

Tokenised Multi-client Provisioning for Dynamic Searchable Encryption with Forward and Backward Privacy

Arnab Bag^{*1}, Sikhar Patranabis^{†2}, and Debdeep Mukhopadhyay^{‡3}

¹Indian Institute of Technology Kharagpur, India

²IBM Research India

³Indian Institute of Technology Kharagpur, India

Abstract

Searchable Symmetric Encryption (SSE) has opened up an attractive avenue for privacy-preserved processing of outsourced data on the untrusted cloud infrastructure. SSE aims to support efficient Boolean query processing with optimal storage and search overhead over large real databases. However, current constructions in the literature lack the support for multi-client search and dynamic updates to the encrypted databases, which are essential requirements for the widespread deployment of SSE on real cloud infrastructures. Trivially extending a state-of-the-art single client dynamic construction, such as ODXT (Patranabis et al., NDSS'21), incurs significant leakage that renders such extension insecure in practice. Currently, no SSE construction in the literature offers efficient multi-client query processing and search with dynamic updates over large real databases while maintaining a benign leakage profile.

This work presents the first dynamic multi-client SSE scheme NOMOS supporting efficient multi-client conjunctive Boolean queries over an encrypted database. Precisely, NOMOS is a multi-reader-single-writer construction that allows only the gate-keeper (or the data-owner) - a trusted entity in the NOMOS framework, to update the encrypted database stored on the adversarial server. NOMOS achieves forward and type-II backward privacy of dynamic SSE constructions while incurring lesser leakage than the trivial extension of ODXT to a multi-client setting. Furthermore, our construction is practically efficient and scalable - attaining linear encrypted storage and sublinear search overhead for conjunctive Boolean queries. We provide an experimental evaluation of software implementation over an extensive real dataset containing millions of records. The results show that NOMOS performance is comparable to the state-of-the-art static conjunctive SSE schemes in practice.

*Email: amiarnabbolchi@gmail.com

†Email: sikharpatranabis@gmail.com

‡Email: debdeep.mukhopadhyay@gmail.com

1 Introduction

Recent advancements in cloud computing have fuelled the development of privacy-preserved processing of sensitive data on third-party cloud servers. Outsourced processing and storing of users' data are becoming standard practices for individuals and organisations. Presently, cloud infrastructures are responsible for handling users' private data obtained from devices/systems used by ordinary citizens, government and industrial establishments. For extended functionalities, the cloud service providers often delegate access to users' data to third-party entities. The involvement of the cloud service providers and other third-party entities - all of whom are considered untrusted, raises deep concern about users' data confidentiality and information privacy. Furthermore, modern cloud applications serve multiple clients, and the data stored on the cloud is frequently updated. In this context, straightforward encryption that provides high confidentiality of data trivially loses the ability to process information in the encrypted form, defeating the advantage of using the cloud.

Several elegant cryptographic primitives such as Fully Homomorphic Encryption (FHE) [Gen09, GSW13], Functional Encryption (FE) [BBG05], Oblivious RAM (ORAM) [Gol87, GO96] and Private Information Retrieval (PIR) [CGKS95, OSI07, BI01] allow implementing diverse functionalities over encrypted outsourced data. However, all of these approaches either incur prohibitively heavy computation/storage overhead or require extremely high communication bandwidth for real applications.

In contrast, SSE offers theoretically robust and implementation-efficient constructions for encrypted data processing, especially searching, while leaking only benign information to untrusted parties. The benign leakage in SSE is formally quantified using precise leakage functions that capture the information leaked. We elaborate more on the general SSE setting and construction below.

1.1 Searchable Symmetric Encryption

Searchable Symmetric Encryption (SSE) [SWP00, Goh03, CM05, CGKO06, CK10, KPR12, KP13, CJJ⁺13, CJJ⁺14, FJK⁺15, KM17, Bos16, BMO17a, EKPE18, LPS⁺18, CPPJ18, PM21] allows querying an encrypted database privately without decryption. In contrast to the abovementioned approaches, SSE offers efficient and low-bandwidth constructions for encrypted data processing, especially searching, while leaking only benign information to untrusted parties. Fundamentally, an SSE scheme offers the following capabilities.

- Allow an (or many) entity (potentially untrusted client) to efficiently search encrypted queries over an encrypted database stored on the cloud (untrusted) without revealing the result to the server.
- Minimise the leakage during query (or update) such that the untrusted entities receive only benign information.

Typical benign leakages include crude statistical information related to the database elements, the query, or the result of the query - but do not include the actual associated (encrypted) data. The database size, *query pattern* (the set of queries corresponding to the same keyword), and the *access pattern* (the set of database records matching a given query) are a few such leakages typically studied in the context of SSE. We present formal syntax of SSE with elaborate details in Section 2, and study these leakages in Appendix C and E.

Dynamic SSE. A dynamic SSE construction [CM05, KPR12, KP13, CJJ⁺14, Bos16, BMO17a, EKPE18] allows dynamic updates (adding or deleting records) to the encrypted database offloaded to the cloud by the client. In contrast, static constructions do not allow updates to the database once it is encrypted. The update capability of dynamic constructions implies two security notions - *forward* privacy and *backward* privacy. Informally, the forward privacy notion states that a current search operation can not be linked to a future update operation, and backward privacy dictates that an insertion operation followed by a deletion does not reveal any information in a future search operation. These two notions are essential for dynamic schemes to prevent a certain class of attacks, specifically, the file injection attack [ZKP16].

Multi-client SSE. A multi-client SSE scheme allows multiple clients to search (or update) the encrypted database. SSE schemes can be classified in the following way according to the different entities involved in the setting.

Single-Reader-Single-Writer (SRSW): Single-reader–single-writer or SRSW setting has a single client and an untrusted cloud server. The single client also acts as the data owner who has permission to update the encrypted database on the server. SRSW constructions [SWP00, Goh03, CM05, CGKO06, CK10, KPR12, KP13, CJJ⁺13, CJJ⁺14, FJK⁺15, KM17, LPS⁺18] have been extensively studied in the literature. A number of dynamic SRSW schemes [CM05, KPR12, KP13, CJJ⁺14, Bos16, BMO17a, EKPE18] have been proposed in recent years and ODXT by Patranabis et al. [PM21] is the only state-of-the-art dynamic scheme with conjunctive query support.

Multi-Reader-Single-Writer (MRSW): In the multi-reader–single-writer or MRSW setting, multiple clients can interact with the untrusted server to search (with individual trapdoors), and a single data owner can generate or update the encrypted database on the untrusted server.

Multi-Reader-Multi-Writer (MRMW): Multi-reader-multi-writer or MRMW is the most generic setting allowing multiple clients to search and update the encrypted database on the cloud server.

Ideally, cloud applications require multi-client access and dynamic update capability to cater to applications over diverse databases involving a potentially large number of users (or clients). Unfortunately, the literature on practically efficient MRSW and MRMW SSE schemes is sparse and entirely restricted to the static setting [BHJP14, JJK⁺13], which prompts us to raise the following question.

Is there a dynamic MRSW/MRMW SSE scheme with strong forward and backward privacy, linear storage requirements and sublinear search overheads – all collectively?

As it turns out, the answer is *no*. ODXT supports dynamic updates and conjunctive queries with sublinear search overhead and linear encrypted storage. However, it is an SRSW construction without support for multiple clients. Furthermore, ODXT is vulnerable to a particular leakage originating from the cross-terms in a conjunctive query (discussed later) that can lead to complete query recovery. Trivially extending ODXT to the multi-client setting by delegating the search token generation phase from multiple clients (assumed to be semi-honest) to the data-owner (a trusted party) retains this leakage. We outline an attack process based on this leakage in Section 3 that leads to complete recovery of the cross-terms. In brief, the existing SSE literature lacks dynamic SSE schemes in MRSW and MRMW settings, which hinders the widespread adoption of SSE in encrypted processing tasks on the cloud.

This work aims to bridge this gap between secure and practically efficient SSE constructions and real multi-client cloud applications. We summarise our goals with the following question.

Can we design an efficient dynamic SSE scheme with forward and backward privacy in the MRSW setting?

We show in this paper that it is possible to design such a scheme, and we present NOMOS construction that achieves the aforementioned practical design and security goals. The following subsection lists the primary contributions of this work. We emphasise that NOMOS is a dynamic multi-keyword construction in the MRSW setting with forward and type-II backward privacy [BMO17a], which is a stepping stone towards building “ideal” MRMW SSE constructions. Extending NOMOS to the MRMW setting is of independent interest requiring separate in-depth exploration, and we leave this as a future work.

1.2 Our Contributions

We summarise our main contributions of this work with brief overview below.

❶ **Multi-client SSE.** We present the first multi-client dynamically updatable SSE construction NOMOS for outsourced encrypted databases. NOMOS supports efficient multi-keyword conjunctive Boolean queries in the MRSW setting, which is essential for practical cloud applications. To the best of our knowledge, this is the first scheme in the literature that can process conjunctive queries from multiple clients with dynamic updates. Clients in NOMOS obtain search tokens from a trusted

entity called gate-keeper, which is allowed to update the encrypted database and holds the keys for token generation. The clients use the search tokens obtained from gate-keeper to query over the encrypted database on the cloud server. We use Oblivious Pseudo-random Function (OPRF)-based mechanism to delegate the search token generation process to the gate-keeper; thus bypassing the need to share the secret keys for token generation among multiple potentially semi-honest clients.

② **Leakage mitigation.** The NOMOS construction mitigates a particular leakage originating from the cross-terms of conjunctive queries. This leakage is inherently present in the state-of-the-art ODXT scheme, and we discuss an attack flow that shows that the trivial extension of ODXT to the multi-client setting remains vulnerable to this leakage that potentially breaks the scheme. NOMOS avoids this specific cross-term-based leakage exploitable from XSet accesses by introducing decorrelated XSet access pattern. We use a variant of Bloom Filter (BF), denoted as the Redundant Bloom Filter (RBF), for decorrelating the repeated memory accesses (for XSet implemented using RBF) to mitigate the cross-term leakage. Using RBF has a minimal impact on the storage, communication and computation overhead of NOMOS compared to non-BF and plain BF versions of our construction.

③ **Tokenised search.** NOMOS allows each client to obtain search tokens from the *gate-keeper* (the data owner holding the keys for token generation) individually and engage in the search protocol with the server to retrieve the query result. The token generation process and search phase work asynchronously, although the search phase requires the tokens to be generated first through the token generation protocol. This tokenised search process for multiple clients allows us to avoid a three-party search protocol involving a client, the gate-keeper, and the server, without blocking other clients from invoking the token generation or the search protocol (whichever is available). This is a desirable capability in multi-user cloud applications where requests arrive asynchronously, and the gate-keeper/cloud needs to serve as early as possible, reducing waiting time (for example, assigning doctors appointments based on patient details in medical applications). It can be easily adopted into general cloud search applications where strong access control is necessary, such as office employee records and bank operations, to name a few. The tokenised functionality allows fine-grained user management and enforcing access permissions for each user individually in a multi-client setting.

④ **Security analysis and implementation.** We provide a concrete security analysis of NOMOS using hybrid arguments of indistinguishability framework. NOMOS setting assumes that gate-keeper (or the data owner) is a trusted entity, and the cloud server is an honest-but-curious polynomial-time adversarial entity. The clients are assumed to be semi-honest entities individually; that is, a semi-honest client follows the specified token generation and search protocol but can obtain/share additional information regarding the queries issued or the tokens received. We provide an overview of the NOMOS leakage profile and concrete security analysis in Appendix C and E.

We implemented the NOMOS framework using C++ (natively multi-threaded) with Redis¹ as the database back-end. We used the Enron dataset² to evaluate NOMOS performance, and we report the results in Section 5. The experimental results show that NOMOS achieves linear storage overhead and sublinear search time, comparable to other state-of-the-art conjunctive SSE constructions.

We provide preliminary notations and syntax in Section 2, which we follow throughout the manuscript. We present an attack in Section 3 on the trivial extension of ODXT to the multi-client setting to demonstrate the devastating effect of cross-term leakage on a multi-client SSE construction. We outline the required security notions and challenges of designing multi-client SSE in Section 3.3. We present our main NOMOS construction in Section 4. Finally, the experimental results and related discussion are provided in Section 5, and we end with a concluding remark on our work. We have diverted the formal security analysis to Appendix E.

¹<https://redis.io/>

²<https://www.cs.cmu.edu/~enron>

2 Basic Notations and Syntax

We provide the basic notations and syntax of SSE below, which we follow throughout this paper. For ease of exposition, we assume the database to be a document collection indexed by keywords and unique document identifiers. More precisely, we assume an inverted index of keywords and corresponding document identifiers for a document collection is available as the plain database.

2.1 Basic Notations

Data. We use w to represent a keyword and id to represent a document identifier in a database. We use Δ to denote the dictionary of all w s in a database and N is number of w s in Δ . We represent the plain database as \mathbf{DB} , and $\mathbf{DB}(q)$ represents the set of ids satisfying a conjunctive query q over \mathbf{DB} . Similarly, for a single w , $\mathbf{DB}(w)$ is the set of ids where w appears. The number of ids in $\mathbf{DB}(q)$ is represented as $|\mathbf{DB}(q)|$. For two values (or strings) v_1 and v_2 , $v_1||v_2$ represents the concatenation of v_1 and v_2 . The cardinality of a set S is denoted by $|S|$, and for a string s (or vector), $|s|$ represents the length of s . We represent the sequence $m, m+1, \dots, n-1, n$ using $[m, n]$ and $1, 2, \dots, n$ using $[n]$. Sampling a value v from a distribution ξ is expressed as $v \xleftarrow{\$} \xi$. We denote a negligible function as $\text{negl}()$. We denote the attribute of a w using $I(w)$, which is essentially encoded as an index. More precisely, I can be considered as a list of (w, a) pairs, where $I(w)$ returns the attribute a , where $w \in \Delta$ and a is a valid keyword attribute. We represent a set of valid keyword attribute combinations using \mathcal{P} , from all unique keyword attributes for all w s in Δ . We denote the number of all unique keyword attributes using d . We also assume that during an update, a complete document (containing multiple keywords) is replaced, and in this process, the existing records are deleted and added again with the modified content. In that case, update operations are usually done in batches of deletions followed by additions of multiple (w, id) pairs.

Entities. We use \mathcal{C} to represent a client and $\mathbf{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_t\}$ to represent a set of t clients. We represent a data-owner by the symbol \mathcal{D} . We denote the server using \mathcal{S} and the gate-keeper using \mathcal{G} (explained in Section 4). In a single client setting, \mathcal{D} and \mathcal{G} serve the same purpose, and we use the \mathcal{D} symbol in the context of single client constructions. However, in an MRSW multi-client setting, \mathcal{G} has the additional responsibility of generating search tokens for \mathcal{C} s. Thus, we denote the single entity responsible for database update in SRSW constructions, such as OXT, ODXT, by \mathcal{D} , and we denote the single entity responsible for database update *and search token generation for other clients*, as in MC-ODXT or NOMOS, by \mathcal{G} . We denote a polynomial-time adversary by \mathcal{A} and a simulator using SIM.

2.2 Cryptographic Primitives

We denote a pseudo-random function (PRF) by $F(K, \cdot)$ and a specific version mapping to \mathbb{F}_p as $F_p(K, \cdot)$. We represent a collision-resistant hash function as CRHF or with the symbol \mathbf{H} , which we assume can be modelled as a random oracle. Additionally, we use an authenticated encryption (AE) [Rog02, BN00] scheme with the routines $\mathbf{AE} = \{\mathbf{AE.Enc}, \mathbf{AE.Dec}\}$ that is IND-CPA and and strongly UF-CMA-secure (unforgeability guarantee) [BN00].

Decisional Diffie-Hellman assumption. Let \mathbb{G} be a cyclic group of prime order q , and let g be any uniformly sampled generator for \mathbb{G} . The decisional Diffie-Hellman (DDH) assumption is that for all PPT algorithms \mathcal{A} , we have,

$$\left| \Pr[\mathcal{A}(g, g^\alpha, g^\beta, g^{\alpha \cdot \beta}) = 1] - \Pr[\mathcal{A}(g, g^\alpha, g^\beta, g^\gamma) = 1] \right| \leq \text{negl}(\lambda),$$

where $\alpha, \beta, \gamma \xleftarrow{\$} \mathbb{Z}_p^*$.

Oblivious pseudo-random function. Oblivious Pseudo-random Function (OPRF) is a cryptographic primitive that allows two parties to jointly evaluate a PRF where party A provides the input plaintext x and party B inputs the key K . At the end of the protocol, A receives the output c , which is indistinguishable from a regular PRF evaluation with the same x and K , and party B receives nothing (or error/nothing symbol \perp).

Hashed Diffie-Hellman OPRF. We use a specific instance of OPRF called hashed Diffie-Hellman (DH) OPRF that works as follows - party A provides input x and a randomly sampled value r . A uses

a hash function H which hashes an input to a group element of \mathbb{G} . A uses H to obtain $H(x) \in \mathbb{G}$ and raises $H(x)$ to generate $s = H(x)^r$. A sends s to B, and B raises s by power K to obtain $t = s^K$. B sends back t to A, and A outputs $y = t^{-1}$. We represent this OPRF evaluation as $y = \text{OPRF}(K, x)$. DH-OPRF is used as a core primitive in our construction to allow multi-client search. Please refer [NR04] for more details on DH-OPRF³. Note that the main NOMOS construction can be instantiated with a generic OPRF. We opted for DH-OPRF for the ease of analytical exposition and implementation.

2.3 SSE Syntax and Security Definition

We denote an SSE algorithm simply using **SSE**. We use **EDB** to denote the encrypted database stored on the encrypted server. We use \mathcal{R}_q to denote the set of encrypted records returned upon searching a conjunctive query q - compactly expressed as $\mathcal{R}_q = \mathbf{EDB}(q)$. We denote a conjunctive query as $q = w_1 \wedge \dots \wedge w_n$ where $w_i \in \Delta$. Without loss of generality, we assume w_1 in q has the least frequency of updates. We call w_1 as the special-term or s-term, and w_2, \dots, w_n are denoted as the cross-terms or x-terms.

A dynamic MRSW SSE scheme **SSE** with tokenised multi-client search is defined by the following ensemble of algorithms.

- **SSE.Setup**(1^λ). This is a PPT algorithm run by \mathcal{G} (who is the data-owner) that samples the master security key \mathbf{sk} , the AE key K_M , and initialises the encrypted database **EDB**.
- **SSE.Update**($\mathbf{sk}, \{\mathbf{op}, (w, \mathbf{id})\}; \mathbf{EDB}$). This is a PPT algorithm jointly executed by \mathcal{G} and \mathcal{S} , where \mathcal{G} 's input is secret key \mathbf{sk} , the update record (w, \mathbf{id}) and the type of the update $\mathbf{op} = \{\text{ADD}, \text{DEL}\}$, and \mathcal{S} 's input is **EDB**. At the end of the protocol, **EDB** stored on \mathcal{S} is updated with the new record.
- **SSE.GenToken**(q, \mathbf{sk}). The **GENTOKEN** routine is executed by a client $\mathcal{C}_j \in \mathcal{C}$ and \mathcal{G} , where \mathcal{C}_j 's input is a query q , and \mathcal{G} 's input is the secret key \mathbf{sk} , and at the end of execution, \mathcal{C}_j receives the search token \mathbf{tk} .
- **SSE.Search**($\mathbf{tk}; \mathbf{EDB}$). In this protocol, a client \mathcal{C}_j sends the search token \mathbf{tk} obtained from \mathcal{G} using **GENTOKEN** method to \mathcal{S} , and \mathcal{S} looks up **EDB** to return the set of matched encrypted records \mathcal{R}_q to \mathcal{C}_j .

An SSE scheme is said to be correct if the **SSE.SEARCH** routine returns all the matching encrypted ids from **EDB** for a query q .

Security of MRSW SSE. A dynamic MRSW SSE scheme **SSE** as described above is said to be secure if it follows the security properties stated below.

Security against a semi-honest client. The following leakage function parameterises the security of a dynamic MRSW SSE construction against a semi-honest client.

$$\mathcal{L}_{\mathcal{C}} = \{\mathcal{L}_{\mathcal{C}}^{\text{SETUP}}, \mathcal{L}_{\mathcal{C}}^{\text{UPDATE}}, \mathcal{L}_{\mathcal{C}}^{\text{GENTOKEN}}, \mathcal{L}_{\mathcal{C}}^{\text{SEARCH}}\}$$

In this ensemble, $\mathcal{L}_{\mathcal{C}}^{\text{SETUP}}$ encapsulates the leakage to a semi-honest client during **SETUP**, $\mathcal{L}_{\mathcal{C}}^{\text{UPDATE}}$ encapsulates the leakage to a semi-honest client during **UPDATE**, $\mathcal{L}_{\mathcal{C}}^{\text{GENTOKEN}}$ encapsulates the leakage to a semi-honest client during **GENTOKEN**, and $\mathcal{L}_{\mathcal{C}}^{\text{SEARCH}}$ encapsulates the leakage to a semi-honest client during **SEARCH**.

