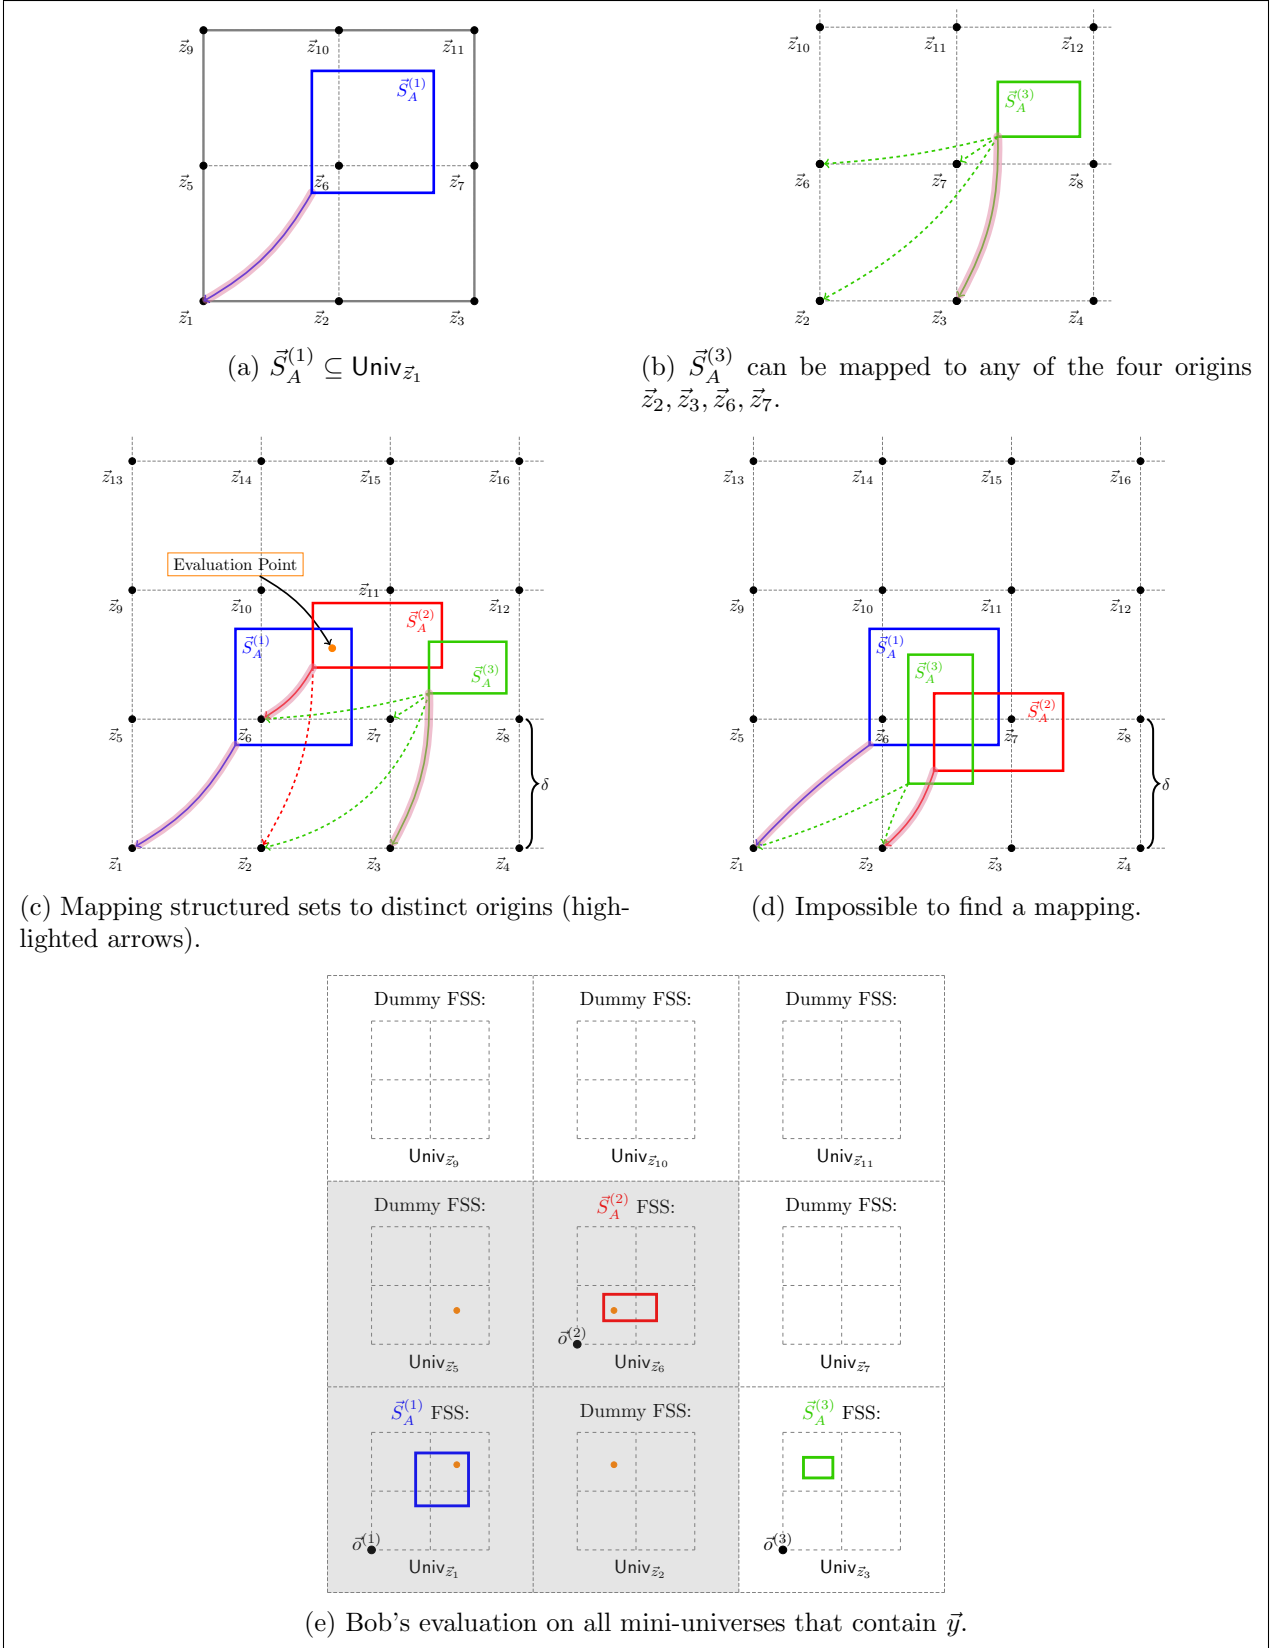


Figure 3: Our new spatial hashing techniques in 2-dimensional space.

communication costs of spatial hashing. In prior works [GRS22, GRS23], Bob needs to perform an expensive sum transformation on the bFSS outputs. This restricts their bFSS to *strong* bFSS that outputs a single bit, for which both computation and communication costs grow exponentially in the dimension d . In contrast, we get rid of the sum transformation, hence allowing for *weak* bFSS that outputs a bit string, with both computation and communication costs linear in d .

Second, when constructing and evaluating our ibFSS in a mini-universe $\text{Univ}_{\vec{o}}$, we *shift* all the points by the origin \vec{o} . Specifically, for $\vec{x} = (x_1, \dots, x_d)$ and $\vec{o} = (o_1, \dots, o_d)$, we consider $\vec{x}.\text{Shift}(\vec{o}) := (x_1 - o_1, \dots, x_d - o_d)$ in the mini-universe with diameter 2δ . When Bob computes the hash value, we include the origin \vec{o} as an index to avoid collisions. This shifting approach eliminates the need for point functions in prior work [GRS23]. On a minor note, we also improve on the domain of the ibFSS (size of *mini-universe*) from 3δ [GRS23] to 2δ .

Finally, we can further relax the restrictions on Alice’s input balls. Instead of requiring all the balls to be disjoint and have the same diameter, the only requirement we need for our protocol to work is that *every structured set $\vec{S}_A^{(i)}$ can be assigned to a distinct origin $\vec{o}^{(i)}$ such that $\vec{S}_A^{(i)} \subseteq \text{Univ}_{\vec{o}^{(i)}}$* . In particular, these balls can have *different diameters* and *overlap* with each other. In fact, they don’t even have to be ℓ_∞ -balls with the same diameter in each dimension; our construction works for any structured set with an ibFSS scheme. We state the theorem below and defer its proof to [Appendix B.1](#). the full version of the paper.

Theorem 6. *Given a two-party $(\log(2\delta), d, t)$ -ibFSS scheme with pseudorandom keys for a family of sets $\mathcal{S} \subseteq \underbrace{2^{\mathcal{U}} \times \dots \times 2^{\mathcal{U}}}_d$ where $\mathcal{U} = 2\delta$ and every set in \mathcal{S} is disjoint union of at most w prefix sets, an oblivious key-value store scheme $(\text{Encode}, \text{Decode})$, and a Hamming-correlation robust hash $H : \{0, 1\}^* \rightarrow \{0, 1\}^{\lambda + \log |N_A| + \log w + \log |N_B| + d(1 + \log \log(2\delta))}$, the protocol in [Figure 4](#) realizes $\mathcal{F}_{\text{sa-PSI}}$ ([Figure 1](#)) in the \mathcal{F}_{OT} -hybrid model secure in the presence of semi-honest adversaries.*

6 ibFSS for d -dimensional ℓ_∞ -balls

In this section, we present a construction an ibFSS for d -dimensional balls with ℓ_∞ -norm. Looking ahead, we will instantiate our structure-aware PSI protocols with it and analyze our computation and communication costs in [Section 7](#). What we need for d -dimensional ℓ_∞ -balls is a weak incremental Boolean Distributed Comparison Function (u, d, t) -ibDCF. It is an instance of (u, d, t) -ibFSS for the specific set family of d -dimensional ℓ_∞ -balls. We formalize the definition in [Section 6.1](#) and give a construction in [Section 6.2](#).

6.1 ibDCF Definition

We consider (u, d, t) -ibFSS for the set family of d -dimensional intervals $\vec{S}_{\text{INT}} = (S_1^{(\alpha_1, \beta_1)} \times \dots \times S_d^{(\alpha_d, \beta_d)})$ represented by interval tuples $((\alpha_1, \beta_1), \dots, (\alpha_d, \beta_d))$, where $\alpha_i \in \{0, 1\}^u$ is a partition point and $\beta_i \in \{0, 1\}$ is an indicator of left or right interval. Specifically, for every $i \in [d]$, the set is either a left interval $S_i^{(\alpha_i, 0)} = \{x \mid x < \alpha_i, x \in \{0, 1\}^u\}$ or a right interval $S_i^{(\alpha_i, 1)} = \{x \mid x > \alpha_i, x \in \{0, 1\}^u\}$. We define the syntax below and the security requirement follows from (u, d, t) -ibFSS ([Definition 4](#)).

Definition 7. *[(u, d, t)-ibDCF: Syntax] A (two-party) d -dimensional weak incremental Boolean distributed comparison function scheme (u, d, t) -ibDCF for the family of d -dimensional interval sets over the universe $(\{0, 1\}^u)^d$ consists of a pair of algorithms $(\text{Share}, \text{Eval})$ with the following syntax:*

Parameters:

- computational security parameter κ and statistical security parameter λ
- set family \mathcal{S} with corresponding $((\log(2\delta), d, t)$ -ibFSS scheme (Share, Eval) (Definition 3)
- oblivious key-value store scheme (Encode, Decode) (Definition 1, Definition 2)
- oblivious transfer ideal functionality \mathcal{F}_{OT}
- ℓ OT instances such that $\Pr[\text{Binomial}(1 - 2^{-t}, \ell) < \kappa] < 2^{-\lambda - |N_B| - d(1 + \log \log(2\delta))}$
- hamming-correlation robust hash H with output length $\lambda + \log |N_A| + \log w + \log |N_B| + d(1 + \log \log(2\delta))$

Inputs:

- Alice has N_A structured sets $S_A = \bigcup_{i \in [N_A]} \vec{S}_A^{(i)}$, where $\vec{S}_A^{(i)} \in \mathcal{S}$.
- Bob has an unstructured set $S_B \subseteq \mathcal{U}^d$ of size N_B .

Notation:

- For $\vec{o} = (o_1, \dots, o_d) \in \mathcal{U}^d$, $\text{Univ}_{\vec{o}} := [o_1, o_1 + 2\delta) \times \dots \times [o_d, o_d + 2\delta)$.
- For $\vec{x} = (x_1, \dots, x_d), \vec{o} = (o_1, \dots, o_d) \in \mathcal{U}^d$, $\vec{x}.\text{Shift}(\vec{o}) := (x_1 - o_1, \dots, x_d - o_d)$.
- For $S \subseteq \mathcal{U}^d$ and $\vec{o} = (o_1, \dots, o_d) \in \mathcal{U}^d$, $S.\text{Shift}(\vec{o}) := \{\vec{x}.\text{Shift}(\vec{o}) \mid \vec{x} \in S\}$.

Protocol:

1. For each $i \in [N_A]$, Alice identifies a distinct origin $\vec{o}^{(i)}$ such that $\vec{S}_A^{(i)} \subseteq \text{Univ}_{\vec{o}^{(i)}}$.
2. For each $\eta \in [\ell]$, Alice does the following:
 - (a) For each $i \in [N_A]$ compute $(k_0^{(i, \eta)}, k_1^{(i, \eta)}) \leftarrow \text{Share}(1^\kappa, \vec{S}_A^{(i)}.\text{Shift}(\vec{o}^{(i)}))$.
 - (b) For $b \in \{0, 1\}$ encode $\mathcal{KV}_b^{(\eta)} \leftarrow \text{Encode}(\{(\vec{o}^{(1)}, k_b^{(1, \eta)}), \dots, (\vec{o}^{(N_A)}, k_b^{(N_A, \eta)})\})$.
3. Bob chooses a random string $s \leftarrow \{0, 1\}^\ell$.
4. The parties invoke ℓ (parallel) instances of \mathcal{F}_{OT} . In the η -th instance:
 - Alice is the sender with input $(\mathcal{KV}_0^{(\eta)}, \mathcal{KV}_1^{(\eta)})$;
 - Bob is the receiver with choice bit $s[\eta]$. He obtains output $\mathcal{KV}_*^{(\eta)} = \mathcal{KV}_{s[\eta]}^{(\eta)}$.
5. Bob initializes a set $Y := \emptyset$ and does the following:
 - (a) For each element $\vec{y} \in S_B$, and for each origin \vec{z} where $\vec{y} \in \text{Univ}_{\vec{z}}$:
 - i. Write $\vec{y}.\text{Shift}(\vec{z}) = (y_1, \dots, y_d)$, where $y_j \in [0, 2\delta)$ for all $j \in [d]$.
 - ii. For each $\eta \in [\ell]$, compute $k_*^{(\eta)} := \text{Decode}(\mathcal{KV}_*^{(\eta)}, \vec{z})$.
 - iii. Compute
$$Y' = \left\{ H \left(\vec{y} \parallel \vec{z}; \text{Eval}(k_*^{(1)}, \vec{y}^1) \parallel \dots \parallel \text{Eval}(k_*^{(\ell)}, \vec{y}^\ell) \right) \mid \begin{array}{l} \ell_1 \in [1, \log(2\delta)], \dots, \ell_d \in [1, \log(2\delta)] \\ \vec{y}^j := ((y_1)_{[1: \ell_1]}, \dots, (y_d)_{[1: \ell_d]}) \end{array} \right\}$$
 - iv. Update $Y := Y \cup Y'$.
 - (b) Pad Y with dummy random strings such that $|Y| = N_B \cdot (2 \cdot \log(2\delta))^d$ and send it to Alice.
6. Alice initializes set $I := \emptyset$. For each set $i \in [N_A]$, Alice does the following:
 - (a) Write $\vec{S}_A^{(i)}.\text{Shift}(\vec{o}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x_j}^{\rightarrow}$ as disjoint union of w' prefix sets (w is the upper bound for w').
 - (b) For each $j \in [w']$, run $\text{INTSEARCH} \left(\{k_0^{(i, \eta)}\}_{\eta \in [\ell]}, \log(2\delta), \vec{o}^{(i)}, \vec{x}^j, Y \right)$, with I being a global variable.
7. **[output]** Alice outputs the intersection I .

Figure 4: Protocol realizing $\mathcal{F}_{\text{sa-PSI}}$ (Figure 1) using spatial hashing techniques.

- $(k_0, k_1) \leftarrow \text{Share}(1^\kappa, ((\alpha_1, \beta_1), \dots, (\alpha_d, \beta_d)))$: The randomized share function takes as input the security parameter κ and the description of a d -dimensional interval, namely a set of partition points and bits (one for each dimension), where $\alpha_i \in \{0, 1\}^u$ and $\beta_i \in \{0, 1\}$, and outputs two key shares.
- $y_{idx} \leftarrow \text{Eval}(idx, k_{idx}, \vec{x})$: The deterministic evaluation function takes as input a party index $idx \in \{0, 1\}$, the corresponding key share k_{idx} , and input $\vec{x} = (x_1, \dots, x_d)$ where each $x_i \in \bigcup_{\ell=1}^u \{0, 1\}^\ell$. It outputs a string $y_{idx} \in \{0, 1\}^t$.

Handling two-sided intervals. While the above definition only considers *single-sided* intervals in each dimension, we can use it to handle *two-sided* intervals in a generic manner. Specifically, a two-dimensional $(u, 2, t)$ -ibDCF for the set family $\vec{S}_{2\text{-INT}} = (S_1^{(L,0)} \times S_2^{(R,1)}) \in 2^u \times 2^u$ can be used as a single-dimensional ibDCF for a two-sided interval (L, R) if evaluating on $\vec{x} = (x, x)$ with the *same* input in both dimensions.

Similarly, a $2d$ -dimensional $(u, 2d, t)$ -ibDCF for the set family $\vec{S}_{2d\text{-INT}} = (S_1^{(L_1,0)} \times S_2^{(R_1,1)} \times \dots \times S_{2d-1}^{(L_{d-1},0)} \times S_{2d}^{(R_d,1)}) \in \underbrace{2^u \times \dots \times 2^u}_{(2d)}$ can be used as an ibDCF for a d -dimensional ℓ_∞ -ball $(L_1, R_1) \times \dots \times (L_d, R_d)$ if evaluating on $\vec{x} = (x_1, x_1, \dots, x_d, x_d)$ with the *same* input along each of the d dimensions of the ball.

6.2 ibDCF Constructions

We first give a construction for a single-dimensional $(u, 1, 1)$ -ibDCF for left-sided intervals in [Figure 6](#). Our construction builds upon the optimized Distributed Point Function (DPF) construction from [\[BGI18\]](#) with the addition of an extra variable. For a left-sided interval $\{x \mid x < \alpha, x \in \{0, 1\}^u\}$, we consider the DPF construction for the point function $f(x) = 1$ if $x = \alpha$ and $f(x) = 0$ otherwise. At a high level, in the GGM-tree based DPF construction [\[BGI18\]](#), each node in the binary tree consists of a PRG seed and an additional output bit. They focus on the *critical path* along α and keep the invariant that the output bit is a secret share of 1 *on* the critical path and secret share of 0 *off* the critical path.

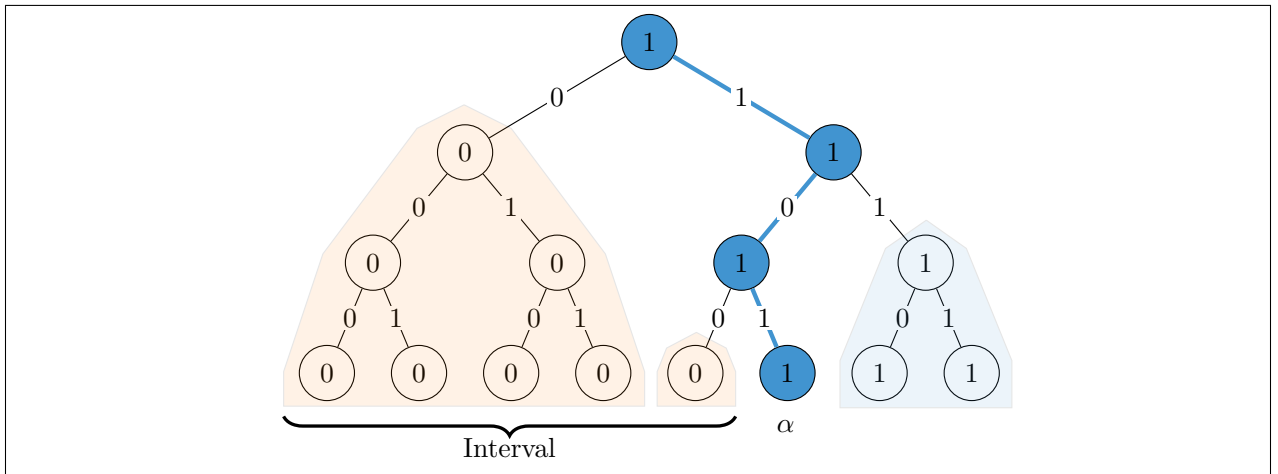


Figure 5: Incremental DCF for a single-dimensional left-sided interval $\{x \mid x < \alpha, x \in \{0, 1\}^u\}$.

$(u, 1, 1)$ -ibDCF (Share, Eval):

Parameters: Let $G : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{2(\kappa+2)}$ be a pseudorandom generator.

Share $(1^\kappa, (\alpha, \beta = 0))$:

Let $\alpha = \alpha^{(1)}\alpha^{(2)} \dots \alpha^{(u)} \in \{0, 1\}^u$ be the bit decomposition.
 Sample $s_0^{(0)} \leftarrow \{0, 1\}^\kappa$ and $s_1^{(0)} \leftarrow \{0, 1\}^\kappa$ uniformly at random.
 Assign $t_0^{(0)} = 0$ and $t_1^{(0)} = 1$.
for $i = 1$ **to** u **do**
 $s_0^L \| t_0^L \| y_0^L \parallel s_0^R \| t_0^R \| y_0^R \leftarrow G(s_0^{(i-1)})$ and $s_1^L \| t_1^L \| y_1^L \parallel s_1^R \| t_1^R \| y_1^R \leftarrow G(s_1^{(i-1)})$
 if $\alpha^{(i)} = 0$ **then**
 $s_{CW} = s_0^R \oplus s_1^R$
 $t_{CW}^L = t_0^L \oplus t_1^L \oplus 1$ and $t_{CW}^R = t_0^R \oplus t_1^R$
 $y_{CW}^L = y_0^L \oplus y_1^L$ and $y_{CW}^R = y_0^R \oplus y_1^R$
 $CW^{(i)} = s_{CW} \| t_{CW}^L \| t_{CW}^R \| y_{CW}^L \| y_{CW}^R$
 $s_b^{(i)} = s_b^L \oplus t_b^{(i-1)} \cdot s_{CW}$ for $b = 0, 1$
 $t_b^{(i)} = t_b^L \oplus t_b^{(i-1)} \cdot t_{CW}^L$ for $b = 0, 1$
 else
 $s_{CW} = s_0^L \oplus s_1^L$
 $t_{CW}^L = t_0^L \oplus t_1^L$ and $t_{CW}^R = t_0^R \oplus t_1^R \oplus 1$
 $y_{CW}^L = y_0^L \oplus y_1^L \oplus 1$ and $y_{CW}^R = y_0^R \oplus y_1^R$
 $CW^{(i)} = s_{CW} \| t_{CW}^L \| t_{CW}^R \| y_{CW}^L \| y_{CW}^R$
 $s_b^{(i)} = s_b^R \oplus t_b^{(i-1)} \cdot s_{CW}$ for $b = 0, 1$
 $t_b^{(i)} = t_b^R \oplus t_b^{(i-1)} \cdot t_{CW}^R$ for $b = 0, 1$
 Let $k_b = s_b^{(0)} \| CW^{(1)} \| CW^{(2)} \| \dots \| CW^{(u)}$
return (k_0, k_1)

Eval $(1^\kappa, b, k_b, x)$:

Parse $k_b = s^{(0)} \| CW^{(1)} \| \dots \| CW^{(u)}$ and $x = x^{(1)}, x^{(2)}, \dots, x^{(\ell)}$
 Assign $t^{(0)} = b$. **Assign** $y^{(0)} = b$.
for $i = 1$ **to** ℓ **do**
 Parse $CW^{(i)} = s_{CW} \| t_{CW}^L \| t_{CW}^R \| y_{CW}^L \| y_{CW}^R$
 $\tau^i = G(s^{(i-1)}) \oplus (t^{(i-1)} \cdot [s_{CW} \| t_{CW}^L \| y_{CW}^L \| s_{CW} \| t_{CW}^R \| y_{CW}^R])$
 Parse $\tau^i = s^L \| t^L \| y^L \parallel s^R \| t^R \| y^R \in \{0, 1\}^{2(\kappa+2)}$
 Update $y^L = y^L \oplus y^{(i-1)}$ and $y^R = y^R \oplus y^{(i-1)}$
 if $x_i = 0$ **then** $s^{(i)} = s^L, t^{(i)} = t^L, y^{(i)} = y^L$
 else $s^{(i)} = s^R, t^{(i)} = t^R, y^{(i)} = y^R$
return $y^{(\ell)}$

Figure 6: Construction of a single-dimensional Incremental Boolean Distributed Comparison Function $(u, 1, 1)$ -ibDCF for left-sided intervals.

In our construction, we use their output bit (t) as an *indicator bit* of whether we are on or off the critical path, and add another bit (y) as the output bit. We leverage the indicator bit and keep the following invariants: (1) if we are on the critical path, then the output bit is a secret share of 1; (2) if we deviate to the right of the critical path (namely $\alpha^{(i)} = 0$), then the output bit remains a secret share of 1; (3) if we deviate to the left of the critical path (namely $\alpha^{(i)} = 1$), then the output bit becomes a secret share of 0; and (4) once we deviate from the critical path, the output remains a secret share of the same bit for the rest of the subtree (1 if deviating to the right and

0 if deviating to the left). The additional output bit y is highlighted in blue in the construction (Figure 6).² We illustrate these invariants in Figure 5. Specifically, for all the nodes highlighted in light orange, the parties obtain secret shares of 0; for all the nodes highlighted in blue and light blue, the parties obtain secret shares of 1.

Note that we present a construction for left-sided intervals. It is straightforward to construct an analogous scheme for right-sided intervals, and we omit the details here. We state the theorem below and defer its security proof to Appendix B.2. the full version of the paper.

Theorem 8. *Given a pseudorandom generator $G : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{2(\kappa+2)}$, the construction (Share, Eval) in Figure 6 is a secure $(u, 1, 1)$ -ibDCF scheme (Definition 7) for the family of single-dimensional one-sided intervals, with key size $|k_0| = |k_1| = u \cdot (\kappa + 4) + \kappa$ bits and output size $|y_0| = |y_1| = 1$.*

Next we present our construction for a d -dimensional (u, d, d) -ibDCF scheme in Figure 7, where we adapt the concat technique from prior work [GRS22]. We state the theorem below and defer its proof to Appendix B.3. the full version of the paper.

Theorem 9. *Given a single-dimensional $(u, 1, 1)$ -ibDCF scheme, the construction (Share*, Eval*) in Figure 7 is a secure (u, d, d) -ibDCF scheme (Definition 7) for the family of d -dimensional intervals, with key size $|k_0| = |k_1| = d \cdot (u \cdot (\kappa + 4) + \kappa)$ bits and output size $|y_0| = |y_1| = d$.*

<p>(u, d, d)-ibDCF (Share*, Eval*):</p> <p>Parameters: Let (Share, Eval) be a single-dimensional $(u, 1, 1)$-ibDCF scheme.</p> <p>Share*$(1^\kappa, ((\alpha_1, \beta_1), \dots, (\alpha_d, \beta_d)))$:</p> <pre style="margin-left: 20px;"> for $i = 1$ to d do $(k_0^i, k_1^i) \leftarrow \text{Share}(1^\kappa, (\alpha_i, \beta_i))$ $k_0 = (k_0^1, \dots, k_0^d)$ and $k_1 = (k_1^1, \dots, k_1^d)$ return (k_0, k_1) </pre> <p>Eval*$(b, k_b = (k_b^1, \dots, k_b^d), \vec{x} = (x_1, \dots, x_d))$:</p> <pre style="margin-left: 20px;"> return $(\text{Eval}(b, k_b^1, x_1) \parallel \dots \parallel \text{Eval}(b, k_b^d, x_d))$ </pre>
--

Figure 7: Construction of Incremental Boolean Distributed Comparison Function (u, d, d) -ibDCF for d -dimensional intervals with output length d .

7 Instantiation and Evaluation

We instantiate our sa-PSI protocols with d -dimensional balls with ℓ_∞ -norm using the ibDCF presented in Section 6, and analyze the computation and communication improvements over prior work [GRS22, GRS23]. Since our protocol is semi-honest secure, we mainly compare with the semi-honest work [GRS22]. Note that [GRS23] follows the same construction framework as [GRS22] for

²Note that similar ideas were presented in [APR⁺22] in to construct DCF from iDPF, where there are also evaluation values at each node (v_i) in their construction (Algorithm 7). However, outputting these values doesn't immediately work for us because we need the critical path to output a secret share of 1 (instead of 0 in their construction). Another critical difference is that v_i in [APR⁺22] remains unchanged when $x_i = 1$, whereas our construction generates fresh randomness for each node, resulting in fresh secret shares at each node and achieving slightly stronger security guarantees that are essential in our sa-PSI protocols.

the most part while focusing on achieving malicious security, except that [GRS23] also introduces a new spatial hashing technique that improves upon the semi-honest construction. Hence we compare our work with both spatial hashing constructions.

7.1 Single Ball with Single Point

We start our comparison with the setting where Alice holds a single ball with ℓ_∞ -norm as the distance metric in the d -dimensional universe $(\{0, 1\}^u)^d$, and Bob holds a single point. The computation cost of the protocol can be broken down into three parts: (1) Alice generates all the FSS keys, (2) Bob evaluates all the FSS on his inputs and computes hash values, and (3) Alice identifies the set intersection from the hash values received from Bob. Our unit of computation cost is one PRG or hash operation.

The communication cost can be broken down into two parts: (1) oblivious transfer where Alice is the OT sender with FSS keys and Bob is the OT receiver with random choice bits, and (2) the hash values sent from Bob to Alice. The computation and communication costs are summarized in Table 1. Note that all the computation only involves efficient symmetric-key operations such as PRGs (which can be instantiated with AES) and hash functions (which can be instantiated with e.g., SHA256). There is also computation cost from OT on both parties, but we can leverage the efficient OT extension [IKNP03], and the computation cost is dominated by the FSS key generation and evaluation, so we omit this cost from the table.