A dynamic MRSW SSE scheme Π is said to be secure against a semi-honest PPT adversary \mathcal{A} , who is allowed make $Q = \text{poly}(\lambda)$ queries, with respect to $\mathcal{L}_{\mathcal{C}}$, there exists a polynomial time simulator **SIM**, such that

$$|\Pr[\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}(\lambda) = 1]| \leq \text{negl}(\lambda)$$

where $\mathbf{Real}_{\mathcal{A}}^{\Pi}$ and $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$ are defined in Algorithm 1 and 2 of in Appendix A, respectively.

Security against a semi-honest server. The following leakage function parameterises the security of a dynamic MRSW SSE construction against a semi-honest server.

$$\mathcal{L}_{\mathcal{S}} = \{\mathcal{L}_{\mathcal{S}}^{\text{SETUP}}, \mathcal{L}_{\mathcal{S}}^{\text{UPDATE}}, \mathcal{L}_{\mathcal{S}}^{\text{GENTOKEN}}, \mathcal{L}_{\mathcal{S}}^{\text{SEARCH}}\}$$

³The constructions in this manuscript can be instantiated with a general syntax of OPRF. We opt for the specific instance of DH-OPRF for the ease of exposition with the existing construction structure of ODXT [PM21] and OSPIR-OXT [JJK⁺13].

In this ensemble, $\mathcal{L}_S^{\text{SETUP}}$ encapsulates the leakage to a semi-honest server during SETUP, $\mathcal{L}_S^{\text{UPDATE}}$ encapsulates the leakage to a semi-honest server during UPDATE, $\mathcal{L}_S^{\text{GENTOKEN}}$ encapsulates the leakage to a semi-honest server during GENTOKEN, and $\mathcal{L}_S^{\text{SEARCH}}$ encapsulates the leakage to a semi-honest server during SEARCH.

A dynamic MRSW SSE scheme Π is said to be secure against a semi-honest PPT adversary \mathcal{A} , who is allowed make $Q = \text{poly}(\lambda)$ queries, with respect to \mathcal{L}_S , there exists a polynomial time simulator SIM, such that

$$|Pr[\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda) = 1] - Pr[\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}(\lambda) = 1]| \leq \text{negl}(\lambda)$$

where $\mathbf{Real}_{\mathcal{A}}^{\Pi}$ and $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$ are defined in Algorithm 3 and 4 in Appendix A, respectively.

SSE data structures. SSE constructions heavily rely on the underlying data structures to store and efficiently search over encrypted data. We consider two widely used SSE-specific data structures in this work, namely TSet and XSet. We assume that **EDB** comprises of both TSet and XSet (as required by the construction discussed later in Section 3 and 4).

TSet. TSet is an encrypted variant of a multi-map data structure that stores the encrypted database in a structured form. Fundamentally, TSet stores and accesses data elements in a uniformly indistinguishable manner that hides the association of an w with respective ids. At a high level, TSet follows the typical syntax of a multi-map.

- Insertion: $\text{TSet}[\text{key}] = \text{val}$
- Retrieval: $\text{val} = \text{TSet}[\text{key}]$

The TSet keys are generated through PRFs (and SKE) such that the probability of an \mathcal{A} distinguishing two different w s from randomly accessed (key, val) pairs is negligible.

XSet. XSet is a data structure typically used in multi-keyword SSE schemes to support conjunctive queries (especially in the cross-tag family of constructions [CJJ⁺13, LPS⁺18, PM21]). XSet stores the cross-term-specific information that is used during conjunctive query search. Note that XSet does not store any encrypted information of individual w s or ids; rather, it stores flags or bits associated with cross-terms that are generated using CRHF or PRFs. At a high level, an XSet has the following syntax,

- Insertion: $\text{XSet}[\text{index}] = b, b \in \{0, 1\}$
- Retrieval: $b = \text{XSet}[\text{index}]$

where the index is typically generated from a combined input of w and id to a CRHF, and $b = 1$ indicates that (w, id) is valid pair (that is w appears in document id). For detailed properties and analysis of TSet and XSet, please refer [CJJ⁺13]. We use a slightly different variant of TSet, as adopted in ODXT [PM21].

We choose ODXT as our base construction for developing the multi-client solution as ODXT is the only state-of-the-art dynamic scheme with an efficient update and conjunctive query search. Unfortunately, ODXT itself does not support multi-client search. We first transform ODXT into a multi-client construction MC-ODXT (see Appendix B) following [JJK⁺13]. However, we show that MC-ODXT is vulnerable to cross-term leakage, and the following attack exploits this leakage to break the scheme.

3 Attack on Multi-client SSE Exploiting Cross-term Leakage

We outline an attack on the trivial multi-client extension of ODXT (or MC-ODXT) following the approach in [JJK⁺13]. This attack demonstrates that the presence of the same cross-terms across different queries leads to severe leakage through XSet access pattern, which can completely break MC-ODXT. More precisely, we show that a semi-honest \mathcal{C} colluding with the semi-honest \mathcal{S} can exactly recover the query keywords of a legitimate \mathcal{C} . Provided that enough query instances,

the colluding \mathcal{C} - \mathcal{S} pair can recover the entire keyword dictionary. This glaring vulnerability puts MC-ODXT (or any construction following the same approach) at risk of a severe data leak. Our main construction NOMOS adopts a computationally “lightweight” redundancy-based approach to reduce this leakage while incurring minimal extra storage overhead without affecting the performance.

We briefly summarise the workflow of MC-ODXT algorithms presented in Appendix B to identify the source of the leakage and subsequently discuss the attack exploiting this leakage. The MC-ODXT workflow below involves a client \mathcal{C} obtaining a search token for a conjunctive query $q = w_1 \wedge \dots \wedge w_n$ from the data owner \mathcal{G} and querying the server \mathcal{S} using those search tokens. Note that, in this MRSW setting, we use the \mathcal{G} notation instead \mathcal{D} as it has the additional responsibility of generating \mathcal{C} s search tokens.

MC-ODXT workflow. MC-ODXT follows the ODXT structure with the SETUP, UPDATE, SEARCH routines and an additional routine GENTOKEN for query token generation. The SETUP routine sets up and initialises the parameters, data structures and generates the secret keys. If there is an initial **DB** present, \mathcal{G} repeatedly invokes UPDATE on **DB** entries to update **EDB** stored on \mathcal{S} . The UPDATE algorithm generates a unique TSet address **addr** for each (w, id) pair by appending a counter value with w and evaluating the resulting value through a PRF. Since this counter value is incremented with each update, an **addr** is never repeated (in other words, a unique **addr** is obtained in each update invocation). Furthermore, the update routine treats an ADD or DEL op identically, as there is no conditional execution based on ADD or DEL. \mathcal{S} stores the encrypted id in TSet (which is a part of **EDB**) along with a w -specific deblinding token (α , this is required during search) received from \mathcal{G} . Additionally, \mathcal{G} generates an **xtag** (or cross-tag) by combining w with id (concatenated with **op**) through PRF and raising to the power of g (such that during a search, **xtag** can be recomputed obliviously using the query tokens and α). \mathcal{G} sends **xtag** to the server, and the server sets a bit 1 at address **xtag** in XSet.

Prior to interacting with \mathcal{S} in the SEARCH, a \mathcal{C} obtains a blinded search token from \mathcal{G} for a conjunctive query comprising of two components. The first one constitutes of **strap** (or trapdoor corresponding to the s -term) and **bstags** (or blinded search tags for TSet look-up) corresponding to the s -term w_1 . The second component consists of **bxtraps** (or blinded trapdoors for x -terms) corresponding to the x -terms w_2, \dots, w_n . \mathcal{C} sends these tokens to \mathcal{S} to look up **EDB**. In this process, \mathcal{S} retrieves the encrypted ids using the de-blinded **bstags**. It also computes the deblinded **xtags** for each x -term and retrieved encrypted id (concatenated with **op**) pair. \mathcal{S} checks whether the **xtags** for *all* cross terms and a particular id is set to 1 in XSet. If all **xtag** locations are set, it returns the encrypted id (concatenated with **op**) to \mathcal{C} . \mathcal{C} locally checks if the id has been added for *all* w s in q , and not deleted even for one w of q . If it is present for all w s, it keeps the id in the final result set, otherwise discards it.

Note that the **xtag** computation process is deterministic as it refers to a physical location of a value in the memory (or storage). For validating a (w, id) pair, the same **xtag** address needs to be generated each time the SEARCH protocol encounters the same (w, id) pair. This association is revealed even across different queries having the same keyword issued by multiple clients and leads to the leakage across multiple clients. The following example expounds on this observation for a clearer understanding.

Table 1: An example of SSE execution sequence.

Entity	Time	Operation	Query/Data
\mathcal{C}_1	T1	Search	$w_1 \wedge w_2 \wedge w_3$
\mathcal{G}	T2	Update	(w_3, id)
\mathcal{C}_1	T3	Search	$w_1 \wedge w_3$
\mathcal{C}_2	T4	Search	$w_2 \wedge w_3$

Consider the sequence of MC-ODXT events shown in Table 1. Assume that w_3 was not present in **EDB** during T1. It is inserted into **EDB** at T2 by \mathcal{G} , and queried again (as an x -term) at T3 by \mathcal{C}_1 followed by \mathcal{C}_2 at T4. Observe that these three instances of queries and update involve w_3 . The second and third instances generate the same **xtag** for (w_3, id) pair following the MC-ODXT construction. Since \mathcal{S} is assumed to be semi-honest, it can “see” that the same **xtag**

is accessed in these instances. \mathcal{S} 's ability to observe these distinct accesses for xtags is the base of this attack below.

This two-phase attack assumes a semi-honest \mathcal{C}_i that colludes with \mathcal{S} in the attack process. In the build phase, \mathcal{C}_i legitimately obtains search tokens for its own conjunctive queries and sends those to \mathcal{S} for searching. \mathcal{S} honestly executes the search routine but at the same time records the xtag access from XSet (as a semi-honest entity, it executes the search according to the protocol). Since \mathcal{C}_i colludes with the \mathcal{S} , \mathcal{C}_i provides the server with the exact query ws it sent the search tokens for (without shuffling). \mathcal{S} associates the recorded xtags with received query ws and stored locally for later references. \mathcal{C}_i can repeat this process multiple times to obtain multiple (w, xtag) mappings, and \mathcal{S} can grow the recorded information covering more ws.

While launching the attack on a benign client \mathcal{C}_j , \mathcal{S} compares the xtags generated for the search tokens of \mathcal{C}_j . If the xtags match, \mathcal{S} can infer the corresponding w from the recorded database with high probability. Observe that if colluding \mathcal{C}_i and \mathcal{S} can cover the complete Δ , \mathcal{S} would be able to recover all query keywords of \mathcal{C}_j with complete certainty. We formalise this attack method below.

3.1 Formalising MC-ODXT Attack Process

We denote a semi-honest client using \mathcal{C}_i that colludes with a semi-honest server \mathcal{S} and a benign client using \mathcal{C}_j . We assume that the colluding client \mathcal{C}_i shares the query keywords for which it received the query tokens tk with \mathcal{S} (without shuffling) as well while executing the SEARCH routine as specified. Upon receiving the (w, tk), \mathcal{S} builds a local database XDB that stores records of the form (w, xtag). We assume that \mathcal{C}_i makes t queries during XDB building phase. We summarise the attack process titled CROSSATTACK formally in Algorithm 1 below.

Algorithm 1 Query recovery attack MC-ODXT in presence of colluding semi-honest client and semi-honest server

Input: Query tokens (tk) and ws from semi-honest \mathcal{C}_i , query tokens of benign \mathcal{C}_j

Output: \mathcal{W} : the set of cross-term ws present in \mathcal{C}_j 's query

```

1: function CROSSATTACK
2:   Build Phase
3:   Server
4:   Initialise empty database XDB
5:   Server + Colluding Client
6:   for  $l = 1$  to  $t$  do
7:     Colluding Client
8:     Generate a random query  $q_l \xleftarrow{\$} \Delta^*$   $\triangleright$  Combination of keywords from  $\Delta$ 
9:     Obtain search tokens  $tk_{q_l}$  for  $q_l$  from  $\mathcal{G}$ 
10:    Send  $q_l$  and  $tk_{q_l}$  (without shuffling) to  $\mathcal{S}$ 
11:    Server
12:    Recover xtags using  $tk_{q_l}$  for  $w \in q_l$  available from  $q_l$  sent by  $\mathcal{C}_i$ 
13:    Set  $XDB[xtag_j] = w_j, \forall w_j \in q_l$  received from  $\mathcal{C}_i$ 
14:   Attack Phase
15:   Benign Client
16:   Obtain search token  $tk_q$  for  $q = w_1 \wedge \dots \wedge w_n$  from  $\mathcal{G}$ 
17:   Send  $tk_q$  to  $\mathcal{S}$ 
18:   Server
19:   Compute the xtags from  $tk_q$ 
20:   Look-up XDB using the computed xtags:  $w_i = XDB[xtag_i]$ 
21:   Repeat this for all xtags to recover  $\mathcal{W} = \{w_2, \dots, w_n\}$ 
22:   Return  $\mathcal{W}$ 

```

The CROSSATTACK attack in Algorithm 1 has two phases - the build phase, where the colluding \mathcal{C}_i engages with \mathcal{S} to build the XDB. In the attack phase, \mathcal{S} obtains the xtags corresponding to

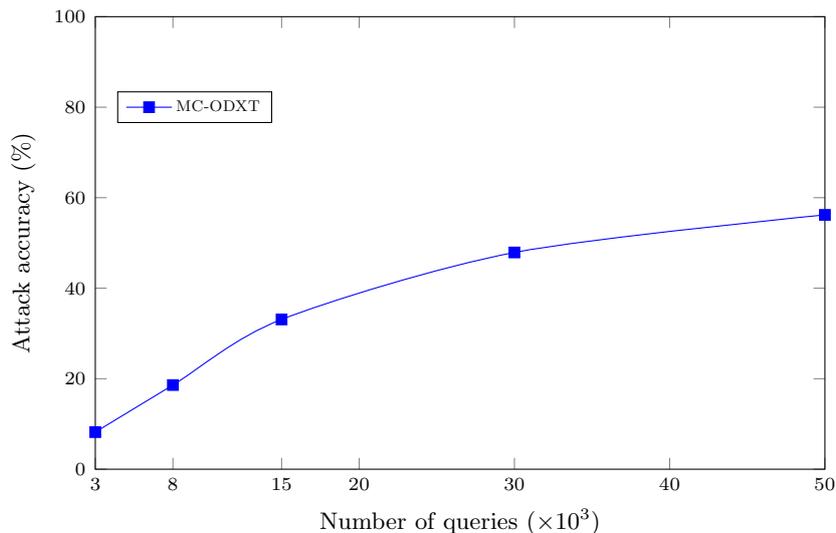


Figure 1: Attack accuracy vs number of queries by the semi-honest client in the building phase. The number of queries by the benign clients was 2000.

benign \mathcal{C}_j 's query q , and looks-up XDB to recover the w s in q . The attack accuracy (probability of exact keyword recovery) improves as more unique w s from Δ are covered in the build phase to populate XDB. Therefore, increasing the number of query iterations t in the build phase leads to higher successful keyword recovery as more w s are covered in XDB. The attack perfectly recovers all cross-terms with probability 1 for the ideal case when XDB contains all w s of Δ .

3.2 Experimental Evaluation

We executed the attack in Algorithm 1 on MC-ODXT to highlight the severity of the leakage discussed above. We used the Enron email dataset for this experiment, and the platform details for this experiment are available in Section 5. The experiment builds XDB from the recorded x tags and the associated w s. Subsequently, in the attack phase, the x tags corresponding to a benign client's queries are recorded and looked up in XDB for successful w recovery. The attack accuracy, defined as the ratio of the number of correct x -term looked up to the total number of x -terms looked up, is plotted in Figure 1 against the number of benign clients queries t in the building phase. The attack accuracy improves as the number of queries in the building phase increases, allowing the \mathcal{S} to cover more keywords in XDB. In this evaluation, the CrossAttack of Algorithm 1 successfully recovered more than 50% of the x -terms of benign clients' queries.

3.3 Challenges in Designing Multi-client SSE

Developing a multi-client SSE construction (MRSW or MRMW) poses several challenges as a multi-client-specific workflow fundamentally differs from an SRSW construction. As illustrated by the attack above, trivial extensions like MC-ODXT suffer from multi-client specific attack(s), which needs to be suitably addressed without compromising functionality or efficiency. At a high level, the following multi-client-specific privacy notions are necessary for a multi-client construction.

Query privacy. A legitimate client needs to share query information with the data owner, but the data owner must not be able to figure out which keywords are being queried.

Preventing token forgery. An adversarial client must not be able to modify or reuse received query tokens with previously obtained (or future) query tokens⁴.

⁴Note that, token forgery implies the presence of a malicious client; whereas the attack in Algorithm 1 requires only a colluding semi-honest client. Nonetheless, our main NOMOS construction considers token forgery prevention as a necessary feature in the multi-client setting.

Token validation. The server should be able to validate that it received a genuine query token from a client generated by the data owner and not by a third party.

The above privacy requirements are provisioned in MC-ODXT setting through a two-party oblivious computation-based mechanism, more precisely through OPRF. In that case, \mathcal{C} does not have to share the actual q with \mathcal{G} . Instead, \mathcal{C} shares hashed values mapped to group elements with \mathcal{G} . However, in a dynamic construction following the ODXT structure, it needs to generate several **bstags** (corresponding to each update operation of the s-term) and additional deblindings for these tokens, which requires modifications of the approach of [JJK⁺13].

Our final construction NOMOS adopts a similar approach to provision multi-client search, and an AE-based authentication method is incorporated into the token generation process to validate query tokens on the server side. However, the leakage from cross-terms in the multi-client settings poses further challenges that need to be addressed without compromising efficiency and security.

Observe that the leakage mentioned in Section 3 appears due to a (w, id) pair validation requiring a valid physical location look-up in the XSet storage, which is deterministic across multiple queries from different clients. Therefore, \mathcal{S} can record this information and exploit later as demonstrated in CROSSATTACK of Algorithm 1. To prevent this leakage, the XSet look-up access pattern needs to be hidden from \mathcal{S} . In our construction NOMOS, we opt for a redundancy-based mechanism for XSet look-up that produces different access patterns for the same (w, id) pair.

Note. We would like to mention that in this work, we specifically focus on cross-term-based leakage. Since the s-terms have a lower frequency, the chance of s-terms colliding across queries from multiple clients compared to x-terms is low. Thus, this type of leakage is more severe for x-terms, and we prioritise mitigating x-term leakage in this development. We leave devising a similar strategy for s-terms as an important future work.

4 Nomos - Dynamic Multi-client SSE Construction

We start by outlining the setting of our main construction with brief details of each entity and how each entity interacts with other entities. We discuss NOMOS construction in two phases - the first phase describes the multi-client provisioning, and the second phase discusses the cross-term leakage mitigation technique incorporated into NOMOS. The core structure of the basic construction follows from ODXT construction, and we encourage the readers to refer [PM21] for more details.

Clients. We assume there are l clients $\{\mathcal{C}_1, \dots, \mathcal{C}_l\}$ who are allowed to obtain search tokens and search over the database. Each \mathcal{C}_i can request a search token following the GENTOKEN routine and engage in the SEARCH protocol with the server to retrieve query results.

Gate-keeper. We denote the data owner using gate-keeper (denoted by \mathcal{G}) who can update **EDB** and issue search tokens to \mathcal{C} (the gate-keeper name signifies additional responsibility to generate search tokens for \mathcal{C} s). \mathcal{G} holds the secret key (**sk**) to generate the search tokens and update the database. Since NOMOS is an MRSW construction, \mathcal{G} is the only entity allowed to update **EDB** and is considered a trusted party. Note that, although \mathcal{G} is considered a trusted entity who can “see” the data, NOMOS ensures query privacy of clients by not revealing query keywords to \mathcal{G} .

Trust in \mathcal{G} We stress here that although \mathcal{G} is a trusted entity, it does not learn the client’s search result, along with not knowing the query keywords. Thus, \mathcal{G} is trusted for updates only. This trust level in the MRSW setting is aptly applicable to real-life examples, where \mathcal{G} can be considered an administrative entity, such as the government or an employer, who can update the database and enforce search policies on the clients accessing the data.