Comp. & Comm. Costs		[GRS22]	Ours	
Comp.	Alice	FSS	$\mathcal{O}(u \cdot d \cdot \ell_{\text{OT}})$	$\mathcal{O}(u \cdot d \cdot \ell_{\text{OT}})$
		Intersection	$\mathcal{O}(\min(S_A \cdot u, 2^u) \cdot d \cdot \ell_{\text{OT}})$	$\mathcal{O}(u^d + u \cdot d \cdot \ell_{\text{OT}})$
	Bob's Eval	$\mathcal{O}(u \cdot d \cdot \ell_{\text{OT}})$	$\mathcal{O}(u^d + u \cdot d \cdot \ell_{\text{OT}})$	
Comm. (bits)	OT	$\mathcal{O}(\kappa \cdot u \cdot d \cdot \ell_{\text{OT}})$	$\mathcal{O}(\kappa \cdot u \cdot d \cdot \ell_{\text{OT}})$	
	Bob's Hashes	$\mathcal{O}(h_{\text{out}})$	$\mathcal{O}(u^d \cdot h_{\text{out}})$	

Table 1: Summary of computation and communication costs for the setting where Alice holds a single ℓ_∞ ball in the d -dimensional universe $(\{0, 1\}^u)^d$ and Bob holds a single point in the universe. $h_{\text{out}} = \mathcal{O}(\lambda)$ is the output length of the final hash function. $\ell_{\text{OT}} = \mathcal{O}(\kappa)$ is the number of OTs. $|S_A|$ is upper bounded by $2^{u \cdot d}$.

Cells in green indicate that our protocol achieves better complexity, and those in grey indicate that our protocol requires higher overhead. Our main improvement over prior work is Alice's computation cost to identify the set intersection. In particular, prior work requires Alice to compute hash values (i.e., perform FSS evaluations) on every element in her set, which could potentially be exponentially large, namely $|S_A|$ is upper bounded by $2^{u \cdot d}$. In our work, by leveraging our new ibFSS construction for d -dimensional ℓ_∞ balls, Alice only needs to compute hash values for up to u^d prefixes. In case there is a match, she can perform an efficient search for Bob's element in time $u \cdot d \cdot \ell_{\text{OT}}$. To enable Alice to perform such a search, we require Bob to compute hash values for all combinations of his prefixes, hence introducing a computation and communication overhead of u^d on Bob's side. Text in red highlights our difference from prior work that attributes to our new ibFSS technique.

7.2 Multiple (Overlapping) Balls with Multiple Points

Next, we consider the setting where Alice holds N_A balls with ℓ_∞ -norm as the distance metric in the d -dimensional universe $(\{0, 1\}^u)^d$, and Bob holds N_B points in the universe. The computation and communication costs can be broken down in a similar way as above, which is summarized in Table 2.

Comp. & Comm. Costs		[GRS22, GRS23]	Ours	
Comp.	Alice	FSS	$\mathcal{O}(N_A \cdot u^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_A \cdot u \cdot d \cdot \ell_{\text{OT}})$
		Intersection	$\mathcal{O}(S_A \cdot N_A \cdot u^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_A \cdot u^d + N_B \cdot u \cdot d \cdot \ell_{\text{OT}})$
	Bob's Eval	$\mathcal{O}(N_A \cdot N_B \cdot u^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_A \cdot N_B \cdot (u^d + u \cdot d \cdot \ell_{\text{OT}}))$	
Comm. (bits)	OT	$\mathcal{O}(\kappa \cdot N_A \cdot u^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(\kappa \cdot N_A \cdot u \cdot d \cdot \ell_{\text{OT}})$	
	Bob's Hashes	$\mathcal{O}(N_B \cdot h_{\text{out}})$	$\mathcal{O}(N_B \cdot N_A \cdot u^d \cdot h_{\text{out}})$	

Table 2: Summary of computation and communication costs for the setting where Alice holds N_A number of ℓ_∞ balls in the d -dimensional universe $(\{0, 1\}^u)^d$ and Bob holds N_B points in the universe. $h_{\text{out}} = \mathcal{O}(\lambda)$ is the output length of the final hash function. $\ell_{\text{OT}} = \mathcal{O}(\kappa)$ is the number of OTs. $|S_A|$ is upper bounded by $2^{u \cdot d}$. Prior work [GRS22, GRS23] only allows Alice to hold *disjoint* balls while our protocol also supports *overlapping* balls.

Besides the improvement from ibFSS (which is highlighted in *red*), we have several new contributions in this setting. The way that prior work deals with multiple balls is via a *sum* transformation to combine multiple FSS into a single one that achieves the *OR* functionality. Since their *sum* transformation only works for FSS that outputs a single bit, they have to use *strong FSS* for d -dimensional balls [BGI16], incurring an overhead of u^d in both computation and communication for every ball and every FSS evaluation. By utilizing our new observation for the *OR* functionality in PSI, we no longer need the *sum* transformation, allowing for the usage of *weak FSS* that outputs multiple bits for each ball. This reduces the exponential overhead (u^d) down to linear in the dimension ($u \cdot d$). Moreover, the *sum* transformation in prior work only works for *disjoint* balls. We eliminate this restriction in our protocol and allow for overlapping balls as well. We highlight our difference from prior work in *blue* in the table.

7.3 Spatial Hashing

In this section, we analyze the multi-ball-multi-point protocols using the spatial hashing technique [GRS22]. In particular, Alice holds N_A number ℓ_∞ -balls of diameter δ in the d -dimensional universe $(\{0, 1\}^u)^d$, and Bob holds N_B points. Note that prior work [GRS22, GRS23] additionally considers more restricted settings where the balls are not only disjoint, but are guaranteed to be far apart ($> 4\delta$ [GRS22] or $> 8\delta$ [GRS23] in distance), which we do not compare with. In fact, our protocol works for a more relaxed setting where Alice's balls do *not* have the same diameter, and they can even overlap. The computation and communication costs can be broken down in a similar way as above, which is summarized in Table 3. All three constructions require an OKVS [GPR⁺21, RR22, BPSY23], typically with linear overhead (and small constants) in the number of keys, so we ignore this overhead.

In addition to the improvements from ibFSS (highlighted in *red*) and those from the new OR observation (highlighted in *blue*), we introduce new techniques to improve upon spatial hashing, which may be of independent interest. The spatial hashing technique is introduced in [GRS22] and

Comp. & Comm. Costs		[GRS22]	[GRS23]	Ours	
Comp.	Alice	FSS	$\mathcal{O}(N_A \cdot (4 \log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_A \cdot u \cdot d \cdot (\log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_A \cdot \log \delta \cdot d \cdot \ell_{\text{OT}})$
		Intersection	$\mathcal{O}(S_A \cdot (2 \log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(S_A \cdot u \cdot d \cdot (2 \log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_A \cdot (\log \delta)^d + N_B \cdot \log \delta \cdot d \cdot \ell_{\text{OT}})$
	Bob's Eval	$\mathcal{O}(N_B \cdot (2 \log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_B \cdot u \cdot d \cdot (2 \log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(N_B \cdot 2^d \cdot ((\log \delta)^d + \log \delta \cdot d \cdot \ell_{\text{OT}}))$	
Comm. (bits)	OT	$\mathcal{O}(\kappa \cdot N_A \cdot (4 \log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(\kappa \cdot N_A \cdot u \cdot d \cdot (\log \delta)^d \cdot \ell_{\text{OT}})$	$\mathcal{O}(\kappa \cdot N_A \cdot \log \delta \cdot d \cdot \ell_{\text{OT}})$	
	Bob's Hashes	$\mathcal{O}(N_B \cdot h_{\text{out}})$	$\mathcal{O}(N_B \cdot h_{\text{out}})$	$\mathcal{O}(N_B \cdot (2 \log \delta)^d \cdot h_{\text{out}})$	

Table 3: Summary of computation and communication costs for spatial hashing. $h_{\text{out}} = \mathcal{O}(\lambda)$ is the output length of the final hash function. $\ell_{\text{OT}} = \mathcal{O}(\kappa)$ is the number of OTs. $|S_A|$ is upper bounded by $2^{u \cdot d}$. Prior work [GRS22, GRS23] only allows Alice to hold *disjoint* balls with *fixed diameter* δ while our protocol also supports *overlapping* balls with *different diameters*.

Comp. & Comm. Costs			$d = 2$		$d = 3$		$d = 5$	
			[GRS23]	Ours	[GRS23]	Ours	[GRS23]	Ours
Comp. (PRGs, SHA256)	Alice	FSS	86M	840K	516M	1.3M	13.8B	2.1M
		Intersection	352B	2.8B	135T	4.2B	1478Q	7.0B
	Bob's Eval	1.2T	11.3B	13.8T	34.6B	1468T	324B	
Comm.	OT	1.3GB	12.5MB	7.7GB	18.8MB	205GB	31.3MB	
	Bob's Hashes	4.77MB	477MB	4.77MB	4.77GB	4.77MB	477GB	

Table 4: Suppose Alice has $N_A = 300$ balls with fixed diameter $\delta = 32$ and Bob has a collection of $N_B = 10^6$ points in $d = \{2, 3, 5\}$ dimensional space over $\mathcal{U} = \{0, 1\}^u$, $u = 32$. Let computational security parameter $\kappa = 128$ and statistical security parameter $\lambda = 40$. We estimate the computation and communication cost of our new, spatial hashing technique and compare against the previous best construction [GRS23]. Our unit of computation cost is one PRG or hash operation. K, M, B, T, Q stand for thousand, million, billion, trillion, quadrillion respectively in computation units. For example, 1K means 1000 AES/SHA256 calls.

improved in [GRS23]. In [GRS22], they consider all the *active grids* that overlap with at least one ball. Each active grid may contain up to 2^d union of disjoint balls, so this problem is reduced to the multi-ball-multi-point setting. Therefore, they need to use the expensive *sum* transformation on strong FSS, inheriting the $(\log \delta)^d$ overhead and introducing a new 2^d overhead as the maximum number of balls per grid. Additionally, each ball may overlap with up to 2^d grids, hence the total number of active grids is upper bounded by $N_A \cdot 2^d$, introducing another 2^d overhead in FSS key generation. The follow-up work [GRS23] improves the spatial hashing technique by mapping every ball into a unique grid, hence reducing the number of active grids. However, for each point in Bob's set, he needs to evaluate FSS for all the adjacent grids and perform the *sum* transformation, which still requires strong FSS and the $(\log \delta)^d$ overhead. Moreover, they apply a distributed point function (DPF) to reduce the domain from the universe to the grid, introducing another overhead of $u \cdot d$.

In our work, we also map every ball into a unique grid, but Bob does not need to perform the *sum* transformation on adjacent cells. Instead, we can leverage our new OR observation and achieve a $\log \delta \cdot d$ overhead similar to the multi-ball-multi-point protocol. Additionally, we eliminate the usage of DPF by shifting both the balls and points to their corresponding mini-universe, saving another factor of $u \cdot d$. Finally, we relax our restriction on Alice's balls in that they can overlap and have different diameters. The only remaining restriction is that every ball can be mapped to a distinct grid that fully contains this ball. We highlight our difference from prior work in *brown* in the table.

To illustrate our improvement, we estimate the concrete computational and communication costs for specific settings in Table 4 comparing our protocol with the previous best spatial hashing

technique [GRS23]. Alice’s computation cost is improved by $125 - 211M\times$, and Bob’s computation cost is improved by $106 - 4,531\times$. In term of communication, we achieve lower communication cost in OT ($104 - 6,550\times$) while incurring higher communication overhead in Bob’s hashes ($100 - 100K\times$). The overhead can outweigh the OT savings in certain scenarios. For instance, when $d = 5$, the total communication cost of [GRS23] is 205GB while ours is 477GB. However, the total computation cost of [GRS23] is 1478Q while ours is 331B, showing a 4.5 million times improvement.

Concretely, we can estimate for the ride sharing application setting where Alice (passengers) has $N_A = 300$ balls with fixed diameter $\delta = 32$ wants to be matched with Bob (drivers) $N_B = 300$ in a two-dimensional space over $\mathcal{U} = \{0,1\}^u$ for $u = 32$. In our protocol (respectively, in prior work [GRS23]), the computation cost of Alice generating FSS is 840K (86M), Alice computing intersection is 848K (352B), Bob evaluating FSS is 3.3M (344M); the communication cost of OT is 12.8MB (1.28GB), Bob’s hashes is 146KB (1.46KB). Our protocol outperforms prior work by orders of magnitudes in both computation and communication. .

8 Extending Functionality

In this section, we present a construction for sa-PSI where Bob learns the output (Section 8.1). With a slight modification, we construct protocols that allow Alice (or Bob) to learn structure-aware PSI cardinality or PSI with associated sum (Section 8.2). One difference from prior constructions is that we only allow Alice’s structured set to be a union of *disjoint* sets.

8.1 sa-PSI with Bob Learning Output

Construction Overview. Following the syntax from spatial hashing, both parties prepare their inputs. In **Step 1**, Alice assigns each of her structures to a distinct origin and prepares a set X of all prefix values along with the origin as an identifier. In **Step 2**, Bob computes a set Y for each input \vec{y} , which consists of all prefixes associated with every origin that contains the input. If $\vec{y} \in S_A \cap S_B$, then there is exactly one matching value (corresponding to Alice’s prefix and origin), namely $|X \cap Y| = 1$, which must be accounted for while computing PSI cardinality between X and Y . If $\vec{y} \notin S_A \cap S_B$, then there is no matching value, hence $|X \cap Y| = 0$. Alice and Bob can use a $\mathcal{F}_{\text{PSI-CA}}$ ideal functionality, for each set Y , for Bob to learn $|X \cap Y|$. If Bob learns that the cardinality is 1 then he includes the input \vec{y} in intersections, if cardinality is 0 then he does not include it in the intersection. We define the ideal functionality for sa-PSI where Bob learns the output in Figure 8. Note that a difference from Figure 1 is that Alice’s structured sets are *disjoint*. We present our protocol in Figure 9, state the theorem below and defer its security proof to Appendix B.4 the full version of the paper.

Theorem 10. *The protocol in Figure 9 securely realizes $\mathcal{F}_{\text{sa-PSI}}$ (Figure 8) in the $\mathcal{F}_{\text{PSI-CA}}$ -hybrid model in the presence of semi-honest adversaries.*

Cost Analysis. For the set family \mathcal{S} as disjoint union of N_A balls with diameter $\leq \delta$ in ℓ_∞ norm, Alice and Bob run N_B instances of $\mathcal{F}_{\text{PSI-CA}}$, between sets of size $n_1 = N_A \cdot w = N_A \cdot (\log(2\delta))^d$ and $n_2 = (2 \log(2\delta))^d$. If we instantiate $\mathcal{F}_{\text{PSI-CA}}$ with a semi-honest construction with linear computation and communication in set sizes [PSTY19, IKN⁺20, KK20, GMR⁺21, RS21, CGS22, HMS21], then we achieve sa-PSI for Bob to learn output with communication and computation costs $O(N_B \cdot (n_1 + n_2))$, which is $\ll |S_A|$.

Parameters: A family of sets $\mathcal{S} \subseteq \underbrace{2^{\mathcal{U}} \times \dots \times 2^{\mathcal{U}}}_d$ where $\mathcal{U} = \{0, 1\}^u$. Number of Alice's structured sets N_A and size of Bob's set N_B .

Functionality:

1. Receive input $S_A = \dot{\bigcup}_{i \in N_A} \vec{S}_A^{(i)}$, where $\vec{S}_A^{(i)} \in \mathcal{S}$ and S_A is a *disjoint* union of structured sets (or a concise representation of S_A) from Alice.
2. Receive input $S_B \subseteq \mathcal{U}^d$ of size N_B from Bob.
3. **[output]** Send $S_A \cap S_B$ to Bob.

Figure 8: Ideal functionality $\mathcal{F}_{\text{sa-PSI}}$ for structure-aware PSI, where Bob learns the output.