Server. The server (denoted by \mathcal{S}) stores the encrypted database **EDB** comprised of TSet and XSet and performs an update or search as requested. \mathcal{S} engages in the UPDATE protocol with \mathcal{G} to update **EDB**. During a search, \mathcal{S} interacts with \mathcal{C} to receive the search tokens and performs the database look-up. At the end of the SEARCH protocol, it returns the retrieved encrypted ids to \mathcal{C} matching the actual query.

We elucidate the MRSW setting here for better clarity and understanding. Consider a genome analysis service provider offering a cloud-based gnome referencing service for clients who can query

Algorithm 2 NOMOS SETUP

Input: Security parameters λ **Output:** sk , UpdateCnt, and **EDB**

- 1: **function** NOMOS.SETUP
 - 2: Gate-keeper
 - 3: Sample a uniformly random key K_S from \mathbb{Z}_p^* for OPRF
 - 3: Sample two sets of uniformly random keys $K_T = \{K_T^1, \dots, K_T^d\}$ and $K_X = \{K_X^1, \dots, K_X^d\}$ from $(\mathbb{Z}_p^*)^d$ for OPRF
 - 4: Sample uniformly random key K_Y from $\{0, 1\}^\lambda$ for PRF F_p
 - 5: Sample uniformly random key K_M from $\{0, 1\}^\lambda$ for AE
 - 6: Initialise UpdateCnt, TSet, XSet to empty maps
 - 7: Gate-keeper keeps $sk = (K_S, K_T, K_X, K_Y)$; UpdateCnt is disclosed to clients when required, and K_M is shared between gate-keeper and the server
 - 8: Set **EDB** = (TSet, XSet)
 - 9: Send **EDB** to server
-

over standard genome data provided by the analysis service. Following the current cloud-based service trend, assume that the genome analysis service has outsourced the infrastructure support to a third-party cloud storage and computing resource provider, such as Amazon AWS. In this example, the genome service provider can be modelled as \mathcal{G} , the cloud server acts as \mathcal{S} , and the clients requesting for look-up can be considered as \mathcal{C} s. Naturally, genome data is considered sensitive private information, which needs to be stored on \mathcal{S} in encrypted form, and a \mathcal{C} needs to query the data for reference information. Such applications can be aptly handled by a scalable MRSW SSE scheme like NOMOS.

NOMOS setting assumes \mathcal{S} as a semi-honest party and treats each client as a semi-honest party individually (who can potentially collude with \mathcal{S} while following the NOMOS description as specified). \mathcal{G} is a trusted entity that can update the encrypted database and generate search tokens for \mathcal{C} 's queries (as \mathcal{G} holds the secret keys). The multi-client provisioning in NOMOS follows the approach [JJK⁺13] and incorporates modifications for dynamic updates. Note that a decentralised trust in \mathcal{G} is ideal for the MRMW setting, where multiple parties update the encrypted database. In contrast, the MRSW setting allows only one party to update the database - typically the data-owner. Thus, the MRSW setting allows provisioning multi-client search using basic OPRF primitive rather than relying on independent system-oriented approaches like trusted execution environments or advanced primitives like MPC, as it handles updates only from the data-owner, unlike the MRMW setting. This model is adopted in existing works like OSPiR-OXT [JJK⁺13], which we have followed here.

4.1 Setup and Update

The NOMOS SETUP routine of Algorithm 2 is executed by \mathcal{G} that initialises the system (including **EDB** on \mathcal{S}). Subsequently, \mathcal{G} and \mathcal{S} can jointly execute NOMOS UPDATE of Algorithm 3 repeatedly on the data records from **DB**.

Update process. The UPDATE algorithm is invoked by \mathcal{G} with (w, id) and op as input. \mathcal{S} receives the encrypted values, along with tags generated by \mathcal{G} and updates the TSet and XSet. The UPDATE process of Algorithm 3 adopts the update routine of ODXT with modifications to support multi-client search. The modifications include the way the TSet and XSet addresses (stags and xtags) are generated, such that in the SEARCH routine, the same addresses can be recomputed from search tokens obtained via OPRF evaluations. Note that, Algorithm 3 incorporates the RBF's (discussed in Section 4.2) redundant xtag generation process. The changes in NOMOS algorithm(s) from MC-ODXT are highlighted in red.

4.2 Mitigating Cross-Term Leakage

Recall from Section 3 that the cross-term leakage arises from repeated xtag accesses (translated to memory location accesses) by \mathcal{S} for the same (w, id) combinations from different queries (and

Algorithm 3 NOMOS UPDATE

Input: $K_T = \{K_T^1, \dots, K_T^d\}$, $K_X = \{K_X^1, \dots, K_X^d\}$, accessed as $K_T[I(\mathbf{w})]$ and $K_X[I(\mathbf{w})]$ for attribute $I(\mathbf{w})$ of \mathbf{w} , ℓ is the number of hash functions for insertion into RBF, (\mathbf{w}, id) pair to be updated, update operation op

Output: Updated **EDB**

```
1: function NOMOS.UPDATE
   Gate-keeper
2:   Parse  $\text{sk} = (K_T, K_X, K_Y)$  and  $\text{UpdateCnt}$ 
3:   Set  $K_Z \leftarrow F((H(\mathbf{w}))^{K_S}, 1)$ 
4:   If  $\text{UpdateCnt}[\mathbf{w}]$  is NULL then set  $\text{UpdateCnt}[\mathbf{w}] = 0$ 
5:   Set  $\text{UpdateCnt}[\mathbf{w}] = \text{UpdateCnt}[\mathbf{w}] + 1$ 
6:   Set  $\text{addr} = (H(\mathbf{w} \parallel \text{UpdateCnt}[\mathbf{w}] \parallel 0))^{K_T[I(\mathbf{w})]}$ 
7:   Set  $\text{val} = (\text{id} \parallel \text{op}) \oplus (H(\mathbf{w} \parallel \text{UpdateCnt}[\mathbf{w}] \parallel 1))^{K_T[I(\mathbf{w})]}$ 
8:   Set  $\alpha = F_p(K_Y, \text{id} \parallel \text{op}) \cdot (F_p(K_Z, \mathbf{w} \parallel \text{UpdateCnt}[\mathbf{w}])^{-1})$ 
9:   Set  $\text{xtag}_i = H(\mathbf{w})^{K_X[I(\mathbf{w})] \cdot F_p(K_Y, \text{id} \parallel \text{op}) \cdot i}$ , where  $i \in [\ell]$ 
10:  Send  $(\text{addr}, \text{val}, \alpha, \{\text{xtag}_i\}_{i \in [\ell]})$  to server
   Server
11:  Parse EDB = (TSet, XSet)
12:  Set  $\text{TSet}[\text{addr}] = (\text{val}, \alpha)$ 
13:  Set  $\text{XSet}[\text{xtag}_i] = 1$ , for  $i \in [\ell]$ 
```

updates). Intuitively, to mitigate this leakage, these memory accesses (to the same address for a particular (\mathbf{w}, id) pair) need to be different for each access without affecting the look-up performance severely. We adopt a simple yet effective way to achieve this through redundant location accesses, where multiple “copies” of the XSet bit value are stored at multiple addresses. A random subset of these locations is looked up in each subsequent access for the same (\mathbf{w}, id) pair.

Randomising XSet access. We opt for a Bloom filter (BF) (which physically stores XSet) based solution to achieve this redundant look-up. At a high level, a BF uses k different hash functions to generate k distinct addresses for an element look-up. However, straightforwardly plugging in BF into MC-ODXT does not hide the repeated access pattern as k addresses for a particular (\mathbf{w}, id) pair are still generated from the same xtag . We modify the BF structure slightly in the following way. Instead of using k hash functions to generate the BF addresses for look-up, we use ℓ hash functions to generate the BF addresses for an input element, where $\ell > k$. During a search, instead of using all ℓ hash functions to generate the BF addresses, a subset of k hash functions out of the ℓ are chosen randomly to generate the BF addresses. Observe that, with this modification, \mathcal{S} receives a different set of BF addresses for each repeated access of a particular (\mathbf{w}, id) pair and hence can not correlate among previously accessed entries. We call this BF variant Redundant Bloom Filter (RBF), and we present elaborate details and analysis of RBF in Appendix D.

Avoiding two rounds. Note that incorporating RBF as a module into NOMOS would incur a two-round solution as the RBF addresses need to be generated from xtags . The xtags must not be revealed to \mathcal{S} , and hence need to be generated on the \mathcal{C} 's side. This is undesirable in a multi-client setting due to communication/computation overhead and increased leakage from additional token exchanges. We avoid this by embedding the RBF address generation phase into the UPDATE and GENTOKEN algorithms in the following way.

$$\begin{aligned} \mathcal{G}(\text{UPDATE}) & : \text{xtag}_i = H(\mathbf{w})^{K_X[I(\mathbf{w})] \cdot F_p(K_Y, \text{id} \parallel \text{op}) \cdot i}, i \in [\ell] \\ \mathcal{G}(\text{GENTOKEN}) & : \overline{\text{bxtrap}'_j} = \overline{\text{bxtrap}'_j} \cup \{(\text{bxtrap}'_j)^{\beta_i}\} \end{aligned}$$

The revised final UPDATE and GENTOKEN algorithms are presented in Algorithm 3 and 4, respectively. The SEARCH routine is modified to compute final k addresses for RBF and is presented in Algorithm 5 in Section 4.3.

Since UPDATE protocol should be executed in batches of multiple deletions and additions involving several (\mathbf{w}, id) pairs (a realistic assumption stated in Section 2), several XSet addresses are generated for inserting multiple (\mathbf{w}, id) pair records into RBF-based XSet. The generated XSet

Algorithm 4 NOMOS GENTOKEN

Input: $q = \{w_1, \dots, w_n\}$ (n keywords in q), \mathcal{P} is the set of allowed attribute sequences, ℓ and k hashes are used for insertion and query in RBF, respectively

Output: $\text{strap}, \text{bstag}_1, \dots, \text{bstag}_m, \delta_1, \dots, \delta_m, \overline{\text{bxtrap}}_1, \dots, \overline{\text{bxtrap}}_n, \text{env}$

```
1: function NOMOS.GENTOKEN
   Client
2:   Set  $m = \text{UpdateCnt}[w_1]$   $\triangleright$  the update count of the least frequently updated  $w$ 
3:   Sample  $r_1, \dots, r_n \xleftarrow{\$} \mathbb{Z}_p^*$ 
4:   Sample  $s_1, \dots, s_m \xleftarrow{\$} \mathbb{Z}_p^*$ 
5:   Set  $a_j = (H(w_j))^{r_j}$ , for  $j = 1, \dots, n$ 
6:   Set  $b_j = (H(w_1 || j || 0))^{s_j}$ , for  $j = 1, \dots, m$ 
7:   Set  $c_j = (H(w_1 || j || 1))^{s_j}$ , for  $j = 1, \dots, m$ 
8:   Set  $\text{av} = (I(w_1), \dots, I(w_n)) = (I_1, \dots, I_n)$ 
   Gate-keeper
9:   Abort if  $\text{av} \notin \mathcal{P}$ 
10:  Sample  $\rho_1, \dots, \rho_n \xleftarrow{\$} \mathbb{Z}_p^*$ 
11:  Sample  $\gamma_1, \dots, \gamma_m \xleftarrow{\$} \mathbb{Z}_p^*$ 
12:  Set  $\text{strap}' = (a_1)^{K_S}$ 
13:  Set  $\text{bstag}'_j = (b_j)^{K_T[I_1] \cdot \gamma_j}$ , for  $j = 1, \dots, m$ 
14:  Set  $\delta'_j = (c_j)^{K_T[I_1]}$ , for  $j = 1, \dots, m$ 
15:  Set  $\text{bxtrap}'_j = (a_j)^{K_X[I_j] \cdot \rho_j}$  for  $j = 2, \dots, n$ 
16:  Sample random indices for RBF  $\beta_i \xleftarrow{\$} [\ell], i \in [k]$ 
17:  for  $j = 2$  to  $n$  do
18:     $\overline{\text{bxtrap}}'_j = \{\}$ 
19:     $\overline{\text{bxtrap}}'_j = \text{bxtrap}'_j \cup \{(\text{bxtrap}'_j)^{\beta_i}\}$ , for  $\beta_i \in \{\beta_1, \dots, \beta_k\}$ 
20:  Set  $\text{env} = \text{AE.Enc}_{K_M}(\rho_1, \dots, \rho_n, \gamma_1, \dots, \gamma_m)$ 
21:  Send  $(\text{strap}', \text{bstag}'_1, \dots, \text{bstag}'_m, \delta'_1, \dots, \delta'_m, \overline{\text{bxtrap}}'_2, \dots, \overline{\text{bxtrap}}'_n, \text{env})$  to client
   Client
22:  Set  $\text{strap} = (\text{strap}')^{r_1^{-1}}$ 
23:  Set  $\text{bstag}_j = (\text{bstag}'_j)^{s_j^{-1}}$ , for  $j = 1, \dots, m$ 
24:  Set  $\delta_j = (\delta'_j)^{s_j^{-1}}$ , for  $j = 1, \dots, m$ 
25:  for  $j = 2$  to  $n$  do
26:    Initialise  $\overline{\text{bxtrap}}_j = \{\}$ 
27:    for  $l = 1$  to  $k$  do
28:       $\overline{\text{bxtrap}}_j = \overline{\text{bxtrap}}_j \cup \{(\overline{\text{bxtrap}}'_j[l])^{r_j^{-1}}\}$ 
29:  Output  $(\text{strap}, \text{bstag}_1, \dots, \text{bstag}_m, \delta_1, \dots, \delta_m, \overline{\text{bxtrap}}_2, \dots, \overline{\text{bxtrap}}_n, \text{env})$  as search token
```

addresses (for all (w, id) pairs) must be shuffled by \mathcal{G} prior to sending to \mathcal{S} in batches. This random shuffling is necessary to prevent \mathcal{S} from associating any XSet access pattern to a probable keyword. In contrast, such shuffling in MC-ODXT is ineffective against \mathcal{S} associating a probable x-term to an observed (or a set of observed) XSet address as only one single xtag match is required for a particular (w, id) pair update and search.

4.3 Token Generation and Search

The multi-client search process starts with the token generation process outlined in Algorithm 4 - a two-party protocol between a \mathcal{C} and \mathcal{G} . We briefly summarise the workflow of the GENTOKEN method of Algorithm 4 that generates the search tokens while maintaining the query privacy of legitimate clients and preventing token forgery by an adversarial client.

Token generation phase. The NOMOS SEARCH algorithm follows the ODXT [PM21] SEARCH process, which generates two types of search tokens - the stags and the xtraps. The stags are

generated from the s-term (the least-frequently updated term or w_1 in Algorithm 4, without loss of generality), which are used to retrieve encrypted ids through TSet look-up. Whereas xtraps are generated from x-terms which are used to check the validity of a (w, id) pair through XSet look-up. The GENTOKEN routine in NOMOS is responsible for generating these tokens for multiple clients. Since the \mathcal{G} holds the keys (K_T, K_X) as a part of the secret key sk to generate the tokens, \mathcal{C}_j needs to send the query to \mathcal{G} to generate tokens without revealing the actual ws. We resort to an OPRF-based computation, allowing \mathcal{C}_i to send query ws in blinded form. The major difference from OSPIR-OXT [JJK⁺13] is that ODXT generates an **stag** for each update count for the s-term, whereas OSPIR-OXT generates a single **stag** (**strap** in GENTOKEN). This is a direct consequence of the dynamic update capability of ODXT, and GENTOKEN routine computes the blinded exponentiations for each **stag**. Similarly, \mathcal{C}_j also computes the blinded xtraps and the set of keyword attributes $av = \{I_1, \dots, I_n\}$ where $I_i = I(w_i)$ for $i \in [n]$, which are sent to \mathcal{G} .

Blinded tokens and query validation. Upon receiving the search tokens, \mathcal{G} first verifies whether av is a valid set of attributes which \mathcal{C}_i is allowed to query by checking $av \in \mathcal{P}$. If not valid, \mathcal{G} aborts the process. Otherwise, \mathcal{G} computes its own part of the OPRF computation (party B’s computation in OPRF as discussed in Section 2) by processing **bstrap**, **bstag** and **bstrap**. The blinded **strap** (**bstrap**) computation is done through OPRF evaluation using K_S , and blinded **stag** (**bstag**) and **xtrap** (**bstrap**) generation are done through OPRF evaluation using K_T and K_X combined with \mathcal{G} ’s own blinding factors $\{\rho_1, \dots, \rho_m\}$ and $\{\gamma_1, \dots, \gamma_m\}$. \mathcal{G} ’s blinding factors ρ_i ’s and γ_i ’s are necessary to prevent a potentially malicious client from modifying the search tokens by replacing the search tokens. Since the blinding factors are randomly generated for each request, a polynomially bound malicious party can not replicate the blinding factors. \mathcal{S} can verify the tokens as \mathcal{G} encrypts $\{\rho_1, \dots, \rho_m\}$ and $\{\gamma_1, \dots, \gamma_m\}$ using AE that \mathcal{S} can authenticate prior to search (in the SEARCH routine).

In the final phase of GENTOKEN routine, \mathcal{C}_i deblinds the doubly-blinded **bstag**’s, δ ’s and **bstrap**’s using its own blinding factors (r_1, \dots, r_n) and (s_1, \dots, s_m) to obtain the \mathcal{G} -blinded tokens (**strap**, **bstag** and **bstrap**), which \mathcal{C} subsequently uses as the search token. Note that \mathcal{S} receives the AE-encrypted blinding factors from \mathcal{G} as a part of the search token, which are used for token validation and deblinding during SEARCH execution. However, \mathcal{S} itself is not involved in the GENTOKEN protocol. The AE decryption key K_M is generated by \mathcal{G} at SETUP and shared with \mathcal{S} .

Search phase. The NOMOS SEARCH of Algorithm 5 is jointly executed by a \mathcal{C} and \mathcal{S} without any involvement of \mathcal{G} . However, \mathcal{C} must have obtained the search tokens from \mathcal{G} prior to invoking the SEARCH routine. In this phase, the client sends the blinded search tokens and encrypted blinding factors to \mathcal{S} that it received from \mathcal{G} at the end of GENTOKEN (blinded with \mathcal{G} ’s blinding factors). At a high level, the SEARCH protocol proceeds in two stages - first, \mathcal{C} computes the final xtraps from the received bstraps. Note that the resulting xtokens are still blinded as \mathcal{C} does not have \mathcal{G} ’s blinding factors. \mathcal{S} receives the bstags and computed xtokens along with \mathcal{G} ’s AE-encrypted blinding factors. \mathcal{S} validates the AE ciphertext using key K_M and proceeds for decryption if the validation is successful. \mathcal{S} deblinds the received bstags to recover the actual stags, and after that, it follows the usual ODXT search routine to retrieve the matching ids. During xtag computation, \mathcal{S} deblinds xtokens using the decrypted \mathcal{G} ’s blinding factors, and follows the usual ODXT search process⁵.

S-term information to clients. In this multi-client setting we assume that \mathcal{C}_i ’s obtain the s-term frequency information from \mathcal{G} via a suitable privacy-preserving mechanism without revealing the keyword to \mathcal{G} . This a reasonable assumption following from the prior works including OXT [CJJ⁺13], OSPIR-OXT [JJK⁺13], HXT [LPS⁺18], and ODXT [PM21] (as the static constructions do not have update capability, they use keyword frequency instead of the update frequency). We elaborate more on possible privacy-preserving mechanisms for frequency information retrieval (such as based on private information retrieval [CGKS95, OSI07, BI01]) in Appendix B.1.

⁵We emphasise that NOMOS follows the vast majority of the SSE literature (including ODXT itself) in its index-only focus and does not incorporate a dedicated mechanism to handle actual document retrieval. We leave extending NOMOS to full SSE scheme with final encrypted document retrieval as an interesting future work.

Algorithm 5 NOMOS SEARCH

Input: $q, \text{strap}, \text{bstag}_1, \dots, \text{bstag}_m, \delta_1, \dots, \delta_m, \overline{\text{bxtrap}}_1, \dots, \overline{\text{bxtrap}}_n, \beta_1, \dots, \beta_n, \text{env}, \text{UpdateCnt}$ **Output:** IdList

```
1: function NOMOS.SEARCH
   Client
2:   Set  $K_Z \leftarrow F(\text{strap}, 1)$ 
3:    $m = \text{UpdateCnt}[w_1]$ 
4:   Initialise  $\text{stokenList}$  to an empty list
5:   Initialise  $\text{xtokenList}_1, \dots, \text{xtokenList}_m$  to empty lists
6:   for  $j = 1$  to  $m$  do
7:      $\text{stokenList} = \text{stokenList} \cup \{\text{bstag}_j\}$ 
8:     for  $i = 2$  to  $n$  do
9:        $\text{xtokenSet}_{i,j} \leftarrow \{\}$ 
10:      for  $t = 1$  to  $k$  do
11:        Set  $\text{xtoken}_{i,j}^t = \overline{\text{bxtrap}}_i[t]^{F_p(K_Z, w_1 || j)}$ 
12:        Set  $\text{xtokenSet}_{i,j} = \text{xtokenSet}_{i,j} \cup \text{xtoken}_{i,j}^t$ 
13:        Randomly permute the tuple-entries of  $\text{xtokenSet}_{i,j}$ 
14:        Set  $\text{xtokenList}_j = \text{xtokenList}_j \cup \text{xtokenSet}_{i,j}$ 
15:   Send  $(\text{stokenList}, \text{xtokenList}_1, \dots, \text{xtokenList}_m)$ 
   Server
16:   Upon receiving  $\text{env}$  from client, verify  $\text{env}$ ; if verification fails, return  $\perp$ ; otherwise decrypt  $\text{env}$ 
17:   Parse  $\text{EDB} = (\text{TSet}, \text{XSet})$ 
18:   Initialise  $\text{sEOpList}$  to an empty list
19:   for  $j = 1$  to  $\text{stokenList.size}$  do
20:     Set  $\text{cnt}_j = 1$ 
21:     Set  $\text{stag}_j \leftarrow (\text{stokenList}[j])^{1/\gamma_j}$ 
22:     Set  $(\text{sval}_j, \alpha_j) = \text{TSet}[\text{stag}_j]$ 
23:     Initialise  $\text{flag} = 1$ 
24:     for  $i = 2$  to  $n$  do
25:       Set  $\text{xtokenSet}_{i,j} = \text{xtokenList}_j[i]$ 
26:       for  $t = 1$  to  $k$  do
27:         Compute  $\text{xtag}_{i,j} = (\text{xtokenSet}_{i,j}[t])^{\alpha_j/\rho_i}$ 
28:         If  $\text{XSet}[\text{xtag}_{i,j}] = 0$ , then set  $\text{flag} = 0$ 
29:         If  $\text{flag} = 1$ , then set  $\text{cnt}_j = \text{cnt}_j + 1$ 
30:       Set  $\text{sEOpList} = \text{sEOpList} \cup \{(j, \text{sval}_j, \text{cnt}_j)\}$ 
31:   Sent  $\text{sEOpList}$  to client
   Client
32:   Initialise  $\text{IdList}$  to an empty list
33:   for  $\ell = 1$  to  $\text{sEOpList.size}$  do
34:     Let  $(j, \text{sval}_j, \text{cnt}_j) = \text{sEOpList}[\ell]$ 
35:     Recover  $(\text{id}_j || \text{op}_j) = \text{sval}_j \oplus \delta_\ell$ 
36:     If  $\text{op}_j$  is DEL and  $\text{cnt}_j = n$  then set  $\text{IdList} = \text{IdList} \setminus \{\text{id}_j\}$ 
37:   Output  $\text{IdList}$ 
```

\triangleright XSet is implemented using RBF

4.4 Computation and Storage Overhead

The following overhead analysis assumes that a single record in TSet or XSet requires constant storage, and the group operations and storage look-ups are the costliest operations in practice.

Computation overhead. The UPDATE routine executes for a (w, id) pair in each invocation. The UPDATE routine computes the TSet addresses along with w -bound deblinding factor, which requires a total of three hash computations, two group operations and field inversion. However, as we use RBF-based XSet, ℓ xtag computations require ℓ group operations that dominates the UPDATE routine with $O(\ell)$ computation overhead. Since ℓ is a constant (which is significantly

small compared to the number of updates) for a specific setting, this $O(\ell)$ can be asymptotically approximated to $O(1)$ per UPDATE invocation for a series of updates.

The GENTOKEN protocol requires $|q| + 2|\text{UpdateCnt}[w_1]|$ hash computations and group operations to generate the client-side values with blinding. The gate-keeper-side processing involves $|q| + 2|\text{UpdateCnt}[w_1]|$ group operations and $|q| + |\text{UpdateCnt}[w_1]|$ field multiplications. The client-side deblinding phase computes $|q| + 2|\text{UpdateCnt}[w_1]|$ group operations. As a result, GENTOKEN incurs $O(|q| + 2|\text{UpdateCnt}[w_1]|)$ computation overhead asymptotically that is sublinear in the total database size $|\mathbf{DB}|$. The communication overhead is also $O(|q| + 2|\text{UpdateCnt}[w_1]|)$ as \mathcal{C} and \mathcal{G} exchange $|q| + 2|\text{UpdateCnt}[w_1]|$ tokens in this process.

The SEARCH protocol computes $k \cdot |q| \cdot |\text{UpdateCnt}[w_1]|$ group operations to compute the blinded xtokens. \mathcal{S} performs $|\text{UpdateCnt}[w_1]|$ TSet look-ups that require $|\text{UpdateCnt}[w_1]|$ group operations for deblinding. Additionally, \mathcal{S} computes a total $k \cdot |q| \cdot |\text{UpdateCnt}[w_1]|$ XSet addresses for look-up. Therefore, NOMOS SEARCH incurs $O(k \cdot |q| \cdot |\text{UpdateCnt}[w_1]|)$ asymptotic computation overhead with all combined. Since k is a small constant, the SEARCH overhead is sublinear in the total database size $|\mathbf{DB}|$. Furthermore, \mathcal{C} needs to send $O(k \cdot |q| \cdot |\text{UpdateCnt}[w_1]|)$ tokens to \mathcal{S} , and it receives $O(|\text{UpdateCnt}[w_1]|)$ encrypted values back as the result. Hence, the asymptotic communication overhead of NOMOS SEARCH routine is $O(k \cdot |q| \cdot |\text{UpdateCnt}[w_1]|)$.

Storage overhead. We analyse the NOMOS EDB storage overhead with respect to the plain database \mathbf{DB} . The storage overhead for EDB in NOMOS is essentially the combined TSet and XSet overhead. The TSet overhead of NOMOS is practically the same as of single client dynamic construction ODXT, which is $O(|\mathbf{DB}|)$ (linear in terms of the number of records in the plain database \mathbf{DB}) as the TSet stores one encrypted value for each entry in \mathbf{DB} . The RBF-backed XSet requires $\ell \cdot O(|\mathbf{DB}|)$ storage. However, compared to TSet, XSet stores only 1/0 for each index and requires lesser storage than TSet that stores encryptions of $O(|\mathbf{DB}|)$ items. As a result, NOMOS has linear $O(|\mathbf{DB}|)$ asymptotic storage overhead in practice.

Security analysis. We divert the elaborate security analysis of NOMOS to Appendix E.

5 Implementation Details and Results

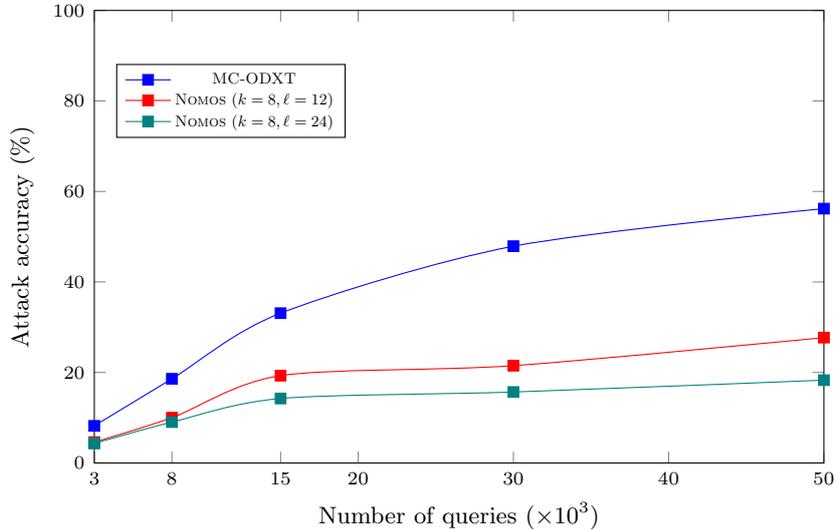


Figure 2: Attack accuracy vs number of queries t by the colluding client in the build phase. The number of queries by the benign clients was set to 2000.

In this section, we describe a prototype implementation of NOMOS and evaluate its overhead and performance over real-world databases.

Data set and platform. We used the Enron email data set⁶ for our experiments. The database

⁶<https://www.cs.cmu.edu/~enron>

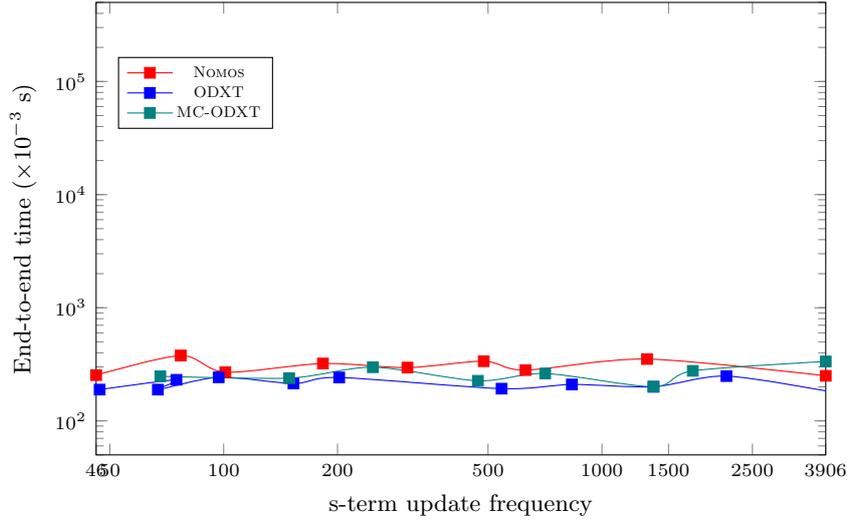


Figure 3: End-to-end query latency for two-keyword queries of the form $q = w_a \wedge w_v$ (s-term w_a update frequency is fixed at 70).

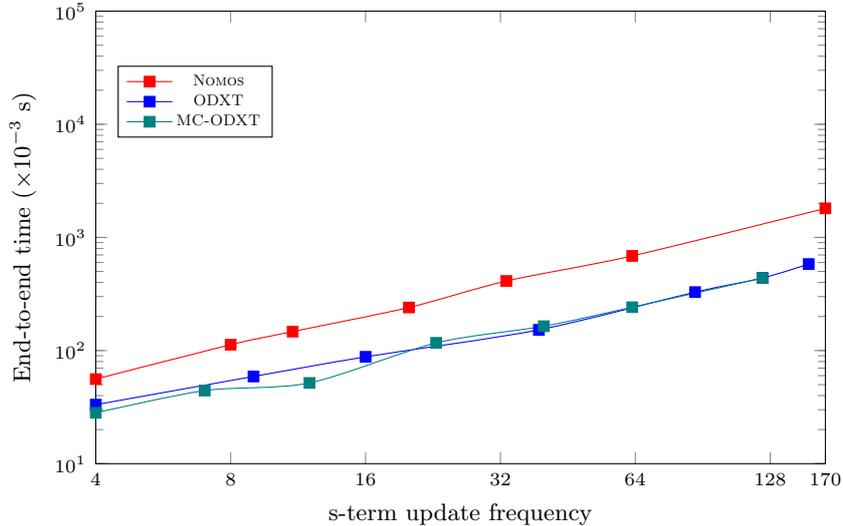


Figure 4: End-to-end query latency for two keyword queries of the form $q = w_v \wedge w_a$ (s-term w_v update frequency is varied).

contains 517,401 documents (emails) and 20 million keyword-document pairs, with a total size 1.9 GB. The final NOMOS algorithm was implemented using C++11 with native multithreading support and was compiled using GCC9. We used Redis as the database backend system to store **EDB** comprising TSet and XSet. We ran the experiments on dual 24-core Intel Xeon E5-2690 2.6GHz CPU with 128GB RAM and 512GB SSD running Ubuntu 20.04 64bit operating system with gigabit network link.

Query processing. We evaluated NOMOS’s performance for two different types of queries - two-keyword and multi-keyword queries. The two-keyword queries are of the form $q = w_1 \wedge w_2$, which we represent as $q = w_a \wedge w_v$ or $q = w_v \wedge w_a$. Here w_a is called the constant term whose frequency is kept fixed, and w_v is the variable term whose frequency is varied during experimentation. For the multi-keyword queries of the form $q = w_1 \wedge \dots \wedge w_n$, $n \in [3, 6]$, the first keyword w_1 is varied and maximum frequency of $\{w_2, \dots, w_n\}$ is fixed in one set of experiments. In another set, the frequency of w_1 is kept fixed and max frequency of $\{w_2, \dots, w_n\}$ is varied. These experiments

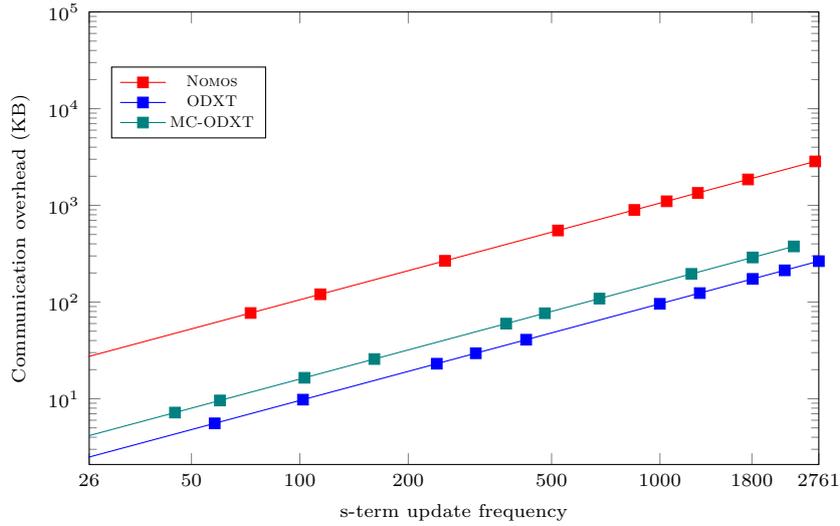


Figure 5: End-to-end communication overhead for two-keyword queries of the form $q = w_v \wedge w_a$ (s-term w_v update frequency is varied).

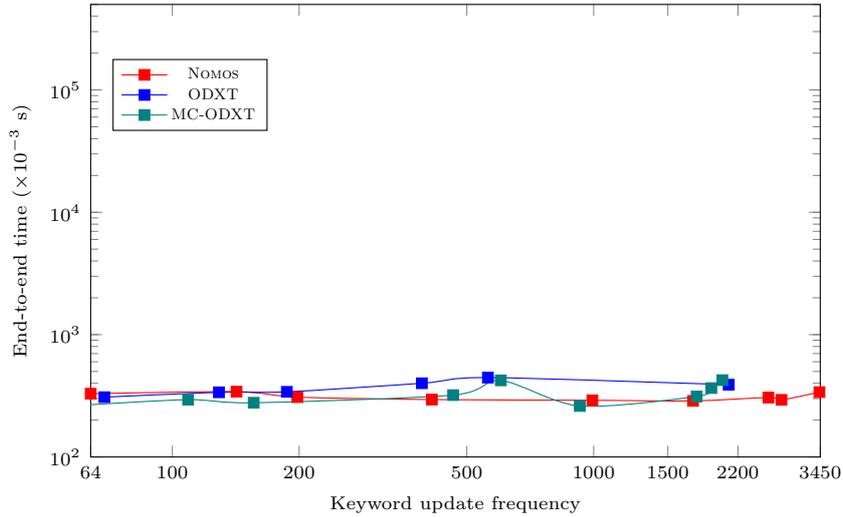


Figure 6: End-to-end query time for multi-keyword queries of the form $q = w_1 \wedge \dots \wedge w_n$ (s-term w_1 update frequency is fixed at 60).

examine the sublinear search and communication overhead of NOMOS for two and multi-keyword queries. The query keywords are sampled randomly from Δ .

5.1 Experiments on Leakage

We executed the `CrossAttack` method of Algorithm 1 on NOMOS to compare the leakage with MC-ODXT. The attack accuracy is plotted in Figure 2 against the number of queries issued by the colluding client in the build phase of the attack. Clearly, NOMOS has a significantly lower attack accuracy compared to MC-ODXT. This low attack accuracy is a direct consequence of the redundancy of RBF during `XSet` accesses. The value of the RBF parameter ℓ was set to 12, while the parameter k was set to 8. Thus, NOMOS achieves more than 50% reduction in attack accuracy with four additional indices for `XSet` insertion. The parameter k needs to be kept at a fixed value to maintain a desired false positive rate for the same database parameters and experiment setting as of the non-RBF version. In contrast, ℓ is responsible for the redundancy in RBF leading to a

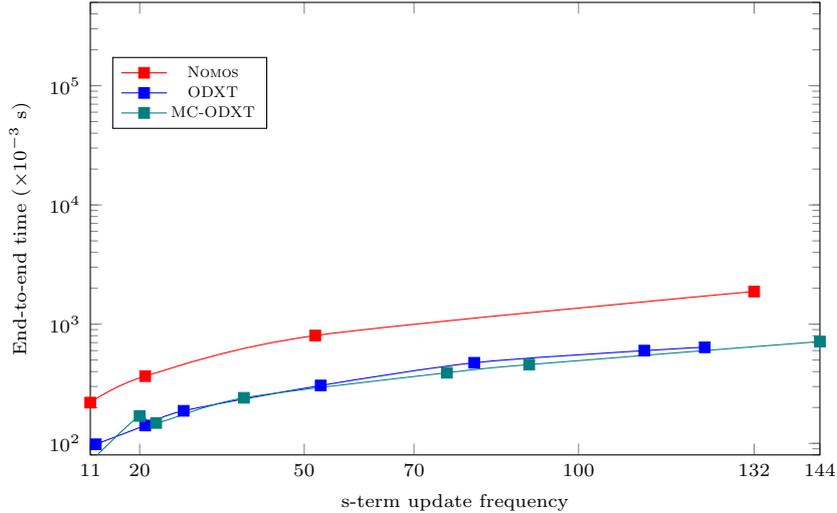


Figure 7: End-to-end query time for multi-keyword queries of the form $q = w_1 \wedge \dots \wedge w_n$ (s-term w_1 update frequency is varied).

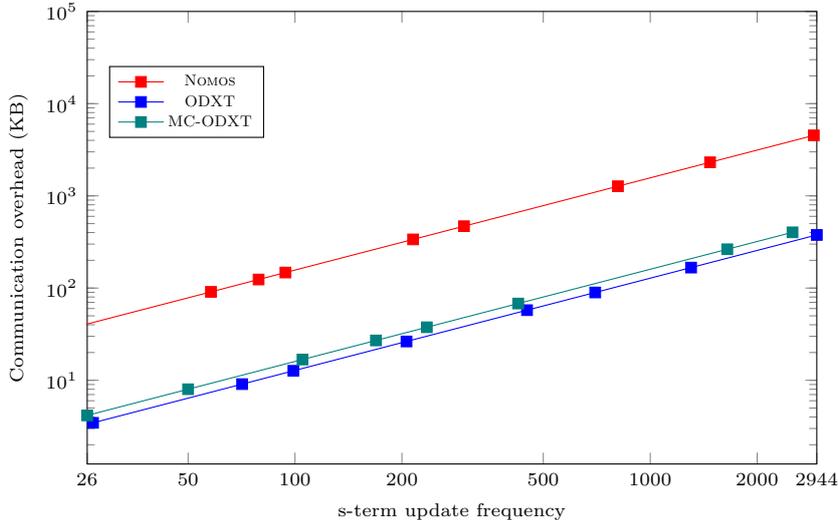


Figure 8: End-to-end communication overhead for multi-keyword queries of the form $q = w_1 \wedge \dots \wedge w_n$ (s-term w_1 update frequency is varied).

lower attack accuracy due to the increased redundant accesses per (w, id) pair. However, a high value of ℓ would incur larger storage to accommodate additional XSet entries, and thus implies a trade-off. We view this as a leakage-versus-storage trade-off in the multi-client setting that is unavoidable in practice. We present comprehensive discussion on the choice of parameters and analysis of RBF in Appendix D.