Parameters:

- private set intersection cardinality ideal functionality $\mathcal{F}_{\text{PSI-CA}}$

Inputs:

- Alice has N_A structured sets $S_A = \dot{\bigcup}_{i \in N_A} \vec{S}_A^{(i)}$, where $\vec{S}_A^{(i)} \in \mathcal{S}$ and the structured sets in S_A are *disjoint*.
- Bob has an unstructured set $S_B \subseteq \mathcal{U}^d$ of size N_B .

Protocol:

1. Alice initializes $X := \emptyset$ and does the following:
 - (a) Identify a distinct origin $\vec{o}^{(i)}$ for each $i \in [N_A]$ such that $\vec{S}_A^{(i)} \subseteq \text{Univ}_{\vec{o}^{(i)}}$.
 - (b) For each $i \in [N_A]$:
 - i. Write $\vec{S}_A^{(i)}. \text{Shift}(\vec{o}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x_j}^{\vec{o}^{(i)}}$ as disjoint union of w' prefix sets (w' is the upper bound for w').
 - ii. Update set $X := X \cup \{(x_j^{\vec{o}^{(i)}} \parallel \vec{o}^{(i)}) \mid j \in [w']\}$.
 - (c) Pad X with dummy random strings such that $|X| = N_A \cdot w'$.
2. Bob initializes $I := \emptyset$ and does the following:
 - (a) For each element $\vec{y} \in S_B$, and for each origin \vec{z} where $\vec{y} \in \text{Univ}_{\vec{z}}$:
 - i. Write $\vec{y}. \text{Shift}(\vec{z}) = (y_1, \dots, y_d)$, where $y_j \in [0, 2\delta)$ for all $j \in [d]$.
 - ii. Compute

$$Y := \left\{ (\vec{y}' \parallel \vec{z}) \mid \begin{array}{l} \ell_1 \in [\log(2\delta)], \dots, \ell_d \in [\log(2\delta)] \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right\}$$
 - (b) Bob sends set Y and Alice sends X to ideal functionality $\mathcal{F}_{\text{PSI-CA}}$.
 - (c) Bob receives output cnt from $\mathcal{F}_{\text{PSI-CA}}$. Update set $I = I \cup \{\vec{y}\}$ if cnt = 1.
3. **[output]** Bob outputs the intersection I .

Figure 9: Protocol realizing $\mathcal{F}_{\text{sa-PSI}}$ (Figure 8), where Bob learns the output $S_A \cap S_B$.

8.2 Structure-Aware PSI Cardinality/Sum

Construction Overview. The above principle can be extended to allow Alice (or Bob) to learn structure-aware PSI cardinality $|S_A \cap S_B|$, except that we accumulate over all the Y sets computed for each of Bob's input. It could be the case that two different inputs have overlapping prefixes within the same origin (for example, points which are "close" to each other). In our protocol,

Bob computes a multi-set with an associated cnt, indicating the number of inputs that any given prefix is associated with. After this, Alice and Bob can use $\mathcal{F}_{\text{PSI-Sum}}$, to learn the cardinality. We define the ideal functionality in [Figure 10](#) and present our construction in [Figure 11](#). We state the theorem below and defer its proof to [Appendix B.5](#) the full version of the paper.

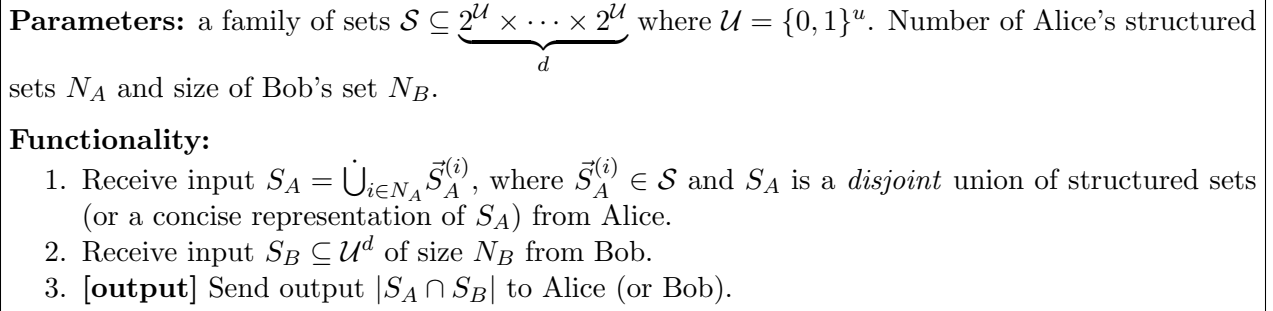


Figure 10: Ideal functionality $\mathcal{F}_{\text{sa-PSI-CA}}$ for structure-aware PSI cardinality.

Theorem 11. *The protocol in [Figure 11](#) securely realizes $\mathcal{F}_{\text{sa-PSI-CA}}$ ([Figure 10](#)) in the $\mathcal{F}_{\text{PSI-Sum}}$ -hybrid model, secure in the presence of semi-honest adversaries.*

Cost Analysis. For the set family \mathcal{S} as disjoint union of N_A balls of diameter $\leq \delta$ in ℓ_∞ norm, Alice and Bob run an instance of $\mathcal{F}_{\text{PSI-Sum}}$ between sets of size $n_1 = N_A \cdot w = N_A \cdot (\log(2\delta))^d$ and $n_2 = |S_B| \cdot (2 \cdot \log(2\delta))^d$. If we realize $\mathcal{F}_{\text{PSI-Sum}}$ with a semi-honest construction with linear computation and communication in set sizes [[PSTY19](#), [KK20](#), [HMS21](#), [RS21](#), [CGS22](#)], then we achieve sa-PSI-CA with communication and computation costs of $O(n_1 + n_2)$, which is $\ll |S_A|$. Note that we require the protocol to only reveal the associated sum, while hiding the intersection cardinality.

Structure-Aware PSI-Sum. The construction in [Figure 11](#) can be extended to realize structure-aware PSI with associated sum, when Bob’s inputs have associated values. The only change is that the associated value cnt, for any prefix $y' || z$ in origin z in set Y_v , is the sum of Bob’s associated values for all inputs that share the prefix y' .

9 Acknowledgments

This project is supported in part by the NSF CNS Award 2247352, Brown DSI Seed Grant, Meta Research Award, Google Research Scholar Award, and Amazon Research Award.

References

[ABD⁺21] Navid Alamati, Pedro Branco, Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, and Sihang Pu. Laconic private set intersection and applications. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part III*, volume 13044 of *LNCS*, pages 94–125. Springer, Heidelberg, November 2021.

[Ads] Private Intersection-Sum Protocols with Applications to Attributing Aggregate Ad Conversions. <https://research.google/pubs/pub51026/>.

Notation: We denote a set of elements as X and a set where every element has an associated value as $X_v = \{(x, v) \mid x \in X, \text{val}(x) = v\}$.

Parameters:

- private set intersection with associated sum ideal functionality $\mathcal{F}_{\text{PSI-Sum}}$

Inputs:

- Alice has N_A structured sets $S_A = \dot{\bigcup}_{i \in [N_A]} \vec{S}_A^{(i)}$, where $\vec{S}_A^{(i)} \in \mathcal{S}$ and sets in S_A are non-overlapping.
- Bob has an unstructured set $S_B \subseteq \mathcal{U}^d$ of size N_B .

Protocol:

1. Alice initializes $X := \emptyset$ and does the following:
 - (a) Identify a distinct origin $\vec{o}^{(i)}$ for each $i \in [N_A]$ such that $\vec{S}_A^{(i)} \subseteq \text{Univ}_{\vec{o}^{(i)}}$.
 - (b) For each $i \in [N_A]$:
 - i. Write $\vec{S}_A^{(i)}.\text{Shift}(\vec{o}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x^j}^{\vec{o}^{(i)}}$ as disjoint union of w' prefix sets (w is the upper bound for w').
 - ii. Update set $X := X \cup \{(x^j \parallel \vec{o}^{(i)}) \mid j \in [w']\}$.
 - (c) Pad X with dummy random strings such that $|X| = N_A \cdot w$.
2. Bob initializes a set with associated value $Y_v := \emptyset$ and does the following:
 - (a) For each element $\vec{y} \in S_B$, and for each origin \vec{z} where $\vec{y} \in \text{Univ}_{\vec{z}}$:
 - i. Write $\vec{y}.\text{Shift}(\vec{z}) = (y_1, \dots, y_d)$, where $y_j \in [0, 2\delta)$ for all $j \in [d]$.
 - ii. Compute
$$Y' = \left\{ (y' \parallel \vec{z}) \mid \begin{array}{l} \ell_1 \in [\log(2\delta)], \dots, \ell_d \in [\log(2\delta)] \\ y' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right\}$$
 - iii. For each $(y' \parallel \vec{z}) \in Y'$: if $(y' \parallel \vec{z}) \notin Y$ then $Y_v = Y_v \cup \{(y' \parallel \vec{z}, 1)\}$, else for $(y' \parallel \vec{z}, \text{cnt}) \in Y_v$, update $\text{cnt} = \text{cnt} + 1$.
 - (b) Pad Y_v with dummy random strings and values such that $|Y_v| = N_B \cdot (2 \cdot \log(2\delta))^d$.
3. Alice sends X and Bob sends Y_v to $\mathcal{F}_{\text{PSI-Sum}}$.
 - **Alice learns output:** functionality $\mathcal{F}_{\text{PSI-Sum}}$ returns output to Alice.
 - **Bob learns output:** functionality $\mathcal{F}_{\text{PSI-Sum}}$ returns output to Bob.

Figure 11: Protocol realizing $\mathcal{F}_{\text{sa-PSI-CA}}$ (Figure 10) in the $\mathcal{F}_{\text{PSI-Sum}}$ -hybrid model.

- [ADT11] Giuseppe Ateniese, Emiliano De Cristofaro, and Gene Tsudik. (If) size matters: Size-hiding private set intersection. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 156–173. Springer, Heidelberg, March 2011.
- [AKB07] Giuseppe Ateniese, Jonathan Kirsch, and Marina Blanton. Secret handshakes with dynamic and fuzzy matching. In *NDSS*, volume 7, pages 43–54, 2007.
- [ALOS22] Diego F. Aranha, Chuanwei Lin, Claudio Orlandi, and Mark Simkin. Laconic private set-intersection from pairings. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 111–124. ACM Press, November 2022.

- [ALP⁺21] Asra Ali, Tancrede Lepoint, Sarvar Patel, Mariana Raykova, Phillipp Schoppmann, Karn Seth, and Kevin Yeo. Communication-computation trade-offs in PIR. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 1811–1828. USENIX Association, August 2021.
- [ALSZ13] Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. More efficient oblivious transfer and extensions for faster secure computation. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 535–548. ACM Press, November 2013.
- [App] Privacy-Preserving Contact Tracing. <https://covid19.apple.com/contacttracing>.
- [APR⁺22] Amit Agarwal, Stanislav Peceny, Mariana Raykova, Phillipp Schoppmann, and Karn Seth. Communication efficient secure logistic regression. Cryptology ePrint Archive, Report 2022/866, 2022. <https://eprint.iacr.org/2022/866>.
- [BBC⁺21] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Lightweight techniques for private heavy hitters. In *2021 IEEE Symposium on Security and Privacy*, pages 762–776. IEEE Computer Society Press, May 2021.
- [BBV⁺20] Alex Berke, Michiel A. Bakker, Praneeth Vepakomma, Ramesh Raskar, Kent Larson, and Alex ‘Sandy’ Pentland. Assessing disease exposure risk with location histories and protecting privacy: A cryptographic approach in response to A global pandemic. *CoRR*, abs/2003.14412, 2020.
- [BCG⁺21] Elette Boyle, Nishanth Chandran, Niv Gilboa, Divya Gupta, Yuval Ishai, Nishant Kumar, and Mayank Rathee. Function secret sharing for mixed-mode and fixed-point secure computation. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 871–900. Springer, Heidelberg, October 2021.
- [BGI15] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 337–367. Springer, Heidelberg, April 2015.
- [BGI16] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1292–1303. ACM Press, October 2016.
- [BGI18] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. Cryptology ePrint Archive, Report 2018/707, 2018. <https://eprint.iacr.org/2018/707>.
- [BGIK22] Elette Boyle, Niv Gilboa, Yuval Ishai, and Victor I. Kolobov. Programmable distributed point functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 121–151. Springer, Heidelberg, August 2022.

- [BPSY23] Alexander Bienstock, Sarvar Patel, Joon Young Seo, and Kevin Yeo. Near-optimal oblivious key-value stores for efficient psi, PSU and volume-hiding multi-maps. In Joseph A. Calandrino and Carmela Troncoso, editors, *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*. USENIX Association, 2023.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
- [CFR21] Anrin Chakraborti, Giulia Fanti, and Michael K. Reiter. Distance-aware private set intersection, 2021.
- [CGS22] Nishanth Chandran, Divya Gupta, and Akash Shah. Circuit-PSI with linear complexity via relaxed batch OPRF. *PoPETs*, 2022(1):353–372, January 2022.
- [CHLR18] Hao Chen, Zhicong Huang, Kim Laine, and Peter Rindal. Labeled PSI from fully homomorphic encryption with malicious security. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1223–1237. ACM Press, October 2018.
- [Chr] Protect your accounts from data breaches with Password Checkup. <https://security.googleblog.com/2019/02/protect-your-accounts-from-data.html>.
- [CILO22] Wutichai Chongchitmate, Yuval Ishai, Steve Lu, and Rafail Ostrovsky. PSI from ring-OLE. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 531–545. ACM Press, November 2022.
- [CLR17] Hao Chen, Kim Laine, and Peter Rindal. Fast private set intersection from homomorphic encryption. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1243–1255. ACM Press, October / November 2017.
- [CM20] Melissa Chase and Peihan Miao. Private set intersection in the internet setting from lightweight oblivious PRF. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 34–63. Springer, Heidelberg, August 2020.
- [CMdG⁺21] Kelong Cong, Radames Cruz Moreno, Mariana Botelho da Gama, Wei Dai, Ilia Iliashenko, Kim Laine, and Michael Rosenberg. Labeled PSI from homomorphic encryption with reduced computation and communication. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1135–1150. ACM Press, November 2021.
- [CRR21] Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 502–534, Virtual Event, August 2021. Springer, Heidelberg.

- [DCW13] Changyu Dong, Liqun Chen, and Zikai Wen. When private set intersection meets big data: an efficient and scalable protocol. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 789–800. ACM Press, November 2013.
- [DKL⁺23] Nico Döttling, Dimitris Kolonelos, Russell W. F. Lai, Chuanwei Lin, Giulio Malavolta, and Ahmadreza Rahimi. Efficient laconic cryptography from learning with errors. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 417–446. Springer, Heidelberg, April 2023.
- [DMRY11] Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Moti Yung. Secure efficient multiparty computing of multivariate polynomials and applications. In Javier Lopez and Gene Tsudik, editors, *ACNS 11*, volume 6715 of *LNCS*, pages 130–146. Springer, Heidelberg, June 2011.
- [DT10] Emiliano De Cristofaro and Gene Tsudik. Practical private set intersection protocols with linear complexity. In Radu Sion, editor, *FC 2010*, volume 6052 of *LNCS*, pages 143–159. Springer, Heidelberg, January 2010.
- [Edg] Password Monitor: Safeguarding passwords in Microsoft Edge. <https://www.microsoft.com/en-us/research/blog/password-monitor-safeguarding-passwords-in-microsoft-edge/>.
- [ElG85] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
- [FNP04] Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 1–19. Springer, Heidelberg, May 2004.
- [GHMM24] Sanjam Garg, Mohammad Hajiabadi, Peihan Miao, and Alice Murphy. Laconic branching programs from the Diffie-Hellman assumption. In Qiang Tang and Vanessa Teague, editors, *PKC 2024, Part II*, volume 14603 of *LNCS*, pages 323–355. Springer, Heidelberg, April 2024.
- [GMR⁺21] Gayathri Garimella, Payman Mohassel, Mike Rosulek, Saeed Sadeghian, and Jaspal Singh. Private set operations from oblivious switching. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 591–617. Springer, Heidelberg, May 2021.
- [GPR⁺21] Gayathri Garimella, Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. Oblivious key-value stores and amplification for private set intersection. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 395–425, Virtual Event, August 2021. Springer, Heidelberg.
- [GRS22] Gayathri Garimella, Mike Rosulek, and Jaspal Singh. Structure-aware private set intersection, with applications to fuzzy matching. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 323–352. Springer, Heidelberg, August 2022.

- [GRS23] Gayathri Garimella, Mike Rosulek, and Jaspal Singh. Malicious secure, structure-aware private set intersection. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part I*, volume 14081 of *LNCS*, pages 577–610. Springer, Heidelberg, August 2023.
- [GS19] Satrajit Ghosh and Mark Simkin. The communication complexity of threshold private set intersection. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 3–29. Springer, Heidelberg, August 2019.
- [GWSW22] Jie Gao, Terrence Wong, Bassant Selim, and Chun Wang. VOMA: A privacy-preserving matching mechanism design for community ride-sharing. *IEEE Trans. Intell. Transp. Syst.*, 23(12):23963–23975, 2022.
- [HEK12] Yan Huang, David Evans, and Jonathan Katz. Private set intersection: Are garbled circuits better than custom protocols? In *NDSS 2012*. The Internet Society, February 2012.
- [HFH99] Bernardo A. Huberman, Matthew K. Franklin, and Tad Hogg. Enhancing privacy and trust in electronic communities. In Stuart I. Feldman and Michael P. Wellman, editors, *Proceedings of the First ACM Conference on Electronic Commerce (EC-99)*, Denver, CO, USA, November 3-5, 1999, pages 78–86. ACM, 1999.
- [HMS21] Kyoohyung Han, Dukjae Moon, and Yongha Son. Improved circuit-based PSI via equality preserving compression. Cryptology ePrint Archive, Report 2021/1440, 2021. <https://eprint.iacr.org/2021/1440>.
- [IKN⁺20] Mihaela Ion, Ben Kreuter, Ahmet Erhan Nergiz, Sarvar Patel, Shobhit Saxena, Karn Seth, Mariana Raykova, David Shanahan, and Moti Yung. On deploying secure computing: Private intersection-sum-with-cardinality. In *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*, pages 370–389. IEEE, 2020.
- [IKNP03] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 145–161. Springer, Heidelberg, August 2003.
- [JL10] Stanislaw Jarecki and Xiaomin Liu. Fast secure computation of set intersection. In Juan A. Garay and Roberto De Prisco, editors, *SCN 10*, volume 6280 of *LNCS*, pages 418–435. Springer, Heidelberg, September 2010.
- [Key] Password Monitoring – Apple Platform Security. <https://support.apple.com/en-al/guide/security/sec78e79fc3b/web>.
- [KK13] Vladimir Kolesnikov and Ranjit Kumaresan. Improved OT extension for transferring short secrets. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 54–70. Springer, Heidelberg, August 2013.

- [KK20] Ferhat Karakoç and Alptekin Küpçü. Linear complexity private set intersection for secure two-party protocols. In Stephan Krenn, Haya Shulman, and Serge Vaudenay, editors, *CANS 20*, volume 12579 of *LNCS*, pages 409–429. Springer, Heidelberg, December 2020.
- [KKRT16] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious PRF with applications to private set intersection. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 818–829. ACM Press, October 2016.
- [KRS⁺19] Daniel Kales, Christian Rechberger, Thomas Schneider, Matthias Senker, and Christian Weinert. Mobile private contact discovery at scale. In Nadia Heninger and Patrick Traynor, editors, *USENIX Security 2019*, pages 1447–1464. USENIX Association, August 2019.
- [KS05] Lea Kissner and Dawn Xiaodong Song. Privacy-preserving set operations. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 241–257. Springer, Heidelberg, August 2005.
- [MPR⁺20] Peihan Miao, Sarvar Patel, Mariana Raykova, Karn Seth, and Moti Yung. Two-sided malicious security for private intersection-sum with cardinality. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 3–33. Springer, Heidelberg, August 2020.
- [OOS17] Michele Orrù, Emmanuela Orsini, and Peter Scholl. Actively secure 1-out-of-N OT extension with application to private set intersection. In Helena Handschuh, editor, *CT-RSA 2017*, volume 10159 of *LNCS*, pages 381–396. Springer, Heidelberg, February 2017.
- [PRTY19] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. SpOT-light: Lightweight private set intersection from sparse OT extension. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 401–431. Springer, Heidelberg, August 2019.
- [PRTY20] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. PSI from PaXoS: Fast, malicious private set intersection. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 739–767. Springer, Heidelberg, May 2020.
- [PSSZ15] Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. Phasing: Private set intersection using permutation-based hashing. In Jaeyeon Jung and Thorsten Holz, editors, *USENIX Security 2015*, pages 515–530. USENIX Association, August 2015.
- [PSTY19] Benny Pinkas, Thomas Schneider, Oleksandr Tkachenko, and Avishay Yanai. Efficient circuit-based PSI with linear communication. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 122–153. Springer, Heidelberg, May 2019.

- [PSWW18] Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. Efficient circuit-based PSI via cuckoo hashing. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 125–157. Springer, Heidelberg, April / May 2018.
- [PSZ14] Benny Pinkas, Thomas Schneider, and Michael Zohner. Faster private set intersection based on OT extension. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 797–812. USENIX Association, August 2014.
- [Rab05] Michael O Rabin. How to exchange secrets with oblivious transfer. *Cryptology ePrint Archive*, 2005.
- [RADN22] Sara Ramezani, Gizem Akman, Mohamed Taoufiq Damir, and Valtteri Niemi. Lightweight privacy-preserving ride-sharing protocols for autonomous cars. In Björn Brücher, Christoph Krauß, Mario Fritz, Hans-Joachim Hof, and Oliver Wasenmüller, editors, *Computer Science in Cars Symposium, CSCS 2022, Ingolstadt, Germany, 8 December 2022*, pages 11:1–11:11. ACM, 2022.
- [RR17] Peter Rindal and Mike Rosulek. Improved private set intersection against malicious adversaries. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 235–259. Springer, Heidelberg, April / May 2017.
- [RR22] Srinivasan Raghuraman and Peter Rindal. Blazing fast PSI from improved OKVS and subfield VOLE. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 2505–2517. ACM Press, November 2022.
- [RS21] Peter Rindal and Phillipp Schoppmann. VOLE-PSI: Fast OPRF and circuit-PSI from vector-OLE. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 901–930. Springer, Heidelberg, October 2021.
- [Sig] Technology preview: Private contact discovery for Signal. <https://signal.org/blog/private-contact-discovery/>.
- [TSS⁺20] Ni Trieu, Kareem Shehata, Prateek Saxena, Reza Shokri, and Dawn Song. Epione: Lightweight contact tracing with strong privacy. *IEEE Data Eng. Bull.*, 43(2):95–107, 2020.
- [UCK⁺21] Erkam Uzun, Simon P. Chung, Vladimir Kolesnikov, Alexandra Boldyreva, and Wenke Lee. Fuzzy labeled private set intersection with applications to private real-time biometric search. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 911–928. USENIX Association, August 2021.
- [vBP24] Aron van Baarsen and Sihang Pu. Fuzzy private set intersection with large hyperballs. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part V*, volume 14655 of *Lecture Notes in Computer Science*, pages 340–369. Springer, 2024.

- [WG14] Yamin Wen and Zheng Gong. Private mutual authentications with fuzzy matching. *International Journal of High Performance Systems Architecture*, 5(1):3–12, 2014.
- [WXL⁺18] Xu An Wang, Fatos Xhafa, Xiaoshuang Luo, Shuaiwei Zhang, and Yong Ding. A privacy-preserving fuzzy interest matching protocol for friends finding in social networks. *Soft Computing*, 22(8):2517–2526, 2018.
- [ZCL21] En Zhang, Jian Chang, and Yu Li. Efficient threshold private set intersection. *IEEE Access*, 9:6560–6570, 2021.

A Cryptographic Tools

A.1 Cryptographic Schemes

Pseudorandom Generator. A pseudorandom generator (PRG) is a deterministic function $G : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{\gamma(\kappa)}$ where $G(x)$ for a randomly sampled $x \leftarrow \{0, 1\}^\kappa$ is computationally indistinguishable from a string randomly sampled from $\{0, 1\}^{\gamma(\kappa)}$ for any PPT adversary.

Hamming Correlation Robustness. Our protocol requires a hash function with a generalization of the “correlation robustness” property from OT-extension [IKNP03].

Definition 12 (Correlation Robust Hash Function [IKNP03]). *Let $H : \{0, 1\}^* \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^v$ be a function and define the related function $F_s(t, x) = H(t; x \oplus s)$, where $s \in \{0, 1\}^\kappa$. We say that H is correlation robust if F_s is computationally indistinguishable from a random function, against distinguishers that never query with repeated t -values. Intuitively, values of the form $H(t_i; x_i \oplus s)$ look jointly pseudorandom, even with known t_i, x_i values and an unknown but common s value.*

The specific property we use is defined below:

Definition 13. *Let $H : \{0, 1\}^* \times (\{0, 1\}^z)^\ell \rightarrow \{0, 1\}^v$ be a function and define the related function $F_s(t, x, \Delta) = H(t; x \oplus (s \odot \Delta))$, where $s \in \{0, 1\}^\ell$; x, Δ are vectors of length ℓ with components in $\{0, 1\}^z$; and \odot is component wise multiplication (of a bit times a string). We say that H is **Hamming correlation robust** if F_s is computationally indistinguishable from a random function, against distinguishers that never query with repeated t -values and always query with Δ having at least κ nonzero components.*

Intuitively, values of the form $H(t_i; x_i \oplus (s \odot \Delta_i))$ look jointly pseudorandom, even with known t_i, x_i, Δ_i values and an unknown but common s value, provided that each Δ_i has high Hamming weight.

This definition generalizes the one from [KKRT16] in that x and Δ are bit strings (vectors of bits) in [KKRT16], whereas in our protocol x and Δ can be vectors with components from $\{0, 1\}^z$.

A.2 Secure Two-Party Computation

Semi-Honest Security. We define secure two-party computation protocols, in the presence of semi-honest adversaries, using the *Universal Composability* (UC) framework [Can01]. Parties P_0 with input x_0 and P_1 with input x_1 run protocol Π to learn the output of a function $f(x_0, x_1)$ where party P_i learns $f_i(x_0, x_1)$. Party P_i 's $\text{View}_i(x_0, x_1)$ during an honest execution consists of her private input x_i , privately chosen randomness and the transcript of the protocol. Let $\text{Out}_i(x_0, x_1)$

be P_i 's output from protocol Π . We say that protocol Π securely realizes a functionality f if there exists a PPT simulator Sim for both parties and for all possible inputs x_0, x_1 such that for $i \in \{0, 1\}$,

$$(\text{Sim}(1^\kappa, i, x_i, f_i(x_0, x_1)), f_{1-i}(x_0, x_1)) \cong_\kappa (\text{View}_i(x_0, x_1), \text{Out}_{1-i}(x_0, x_1)).$$

Oblivious Transfer. Oblivious transfer (OT) [Rab05] is a special secure two-party computation protocol. In a single OT instance, the sender submits a pair of strings (m_0, m_1) and the receiver learns exactly one of the messages m_b , based on his choice bit $b \in \{0, 1\}$. The sender is oblivious of the choice bit. In Figure 12 we formally define the ideal functionality for ℓ parallel instances of oblivious transfer. A single OT instances requires public-key operations, while a large number of OT instances can be efficiently realized from κ number of public-key operations using *OT extension* techniques [IKNP03, KK13, ALSZ13].

Parameters: number of OTs ℓ , payload length r .

- Receive $(m_{1,0}, m_{1,1}), \dots, (m_{\ell,0}, m_{\ell,1})$ from the sender, where each $m_{i,b} \in \{0, 1\}^r$.
- Receive $\vec{b} \in \{0, 1\}^\ell$ from the receiver.
- Send output $(m_{1,b_1}, \dots, m_{\ell,b_\ell})$ to the receiver.

Figure 12: Ideal functionality \mathcal{F}_{OT} for ℓ oblivious transfers.

PSI Cardinality and PSI with Associated Sum. We formalize the ideal functionality $\mathcal{F}_{\text{PSI-CA}}$ for two-party PSI cardinality in Figure 13, where one party (either Alice or Bob), learns the output. We define the ideal functionality $\mathcal{F}_{\text{PSI-Sum}}$ for two-party PSI with associated sum in Figure 14, where Alice (or Bob) learns *only* the sum of the associated values for items in the intersection.

Parameters: Size of sets n_1, n_2 .

- Receive input $A = \{a_1, \dots, a_{n_1}\}$ where $a_i \in \{0, 1\}^*$ from Alice.
- Receive input $B = \{b_1, \dots, b_{n_2}\}$ where $b_i \in \{0, 1\}^*$ from Bob.
- Alice (or Bob) learns $|A \cap B|$, the cardinality of intersection.

Figure 13: Ideal functionality $\mathcal{F}_{\text{PSI-CA}}$ for computing cardinality of intersection.

Parameters: Size of sets n_1, n_2 .

- Receive input $A = \{a_1, \dots, a_n\}$ where $a_i \in \{0, 1\}^*$ from Alice.
- Receive input $B = \{(b_1, v_1), \dots, (b_{n_2}, v_{n_2})\}$ where $b_i \in \{0, 1\}^*$ and $v_i \in \mathbb{Z}$ from Bob.
- Alice learns $s = \sum_{b_i \in A} v_i$, the sum of associated values for intersection elements.

Figure 14: Ideal functionality $\mathcal{F}_{\text{PSI-Sum}}$ for computing sum of associated values in intersection.

B Security Proofs

B.1 Proof of Theorem 6

Bob is corrupt. The goal is to simulate Bob's view in the protocol (independent of Alice's input) which consists of ℓ OKVS encodings received from the \mathcal{F}_{OT} .

Hybrid 0. We begin with the honest execution of the protocol where Bob's view is generated using Alice's private input S_A .

Hybrid 1. In the protocol, based on Bob's choice $s[j]$ for each of his ℓ OT instances, he learns one $\mathcal{KV}_{s[j]}^{(j)}$ out of the two encoding sets $(\mathcal{KV}_0^{(j)}, \mathcal{KV}_1^{(j)})$. Recall that, an encoding set $\mathcal{KV}_b^{(j)}$ with index b only encodes ibFSS shares with the same index, where $b \in \{0, 1\}$. As a result, Bob is guaranteed to learn only one of two ibFSS shares for each of the ℓ independent sharings of her N_A inputs. In this hybrid, we modify how each of the sets $\{\mathcal{KV}_{s[1]}^{(1)}, \dots, \mathcal{KV}_{s[\ell]}^{(\ell)}\}$ to be computed as:

$$\mathcal{KV}_{s[j]}^{(j)} \leftarrow \text{Encode}(\{(\vec{\sigma}^{(1)}, k_{s[j]}^{\prime(1,j)}), \dots, (\vec{\sigma}^{(N_A)}, k_{s[j]}^{\prime(N_A,j)})\})$$

For every structure $i \in [N_A]$, the ibFSS key share is generated $k_{s[j]}^{\prime(i,j)} \leftarrow \text{Sim}(1^\kappa, s[j])$ by invoking the underlying ibFSS scheme privacy simulator ([Definition 4](#)). These shares are simulated independent of Alice's input and just using the security parameter and Bob's choice bit $s[j]$. This hybrid is indistinguishable from the honest execution in the \mathcal{F}_{OT} -hybrid model and by the privacy of the ibFSS scheme.