5.2 Experiments on the Search Latency

We considered two types of queries to evaluate the search performance of NOMOS as stated earlier in this section. For two-keyword queries, we fix the frequency of the constant term w_a and vary the frequency of the variable term w_v from 10 to 5000. The end-to-end search latency for the queries of the form $q = w_a \wedge w_v$ in Figure 3 and queries of the form $q = w_v \wedge w_a$ in Figure 4. Observe that, in Figure 3, the end-to-end search latency remains almost constant. Whereas in Figure 4, the search latency varies linearly with the frequency of the variable term. This behaviour validates

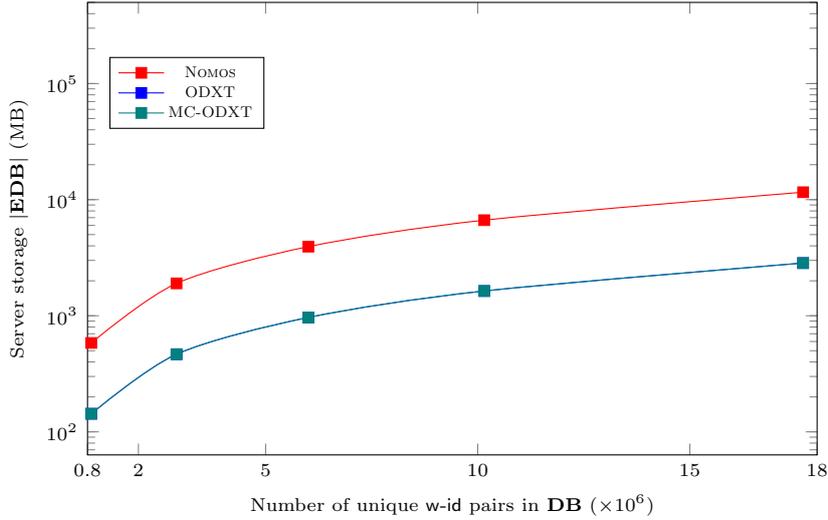


Figure 9: Server-side storage overhead (**EDB** size) for NOMOS, MC-ODXT and ODXT for same plain input database size.

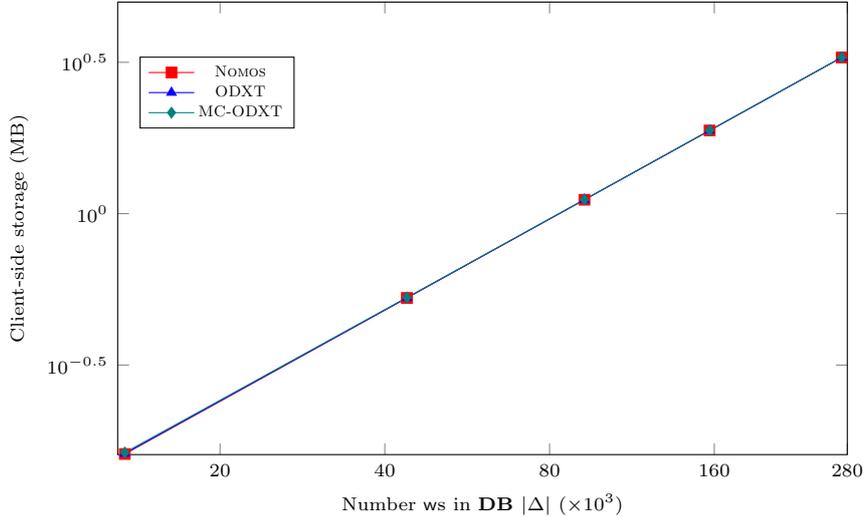


Figure 10: Client-side storage overhead for NOMOS, ODXT and MC-ODXT for same plain input database size.

the sublinearity of the NOMOS search algorithm where the search latency linearly depends on the frequency of the s -term of $|\mathbf{DB}(w_1)|$, which is sublinear in terms of total number of records in the database.

The communication overheads of NOMOS **GENTOKEN** and **SEARCH** are plotted in Figure 5 for two-keyword queries. NOMOS incurs sublinear communication overhead (linear in terms of s -term update frequency) for **GENTOKEN** and **SEARCH** as shown in the plot. However, in comparison with ODXT, NOMOS has the additional *necessary* overhead of search tokens generated by **GENTOKEN**, and the **SEARCH** communication overhead increases due to RBF-based **XSet**.

We also report experimental results for multi-keyword queries of the form $q = w_1 \wedge \dots \wedge w_n$, where $n \in [3, 6]$, and we denote w_1 as the s -term in the multi-keyword queries. The end-to-end search latency for two sets of experiments is plotted for fixed and variable s -term update frequencies in Figure 6 and Figure 7, respectively. Observe that the search overhead remains sublinear (proportional to the frequency of the s -term) in terms of the number of total records in **EDB**. Similarly, the communication overhead remains sublinear in the total database size scaled

by the number of cross-terms, as illustrated in Figure 8.

5.3 Evaluation of the Storage Overhead

We varied the number of w s in **DB** and generated the corresponding **EDB** by executing NOMOS UPDATE. We compare the NOMOS storage overhead with ODXT to illustrate the storage overhead as a trade-off with lesser leakage. The **EDB** overhead for both NOMOS and ODXT are plotted in Figure 9, and the client-side storage overhead is plotted in Figure 10. Observe that the storage overhead profile for both NOMOS and ODXT remains linear with **DB** size, and NOMOS **EDB** overhead is approximately 2.5 times of ODXT which is manageable on the cloud in practice. All three constructions require $O(|\Delta|)$ client-side storage as illustrated in Figure 10.

The performance overhead of NOMOS is slightly higher compared to MC-ODXT as NOMOS generates more tokens due to RBF-based XSet look-up. The increased performance and storage overhead of NOMOS is a necessary leakage-versus-overhead trade-off to allow secure searches in the multi-client setting of NOMOS. The storage overhead of NOMOS is less than $2\times$ - $2.5\times$ of ODXT (and MC-ODXT) for storage that is practically manageable in a cloud infrastructure. Compared with the plain version of search, encrypted search incurs higher overhead which varies depending upon the cryptographic parameters/algorithms used, data structure, and the database system used. However, it provides strong confidentiality against unauthorised access compared to the typical unencrypted search process and incurs less leakage and communication overhead than the naïve symmetrically encrypted database with trivial decryption.

In summary, NOMOS provides secure multi-client search with dynamic updates at the expense of minimal additional storage overhead. Modern cloud services can easily manage this additional storage to provide a secure environment to process encrypted data which multiple users can access. Naturally, depending upon the requirement, a service provider needs to make necessary modifications to the core algorithms. These modifications mainly include parsing the data into a suitable format to store and process as a multi-map. The data owner is responsible for this pre-processing of unstructured data to be updated into the encrypted database on the remote server. Thus, in general, NOMOS can be adopted in a broad set of practical applications with minimal modification, such as healthcare, government records, and banks, to name a few.

6 Conclusion and Future Directions

We introduced the first forward and backward secure dynamic multi-client SSE scheme NOMOS supporting conjunctive Boolean queries. NOMOS is an MRSW construction that builds upon the state-of-the-art SRSW dynamic construction ODXT [PM21]. We showed that the straight-forward extension of ODXT to the multi-client setting is completely insecure against collusion between a semi-honest client and a semi-honest server due to a cross-term-based leakage. Our NOMOS construction mitigates this leakage by adopting a customised Bloom filter called redundant Bloom filter while supporting efficient single-round multi-client queries. We presented extensive experimentation to demonstrate the practical performance of NOMOS over real-world databases. We leave extending NOMOS to the MRMW setting as an interesting direction of future research.

References

- [AG22] Megumi Ando and Marilyn George. On the cost of suppressing volume for encrypted multi-maps. *Proceedings on Privacy Enhancing Technologies*, 4:44–65, 2022.
- [BBG05] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical identity based encryption with constant size ciphertext. In *EUROCRYPT*, 2005.
- [BHJP14] Christoph Bösch, Pieter H. Hartel, Willem Jonker, and Andreas Peter. A survey of provably secure searchable encryption. *ACM Comput. Surv.*, 47, 2014.
- [BI01] Amos Beimel and Yuval Ishai. Information-theoretic private information retrieval: A unified construction. In *ICALP*, 2001.

- [BMO17a] Raphaël Bost, Brice Minaud, and Olga Ohrimenko. Forward and backward private searchable encryption from constrained cryptographic primitives. In *CCS*, 2017.
- [BMO17b] Raphaël Bost, Brice Minaud, and Olga Ohrimenko. Forward and backward private searchable encryption from constrained cryptographic primitives. In *ACM CCS 2017*, pages 1465–1482, 2017.
- [BN00] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In *ASIACRYPT*, 2000.
- [Bos16] Raphael Bost. $\sigma\phi\phi\sigma$: Forward secure searchable encryption. In *CCS*, 2016.
- [CGKO06] Reza Curtmola, Juan A. Garay, Seny Kamara, and Rafail Ostrovsky. Searchable symmetric encryption: improved definitions and efficient constructions. In *ACM CCS*, 2006.
- [CGKS95] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *FOCS*, 1995.
- [CJJ⁺13] David Cash, Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Catalin Rosu, and Michael Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *CRYPTO*, 2013.
- [CJJ⁺14] David Cash, Joseph Jaeger, Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Catalin Rosu, and Michael Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *NDSS*, 2014.
- [CK10] Melissa Chase and Seny Kamara. Structured encryption and controlled disclosure. In *ASIACRYPT*, 2010.
- [CM05] Yan-Cheng Chang and Michael Mitzenmacher. Privacy preserving keyword searches on remote encrypted data. In *ACNS*, 2005.
- [CPPJ18] Javad Ghareh Chamani, Dimitrios Papadopoulos, Charalampos Papamanthou, and Rasool Jalili. New constructions for forward and backward private symmetric searchable encryption. In *CCS*, 2018.
- [EKPE18] Mohammad Etemad, Alptekin Küpçü, Charalampos Papamanthou, and David Evans. Efficient dynamic searchable encryption with forward privacy. *PoPETS*, 2018.
- [FJK⁺15] Sky Faber, Stanislaw Jarecki, Hugo Krawczyk, Quan Nguyen, Marcel-Catalin Rosu, and Michael Steiner. Rich queries on encrypted data: Beyond exact matches. In *ESORICS*, 2015.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *ACM STOC*, 2009.
- [GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *J. ACM*, 43, 1996.
- [Goh03] Eu-Jin Goh. Secure indexes. *IACR Cryptology ePrint Archive*, 2003.
- [Gol87] Oded Goldreich. Towards a theory of software protection and simulation by oblivious rams. In *STOC*, 1987.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO*, 2013.
- [HSWW18] Ariel Hamlin, Abhi Shelat, Mor Weiss, and Daniel Wichs. Multi-key searchable encryption, revisited. In *Public-Key Cryptography - PKC 2018*, Lecture Notes in Computer Science. Springer, 2018.

- [JJK⁺13] Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Catalin Rosu, and Michael Steiner. Outsourced symmetric private information retrieval. In *CCS*, 2013.
- [JT20] Joseph Jaeger and Nirvan Tyagi. Handling adaptive compromise for practical encryption schemes. In *Advances in Cryptology - CRYPTO*, 2020.
- [KM17] Seny Kamara and Tarik Moataz. Boolean searchable symmetric encryption with worst-case sub-linear complexity. In *EUROCRYPT*, 2017.
- [KP13] Seny Kamara and Charalampos Papamanthou. Parallel and dynamic searchable symmetric encryption. In *FC*, 2013.
- [KPR12] Seny Kamara, Charalampos Papamanthou, and Tom Roeder. Dynamic searchable symmetric encryption. In *ACM CCS*, 2012.
- [LPS⁺18] Shangqi Lai, Sikhar Patranabis, Amin Sakzad, Joseph K. Liu, Debdeep Mukhopadhyay, Ron Steinfeld, Shifeng Sun, Dongxi Liu, and Cong Zuo. Result pattern hiding searchable encryption for conjunctive queries. In *CCS*, 2018.
- [NR04] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. *JACM*, 51, 2004.
- [OSI07] Rafail Ostrovsky and William E Skeith III. A survey of single-database private information retrieval: Techniques and applications. In *PKC*, 2007.
- [PM21] Sikhar Patranabis and Debdeep Mukhopadhyay. Forward and backward private conjunctive searchable symmetric encryption. In *NDSS*, 2021.
- [Rog02] Phillip Rogaway. Authenticated-encryption with associated-data. In *CCS*, 2002.
- [Sta90] Wolfgang Stadje. The collector’s problem with group drawings. *Advances in Applied Probability*, 22(4):866–882, 1990.
- [SWP00] Dawn Xiaodong Song, David A. Wagner, and Adrian Perrig. Practical techniques for searches on encrypted data. In *IEEE S&P*, 2000.
- [ZKP16] Yupeng Zhang, Jonathan Katz, and Charalampos Papamanthou. All your queries are belong to us: The power of file-injection attacks on searchable encryption. In *USENIX Security Symposium*, 2016.

A Security Games for Dynamic Multi-client SSE

We present the **Real** and **Ideal** security games for dynamic multi-client SSE as stated in Section V of the main manuscript.

Games for adversarial client The **Real** and **Ideal** games for the security of a dynamic MRSW SSE scheme Π from an adversarial client are presented in Algorithms 6 and 7.

Games for adversarial server The **Real** and **Ideal** games for the security of a dynamic MRSW SSE scheme Π from an adversarial server are presented in Algorithms 8 and 9.

B MC-ODXT Construction Details

We present MC-ODXT construction algorithm in this Appendix. MC-ODXT is an extension of ODXT [PM21] to multi-client MRSW setting following the OPRF-based approach of OSPIROXT [JJK⁺13]. The MC-ODXT SETUP algorithm (Algorithm 10) executed by \mathcal{G} samples OPRF, PRF and AE keys. A client \mathcal{C} engages with \mathcal{G} in the GENTOKEN protocol (Algorithm 12) to obtain the blinded search tokens for a conjunctive query to search over the **EDB**. After obtaining the search tokens, \mathcal{C} engages with \mathcal{S} in the SEARCH protocol (Algorithm 13) to retrieve the set of ids matching the conjunctive query. \mathcal{G} jointly executes the UPDATE protocol (Algorithm 11) with \mathcal{S} to insert contents into or delete contents from **EDB**.

Algorithm 6 Experiment $\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda)$ adversarial client

```
1: function  $\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda)$ 
2:    $N \leftarrow \mathcal{A}(\lambda)$ 
3:    $(\mathbf{sk}, \mathbf{st}_0, \mathbf{EDB}_0) \leftarrow \Pi.\text{SETUP}(\lambda, N)$ 
4:   for  $k \leftarrow 1$  to  $Q$  do
5:     if query-type == gentoken then
6:        $q_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
7:        $(\mathbf{st}_k, \mathbf{t}_k) \leftarrow \text{GENTOKEN}(\mathbf{sk}, \mathbf{st}_{k-1}, q_k)$ 
8:     else if query-type == search then
9:        $t_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
10:       $(\mathbf{st}_k, \mathbf{EDB}_k, \mathbf{DB}(q_k)) \leftarrow \text{SEARCH}(\mathbf{sk}, \mathbf{st}_{k-1}, \mathbf{t}_k; \mathbf{EDB}_{k-1})$ 
11:       $\tau_k$  denotes the view of the adversary after the  $k^{\text{th}}$  query
12:    $b \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_Q, \tau_1, \dots, \tau_Q)$ 
13:   return  $b$ 
```

Algorithm 7 Experiment $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}(\lambda)$ for adversarial client

```
1: function  $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}(\lambda)$ 
2:   Parse the  $\mathcal{L}$  as:  $\{\mathcal{L}^{\text{SETUP}}, \mathcal{L}^{\text{UPDATE}}, \mathcal{L}^{\text{GENTOKEN}}, \mathcal{L}^{\text{SEARCH}}\}$ .
3:    $N \leftarrow \mathcal{A}(\lambda)$ 
4:    $(\mathbf{st}_{\text{SIM}}, \mathbf{EDB}_0) \leftarrow \text{SIM}_0(\mathcal{L}^{\text{SETUP}}(\lambda, N))$ 
5:   for  $k \leftarrow 1$  to  $Q$  do
6:     if query-type == gentoken then
7:        $q_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
8:        $(\mathbf{st}_{\text{SIM}}, \mathbf{t}_k) \leftarrow \text{SIM}_1(\mathbf{st}_{\text{SIM}}, \mathcal{L}^{\text{GENTOKEN}}(q_k))$ 
9:     else if query-type == search then
10:       $t_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
11:       $(\mathbf{st}_{\text{SIM}}, \mathbf{EDB}_k, \mathbf{DB}(q_k)) \leftarrow \text{SIM}_{\text{SEARCH}}(\mathbf{st}_{\text{SIM}}, \mathcal{L}^{\text{SEARCH}}(\mathbf{t}_k); \mathbf{EDB}_{k-1})$ 
12:       $\tau_k$  denote the view of the adversary after the  $k^{\text{th}}$  query
13:    $b \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_Q, \tau_1, \dots, \tau_Q)$ 
14:   return  $b$ 
```

Algorithm 8 Experiment $\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda)$ adversarial server

```
1: function  $\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda)$ 
2:    $N \leftarrow \mathcal{A}(\lambda)$ 
3:    $(\mathbf{sk}, \mathbf{st}_0, \mathbf{EDB}_0) \leftarrow \Pi.\text{SETUP}(\lambda, N)$ 
4:   for  $k \leftarrow 1$  to  $Q$  do
5:     if query-type == search then
6:        $t_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
7:        $(\mathbf{st}_k, \mathbf{EDB}_k, \mathbf{DB}(q_k)) \leftarrow \Pi.\text{SEARCH}(\mathbf{sk}, \mathbf{st}_{k-1}, \mathbf{t}_k; \mathbf{EDB}_{k-1})$ 
8:     else if query-type == update then
9:        $(\text{op}_k, \{\mathbf{w}_k, \text{id}_k\}) \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
10:       $(\mathbf{st}_k, \mathbf{EDB}_k) \leftarrow$ 
11:       $\Pi.\text{UPDATE}(\mathbf{sk}, \mathbf{st}_{k-1}, (\text{op}_k, \{\mathbf{w}_k, \text{id}_k\}); \mathbf{EDB}_{k-1})$ 
12:       $\tau_k$  denotes the view of the adversary after the  $k^{\text{th}}$  query
13:    $b \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_Q, \tau_1, \dots, \tau_Q)$ 
14:   return  $b$ 
```

B.1 Private retrieval of s-term information

As stated in Section IV.C of the main body, the s-term update frequency requires a commonly accessible record in this setting. This necessitates a secure mechanism to disclose the keyword frequencies to prevent \mathcal{G} from knowing \mathcal{C} 's query keywords. Naively, this can be achieved via

Algorithm 9 Experiment $\text{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}(\lambda)$ for adversarial server

```
1: function  $\text{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}(\lambda)$ 
2:   Parse  $\mathcal{L}_S$  as:  $\{\text{left}(\mathcal{L}^{\text{SETUP}}, \mathcal{L}^{\text{UPDATE}}, \mathcal{L}^{\text{GENTOKEN}}, \mathcal{L}^{\text{SEARCH}})\}$ .
3:    $N \leftarrow \mathcal{A}(\lambda)$ 
4:    $(\text{st}_{\text{SIM}}, \mathbf{EDB}_0) \leftarrow \text{SIM}_0(\mathcal{L}^{\text{SETUP}}(\lambda, N))$ 
5:   for  $k \leftarrow 1$  to  $Q$  do
6:     if query-type == search then
7:        $t_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
8:        $(\text{st}_{\text{SIM}}, \mathbf{EDB}_k, \mathbf{DB}(q_k)) \leftarrow \text{SIM}_2(\text{st}_{\text{SIM}}, \mathcal{L}^{\text{SEARCH}}(t_k); \mathbf{EDB}_{k-1})$ 
9:     else if query-type == update then
10:       $(\text{op}_k, \{\mathbf{w}_k, \text{id}_k\}) \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
11:       $(\text{st}_k, \mathbf{EDB}_k) \leftarrow$ 
12:       $\text{SIM}_1(\text{st}_{\text{SIM}}, \mathcal{L}^{\text{UPDATE}}((\text{op}_k, \{\mathbf{w}_k, \text{id}_k\})); \mathbf{EDB}_{k-1})$ 
13:       $\tau_k$  denote the view of the adversary after the  $k^{\text{th}}$  query
14:    $b \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_Q, \tau_1, \dots, \tau_Q)$ 
15:   return  $b$ 
```

Algorithm 10 MC-ODXT SETUP

Input: λ **Output:** $\text{sk}, \text{UpdateCnt}, \mathbf{EDB}$

```
1: function MC-ODXT.SETUP
2:   Gate-keeper
3:   Sample a uniformly random key  $K_S$  from  $\mathbb{Z}_p^*$  for OPRF
4:   Sample two sets of uniformly random keys  $K_T = \{K_T^1, \dots, K_T^d\}$  and  $K_X = \{K_X^1, \dots, K_X^d\}$ 
5:   from  $(\mathbb{Z}_p^*)^d$  for OPRF
6:   Sample uniformly random key  $K_Y$  from  $\{0, 1\}^\lambda$  for PRF  $F_p$ 
7:   Sample shared uniformly random key  $K_M$  from  $\{0, 1\}^\lambda$  for AE
8:   Initialise UpdateCnt, TSet, XSet to empty maps
9:   Gate-keeper keeps  $\text{sk} = (K_S, K_T, K_X, K_Y)$ ; UpdateCnt is disclosed to clients when required,
10:  and  $K_M$  is shared between gate-keeper and the server
11:  Set  $\mathbf{EDB} = (\text{TSet}, \text{XSet})$ 
12:  Send  $\mathbf{EDB}$  to server
```

a publicly hosted website or storage with \mathcal{G} having the write permission (implicitly adopted by OXT [CJJ⁺13], OSPIR-OXT [JJK⁺13], HXT [LPS⁺18], and ODXT [PM21]). In this work we assumed that \mathcal{C} 's have some suitable mechanism to retrieve the frequency information from \mathcal{G} , and focused on the core construction to allow multi-client search capability.