Hybrid 2. This hybrid is identical to the previous hybrid, except that we replace the keys $\{\vec{\sigma}^{(1)}, \dots, \vec{\sigma}^{(N_A)}\}$ with an arbitrary set of distinct keys $\{\vec{\sigma}^{\prime(1)}, \dots, \vec{\sigma}^{\prime(N_A)}\}$ in [Step 1](#). Now, the encoding sets $\{\mathcal{KV}_{s[1]}^{(1)}, \dots, \mathcal{KV}_{s[\ell]}^{(\ell)}\}$ from previous hybrid are computed as follows:

$$\mathcal{KV}_{s[j]}^{(j)} \leftarrow \text{Encode}(\{(\vec{\sigma}^{\prime(1)}, k_{s[j]}^{\prime(1,j)}), \dots, (\vec{\sigma}^{\prime(N_A)}, k_{s[j]}^{\prime(N_A,j)})\})$$

By the obliviousness property of the OKVS scheme ([Definition 1](#)), the original and modified encoding sets are computationally indistinguishable when associated values (ibFSS key shares) in the key-value mapping are pseudorandom. Since we instantiate the protocol with an ibFSS scheme that has pseudorandom keys property, this hybrid is indistinguishable from previous hybrid.

Simulator. Hybrid 2 defines a valid simulation in the ideal world where Bob's view is simulated independent of Alice's private input S_A .

Alice is corrupt. We define the following helper functions to evaluate on any input \vec{y} contained in the mini-universe with origin \vec{z} , namely $\vec{y} \in \text{Univ}_{\vec{z}}$, where the function output $\{\{0, 1\}^t\}^\ell$ is a vector of length ℓ with components of string length t corresponding to ibFSS evaluations:

$$E_*(\vec{y}||\vec{z}) = \left(\text{Eval}(\text{Decode}(\mathcal{KV}_*^{(1)}, \vec{z}), \vec{y}) || \dots || \text{Eval}(\text{Decode}(\mathcal{KV}_*^{(\ell)}, \vec{z}), \vec{y}) \right)$$

$$E_b(\vec{y}||\vec{z}) = \left(\text{Eval}(\text{Decode}(\mathcal{KV}_b^{(1)}, \vec{z}), \vec{y}) || \dots || \text{Eval}(\text{Decode}(\mathcal{KV}_b^{(\ell)}, \vec{z}), \vec{y}) \right) \text{ where } b \in \{0, 1\}$$

$$\Delta(\vec{y}||\vec{z}) = E_0(\vec{y}||\vec{z}) \oplus E_1(\vec{y}||\vec{z})$$

Since $\text{Decode}(\mathcal{KV}_*^{(\eta)}, \vec{z}) = \text{Decode}(\mathcal{KV}_0^{(\eta)}, \vec{z}) \oplus s[\eta] \odot (\text{Decode}(\mathcal{KV}_0^{(\eta)}, \vec{z}) \oplus \text{Decode}(\mathcal{KV}_1^{(\eta)}, \vec{z}))$, we get

$$E_*(\vec{y}||\vec{z}) = E_0(\vec{y}||\vec{z}) \oplus s \odot (E_0(\vec{y}||\vec{z}) \oplus E_1(\vec{y}||\vec{z})) \quad (1)$$

Now, we will show through a sequence of hybrids how Alice's view can be simulated independent of Bob's input.

Hybrid 0. In the protocol, Alice's receives a set of values Y from Bob.

$$Y = \left\{ \text{H} \left(\vec{y}' \| \vec{z}; E_*(\vec{y}' \| \vec{z}) \right) \left| \begin{array}{l} \vec{y}' \in S_B, \vec{y}' \in \text{Univ}_{\vec{z}}, \\ \vec{y}' \cdot \text{Shift}(\vec{z}) = (y_1, \dots, y_d), \\ \ell_1 \in [\log(2\delta)], \dots, \ell_d \in [\log(2\delta)], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\}$$

We write the set $Y = Y_1 \dot{\cup} Y_2 \dot{\cup} Y_3 \dot{\cup} Y_4 \dot{\cup} Y_5$ as the disjoint union of sets, where each term can be expressed as follows.

$$Y_1 = \left\{ \text{H} \left(\vec{y}' \| \vec{\sigma}^{(i)}; E_0(\vec{y}' \| \vec{\sigma}^{(i)}) \right) \left| \begin{array}{l} i \in [N_A], \vec{S}_A^{(i)} \cdot \text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x_j^{\vec{\sigma}^{(i)}}}, \vec{x}^j = (x_1^j, \dots, x_d^j), \\ \vec{y}' \in S_B, \vec{y}' \in \text{Univ}_{\vec{\sigma}^{(i)}}, \vec{y}' \cdot \text{Shift}(\vec{\sigma}^{(i)}) = (y_1, \dots, y_d) \in \text{PreS}_{x_j^{\vec{\sigma}^{(i)}}}, \\ \ell_1 \in [x_1^j : \log(2\delta)], \dots, \ell_d \in [x_d^j : \log(2\delta)], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\}$$

$$Y_2 = \left\{ \text{H} \left(\vec{y}' \| \vec{\sigma}^{(i)}; E_*(\vec{y}' \| \vec{\sigma}^{(i)}) \right) \left| \begin{array}{l} i \in [N_A], \vec{S}_A^{(i)} \cdot \text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x_j^{\vec{\sigma}^{(i)}}}, \vec{x}^j = (x_1^j, \dots, x_d^j), \\ \vec{y}' \in S_B, \vec{y}' \in \text{Univ}_{\vec{\sigma}^{(i)}}, \vec{y}' \cdot \text{Shift}(\vec{\sigma}^{(i)}) = (y_1, \dots, y_d) \in \text{PreS}_{x_j^{\vec{\sigma}^{(i)}}}, \\ (\ell_1, \dots, \ell_d) \in [\log(2\delta)]^d, \text{ where } \ell_1 \in [x_1^j - 1] \vee \dots \vee \ell_d \in [x_d^j - 1], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\}$$

$$Y_3 = \left\{ \text{H} \left(\vec{y}' \| \vec{\sigma}^{(i)}; E_*(\vec{y}' \| \vec{\sigma}^{(i)}) \right) \left| \begin{array}{l} i \in [N_A], \vec{y}' \in S_B, \vec{y}' \in \text{Univ}_{\vec{\sigma}^{(i)}}, \\ \vec{y}' \cdot \text{Shift}(\vec{\sigma}^{(i)}) = (y_1, \dots, y_d) \notin \vec{S}_A^{(i)} \cdot \text{Shift}(\vec{\sigma}^{(i)}), \\ \ell_1 \in [\log(2\delta)], \dots, \ell_d \in [\log(2\delta)], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\}$$

$$Y_4 = \left\{ \text{H} \left(\vec{y}' \| \vec{z}; E_*(\vec{y}' \| \vec{z}) \right) \left| \begin{array}{l} \vec{y}' \in S_B, \vec{y}' \in \text{Univ}_{\vec{z}}, \vec{y}' \cdot \text{Shift}(\vec{z}) = (y_1, \dots, y_d), \\ \vec{z} \notin \{\vec{\sigma}^{(1)}, \dots, \vec{\sigma}^{(N_A)}\} \\ \ell_1 \in [\log(2\delta)], \dots, \ell_d \in [\log(2\delta)], \\ \vec{y}' := ((y_1)_{[1:\ell_1]}, \dots, (y_d)_{[1:\ell_d]}) \end{array} \right. \right\}$$

and Y_5 consists of dummy pseudorandom strings of length equal to the hash outputs, such that $|Y| = N_B \cdot (2 \cdot \log(2\delta))^d$. In the protocol, every set $S_A^{(i)}$ in Alice's input is uniquely mapped to the origin $\vec{\sigma}^{(i)}$ of a mini-universe, such that $\vec{S}_A^{(i)} \cdot \text{Shift}(\vec{\sigma}^{(i)}) \subseteq \text{Univ}_{\vec{\sigma}^{(i)}}$. Therefore, Alice's input defines a set of origins $O = \{\vec{\sigma}^{(1)}, \dots, \vec{\sigma}^{(N_A)}\}$. When Bob evaluates on a point $\vec{y}' \in S_B$, he considers all the (bounded set of) origins (and in turn mini-universes) that contain his input. For each of these origins, he sends hash evaluations on all prefix combinations of his input vector.

If Bob's input $\vec{y}' \in \text{Univ}_{\vec{\sigma}^{(i)}}$ belongs to a mini-universe defined by one of Alice's origins $\vec{\sigma}^{(i)}$, then $\text{Decode}(\mathcal{KV}_*^{(\eta)}, \vec{\sigma}^{(i)})$ is valid **ibFSS** share of Alice's input $\vec{S}_A^{(i)}$ (by correctness property of **OKVS** scheme), for $\eta \in [\ell]$. For all such inputs $\vec{y}' \in S_B, \vec{y}' \in \text{Univ}_{\vec{\sigma}^{(i)}}$, we get the following cases:

1. If Bob's input $\vec{y}' \cdot \text{Shift}(\vec{\sigma}^{(i)}) \in \vec{S}_A^{(i)} \cdot \text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x_j^{\vec{\sigma}^{(i)}}}$ belongs to some prefix set: set Y_1 consists of hash outputs $\text{H}(\vec{y}' \| \vec{\sigma}^{(i)}; E_*(\vec{y}' \| \vec{\sigma}^{(i)}))$ of every input prefix $\vec{y}' \in \vec{S}_A^{(i)} \cdot \text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x_j^{\vec{\sigma}^{(i)}}}$ that also *belongs to the same* prefix set. In this case, $(E_0(\vec{y}' \| \vec{\sigma}^{(i)}) \oplus E_1(\vec{y}' \| \vec{\sigma}^{(i)})) = 0^t, \Delta(\vec{y}' \| \vec{\sigma}^{(i)})$ has all zero components, by the correctness of **ibFSS** scheme. Thus, we can simplify [Equation 1](#) and hash outputs in Y_1 as $\text{H}(\vec{y}' \| \vec{\sigma}^{(i)}; E_0(\vec{y}' \| \vec{\sigma}^{(i)}))$.

2. If Bob's input $\vec{y}.\text{Shift}(\vec{\sigma}^{(i)}) \in \vec{S}_A^{(i)}.\text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x^j} \rightarrow$ belongs to some prefix set: set Y_2 consists of hash outputs $H(\vec{y}' \parallel \vec{\sigma}^{(i)}; E_*(\vec{y}' \parallel \vec{\sigma}^{(i)}))$ of every input prefix $\vec{y}' \notin \vec{S}_A^{(i)}.\text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x^j} \rightarrow$ outside the prefix set. In this case, $(E_0(\vec{y}' \parallel \vec{\sigma}^{(i)}) \oplus E_1(\vec{y}' \parallel \vec{\sigma}^{(i)})) \neq 0^t$, $\Delta(\vec{y}' \parallel \vec{\sigma}^{(i)})$ consists of $\ell = \mathcal{O}(\kappa)$ non-zero components, by the correctness of ibFSS scheme.
3. If Bob's input $\vec{y}.\text{Shift}(\vec{\sigma}^{(i)}) \notin \vec{S}_A^{(i)}.\text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x^j} \rightarrow$ is outside Alice's input: set Y_3 consists of hash outputs $H(\vec{y}' \parallel \vec{\sigma}^{(i)}; E_*(\vec{y}' \parallel \vec{\sigma}^{(i)}))$ of every input prefix \vec{y}' of input \vec{y} . Again, $(E_0(\vec{y}' \parallel \vec{\sigma}^{(i)}) \oplus E_1(\vec{y}' \parallel \vec{\sigma}^{(i)})) \neq 0^t$, $\Delta(\vec{y}' \parallel \vec{\sigma}^{(i)})$ consists of $\ell = \mathcal{O}(\kappa)$ non-zero components, by the correctness of ibFSS scheme.

Finally, consider the case when Bob's input $\vec{y}.\text{Shift}(\vec{z}) \in \text{Univ}_{\vec{z}}$ is evaluated on a mini-universe defined by $\vec{z} \notin \{\vec{\sigma}^{(1)}, \dots, \vec{\sigma}^{(N_A)}\}$, then $\text{Decode}(\mathcal{KV}_*^{(\eta)}, \vec{z})$ is not a valid ibFSS share (by correctness property of OKVS scheme), for $\eta \in [\ell]$. Instead, $\text{Decode}(\mathcal{KV}_b^{(\eta)}, \vec{z}), \text{Eval}(\text{Decode}(\mathcal{KV}_b^{(\eta)}, \vec{z}), \vec{y}')$ are uniform strings for every $b \in \{0, 1\}, \eta \in [\ell]$, by the independence property of OKVS scheme (**Definition 2**). Set Y_4 consists of hash outputs $H(\vec{y}' \parallel \vec{z}; E_*(\vec{y}' \parallel \vec{z}))$ of every input prefix \vec{y}' of such inputs \vec{y} and $\Delta(\vec{y}' \parallel \vec{z})$ is a uniform string of length $d \cdot \ell$ (by OKVS independence property).

Hybrid 1. This hybrid is identical to the previous one, except that we abort if there exists any hash in set Y_4 whose $\Delta(\vec{y}' \parallel \vec{z})$ has less than κ non-zero components. From our previous hybrid, we know that $\Delta(\vec{y}' \parallel \vec{z})$ is a uniform string for each hash evaluations in Y_4 . Further, protocol specifies that parameter ℓ (number of components) is chosen so that $\Pr[\text{Binomial}(1 - 2^{-d}, \ell) < \kappa] < 2^{-\lambda - |N_B| - d(1 + \log \log(2\delta))}$. We use the fact that $|Y_4| \leq |Y|$ and by union bound, the abort probability is upper bounded by $2^{-\lambda - |N_B| - d(1 + \log \log(2\delta))} \cdot |N_B| \cdot (2 \cdot \log(2\delta))^d = 2^{-\lambda}$, which is negligible. Thus, this hybrid is indistinguishable from the honest execution.

Hybrid 2. Only change is that, we replace the hash outputs in Y_2, Y_3 and Y_4 with uniform strings of length $\lambda + \log |N_A| + \log w + \log |N_B| + d(1 + \log \log(2\delta))$ as shown below:

$$Y = Y_1 \dot{\cup} \{h_1, \dots, h_{|Y_2|}\} \dot{\cup} \{h'_1, \dots, h'_{|Y_3|}\} \dot{\cup} \{h''_1, \dots, h''_{|Y_4|}\} \dot{\cup} Y_5$$

where each $h_i, h'_i, h''_i \leftarrow \{0, 1\}^{\lambda + \log |N_A| + \log w + \log |N_B| + d(1 + \log \log(2\delta))}$ are sampled uniformly. This hybrid is indistinguishable from previous hybrid by the hamming correlation robust property (**Definition 13**) which states that hash values of the form $H(\vec{y}' \parallel \vec{z}; t_i \oplus s \odot \Delta(\vec{y}' \parallel \vec{z}))$ are jointly pseudorandom, if they are queried on unique $\vec{y}' \parallel \vec{z}$ values and $\Delta(\vec{y}' \parallel \vec{z})$ has at least κ non-zero components. These properties hold for sets Y_2, Y_3, Y_4 since they are disjoint sets with hash outputs computed on unique $\vec{y}' \parallel \vec{z}$ values and $\Delta(\vec{y}' \parallel \vec{z})$ evaluation has at least κ non-zero components (from discussion in Hybrid 0 and Hybrid 1).

Simulator. Hybrid 2 defines a valid simulation in the ideal world computed just using Alice's input S_A and output $S_A \cap S_B$. Note that, the output can be used to determine $|Y_2|, |Y_3| + |Y_4|$, and $|Y_5|$.

Correctness. We will prove that Alice indeed outputs the intersection $S_A \cap S_B$ as her output, except with negligible probability. Let's consider an input \vec{a} from Alice's set $\vec{S}_A^{(i)}$ which belongs to prefix set $\vec{a} \in \text{PreS}_{x^j} \rightarrow, x^j = (x_1^j, \dots, x_d^j)$, that is, $\vec{a}.\text{Shift}(\vec{\sigma}^{(i)}) \in \vec{S}_A^{(i)}.\text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x^j} \rightarrow$. (The ibFSS scheme guarantees $\Delta(\vec{a}' \parallel \vec{\sigma}^{(i)}) = 0^t$, where $\vec{a}' \in \text{PreS}_{x^j} \rightarrow$). For every such input $\vec{a} = (a_1, \dots, a_d)$ we get the following cases:

- **When $\vec{a} \in S_B$:** set Y_1 will include the following hash values

$$\left\{ \mathbf{H} \left(\vec{a}' \| \vec{\sigma}^{(i)}; E_0(\vec{a}' \| \vec{\sigma}^{(i)}) \right) \left| \begin{array}{l} \ell_1 \in [|x_1^j| : u], \dots, \ell_d \in [|x_d^j| : u] \\ \vec{a}' := ((a_1)_{[1:\ell_1]}, \dots, (a_d)_{[1:\ell_d]}) \end{array} \right. \right\}$$

By collision resistant hash property, Alice can uniquely recognize these hash values and correctly includes \vec{a} in her output.

- **When $\vec{a} \notin S_B$:** Alice wrongly includes \vec{a} in her output only when she finds hash of her prefix set $H(\vec{x}^j \| \vec{\sigma}^{(i)}; E_0(\vec{x}^j \| \vec{\sigma}^{(i)}))$ in Bob's message. By union bound, the probability of finding a matching value for a specific prefix is $2^{-\lambda - \log |N_A| - \log w - \log |N_B| - d(1 + \log \log(2\delta))} \cdot |Y| = 2^{-\lambda - \log |N_A| - \log w}$. By union bound over total number of prefixes, the probability of including a wrong value is less than $w \cdot |N_A| \cdot 2^{-\lambda - \log |N_A| - \log w} = 2^{-\lambda}$, is negligible.

B.2 Proof of Theorem 8

Notation. We begin by presenting notation that is helpful for both the correctness and security proofs. Let $\text{Lose} \in \{L, R\}$ refers the path that exits the special path determined by α . If $\alpha^{(j)} = 0$ then $\text{Lose} = R$ is the right subtree and if $\alpha^{(j)} = 1$ then $\text{Lose} = L$ is the left subtree. We define $\text{Keep} \neq \text{Lose}$ as the path determined by α . The correction word $CW^{(j)}$ consists of five part computed as follows:

- $s_{CW} = s_b^{\text{Lose}} \oplus s_{1-b}^{\text{Lose}}$
- $t_{CW}^L = t_b^L \oplus t_{1-b}^L \oplus \alpha^{(j)} \oplus 1$
- $t_{CW}^R = t_b^R \oplus t_{1-b}^R \oplus \alpha^{(j)}$
- $y_{CW}^L = y_b^L \oplus y_{1-b}^L \oplus \alpha^{(j)}$
- $y_{CW}^R = y_b^R \oplus y_{1-b}^R$

This entails the following:

- $t_{CW}^{\text{Keep}} = t_b^{\text{Keep}} \oplus t_{1-b}^{\text{Keep}} \oplus 1$
- $t_{CW}^{\text{Lose}} = t_b^{\text{Lose}} \oplus t_{1-b}^{\text{Lose}}$
- $y_{CW}^{\text{Keep}} = y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}$
- $y_{CW}^{\text{Lose}} = y_b^{\text{Lose}} \oplus y_{1-b}^{\text{Lose}} \oplus \alpha$

Given these, the correction word is constructed as $CW^{(j)} = s_{CW} \| t_{CW}^L \| y_{CW}^L \| s_{CW} \| t_{CW}^R \| y_{CW}^R$.

Correctness. We will show that for each pairing of interval boundary prefix $\alpha_{[1:j]}$ and evaluation prefix $x_{[1:j]}$, our ibDCF protocol outputs the correct secret shared value $y^{(j)} := y_0^{(j)} \oplus y_1^{(j)}$, given that it already determined the correct $y^{(j-1)}$ as well as $t^{(j-1)}$ and $s^{(j-1)}$. $t^{(j-1)} := t_0^{(j-1)} \oplus t_1^{(j-1)}$ is the indicator for if $\alpha_{[1:j-1]} = x_{[1:j-1]}$, and $s^{(j-1)} := s_0^{(j-1)} \oplus s_1^{(j-1)}$ is the random seed used to calculate correction words.

Proof. We proceed by cases.

Case 0: $j = 0$:

For a zero bit input, correctness relies only on initialization of values. $y_b^{(0)} = b$ and $t_b^{(0)} = b$ by assignment, so $y^{(0)} = 1$ and $t^{(0)} = 1$. Further, $s^{(0)}$ is generated uniformly at random.

Case 1: $x_{[1:j]} = \alpha_{[1:j]}$:

$$\begin{aligned}
t_b^{(j)} \oplus t_{1-b}^{(j)} &= (t_b^{\text{Keep}} \oplus t_b^{(j-1)} \cdot t_{CW}^{\text{Keep}}) \oplus (t_{1-b}^{\text{Keep}} \oplus t_{1-b}^{(j-1)} \cdot t_{CW}^{\text{Keep}}) \\
&= (t_b^{\text{Keep}} \oplus t_{1-b}^{\text{Keep}}) \oplus t_{CW}^{\text{Keep}} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \\
&= (t_b^{\text{Keep}} \oplus t_{1-b}^{\text{Keep}}) \oplus (t_b^{\text{Keep}} \oplus t_{1-b}^{\text{Keep}} \oplus 1) \cdot 1 \\
&= 1 \\
s_b^{(j)} \oplus s_{1-b}^{(j)} &= (s_b^{\text{Keep}} \oplus t_b^{(j-1)} \cdot s_{CW}) \oplus (s_{1-b}^{\text{Keep}} \oplus t_{1-b}^{(j-1)} \cdot s_{CW}) \\
&= (s_b^{\text{Keep}} \oplus s_{1-b}^{\text{Keep}}) \oplus s_{CW} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \\
&= (s_b^{\text{Keep}} \oplus s_{1-b}^{\text{Keep}}) \oplus (s_b^{\text{Lose}} \oplus s_{1-b}^{\text{Lose}}) \cdot 1 \\
&= (s_0^L \oplus s_0^R) \oplus (s_1^L \oplus s_1^R)
\end{aligned}$$

These two pairs of values are substrings of outputs from our pseudorandom generator G with pseudorandom inputs $s_0^{(j-1)}$ and $s_1^{(j-1)}$, respectively. Thus, $s^{(j)}$ is also pseudorandom.

$$\begin{aligned}
y_b^{(j)} \oplus y_{1-b}^{(j)} &= (y_b^{\text{Keep}} \oplus t_b^{(j-1)} \cdot y_{CW}^{\text{Keep}} \oplus y_b^{(j-1)}) \oplus (y_{1-b}^{\text{Keep}} \oplus t_{1-b}^{(j-1)} \cdot y_{CW}^{\text{Keep}} \oplus y_{1-b}^{(j-1)}) \\
&= (y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}) \oplus y_{CW}^{\text{Keep}} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \oplus (y_b^{(j-1)} \oplus y_{1-b}^{(j-1)}) \\
&= (y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}) \oplus y_{CW}^{\text{Keep}} \cdot 1 \oplus 1 \\
&= (y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}) \oplus (y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}) \oplus 1 \\
&= 1
\end{aligned}$$

Case 2: $x_{[1:j-1]} = \alpha_{[1:j-1]}$, $x^{(j)} \neq \alpha^{(j)}$:

$$\begin{aligned}
t_b^{(j)} \oplus t_{1-b}^{(j)} &= (t_b^{\text{Lose}} \oplus t_b^{(j-1)} \cdot t_{CW}^{\text{Lose}}) \oplus (t_{1-b}^{\text{Lose}} \oplus t_{1-b}^{(j-1)} \cdot t_{CW}^{\text{Lose}}) \\
&= (t_b^{\text{Lose}} \oplus t_{1-b}^{\text{Lose}}) \oplus t_{CW}^{\text{Lose}} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \\
&= (t_b^{\text{Lose}} \oplus t_{1-b}^{\text{Lose}}) \oplus (t_b^{\text{Lose}} \oplus t_{1-b}^{\text{Lose}}) \cdot 1 \\
&= 0 \\
s_b^{(j)} \oplus s_{1-b}^{(j)} &= (s_b^{\text{Lose}} \oplus t_b^{(j-1)} \cdot s_{CW}) \oplus (s_{1-b}^{\text{Lose}} \oplus t_{1-b}^{(j-1)} \cdot s_{CW}) \\
&= (s_b^{\text{Lose}} \oplus s_{1-b}^{\text{Lose}}) \oplus s_{CW} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \\
&= (s_b^{\text{Lose}} \oplus s_{1-b}^{\text{Lose}}) \oplus (s_b^{\text{Lose}} \oplus s_{1-b}^{\text{Lose}}) \cdot 1 \\
&= 0 \\
y_b^{(j)} \oplus y_{1-b}^{(j)} &= (y_b^{\text{Lose}} \oplus t_b^{(j-1)} \cdot y_{CW}^{\text{Lose}} \oplus y_b^{(j-1)}) \oplus (y_{1-b}^{\text{Lose}} \oplus t_{1-b}^{(j-1)} \cdot y_{CW}^{\text{Lose}} \oplus y_{1-b}^{(j-1)}) \\
&= (y_b^{\text{Lose}} \oplus y_{1-b}^{\text{Lose}}) \oplus y_{CW}^{\text{Lose}} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \oplus (y_b^{(j-1)} \oplus y_{1-b}^{(j-1)}) \\
&= (y_b^{\text{Lose}} \oplus y_{1-b}^{\text{Lose}}) \oplus y_{CW}^{\text{Lose}} \cdot 1 \oplus 1 \\
&= (y_b^{\text{Lose}} \oplus y_{1-b}^{\text{Lose}}) \oplus (y_b^{\text{Lose}} \oplus y_{1-b}^{\text{Lose}} \oplus \alpha^{(j)}) \oplus 1 \\
&= \alpha^{(j)}
\end{aligned}$$

When $\alpha^{(j)} = 1$, and $x^{(j)} = 0$, $x_{[1:j]} < \alpha_{[1:j]}$ and $y^{(j)} = 0$ as expected. Conversely, when $\alpha^{(j)} = 0$, and $x^{(j)} = 1$, $x_{[1:j]} > \alpha_{[1:j]}$ and $y^{(j)} = 1$.

Case 3: $x_{[1:j-1]} \neq \alpha_{[1:j-1]}$:

From case 2, $s_b^{(j-1)} = s_{1-b}^{(j-1)}$. Therefore, $t_b^L \| y_b^L \| t_b^R \| y_b^R = t_{1-b}^L \| y_{1-b}^L \| t_{1-b}^R \| y_{1-b}^R$, and we can disregard the party index subscripts on these terms. Further, we will see that the **Keep/Lose** terms cancel out, so we will handle both cases with **Choice** for succinctness.

$$\begin{aligned}
t_b^{(j)} \oplus t_{1-b}^{(j)} &= (t_b^{\text{Choice}} \oplus t_b^{(j-1)} \cdot t_{CW}^{\text{Choice}}) \oplus (t_{1-b}^{\text{Choice}} \oplus t_{1-b}^{(j-1)} \cdot t_{CW}^{\text{Choice}}) \\
&= (t^{\text{Choice}} \oplus t_b^{(j-1)} \cdot t_{CW}^{\text{Choice}}) \oplus (t^{\text{Choice}} \oplus t_{1-b}^{(j-1)} \cdot t_{CW}^{\text{Choice}}) \\
&= (t^{\text{Choice}} \oplus t^{\text{Choice}}) \oplus (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \cdot t_{CW}^{\text{Choice}} \\
&= 0 \oplus 0 \cdot t_{CW}^{\text{Choice}} \\
&= 0
\end{aligned}$$

$$\begin{aligned}
s_b^{(j)} \oplus s_{1-b}^{(j)} &= (s_b^{\text{Choice}} \oplus t_b^{(j-1)} \cdot s_{CW}) \oplus (s_{1-b}^{\text{Choice}} \oplus t_{1-b}^{(j-1)} \cdot s_{CW}) \\
&= (s^{\text{Choice}} \oplus t_b^{(j-1)} \cdot s_{CW}) \oplus (s^{\text{Choice}} \oplus t_{1-b}^{(j-1)} \cdot s_{CW}) \\
&= (s^{\text{Choice}} \oplus s^{\text{Choice}}) \oplus s_{CW} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \\
&= 0 \oplus s_{CW} \cdot 0 \\
&= 0
\end{aligned}$$

$$\begin{aligned}
y_b^{(j)} \oplus y_{1-b}^{(j)} &= (y_b^{\text{Choice}} \oplus t_b^{(j-1)} \cdot y_{CW}^{\text{Choice}} \oplus y_b^{(j-1)}) \oplus (y_{1-b}^{\text{Choice}} \oplus t_{1-b}^{(j-1)} \cdot y_{CW}^{\text{Choice}} \oplus y_{1-b}^{(j-1)}) \\
&= (y^{\text{Choice}} \oplus t_b^{(j-1)} \cdot y_{CW}^{\text{Choice}} \oplus y_b^{(j-1)}) \oplus (y^{\text{Choice}} \oplus t_{1-b}^{(j-1)} \cdot y_{CW}^{\text{Choice}} \oplus y_{1-b}^{(j-1)}) \\
&= (y^{\text{Choice}} \oplus y^{\text{Choice}}) \oplus y_{CW}^{\text{Choice}} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \oplus (y_b^{(j-1)} \oplus y_{1-b}^{(j-1)}) \\
&= 0 \oplus y_{CW}^{\text{Choice}} \cdot 0 \oplus (y_b^{(j-1)} \oplus y_{1-b}^{(j-1)}) \\
&= y_b^{(j-1)} \oplus y_{1-b}^{(j-1)} \\
&= y^{(j-1)}
\end{aligned}$$

The shortest differing equal-length prefixes of x and α , if they exist, are enough to determine the output value y of every longer evaluation prefix, because the difference between prefixes can not change sign as the prefixes grow in length. Indeed, we echo the output of the previous prefix in this case.

Therefore, by induction, our **ibDCF** protocol outputs the correct $y := y_0 \oplus y_1$ for any interval boundary α and evaluation input x . \square

Security. Through a sequence of hybrids, we show that for each party $b \in \{0, 1\}$, the key $k_b = s_b^{(0)} \| CW^{(1)} \| CW^{(2)} \| \dots \| CW^{(u)}$ is pseudorandom. First, let us define hybrid $\text{Hyb}_j(1^\kappa, \alpha, b)$ for $j \in \{0, 1, \dots, u\}$ as follows:

- Let $\alpha = \overline{\alpha^{(1)}\alpha^{(2)} \dots \alpha^{(u)}} \in \{0, 1\}^u$ be the bit decomposition.
- Sample uniformly $s_b^{(0)} \leftarrow \{0, 1\}^\kappa$, $t_b^{(0)} = b$ and $y_b^{(0)} = b$.
- For $i \leq j$, sample the correction words $CW^{(1)}, CW^{(2)}, \dots, CW^{(j)} \leftarrow \{0, 1\}^{\kappa+4}$ uniformly at random and use them to compute $s_b^{(i)} \| t_b^{(i)} \| y_b^{(i)}$ (according to the pseudocode).