A plausible way to allow such frequency information retrieval in a privacy-preserving manner is to employ a private information retrieval (PIR) [CGKS95, OSI07, BI01] based mechanism involving a \mathcal{C} and the \mathcal{G} . The idea is (at a high level) that the \mathcal{G} holds an array of keyword frequency values, indexed by keywords, and the \mathcal{C} wants to look up the frequency of a particular keyword without revealing to the \mathcal{G} what is the keyword being looked up. This lends itself to a classic single-server PIR-based solution. In fact, one could also use a multi-server PIR solution [BI01] where the \mathcal{G} simply distributes the frequency information across multiple (mutually distrusting) servers, thus decentralising the trust assumptions involved.

We note, however, that using a (multi-server) PIR-based frequency information retrieval in the search token generation phase brings in additional computation overhead, and possibly, additional rounds. We leave incorporating a dedicated s-term update frequency retrieval mechanism as a part of the MRMW extension of NOMOS.

Algorithm 11 MC-ODXT UPDATE

Input: $K_S, K_T = \{K_T^1, \dots, K_T^d\}, K_X = \{K_X^1, \dots, K_X^d\}$, accessed as $K_T[I(w)]$ and $K_X[I(w)]$ for attribute $I(w)$ of w , $K_Y, (w, \text{id})$ pair to be updates, update operation op

Output: Updated **EDB**

```
1: function MC-ODXT.UPDATE
   Gate-keeper
2:   Parse  $(K_T, K_X, K_Y)$  and UpdateCnt
3:   Set  $K_Z \leftarrow F((H(w))^{K_S}, 1)$ 
4:   If UpdateCnt[w] is NULL then set UpdateCnt[w] = 0
5:   Set UpdateCnt[w] = UpdateCnt[w] + 1
6:   Set  $\text{addr} = (H(w || \text{UpdateCnt}[w] || 0))^{K_T[I(w)]}$ 
7:   Set  $\text{val} = (\text{id} || \text{op}) \oplus (H(w || \text{UpdateCnt}[w] || 1))^{K_T[I(w)]}$ 
8:   Set  $\alpha = F_p(K_Y, \text{id} || \text{op}) \cdot (F_p(K_Z, w || \text{UpdateCnt}[w])^{-1})$ 
9:   Set  $\text{xtag} = H(w)^{K_X[I(w)]} \cdot F_p(K_Y, \text{id} || \text{op})$ 
10:  Send (addr, val,  $\alpha$ , xtag) to server
   Server
11:  Parse EDB = (TSet, XSet)
12:  Set TSet[addr] = (val,  $\alpha$ )
13:  Set XSet[xtag] = 1
```

Algorithm 12 MC-ODXT GENTOKEN

Input: $q = w_1 \wedge \dots \wedge w_n$. \mathcal{P} is the set of allowable attribute sequences, K_S, K_T, K_X, K_M

Output: strap, bstag₁, \dots , bstag_m, $\delta_1, \dots, \delta_m$, bstrap₂, \dots , bstrap_n, env

```
1: function MC-ODXT.GENTOKEN
   Client
2:   Pick  $r_1, \dots, r_n \xleftarrow{\$} \mathbb{Z}_p^*$ 
3:   Set  $m = \text{UpdateCnt}[w_1]$ 
4:   Pick  $s_1, \dots, s_m \xleftarrow{\$} \mathbb{Z}_p^*$ 
5:   Set  $a_j \leftarrow (H(w_j))^{r_j}$ , for  $j = 1, \dots, n$ 
6:   Set  $b_j \leftarrow (H(w_1 || j || 0))^{s_j}$ , for  $j = 1, \dots, m$ 
7:   Set  $c_j \leftarrow (H(w_1 || j || 1))^{s_j}$ , for  $j = 1, \dots, m$ 
8:   Set  $\text{av} = (I(w_1), \dots, I(w_n)) = (I_1, \dots, I_n)$ 
   Gate-keeper
9:   Abort if  $\text{av} \notin \mathcal{P}$   $\triangleright$  Abort if attributes do not match
10:  Pick  $\rho_1, \dots, \rho_n \xleftarrow{\$} \mathbb{Z}_p^*$ 
11:  Pick  $\gamma_1, \dots, \gamma_m \xleftarrow{\$} \mathbb{Z}_p^*$ 
12:  Set  $\text{strap}' \leftarrow (a_1)^{K_S}$ 
13:  Set  $\text{bstag}'_j \leftarrow (b_j)^{K_T[I_1] \cdot \gamma_j}$ , for  $j = 1, \dots, m$ 
14:  Set  $\delta'_j \leftarrow (c_j)^{K_T[I_1]}$ , for  $j = 1, \dots, m$ 
15:  Set  $\text{bstrap}'_j \leftarrow (a_j)^{K_X[I_j] \cdot \rho_j}$  for  $j = 2, \dots, n$ 
16:  Set  $\text{env} = \text{AE.Enc}_{K_M}(\rho_1, \dots, \rho_n, \gamma_1, \dots, \gamma_m)$ 
17:  Send  $\text{strap}', \text{bstag}'_1, \dots, \text{bstag}'_m, \delta'_1, \dots, \delta'_m, \text{bstrap}'_2, \dots, \text{bstrap}'_n, \text{env}$  to Client
   Client
18:  Set  $\text{strap} \leftarrow (\text{strap}')^{r_1^{-1}}$ 
19:  Set  $\text{bstag}_j \leftarrow (\text{bstag}'_j)^{s_j^{-1}}$ , for  $j = 1, \dots, m$ 
20:  Set  $\delta_j \leftarrow (\delta'_j)^{s_j^{-1}}$ , for  $j = 1, \dots, m$ 
21:  Set  $\text{bstrap}_j \leftarrow (\text{bstrap}'_j)^{r_j^{-1}}$ , for  $j = 2, \dots, n$ 
22:  Output (strap, bstag1,  $\dots$ , bstagm,  $\delta_1, \dots, \delta_m$ , bstrap2,  $\dots$ , bstrapn, env) as search token
```

Algorithm 13 MC-ODXT SEARCH

Input: strap, bstag₁, …, bstag_m, δ₁, …, δ_m, bstrap₂, …, bstrap_n, env, UpdateCnt**Output:** IdList

```
1: function MC-ODXT.SEARCH
   Client
2:   Set  $K_Z \leftarrow F(\text{strap}, 1)$ 
3:    $m = \text{UpdateCnt}[w_1]$ 
4:   Initialise stokenList to an empty list
5:   Initialise xtokenList1, …, xtokenListm to empty lists
6:   for  $j = 1$  to  $m$  do
7:     stokenList = stokenList  $\cup$  {bstrapj}
8:     for  $i = 2$  to  $n$  do
9:       Set  $\text{xtoken}_{i,j} = \text{bstrap}_i^{F_p(K_Z, w_1 || j)}$ 
10:      Set  $\text{xtokenList}_j = \text{xtokenList}_j \cup \text{xtoken}_{i,j}$ 
11:      Randomly permute the tuple-entries of xtokenListj
12:   Send (stokenList, xtokenList1, …, xtokenListm)
   Server
13:   Upon receiving env from client, verify env; if verification fails, return  $\perp$ ; otherwise decrypt
   env
14:   Parse EDB = (TSet, XSet)
15:   Initialise sEOpList to an empty list
16:   for  $j = 1$  to stokenList.size do
17:     Set cntj = 1
18:     Set stagj  $\leftarrow$  (stokenList[j])1/γj
19:     Set (svalj, αj) = TSet[stagj]
20:     for  $i = 2$  to  $n$  do
21:       Set  $\text{xtoken}_{i,j} = \text{xtokenList}_j[i]$ 
22:       Compute  $\text{xtag}_{i,j} = (\text{xtoken}_{i,j})^{\alpha_j / \rho_i}$ 
23:       If XSet[xtagi,j] == 1, then set cntj = cntj + 1
24:     Set sEOpList = sEOpList  $\cup$  {(j, svalj, cntj)}
25:   Sent sEOpList to client
   Client
26:   Initialise IdList to an empty list
27:   for  $\ell = 1$  to sEOpList.size do
28:     Let (j, svalj, cntj) = sEOpList[ℓ]
29:     Recover (idj || opj) = svalj  $\oplus$  δℓ
30:     If opj is ADD and cntj = n then set IdList = IdList  $\setminus$  {idj}
31:   Output IdList
```

C Leakage Analysis of Nomos

We outline the leakage profile and security analysis (informal) of NOMOS in a concise manner in this section. Since NOMOS adopts the ODXT structure, we closely follow the analysis of ODXT for NOMOS. As our multi-client extension follows the approach of OSPiR-OXT, we adopt a similar approach (and setting) to prove the security of NOMOS in the multi-client setting.

For the ease of analysis, we first categorise the leakages of NOMOS according to the leakage to the clients and to the server. NOMOS incurs the following class of leakages for the respective entities as summarised below.

Leakage to clients. A client \mathcal{C} learns the following information apart from the data/information trivially available from the execution of respective protocols GENTOKEN and SEARCH. We denote this ensemble of leakages as $\mathcal{L}_{\mathcal{C}}$.

Token generation leakage. The GENTOKEN protocol of NOMOS involves both s-terms and x-terms in blinded form. However, **EDB** is not accessed in this phase and \mathcal{C} receives only the following information.

s-term leakage. \mathcal{C} receives back the doubly blinded strap, bstag and δ values that are obtained from the s-term. However, it does not learn anything about the actual blinding factors or the OPRF key. Hence, the s-term leakage to \mathcal{C} in GENTOKEN is $\mathcal{L}_{\mathcal{C}}^{\text{GENTOKEN},s\text{-term}} = \perp$.

x-term Leakage. Similarly, \mathcal{C} does not learn anything about the x-term blinding factors or the OPRF evaluation keys. Hence, the x-term leakage to \mathcal{C} from GENTOKEN is $\mathcal{L}_{\mathcal{C}}^{\text{GENTOKEN},x\text{-term}} = \perp$.

Note that, \mathcal{C} is not involved in the UPDATE protocol, and hence, $\mathcal{L}_{\mathcal{C}}^{\text{UPDATE}} = \perp$. In the SEARCH protocol, \mathcal{C} inputs the search tokens and receives the query result. It does not receive anything beyond the query result in encrypted form, which is part of the actual protocol execution. Therefore, in SEARCH, \mathcal{C} learns nothing about the s-term and x-terms and the SEARCH leakage to \mathcal{C} can be expressed as $\mathcal{L}_{\mathcal{C}}^{\text{SEARCH},s\text{-term}} = |\text{DB}(q)|$ (the volume of the result). Thus, the complete leakage profile for an adversarial \mathcal{C} therefore is $\mathcal{L}_{\mathcal{C}} = \{|\text{DB}(q)|\}$.

Leakage to server. \mathcal{S} engages with \mathcal{G} and a \mathcal{C} in UPDATE and SEARCH respectively. In these routines, \mathcal{S} learns the following information about the query (tokens) or the encrypted data in addition to trivially received data/information as the output of the interaction during UPDATE and SEARCH with the \mathcal{G} and \mathcal{C} , respectively. We classify these leakages as below.

Update leakage. The UPDATE process practically incurs zero leakage to \mathcal{S} as the server receives (addr, val) to insert into TSet and a set of XSet addresses to set to 1. The addr and val are obtained using PRF/OPRF evaluation where none of the values is repeated. Similarly, the XSet addresses are also uniquely generated each time and are never repeated. Moreover, ADD and DEL operations are treated identically during an update to EDB. Hence, \mathcal{S} learns no information about the w or the id, or the op involved in the UPDATE process and the leakage to \mathcal{S} can be expressed as $\mathcal{L}_{\mathcal{S}}^{\text{UPDATE}} = \perp$.

Search leakage. We define a few notations prior to analysing the SEARCH leakage to \mathcal{S} . We state the forward and backward privacy notions briefly here (refer [BMO17a] for more details). Forward privacy dictates that an update involving a w does not reveal any information of a prior search involving w. Backward privacy states that if a sequence of updates involve adding w and deleting later, a subsequent search involving w does not reveal that w was involved in these updates. We provide more details in Appendix F of this document.

Since NOMOS is an MRSW construction, \mathcal{S} and \mathcal{G} treat each \mathcal{C} identically - the views to \mathcal{G} and \mathcal{S} remain the same for all clients. Thus, the forward privacy and backward privacy analysis can be followed, considering that the views remain identical for all \mathcal{C} s. For this analysis, we assume a list \mathcal{Q} that stores the following information.

1. (t, w) : w searched at time t .
2. $(t, \text{op}, (w, \text{id}))$: (w, id) pair was updated with update type op at time t .

Let TimeDB be a function that takes a w as input and returns the respective ids along with the timestamp t . We express this as below.

$$\begin{aligned} \text{TimeDB}(w) = \{ & (t, \text{id}) \mid (t, \text{ADD}, (w, \text{id})) \in \mathcal{Q} \\ & \text{and } \forall t' : (t', \text{DEL}, (w, \text{id})) \notin \mathcal{Q} \} \end{aligned}$$

For a conjunctive query $q = w_1 \wedge \dots \wedge w_n$, the TimeDB notion is extended as follows.

$$\begin{aligned} \text{TimeDB}(q) = \{ & (\{t_i\}_{i \in [n]}, \text{id}) \mid (t, \text{ADD}, (w_i, \text{id})) \in \mathcal{Q} \\ & \text{and } \forall t' : (t', \text{DEL}, (w_i, \text{id})) \notin \mathcal{Q} \} \end{aligned}$$

This essentially corresponds to the ids along with timestamps which satisfies q and have not been deleted from EDB. We denote the following s-term and x-term leakages as below.

s-term leakage. Let UPD(w) for a w be defined as follows.

$$\text{UPD}(w) = \{t \mid \exists (\text{op}, \text{id}) : (t, \text{op}, (w, \text{id})) \in \mathcal{Q}\}$$

In summary, UPD(w) captures the s-term (w without loss of generality) leakages of q .

x-term leakage. We modify the UPD for a pair (w_1, w_2) in the following way.

$$\text{UPD}(w_1, w_2) = \{(t_1, t_2) \mid \exists(\text{op}, \text{id}) : (t_1, \text{op}, (w_1, \text{id})) \in \mathcal{Q} \\ \text{and } (t_2, \text{op}, (w_2, \text{id})) \in \mathcal{Q}\}$$

The above expression implies that $\text{UPD}(w_1, w_2)$ returns the timestamps of the update operations on w_1 and w_2 involving the same id . For a conjunctive query q , this essentially encapsulates the x -term leakage as $\{\text{UPD}(w_1, w_j)\}_{j \in [2, n]}$.

The s -term and x -term leakages are combined to obtain the total search leakage to the server as follows.

$$\mathcal{L}_S^{\text{SEARCH}} = \text{UPD}(q) = \text{UPD}(w_1) \cup \left(\bigcup_{j=2}^n \text{UPD}(w_1, w_j) \right)$$

Note that, \mathcal{S} is not involved in GENTOKEN protocol, and hence, $\mathcal{L}_S^{\text{GENTOKEN}} = \perp$. Finally, combining all we have

$$\mathcal{L}_S = \{\text{TimeDB}(q), \text{UPD}(q)\}.$$

D Redundant Bloom Filter

The plain Bloom filter (denoted by BF) is a probabilistic data structure suitable for membership checking within a set encoded in the BF. At a high level, to index an element (or insert), a BF uses k different hash functions to obtain k addresses, which are set to 1. During membership check (or look-up), k indices are obtained from the queried element and checked for 1's. These BF insertion and look-up/query routines are outlined in Algorithm 14 and Algorithm 15 respectively.

Note that if we directly plug BF into MC-ODXT, replacing the XSet insertion with BF.INSERT and XSet retrieval with BF.QUERY, the construction essentially works the same. However, the leakage is not mitigated as the bfidx_s (or xtags in the context of SSE) generated during BF query are deterministically generated using k hash functions. Hence, the server still can associate a w using the observed BF (or XSet) addresses.

Algorithm 14 Bloom Filter Insert

Input: Input parameters: x - element to be inserted into BF

Output: Output parameters:

- 1: **function** BF.INSERT(x)
 - 2: [Gate-keeper](#)
 - 3: Select k hash functions $\{h_1, \dots, h_k\}$ for BF indices
 - 4: Initialise empty index set BFIdxSet
 - 5: **for** $i \leftarrow 1$ to k **do**
 - 6: $\text{bfidx}_i \leftarrow h_i(x)$
 - 7: BFIdxSet = BFIdxSet \cup $\{\text{bfidx}_i\}$
 - 8: Shuffle elements in BFIdxSet
 - 9: Send BFIdxSet to the server
 - 10: [Server](#)
 - 11: **for** Each $\text{idx} \in \text{BFIdxSet}$ **do**
 - 12: Set BF[idx] = 1
-

We modify this basic BF construction to allow storing redundant elements (total ℓ indices generated from ℓ hash functions) to be stored for each xtag (corresponding to each (w, id) pair). During a query, BF accesses only a random subset of size k of these ℓ indices. Such that the server “sees” that each time different k indices are being accessed and can not correlate with the recorded information. We call this redundant element-based Bloom filter construction as a Redundant Bloom Filter (RBF). The updated RBF.INSERT and RBF.QUERY routines are provided in Algorithm 16 and Algorithm 17 respectively.

Algorithm 15 Bloom Filter Query

Input: Input parameters: x - element to be queried in **BF****Output:** Output parameters: True/False

```
1: function BF.QUERY( $x$ )
   Client
2:   Select  $k$  hash functions  $\{h_{i_1}, \dots, h_{i_k}\}$  for BF indices
3:   Initialise empty index set BFIdxSet
4:   for  $j \in \{i_1, \dots, i_k\}$  do
5:     bfid $x_j \leftarrow h_j(x)$ 
6:     BFIdxSet = BFIdxSet  $\cup \{bfidx_j\}$ 
7:   Shuffle elements in BFIdxSet
8:   Send BFIdxSet to the server
   Server
9:   for Each idx  $\in$  BFIdxSet do
10:    if BF[idx]  $\neq 1$  then
11:      Return False
12:   Return True
```

Algorithm 16 Redundant Bloom Filter Insert

Input: Input parameters: x - element to be inserted into RBF**Output:** Output parameters:

```
1: function RBF.INSERT( $x$ )
   Gate-keeper
2:   Select  $\ell$  hash functions  $\{h_1, \dots, h_\ell\}$  for RBF indices
3:   Initialise empty index set RIdxSet
4:   for  $i \leftarrow 1$  to  $\ell$  do
5:     rbfidx $_i \leftarrow h_i(x)$ 
6:     RIdxSet = RIdxSet  $\cup \{rbfidx_i\}$ 
7:   Shuffle elements in RIdxSet
8:   Send RIdxSet to the server
   Server
9:   for Each idx  $\in$  RIdxSet do
10:    Set RBF[idx] = 1
```

Algorithm 17 Redundant Bloom Filter Query

Input: Input parameters: x - element to be queried in **RBF****Output:** Output parameters: True/False

```
1: function RBF.QUERY( $x$ )
   Client
2:   Select  $k$  hash functions  $\{h_{i_1}, \dots, h_{i_k}\}$  for RBF indices
3:   Initialise empty index set RIdxSet
4:   for  $j \in \{i_1, \dots, i_k\}$  do
5:     rbfidx $_j \leftarrow h_j(x)$ 
6:     RIdxSet = RIdxSet  $\cup \{rbfidx_j\}$ 
7:   Shuffle elements in RIdxSet
8:   Send RIdxSet to the server
   Server
9:   for Each idx  $\in$  RIdxSet do
10:    if RBF[idx]  $\neq 1$  then
11:      Return False
12:   Return True
```

D.0.1 RBF Overhead

In RBF, the server sets ℓ locations, and accesses k locations during query, where $\ell > k$. The value of k needs to be chosen suitably to have negligible false positive probability (similar to normal BF) without blowing up storage.

Storage overhead. A conventional BF requires

$$k \cdot \sum_{w \in \Delta} |\mathbf{DB}(w)|$$

storage for BF with k hashes. Here, we have k hashes during insert, and ℓ hashes during queries. Hence, the storage requirement of RBF is

$$\ell \cdot \sum_{w \in \Delta} |\mathbf{DB}(w)|$$

RBF storage overhead is $\frac{\ell}{k}$ times (greater than one as $\ell > k$) than BF for the same database.

Communication overhead. An RBF sends k indices for a single element while inserting into and querying on BF. Thus, the communication overhead can be expressed as $O(1)$ for each inserted element (from $O(k)$, as k remains constant for a particular database). For a complete query $q = w_1 \wedge \dots \wedge w_n$, the query overhead can be expressed as follows.

$$k \cdot \sum_{w \in \{w_2, \dots, w_n\}} |\mathbf{DB}(w_1)|$$

With RBF, the communication overhead during the insertion of a single element is $O(1)$ (from $O(\ell)$) as ℓ remains constant for a database. For a conjunctive query of the form $q = w_1 \wedge \dots \wedge w_n$, the communication overhead can be estimated as follows,

$$k \cdot \sum_{w \in \{w_2, \dots, w_n\}} |\mathbf{DB}(w_1)|$$

since k indices are used for query (instead of all ℓ indices). Clearly, the communication overhead of RBF is $\frac{\ell}{k}$ times than BF (greater than one as $\ell > k$). However, if the same k values are chosen for RBF and BF, the communication overhead essentially remains the same for both RBF and BF during look-up.

D.1 Security Analysis of RBF

The redundant access in RBF aims to mitigate the cross-term leakage in a practical yet finite manner. We analyse the effect of repeated RBF look-ups to evaluate the leakage profile of RBF. Since RBF relies on plain redundancy rather than standard cryptographic hardness assumptions, the analysis primarily assesses the advantage of RBF compared to BF (for which the CROSSATTACK of Section III in the main manuscript works) with respect to the redundancy parameters ℓ and k .

Recall from Section IV.B of the main manuscript that RBF uses ℓ hashes to index a value during insertion, while it uses a random k -subset of these ℓ hashes during look-up to hide the access pattern. Since it requires more number of RBF accesses to deterministically relate a specific (w, id) pair to the ℓ indices in RBF, compared to a single one in BF, in practice it results in less leakage.

Previous works [AG22, HSWW18] explored that redundancy-based approaches would require a linear amount of storage on the client side to eliminate access pattern leakage (which indicates that pattern leakage may not be possible to eliminate in practice). Hence, RBF focuses on reducing the attack probability compared to a plain BF up to a limit. We argue in this discussion that, in practice, it is not feasible for the adversary to build the required (w, id) -xtag association within a reasonable time. Hence, this brings a notion of access upper-bound in RBF beyond which the advantage would reduce to a traditional BF. Analysing this RBF access upper bound is equivalent to addressing the following question.

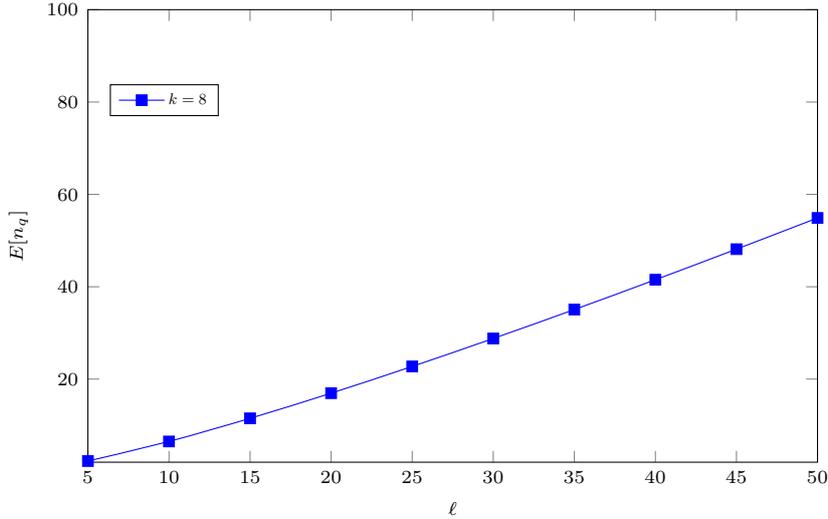


Figure 11: Growth of the RBF access upper bound with redundancy. The look-up indexing hashes were fixed at random 8-sized subset.

Maximum how many RBF access does the adversarial server needs to “see” for the same (w, id) pair from a benign client’s queries before it can figure out that the same (w, id) pair is being accessed?

As it turns out, this problem is identical to the well-known *coupon collector’s problem* with coupons collected in batches. When translated to the coupon collector’s problem, the problem statement can be expressed in the following way.

Maximum how many draws are necessary to collect all ℓ coupons if in each draw uniformly random k -subset of the ℓ coupons is collected with replacement?

The closed-form expression of this upper bound can be expressed as below.

$$E[n_q] = \sum_{j=1}^{\ell} (-1)^{j+1} \cdot \binom{\ell}{k} \cdot \frac{\binom{\ell}{j}}{1 - \frac{\binom{\ell-j}{k}}{\binom{\ell}{j}}} \quad (1)$$

The complete proof can be found in [Sta90] (albeit in a slightly different form). We plot the growth of this upper bound $E[n_q]$ with the number of indexing hashes ℓ in Figure 11, which shows the monotonically increasing behaviour. The parameter k remains constant to maintain the same desired false positive probability of a BF. This bound grows slower than the exponential that is ideally required to eliminate the access pattern completely in practice. We argue that if practical client search policies are formed based on the query semantics and the upper bound, this approach prevents cross-term-based leakage.

In reality, a client’s queries are concentrated on a particular topic, and a client typically issues only a limited number of queries in a search session. Therefore, the essential requirement of encountering the same (w, id) pair across these queries is limited to a query session only. It is unlikely that a client would issue more queries than the upper bound $E[n_q]$ in a session, as plotted in Figure 11. Thus, the administrator can enforce the search policy such that a client cannot issue more queries than the upper bound $E[n_q]$. We state the following assumption to reflect the RBF’s influence in NOMOS security analysis.

Assumption. Given an RBF with ℓ hashes for insertion and a uniformly random k -subset of these ℓ hashes for look-up, τ number of RBF look-ups for the same (w, id) pair are indistinguishable from a random RBF element (corresponding to a random (w, id) pair) look-up using a uniformly random k -subset of these ℓ hashes up to $E[n_q]$ accesses, where $\tau < E[n_q]$.

Remarks. The current analysis of access bound assumes the server would definitely be able to figure out after the expected number of queries given by the closed-form expression in Equation (1). In that case, the problem is the same as the Coupon collector’s problem, as stated above and followed in the abovementioned analysis. In case of an approximation or relatedness measure is considered in the attack process instead of an exact match for x_{tags} to figure out if two sets of k x_{tags} are “related” rather than knowing that these are from exactly the same (w, id) pair. In that case, a different analysis approach is necessary. However, the exact-match base analysis suffices to provide an in-depth overview of the RBF security, and we resorted to the upper-bound-based analysis to choose RBF parameters and simplify the analysis, compared to a “relation”-based analysis which we leave for as an extension of this work.

E Security of Nomos

We follow the formal leakage profile outlined in Section C of this document for the proof sketch of Theorem E.1 and E.2 of the main manuscript in this section. We use two additional cryptographic assumptions for the formal security proofs - the One-More Gap Diffie-Hellman (OM-GDH) assumption and the Extended Decisional Diffie-Hellman (EDDH) assumption. We first state the additional cryptographic assumptions that are necessary for the proofs below.

One-More Gap Diffie-Hellman Assumption. Denote a prime order cyclic group \mathbb{G} with order p (polynomially large in the security parameter λ) and its generator by g . The One-More Gap Diffie-Hellman (OM-GDH) assumption holds in \mathbb{G} if the advantage $\text{Adv}_{\mathbb{G}, \mathcal{A}}^{\text{DDH}}(\lambda)$ is negligible for all adversaries \mathcal{A} . $\text{Adv}_{\mathbb{G}, \mathcal{A}}^{\text{DDH}}(\lambda)$ is the probability of \mathcal{A} winning the following game.

The game samples a $r \leftarrow \mathbb{Z}_p^*$ at random, and samples two other random elements of \mathbb{G} - (h_1, h_2) . The values (h_1, h_2) are shared with \mathcal{A} and it makes a query to the Decisional Diffie-Hellman (DDH) oracle that returns $b \leftarrow a^r$ upon receiving a . \mathcal{A} is allowed make any number of queries to a DDH oracle $\text{DDH}_t(\cdot, \cdot)$, which takes input as (h, v) and returns 1 if $v = h^r$ or 0 otherwise. At the end of game, \mathcal{A} outputs (v_1, v_2) and it wins the game if $v_1 = (h_1)^r$ and $v_2 = (h_2)^r$.

Extended Decisional Diffie-Hellman Assumption. For a prime order cyclic group \mathbb{G} and its generator g , and two arbitrary integers $m, n \in \mathbb{N}$, define the following matrix

$$M := \begin{bmatrix} g^{\alpha_1 \cdot \beta_1} & \dots & g^{\alpha_1 \cdot \beta_n} \\ \vdots & \ddots & \vdots \\ g^{\alpha_m \cdot \beta_1} & \dots & g^{\alpha_m \cdot \beta_n} \end{bmatrix}$$

where $\alpha_i \leftarrow \mathbb{Z}_p^*$, $i \in [m]$ and $\beta_j \leftarrow \mathbb{Z}_p^*$, $j \in [n]$. The extended DDH assumption states that

$$|\text{Pr}[\mathcal{A}(g, M) = 1] - \text{Pr}[\mathcal{A}(g, M') = 1]| \leq \text{negl}(\lambda)$$

where M' is distributed as follows,

$$M := \begin{bmatrix} g^{\gamma_{1,1}} & \dots & g^{\gamma_{1,n}} \\ \vdots & \ddots & \vdots \\ g^{\gamma_{m,1}} & \dots & g^{\gamma_{m,n}} \end{bmatrix}$$

where $\gamma_{i,j} \leftarrow \mathbb{Z}_p^*$, $i \in [m]$, $j \in [n]$.

Search leakages in SSE schemes. An SSE scheme typically incurs the following leakages to the server during search execution - keyword frequency, equality pattern, size pattern, result pattern, and conditional intersection pattern.

Search leakages in SSE schemes. The following leakages are typically incurred by an SSE scheme to the server during SEARCH execution. Assume that Q is a sequence of conjunctive queries issued over time.

Keyword frequency. The total number of time the keywords appear in documents: $N = \sum_{i=1}^d |\Delta(w_i)|$.

Equality pattern. Equality pattern corresponds to the queries that have equal s-terms. Typically, equality pattern $\bar{s} \in [|\Delta|]^Q$ is expressed as a sequence of s-term, where each s-term is assigned an integer with repetitions for each s-term.

Size pattern. Size pattern refers to the number of documents retrieved from encrypted database for the s-term of each query.

Result pattern. Result pattern is the set of ids matching the conjunction of the keywords in the query.

Conditional intersection pattern. Conditional intersection pattern refers to the ids matching a common cross-term for queries with different s-terms.

E.1 Security notions of Nomos

Our security analysis of NOMOS follows the standard ideal/real indistinguishability framework parameterised by a leakage function \mathcal{L} capturing the information leaked to an adversarial entity, as outlined in Section II.C of the main manuscript. We quantify the leakage to a semi-honest client using \mathcal{L}_C and the leakage to a semi-honest server using \mathcal{L}_S . We defer the detailed discussion of \mathcal{L}_C and \mathcal{L}_S to Appendix C due to space constraint.

Security against adversarial \mathcal{C} . We follow the analysis approach from [JJK⁺13] for analysing the security of NOMOS against adversarial clients. Concretely, we present the following definition to analyse the security of NOMOS against adversarial clients following the generic definition of Section II.C of the main manuscript. Definition E.1 compares the real execution to an emulated interaction of $\text{SSE}_{\mathcal{L}_C}$ that models the functionality of NOMOS instantiated from \mathcal{L}_C (see Appendix C). $\text{SSE}_{\mathcal{L}_C}$ takes $(\mathbf{DB}, \mathcal{P})$ as input and process queries q if $I(q) \in \mathcal{P}$. If $I(q) \in \mathcal{P}$, it replies with $\mathbf{DB}(q), \mathcal{L}_C$; otherwise it returns error symbol \perp .

Definition E.1 (Security against an adversarial client). *Let Π be a NOMOS SSE scheme, where*

$$\Pi = \{\text{SETUP}, \text{UPDATE}, \text{GENTOKEN}, \text{SEARCH}\}.$$

Define the following $\mathbf{Real}_{\mathcal{A}}^{\Pi}$ and $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$ experiments (algorithms with running time in 1^λ) as below, provided \mathcal{L}_C , \mathcal{A} , and $\text{SIM} = \{\text{SIM}_0, \text{SIM}_1, \text{SIM}_2\}$.

$\mathbf{Real}_{\mathcal{A}}^{\Pi}$: *\mathcal{A} chooses \mathbf{DB} and the experiment executes SETUP to receive sk and the repeatedly runs UPDATE to obtain \mathbf{EDB} . After that, \mathcal{A} adaptively invokes GENTOKEN and SEARCH with input sk , where \mathcal{A} interacts with \mathcal{G} and \mathcal{S} , respectively. Let t be the number of GENTOKEN+SEARCH instances invoked and τ_i be the local output of \mathcal{G} in i 'th instance. As above, if at any point \mathcal{A} halts and output a bit b , the game outputs $(b, \mathbf{EDB}, \tau_1, \dots, \tau_t)$.*

$\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$: *\mathcal{A} chooses \mathbf{DB} and the experiment initialises*

$$\text{SIM} = \{\text{SIM}_0, \text{SIM}_1, \text{SIM}_2\}$$

by executing $st \leftarrow \text{SIM}_0(\lambda)$. After that, each time \mathcal{A} invokes GENTOKEN, it interacts with $\text{SIM}_1(st)$ and each time \mathcal{A} invokes SEARCH, it interacts with $\text{SIM}_2(st)$. Both SIM_1 and SIM_2 are allowed to update global state st while interacting with \mathcal{A} . Both can issue queries to an ideal emulation of NOMOS. Let t be the number of such queries and let $\tau_i = I(q)$. As above, if at any point \mathcal{A} halts and output a bit b , the game outputs $(b, \mathbf{EDB}, \tau_1, \dots, \tau_t)$.

Π is called \mathcal{L}_C -semantically-secure against an adversarial \mathcal{C} s if for any efficient \mathcal{A} , there is an efficient algorithm SIM , such that statistical difference between $(b, \mathbf{EDB}, \tau_1, \dots, \tau_t)$ outputs from $\mathbf{Real}_{\mathcal{A}}^{\Pi}$ and $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$ experiments is negligible in λ .

Based on the above definition, we present the following theorem about NOMOS security against adversarial clients.

Theorem E.1. *A NOMOS scheme instantiated with Hashed Diffie-Hellman OPRF is adaptively \mathcal{L}_C -semantically-secure against adversarial semi-honest clients provided that the DH assumption holds in \mathbb{G} , F_p and F are secure PRFs, and $(\text{AE.Enc}, \text{AE.Dec})$ is an IND-CPA and strongly UF-CMA-secure AE scheme, and all hash functions are modelled as random oracles.*

Proof. The proof is given in Appendix E.2. □

Security against adversarial \mathcal{S} . The adaptive security analysis against an adversarial \mathcal{S} follows from ODXT, which is $\mathcal{L}_{\text{ODXT}}$ -semantically-secure. NOMOS is $\mathcal{L}_{\mathcal{S}}$ -semantically-secure that absorbs $\mathcal{L}_{\text{ODXT}}$. The definition below follows from the generic security state of Section II.C of the main manuscript.

Definition E.2 (Security against an adversarial server). *Let Π be a NOMOS SSE scheme, where*

$$\Pi = \{\text{SETUP}, \text{UPDATE}, \text{GENTOKEN}, \text{SEARCH}\}.$$

Define the following $\mathbf{Real}_{\mathcal{A}}^{\Pi}$ and $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$ experiments (algorithms with running time in 1^{λ}) as below, provided $\mathcal{L}_{\mathcal{S}}$, \mathcal{A} , and $\text{SIM} = \{\text{SIM}_0, \text{SIM}_1, \text{SIM}_2\}$.

$\mathbf{Real}_{\mathcal{A}}^{\Pi}$: *\mathcal{A} chooses DB and the experiment executes SETUP to receive sk and the repeatedly runs UPDATE to obtain EDB . After that, \mathcal{A} adaptively invokes UPDATE and SEARCH with input sk , where \mathcal{A} interacts with \mathcal{G} and \mathcal{C} , respectively. Let t be the number of $\text{UPDATE}+\text{SEARCH}$ instances invoked and τ_i be the local output of \mathcal{G} or \mathcal{C} in i 'th instance. As above, if at any point \mathcal{A} halts and output a bit b , the game outputs $(b, \text{EDB}, \tau_1, \dots, \tau_t)$.*

$\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$: *\mathcal{A} chooses DB and the experiment initialises*

$$\text{SIM} = \{\text{SIM}_0, \text{SIM}_1, \text{SIM}_2\}$$

by executing $st \leftarrow \text{SIM}_0(\lambda)$. After that, each time \mathcal{A} invokes UPDATE , it interacts with $\text{SIM}_1(st)$ and each time \mathcal{A} invokes SEARCH , it interacts with $\text{SIM}_2(st)$. Both SIM_1 and SIM_2 are allowed to update global state st while interacting with \mathcal{A} . Both can issue queries to an ideal emulation of NOMOS. Let t be the number of $\text{UPDATE}+\text{SEARCH}$ queries and let τ_i be the local output SIM_1 or SIM_2 . As above, if at any point \mathcal{A} halts and output a bit b , the game outputs $(b, \text{EDB}, \tau_1, \dots, \tau_t)$.

Π is called $\mathcal{L}_{\mathcal{S}}$ -semantically-secure against adversarial \mathcal{S} if for any efficient \mathcal{A} , there is an efficient algorithm SIM , such that statistical difference between $(b, \text{EDB}, \tau_1, \dots, \tau_t)$ outputs from $\mathbf{Real}_{\mathcal{A}}^{\Pi}$ and $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$ experiments is negligible in λ .

Specifically, the PRF instances in ODXT are replaced with OPRF, and \mathcal{S} 's view of SEARCH protocol in NOMOS can be generated from SEARCH of ODXT with minor modifications due to the group exponentiations by ρ and γ values that are randomly sampled from \mathbb{Z}_p^* . We state the theorem below where $\mathcal{L}_{\mathcal{S}}$ is defined as in Appendix C.

Theorem E.2. *A NOMOS scheme instantiated with DH OPRF is adaptively $\mathcal{L}_{\mathcal{S}}$ -semantically-secure against adversarial server provided that the DH assumption holds in \mathbb{G} , F_p and F are secure PRFs, and $(\text{AE.Enc}, \text{AE.Dec})$ is an IND-CPA and strongly UF-CMA-secure AE scheme, and all hash functions are modelled as random oracles.*

Proof. The proof is given in Appendix E.3. □

E.2 Security against adversarial client

We build the simulators $\text{SIM} = (\text{SIM}_0, \text{SIM}_1, \text{SIM}_2)$ which are described Algorithm 18, 19 and 20. We prove Theorem E.1 via a sequence of games, where $\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda)$ models the interaction of \mathcal{A} with real instance of NOMOS, and $\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}(\lambda)$ models the interaction of \mathcal{A} with ideal instance of NOMOS using SIM . We outline the sequence of games below along with concise proofs of the associated lemmas based on the games.

Game G_1 . We modify game G_0 by adding an abort if the number of requested tokens in a session exceeds $E[n_q]$.

Lemma E.1. *Game G_1 is indistinguishable from game G_0 .*

Proof. By the indistinguishability property of RBF (within the limit), G_1 is indistinguishable from G_0 . □

Game G_2 . We modify G_1 to add an abort if env is accepted by \mathcal{S} within SEARCH is not genuinely generated by \mathcal{G} in GENTOKEN .