- For j , uniformly sample the other party's PRG seed $s_{1-b}^{(j)} \leftarrow \{0, 1\}^\kappa$ and let $t_{1-b}^{(j)} = 1 - t_b^{(j)}$ and $y_{1-b}^{(j)} = 1 - y_b^{(j)}$
- For $i > j$, compute the values honestly $CW^{(i)}$, $s_b^{(i)} \| t_b^{(i)} \| y_b^{(i)}$, $s_{1-b}^{(i)} \| t_{1-b}^{(i)} \| y_{1-b}^{(i)}$ (according to the pseudocode) as a function of the values sampled in the previous step.
- The output of the hybrid is the key $k_b = s_b^{(0)} \| CW^{(1)} \| CW^{(2)} \| \dots \| CW^{(u)}$.

By design, we can see that Hyb_0 corresponds to honest key generation process from the pseudocode. On the other hand, Hyb_u yields a key that is truly random. Now, we will show that every pair of adjacent hybrids are indistinguishable by the security of the pseudorandom generator. The following claims and theorem complete the formal proof.

Claim 14. For both parties $b \in \{0, 1\}$, any choice of $\alpha \in \{0, 1\}^u$:

$$\{k_b \leftarrow \text{Hyb}_0(1^\kappa, \alpha, b)\} \equiv \{k_b : (k_0, k_1) \leftarrow \text{Gen}(1^\kappa, \alpha)\}$$

Claim 15. For both parties $b \in \{0, 1\}$, any choice of $\alpha \in \{0, 1\}^u$:

$$\{k_b \leftarrow \text{Hyb}_u(1^\kappa, \alpha, b)\} \equiv \{k_b \leftarrow \{0, 1\}^{u \cdot (\kappa+4) + \kappa}\}$$

Theorem 16. Given any secure pseudorandom generator G , such that, any adversary \mathcal{B} has distinguishing advantage ϵ in time T . There exists a polynomial p , so that for every $j \in [u]$, $b \in \{0, 1\}$, $\alpha \in \{0, 1\}^u$ and every adversary \mathcal{A} running in time $T' \leq T - p(\kappa)$ can distinguish between Hyb_{j-1} , Hyb_j with ϵ advantage as follows:

$$|\Pr[k_b \leftarrow \text{Hyb}_{j-1}(1^\kappa, \alpha, b); c \leftarrow \mathcal{A}(1^\kappa, k_b) : c = 1] - \Pr[k_b \leftarrow \text{Hyb}_j(1^\kappa, \alpha, b); c \leftarrow \mathcal{A}(1^\kappa, k_b) : c = 1]| < \epsilon$$

Proof. We present a PRG adversary \mathcal{B} in [Figure 15](#) with the same advantage as an adversary \mathcal{A} that can distinguish between $\text{Hyb}_{j-1}(1^\kappa, \alpha, b)$ and $\text{Hyb}_j(1^\kappa, \alpha, b)$. If r is sampled pseudorandomly ($s \leftarrow \{0, 1\}^\kappa$ and $G(s) = r$), then the output k_b of \mathcal{B} is identically distributed to $\text{Hyb}_{j-1}(1^\kappa, \alpha, b)$. If r is sampled uniformly at random $r \leftarrow \{0, 1\}^{2(\kappa+2)}$ then output of \mathcal{B} is identically distributed to $\text{Hyb}_j(1^\kappa, \alpha, b)$ by the following argument (utilizing notation defined above).

If r is random, then $s_{1-b}^{\text{Lose}}, t_{1-b}^L, t_{1-b}^R, y_{1-b}^L, y_{1-b}^R$ acts as one-time pad so that $CW^{(j)}$ is uniformly distributed. The seed for further computation $s_{1-b}^{(j)}$ is also uniformly distributed since $s_{1-b}^{(j)}$ is either s_{1-b}^{Keep} or $s_{1-b}^{\text{Keep}} \oplus s_{CW}$. Next, we verify the values $t_b^{(j)}, t_{1-b}^{(j)}, y_b^{(j)}, y_{1-b}^{(j)}$ following using the fact that ($t_b^{(j-1)} \oplus t_{1-b}^{(j-1)} = 1$) and ($y_b^{(j-1)} \oplus y_{1-b}^{(j-1)} = 1$):

$$\begin{aligned} t_b^{(j)} \oplus t_{1-b}^{(j)} &= (t_b^{\text{Keep}} \oplus t_b^{(j-1)} \cdot t_{CW}^{\text{Keep}}) \oplus (t_{1-b}^{\text{Keep}} \oplus t_{1-b}^{(j-1)} \cdot t_{CW}^{\text{Keep}}) \\ &= (t_b^{\text{Keep}} \oplus t_{1-b}^{\text{Keep}}) \oplus t_{CW}^{\text{Keep}} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \\ &= (t_b^{\text{Keep}} \oplus t_{1-b}^{\text{Keep}}) \oplus 1 \cdot (t_0^{\text{Keep}} \oplus t_1^{\text{Keep}} \oplus 1) \\ &= 1 \end{aligned}$$

$$\begin{aligned} y_b^{(j)} \oplus y_{1-b}^{(j)} &= (y_b^{\text{Keep}} \oplus t_b^{(j-1)} \cdot y_{CW}^{\text{Keep}} \oplus y_b^{(j-1)}) \oplus (y_{1-b}^{\text{Keep}} \oplus t_{1-b}^{(j-1)} \cdot y_{CW}^{\text{Keep}} \oplus y_{1-b}^{(j-1)}) \\ &= (y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}) \oplus y_{CW}^{\text{Keep}} \cdot (t_b^{(j-1)} \oplus t_{1-b}^{(j-1)}) \oplus (y_b^{(j-1)} \oplus y_{1-b}^{(j-1)}) \\ &= (y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}) \oplus y_{CW}^{\text{Keep}} \cdot 1 \oplus 1 \\ &= (y_b^{\text{Keep}} \oplus y_{1-b}^{\text{Keep}}) \oplus (y_0^{\text{Keep}} \oplus y_1^{\text{Keep}}) \oplus 1 \\ &= 1 \end{aligned}$$

PRG adversary $\mathcal{B}^*(1^\kappa, (j, b, \alpha), r)$:

Let $\alpha = \overline{\alpha^{(1)}\alpha^{(2)} \dots \alpha^{(u)}} \in \{0, 1\}^u$ be the bit decomposition.

Sample $s_b^{(0)} \leftarrow \{0, 1\}^\kappa$ uniformly at random, $t_b^{(0)} = b$ and $y_b^{(0)} = b$.

for $i = 1$ **to** $(j - 1)$ **do**

$CW^{(i)} \leftarrow \{0, 1\}^{\kappa+4}$. Parse $CW^{(i)} = s_{CW} \| t_{CW}^L \| t_{CW}^R \| y_{CW}^L \| y_{CW}^R$

$s_b^L \| t_b^L \| y_b^L \parallel s_b^R \| t_b^R \| y_b^R \leftarrow G(s_b^{(i-1)})$

if $\alpha^{(i)} = 0$:

$s_b^{(i)} = s_b^L \oplus t_b^{(i-1)} \cdot s_{CW}$

$t_b^{(i)} = t_b^L \oplus t_b^{(i-1)} \cdot t_{CW}^L$

if $\alpha^{(i)} = 1$:

$s_b^{(i)} = s_b^R \oplus t_b^{(i-1)} \cdot s_{CW}$

$t_b^{(i)} = t_b^R \oplus t_b^{(i-1)} \cdot t_{CW}^R$

Assign $t_{1-b}^{(j-1)} = 1 - t_b^{(j-1)}$

for $i = j$ **to** u **do**

Let $s_b^L \| t_b^L \| y_b^L \parallel s_b^R \| t_b^R \| y_b^R \leftarrow G(s_b^{(i-1)})$

if $i = j$ **then**

$s_{1-b}^L \| t_{1-b}^L \| y_{1-b}^L \parallel s_{1-b}^R \| t_{1-b}^R \| y_{1-b}^R = r$ (prg challenge)

else

$s_{1-b}^L \| t_{1-b}^L \| y_{1-b}^L \parallel s_{1-b}^R \| t_{1-b}^R \| y_{1-b}^R \leftarrow G(s_{1-b}^{(i-1)})$

if $\alpha^{(i)} = 0$ **then**

$s_{CW} = s_0^R \oplus s_1^R$

$t_{CW}^L = t_0^L \oplus t_1^L \oplus 1$ and $t_{CW}^R = t_0^R \oplus t_1^R$

$y_{CW}^L = y_0^L \oplus y_1^L$ and $y_{CW}^R = y_0^R \oplus y_1^R$

$CW^{(i)} = s_{CW} \| t_{CW}^L \| t_{CW}^R \| y_{CW}^L \| y_{CW}^R$

$s_b^{(i)} = s_b^L \oplus t_b^{(i-1)} \cdot s_{CW}$ for $b = 0, 1$

$t_b^{(i)} = t_b^L \oplus t_b^{(i-1)} \cdot t_{CW}^L$ for $b = 0, 1$

else

$s_{CW} = s_0^L \oplus s_1^L$

$t_{CW}^L = t_0^L \oplus t_1^L$ and $t_{CW}^R = t_0^R \oplus t_1^R \oplus 1$

$y_{CW}^L = y_0^L \oplus y_1^L \oplus 1$ and $y_{CW}^R = y_0^R \oplus y_1^R$

$CW^{(i)} = s_{CW} \| t_{CW}^L \| t_{CW}^R \| y_{CW}^L \| y_{CW}^R$

$s_b^{(i)} = s_b^R \oplus t_b^{(i-1)} \cdot s_{CW}$ for $b = 0, 1$

$t_b^{(i)} = t_b^R \oplus t_b^{(i-1)} \cdot t_{CW}^R$ for $b = 0, 1$

return $k_b = s_b^{(0)} \| CW^{(1)} \| CW^{(2)} \| \dots \| CW^{(u)}$

Figure 15: PRG adversary \mathcal{B} that can distinguish between r that is either sampled randomly $r \leftarrow \{0, 1\}^{2(\kappa+2)}$ or the output of a PRG function $s \leftarrow \{0, 1\}^\kappa$ and $r = G(s)$ from a hybrid distinguishing adversary \mathcal{A} with advantage ϵ for any choice of $j \in [u]$, $b \in \{0, 1\}$, $\alpha \in \{0, 1\}^u$.

□

B.3 Proof of Theorem 9

Correctness. Given $y_0 \leftarrow \text{Eval}(0, k_0, \vec{x})$, $y_1 \leftarrow \text{Eval}(1, k_1, \vec{x})$ we need to show that $y_0 \oplus y_1 = 0^d$ for yes-instances and $y_0 \oplus y_1 \neq 0^d$ for no-instances. According to the scheme, $y_0 \oplus y_1 = y_0^{(1)} \oplus y_1^{(1)} \parallel \dots \parallel y_0^{(d)} \oplus y_1^{(d)}$, where for $i \in [d]$: $y_0^{(i)}, y_1^{(i)}$ are result of evaluation of $(u, 1, 1)$ -ibDCF scheme on keys $k_0^{(i)}, k_1^{(i)}$ respectively.

- For every input $\vec{x} = (x_1, \dots, x_d)$ for which $\text{PreS}_{x_1} \subseteq S_1, \dots, \text{PreS}_{x_d} \subseteq S_d$, $y_0 \oplus y_1 = 0^d$ since $y_0^{(i)} \oplus y_1^{(i)} = 0$ for every $i \in [d]$.
- For every input $\vec{x} = (x_1, \dots, x_d)$ for which $\text{PreS}_{x_1} \not\subseteq S_1$ or \dots or $\text{PreS}_{x_d} \not\subseteq S_d$, $y_0 \oplus y_1 \neq 0^d$ since $y_0^{(i)} \oplus y_1^{(i)} \neq 0$ for at least one index $i \in [d]$.

Privacy. To show that the scheme is secure, the simulator must generate a key k_{idX} that is indistinguishable from output of the Share^* algorithm. For each of the d dimensions, the simulator invokes the simulator $k^{(i)} \leftarrow \text{Sim}(1^\kappa, \text{idX})$ of secure $(u, 1, 1)$ -ibDCF scheme, concatenates outputs and returns the key $k_{\text{idX}} = (k^{(1)}, \dots, k^{(d)})$.

B.4 Proof of Theorem 10

Simulator. Simply forwards the output from ideal functionality to Bob.

Correctness. For every input $\vec{y} \in S_A \cap S_B$, there is exactly one set $\vec{S}_A^{(i)}$, such that $\vec{y} \in \vec{S}_A^{(i)}$, since Alice's input is a disjoint union of sets. This set is uniquely mapped to a single origin $\vec{o}^{(i)}$, by construction. Input prefix \vec{y}^j matches with exactly one critical prefix \vec{x}^j , since $\vec{S}_A^{(i)}.\text{Shift}(\vec{o}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{\vec{x}^j}$ is the disjoint union of prefix sets. Therefore, $X \cap Y = \{(\vec{x}^j \parallel \vec{o}^{(i)})\}$, $|X \cap Y| = 1$ (where Y corresponds to input \vec{y}) and Bob is guaranteed to include \vec{y} in the output.

For every $\vec{y} \notin S_A \cap S_B$, the input is outside every set $\vec{S}_A^{(i)}$ in Alice's input. For every $i \in [N_A], j \in [w']$, $\vec{S}_A^{(i)}.\text{Shift}(\vec{o}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{\vec{x}^j}$, Alice's prefix \vec{x}^j is not a prefix of \vec{y} . Therefore, $X \cap Y = \emptyset$ and Bob will not include \vec{y} in the intersection.

B.5 Proof of Theorem 11

We sketch a proof for the case when Alice learns the output. Analogous arguments hold for the case when Bob learns the output.

Simulator. Simply forwards the output from ideal functionality to Alice.

Correctness. First, for any input $\vec{y} \in S_A \cap S_B$, there is exactly one set $\vec{S}_A^{(i)}$ such that $\vec{y} \in \vec{S}_A^{(i)}$, since Alice's input is a disjoint union of sets. This set is uniquely mapped to a single origin $\vec{o}^{(i)}$, by construction. Some prefix \vec{y}^j of \vec{y} matches with exactly one prefix \vec{x}^j , since $\vec{S}_A^{(i)}.\text{Shift}(\vec{o}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{\vec{x}^j}$ is the disjoint union of prefix sets. Therefore, $X \cap Y' = \{(\vec{x}^j \parallel \vec{o}^{(i)})\}$, $|X \cap Y'| = 1$ (where Y' corresponds to input \vec{y}), any input in Bob's set matches with *exactly one* prefix in X .

Second, for every set $\vec{S}_A^{(i)}$, mapped to origin $\vec{o}^{(i)}$, with prefix sets $\vec{S}_A^{(i)}.\text{Shift}(\vec{o}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{\vec{x}^j}$, the ideal functionality $\mathcal{F}_{\text{PSI-Sum}}$ counts the number of times $(\vec{x}^j \parallel \vec{o}^{(i)})$ value appears in Bob's multi-set Y_v (each count is associated with a distinct output in Bob's set). Thus, Alice accounts for every matching input in Bob's set.

For every $\vec{y} \notin S_A \cap S_B$, the input is outside every set $\vec{S}_A^{(i)}$ in Alice's input. For every $i \in [N_A], j \in [w'], \vec{S}_A^{(i)} \cdot \text{Shift}(\vec{\sigma}^{(i)}) = \dot{\bigcup}_{j \in [w']} \text{PreS}_{x_j^{\vec{y}}}$, Alice's prefix $x_j^{\vec{y}}$ is not a prefix of \vec{y} . Therefore, $X \cap Y' = \emptyset$ and input \vec{y} is not counted in the output from $\mathcal{F}_{\text{PSI-Sum}}$.