Algorithm 18 Simulator SIM_0

Input: λ **Output:** $\{K_S, K_Y, K_X, K_T, K_M\}$, and QList, TList

- 1: **function** $\text{SIM}_0(\lambda)$
 - 2: Select keys K_S, K_X and K_T for OPRF, K_Y for F_p and K_M for AE.
 - 3: Initialise an empty table QList, to be indexed by AE ciphertexts env
 - 4: Initialise an empty table TList, to be indexed by ws in Δ (TList initially holds and empty set for each w)
 - 5: Return keys K_S, K_Y, K_X, K_T, K_M and tables QList, TList
-

Algorithm 19 Simulator SIM_1

Input: st, \mathcal{P} **Output:** Updated st

- 1: **function** $\text{SIM}_1(\text{st})$
 - 2: Get tuples $(a_1, \dots, a_n), (b_1, \dots, b_m), (c_1, \dots, c_m)$, and tuple (I_1, \dots, I_n) as input from \mathcal{A}
 - 3: Abort if $(I_1, \dots, I_n) \notin \mathcal{P}$
 - 4: Sample $\rho_i \xleftarrow{\$} \mathbb{Z}_p^*$ for $i \in [n]$
 - 5: Pick $\gamma_i \xleftarrow{\$} \mathbb{Z}_p^*$ for $i \in [m]$
 - 6: Compute $\tau_s = (a_1)^{K_S}, \tau_1 = (a_1)^{K_T[I_1] \cdot \rho_1}$, and $\tau_i = (a_i)^{(K_X[I_i] \cdot \rho_i)}$ for $i \in [2, n]$
 - 7: Compute $\epsilon_i = (b_i)^{(K_T[I_1] \cdot \gamma_i)}$ for $i \in [m]$
 - 8: Compute $\mu_i = (c_i)^{(K_T[I_1])}$ for $i \in [m]$
 - 9: **for** $i = 2$ to n **do**
 - 10: Initialise $\bar{\tau}_i$
 - 11: **for** $j = 1$ to k **do**
 - 12: Sample $\beta_j \xleftarrow{\$} \mathbb{Z}_p^*$
 - 13: Compute $\tau'_{i,j} = \tau_i^{\beta_j}$
 - 14: Compute $\bar{\tau}_i = \bar{\tau}_i \cup \{\tau'_{i,j}\}$
 - 15: Update QList within st: set $\text{QList}(\text{env}) \leftarrow (I_1, \dots, I_n; \rho_1, \dots, \rho_n; \gamma_1, \dots, \gamma_m)$
 - 16: Encrypt $\text{env} = \text{AE.Enc}_{K_M}(\rho_1, \dots, \rho_n, \gamma_1, \dots, \gamma_m)$
 - 17: Output $(\text{env}, \tau_s, \tau_1, \bar{\tau}_2, \dots, \bar{\tau}_n, \epsilon_1, \dots, \epsilon_m, \mu_1, \dots, \mu_m)$
-

Lemma E.2. Game G_2 is indistinguishable from game G_1 .

Proof. By Strong-UF-CMA unforgeability of AE scheme, G_2 is indistinguishable from G_1 . \square

Game G_3 . We abort if two different GENTOKEN instances generate the same env.

Lemma E.3. Game G_3 is indistinguishable from G_2 .

Proof. The blinding factors ρ_i 's are randomly sampled from \mathbb{Z}_p^* . A collision in ciphertext implies the same in the plaintext, hence; it will be a contradiction to have such an occurrence. Therefore, G_3 is indistinguishable from G_2 . \square

Game G_4 . Abort if for any two distinct w_1 and w_2 in Δ , two OPRF instances from different keys K_T or K_X collide – two distinct OPRF evaluations involving w_1 and w_2 using K_T or K_X output the same value.

Lemma E.4. G_4 is indistinguishable from G_3 .

Proof. Since OPRFs using K_T and K_X are essentially secure PRF instances, probability of such collisions is negligible, and hence; G_4 is indistinguishable from G_3 . \square

Game G_5 . In this game, we change the way of interaction with the SEARCH protocol. The game generates $\text{stag} = \text{bstag}^{1/\rho_1}$ in G_0 . After that, it searches within Δ to find a w such that corresponding stag generation results in stag. If such stag is found, it obtains corresponding value from from TSet as a tuple (sval, α) ; otherwise aborts.

Algorithm 20 Simulator SIM_2

Input: st , upon receiving $(\text{env}, \text{stokenList}, \text{xtokenList}_1, \dots, \text{xtokenList}[m])$ from \mathcal{A}

Output: Updated st

```
1: function  $\text{SIM}_2(\text{st})$ 
2:   Retrieve  $(I_1, \dots, I_n; \rho_1, \dots, \rho_n, \gamma_1, \dots, \gamma_m)$  from  $\text{QList}(\text{env})$ , abort if  $\text{QList}(\text{env}) = \perp$ 
3:   For a  $w_1 \in \Delta$ , s.t.  $\text{stag} = (H(w_1))^{K_T[I_1]}$ , and  $I(w_1) = I_1$ ; set  $\text{strap} = (H(w_1))^{K_S}$  and
    $K_Z = F(\text{strap}, 1)$ . Abort if there is no such  $w_1$ .
4:   for  $i = 1$  to  $\text{stokenList.size}$  do
5:      $\text{stag}_s = \text{stokenList}[i]$ 
6:     For a  $w_s \in \Delta$ , s.t.  $\text{stag}_s = (H(w_s))^{K_T[I_{i+1}]}$ , and  $I(w_{i+1}) = I_{i+1}$ . Abort if there is no
     such  $w_s$ .
7:   Initialise  $c \leftarrow 0$  and  $\text{found} \leftarrow \text{False}$ . Execute the following loop till  $\text{found} == \text{False}$ 
   • Increment  $c = c + 1$ , and  $z_c = F_p(K_Z, w_1 || c)$ . For an  $\text{xtokenList}[c] =$ 
    $(\text{xtokenList}[c, 2], \dots, \text{xtokenList}[c, n])$ , if  $\exists (w_2, \dots, w_n) \in \Delta^{n-1}$ , s.t.  $\text{xtokenList}[c, i] =$ 
    $(H(w_i))^{K_X[I(w_i)] \cdot z_c \cdot \alpha_c \cdot \rho_i}$  for  $i = 2, \dots, n$ , then set  $\text{found} = \text{True}$ .
   • Abort if  $\text{found} = \text{False}$  and  $\text{xtoken}[c]$  is the last element in  $\mathcal{A}$ 's message.
8:   Send  $q = (w_1, \dots, w_n)$  found in the last step to  $\text{SSE}_{\mathcal{L}}$ . As  $\mathcal{P}$  contains  $\text{av}$  with certainty,
    $\text{SIM}_2$  gets back  $\text{DB}(q)$  and  $\text{TSetL}$ .
9:   Initialise  $S' \leftarrow \emptyset$  and retrieve  $D \leftarrow \text{DB}(q)$ . Check  $\forall \text{id} \in D$  s.t.  $(c, \text{id}, e) \in \text{TList}(w_1)$ . Add
    $c$  to  $S'$  and delete  $\text{id}$  from  $D$ .
10:  Sample  $S$  as a  $|D|$ -sized subset in  $\{1, \dots, \text{TSetL}\} \setminus S'$ 
11:  While  $S$  is non-empty do:
   • Remove an element  $c$  from  $S$  and an  $\text{id}$  from  $D$  at random
   • Compute  $e \leftarrow (\text{id} || \text{op}) \oplus (H(w_i || c || 1))^{K_T[I(w_i)]}$ .
   • Insert  $(c, \text{id}, e)$  to  $\text{TList}(w_1)$  into  $\text{TList}$ .
12:  Execute loop while  $c \leq \text{TSetL}$ , without initialising  $c$  (retaining the last value):
   • Compute  $z_c = F_p(K_Z, w_i || c)$ . If  $\text{xtokenList}[c, i] = (H(w_i))^{K_T[I(w_i)] \cdot z_c \cdot \alpha_c \cdot \rho_i}$  for  $i =$ 
    $2, \dots, n$  and if  $\exists (c, \text{id}, e)$  in  $\text{TList}(w_1)$  s.t.  $\text{id} \in \text{DB}(q)$ , then send  $e$  to  $\mathcal{A}$ . Increment
    $c = c + 1$ .
13:  halt after sending  $\text{stop}$  to  $\mathcal{A}$ .
```

Lemma E.5. G_5 is indistinguishable from G_4 .

Proof. Since collisions in stag have been eliminated in previous games and each stag is uniquely generate from w , the retrieval in G_5 is essentially the same as from the stag received as in G_4 with overwhelming probability. In contrast, if two stags do not match, the TSet has negligible probability of returning a non-empty result. Therefore, in this case G_5 is indistinguishable from G_4 . \square

Game G_6 . In this game, instead of encrypting $(\rho_1, \dots, \rho_n, \gamma_1, \dots, \gamma_m)$, the game encrypts a set of randomly generated values $(\rho'_1, \dots, \rho'_n, \gamma'_1, \dots, \gamma'_m)$. The game retains QList indexed by env . When responding to a GENTOKEN request, it also stores the attributes $(I(w_1), \dots, I(w_n))$ and actual blinding factors $(\rho_1, \dots, \rho_n, \gamma_1, \dots, \gamma_m)$. During SEARCH , the game retrieves $(\rho_1, \dots, \rho_n, \gamma_1, \dots, \gamma_m)$ from QList .

Lemma E.6. G_6 is indistinguishable from G_5 .

Proof. In game G_3 , only envs that are uniquely generated in GENTOKEN invocations are accepted. Hence, by IND-CCA guarantee of AE , G_6 is indistinguishable from G_5 . \square

Game G_7 . We consider the case where the game identifies w_1 from bstag and ρ_1 recovered from $\text{QList}(\text{env})$, such that $(\text{bstag})^{1/\rho_1} = (H(w_1))^{K_T[I_1]}$ where $I_1 = I(w_1)$. This is identical to G_5 , except G_7 ignores $\text{stag} = \text{bstag}^{1/\rho_1} = (H(w_1))^{K_T[I(w_1)]}$ but $I_1 \neq I(w_1)$.

Lemma E.7. The probability of finding a w_1 satisfying the above conditions is negligible and thus G_7 and G_6 are indistinguishable.

Proof. Denote $K_T[I_1] \cdot \rho_1$ as e_j exponent for obtaining b_1 from a_1 during j 'th GENTOKEN invocation. By construction, e_j s are randomly uniformly distributed in \mathbb{Z}_p^* . The process of validating whether w_1 is associated with **bstag** in SEARCH can be alternatively written as validating $\text{bstag} = (H(w_1))^{(K_T[I(w_1)]/K_T[I_1]) \cdot e_j}$ with e_j used to generate **env** via GENTOKEN to be used in the subsequent SEARCH instance. In that case, $(H(w_1))^{e_j}$ and **bstag** must have same discrete logarithm equal to $K_T[I(w_1)]/K_T[I_1]$. As K_T is not used in any other way in G_7 and K_T sub-keys are all randomly uniformly sampled from \mathbb{Z}_p^* and the number of validations is polynomially bounded, the success probability of $I(w_1) = I_1$ is negligible. Thus G_7 and G_6 are indistinguishable from each other. \square

Game G_8 . We add an abort two different SEARCH invocations by \mathcal{A} have the same **env** but two different **bstags** - $\{\text{bstag}, \text{bstag}'\}$, such that $\text{bstag}^{1/\rho_1} = (H(w_1))^{K_T[I(w_1)]}$ and $\text{bstag}'^{1/\rho_1} = (H(w_1'))^{K_T[I(w_1)]}$. Denote by e_j the operation $K_T[I(w_1)]$ done in j 'th instance of GENTOKEN.

Lemma E.8. *Two pairs (w_1, bstag) and (w_1', bstag') having the same **env** is negligible and therefore G_8 is indistinguishable from G_7 .*

Proof. The game emulates G_4 by sampling r_1 , and r_2 from \mathbb{Z}_p^* and computes $H(w)$ as $(h_1)^{r_1}(h_2)^{r_2}$ where h_1, h_2 DH challenge inputs. Also, it picks a random index j in $[1, \eta]$ where η is the number of maximum invocations of GENTOKEN by \mathcal{A} allowed by \mathcal{A} . The experiment sends a_1 to the DH challenger that outputs $b_1 = (a_1)^t$ where t is chosen by DH challenger, and b_1 is passed to \mathcal{A} . Subsequently, for each SEARCH invocation where \mathcal{A} inputs **env**, the experiment takes **bstag** as input from \mathcal{A} and **env**. For \mathcal{A} 's each query q to OPRF, the DDH oracle is consulted with $(a_1, b_1, H)(w), \text{bstag}$ to figure out if this is a valid DDH tuple. In case two invocations of SEARCH the checks are verified with (w, bstag) and $(w', \text{bstag}') \neq (w, \text{bstag})$, the following computation is performed - $(\text{bstag}^{r_2}(\text{bstag}')^{-r_2})^{1/(r_1r_2 - r_1'r_2')}$ and $(\text{bstag}^{r_1}(\text{bstag}')^{-r_1})^{1/(r_1'r_2' - r_1r_2)}$, where $H(w) = (h_1)^{r_1}(h_2)^{r_2}$ and $H(w') = (h_1)^{r_1'}(h_2)^{r_2'}$. The success probability of this event is $(1/m_c)$ time the probability of the event G_7 . Here, m_c is the maximum number of GENTOKEN instances (an upper bound) excluding the probability of $r_1r_2' = r_1'r_2$. \square

Game G_9 . In this game, we replace the PRF instances of the form $F_p(K, \cdot)$ with a random function $F_R(\cdot)$ with a range onto \mathbb{Z}_p^* .

Lemma E.9. *G_9 and G_8 are indistinguishable.*

Proof. By the indistinguishability property of a PRF G_9 and G_8 are indistinguishable. \square

Game G_{10} . In this game the SEARCH process of G_9 is modified in the following way. G_{10} finds w_1 using **bstag** and (I_1, ρ_1) from QList(**env**) as stated earlier and computes **strap** \leftarrow OPRF(K_S, w_1) and $(K_Z, K_e) \leftarrow (F_\tau(\text{strap}, 1), F_\tau(\text{strap}, 2))$. After that, it computes $z_c = F_p(K_Z, c)$ for each c . Provided

$$\text{xtokenList}[c] = (\text{xtokenList}[c, 2], \dots, \text{xtokenList}[c, n]),$$

it looks for a w_i in Δ such that $\text{xtokenList}[c, i] = (\text{OPRF}(K_X, w_i))^{z_c \cdot \rho_i}$ and $\text{id}_c \in \mathbf{DB}(w_i)$ where id_c is the c 'th (e, y) entry in $\mathbf{t} = \mathbf{T}[w_1]$. Any such w_i is found for all i is sent to \mathcal{A} .

Let $\Delta(\text{id})$ denote the set of w s which appear in id . For G_{10} to succeed for any c, i in the above check, G_9 must be successful as $\text{xtokenList}[c, i] = (\text{OPRF}(K_X, w_i))^{z_c \cdot \rho_i}$ would imply $(\text{xtokenList}[c, i])^{y_c/\rho_i}$ is equal to $(\text{OPRF}(K_X, w_i))^{\text{id}_c}$. If $\text{id}_c \in \mathbf{DB}(w_i)$ then this particular **xtrap** entry is present in **XSet**. G_{10} and G_9 can only differ if $(\text{xtokenList}[c, i])^{y_c/\rho_i} \in \mathbf{XSet}$ and $\text{xtokenList}[c, i] \neq (\text{OPRF}(K_X, w))^{z_c \cdot \rho_i}$ for any (c, i) and for all $w \in \Delta(\text{id}_c)$.

Denote the exponent $K_X[I_i] \cdot \rho_i$ with a_j in GENTOKEN as $e_{i,j}$. G_9 can choose $e_{i,j}$ a random value in \mathbb{Z}_p^* and compute a ρ_i as $e_{i,j}/K_X[I_i]$. Therefore, $(\text{xtokenList}[c, i])^{y_c/\rho_i}$ finds a match in **XSet** as the existence of an id and $\bar{w} \in \Delta(\text{id})$ such that following expression holds.

$$(\text{xtokenList}[c, i])^{1/z_c} = (H(\bar{w}))^{e_{i,j} \cdot \frac{K_X[I_i(\bar{w})]}{K_X[I_i]} \cdot \frac{F_R(\bar{\text{id}})}{F_R(\text{id}_c)}} \quad (2)$$

Also, as $\text{xtokenList}[c, i] \neq (\text{OPRF}(K_X, w))^{z_c \cdot \rho_i}$ for $w \in \Delta(\text{id}_c)$, the following expression holds.

$$(\text{xtokenList}[c, i])^{1/z_c} \neq (H(w))^{e_{i,j} \cdot \frac{K_X[I(w)]}{K_X[I_i]}} \quad (3)$$

Lemma E.10. *The probability of that Equation (2) holds and Equation (3) does not hold for every $w \in \Delta$ is negligible and G_{10} is indistinguishable from G_9 .*

Proof. If Equation (2) holds for $\bar{\text{id}} = \text{id}_c$, and $I(\bar{w}) = I_i$, then $(\text{xtokenList}[c, i])^{1/z_c} = (H(\bar{w}))^{e_{i,j}}$ for $\bar{w} \in \Delta(\text{id}_c)$. However, this is a contradiction with Equation (3) where $I(\bar{w}) = I_i$ implies $(\text{xtokenList}[c, i])^{1/z_c} \neq (H(\bar{w}))^{e_{i,j}}$. We consider the following two cases below for $\bar{\text{id}} \neq \text{id}_c$ (case 1) and $I(\bar{w}) \neq I_i$ (case 2) to show that these two cases occur with only negligible probability.

Case 1. This argument essentially shows that $\bar{\text{id}} \neq \text{id}$ holds for negligible probability. In order to show that, G_9 is slightly modified to G'_9 where (e, y) is not appended to $T[w]$ tuples. Instead, it appends (e, id, z) such that does not have to query $F_I(\cdot)$ while creating \mathbf{T} . Furthermore, it does not create \mathbf{XSet} during SETUP . The test procedure of G_9 is modified for each c and i in SEARCH in the following way. G'_9 searches for $\bar{w} \in \Delta$ and $\bar{\text{id}} \in \mathbf{DB}(w)$ rather than checking $(\text{xtokenList}[c, i])^{y_c/\rho_i}$ in \mathbf{XSet} , using id_c, z_c stored in c 'th entry in $\mathbf{T}[w_1]$. These two are equivalent cases and G'_9 essentially is an identical view of G_9 . Additionally, test of Equation (2) can be done by checking $a = b^{F_I(\bar{\text{id}})/F_I(\text{id}_c)}$ where $a = (\text{xtokenList}[c, i])^{1/z_c}$ and $b = (H(\bar{w}))^{e_{i,j} \cdot (K_X[I(\bar{w})]/K_X[I_i])}$. In summary, G'_9 can work with $F_I(\cdot)$ as follows. G'_9 presents a tuple (a, b, x, \bar{x}) to the DDH oracle for 1 if $a^{F_I(x)} = b^{F_I(\bar{x})}$, otherwise 0. As H maps to \mathbb{G} , all (a, b, x, \bar{x}) invocation in G'_9 has $b \neq 1$ (except for negligible probability). Also, as the random function $F_I(\cdot)$ maps to \mathbb{Z}_p^* which is invoked polynomially bounded number of times by G'_9 , there is negligible probability of that the oracle returns $x \neq \bar{x}$. This implies that Equation (2) holds with negligible probability.

Case 2. In this case, we show that Equation (3) holds with negligible probability for $I(\bar{w}) \neq I_i$. Recall that, G'_9 uses K_X for testing Equation (2) that is equivalent to presenting a query of the form (a, b, I, \bar{I}) where $a = (\text{xtokenList}[c, i])^{1/z_c \cdot F_I(\text{id}_c)}$, $b = (H(\bar{w}))^{e_{i,j} \cdot F_I(\bar{\text{id}})}$, $I = I_i$, and $\bar{I} = I(\bar{w})$ and gets back 1 if $a^{K_X[I]} = b^{K_X[\bar{I}]}$. Since H randomly maps to \mathbb{G} , $b = 1$ occurs with negligible probability. Furthermore, as K_X is randomly constructed by choosing from \mathbb{Z}_p and G'_9 makes polynomially many queries, the probability of such successful queries is negligible. \square

Game G_{11} . In this game, G_{10} is modified to select w_i from $\text{xtokenList}[c, i]$ and (z_c, ρ_i) with verification using $\text{xtoken}[c, i] = (\text{OPRF}(K_X, w_i))^{z_c \cdot \rho_i}$ and $\text{id}_c \in \mathbf{DB}(w_i)$ and $I(w_i) = I_i$ retrieved from $\mathbf{QList}(\text{env})$. The differentiating part of the games - $\text{xtokenList}[c, i] = (H(w_i))^{K_X[I(w_i)] \cdot z_c \cdot \rho_i}$ given $\text{id}_c \in \mathbf{DB}(w_i)$ and $I(w_i) \neq I_i$.

Lemma E.11. *The above event occurs with negligible probability and therefore, G_{11} is indistinguishable from G_{10} .*

Proof. Recall the notation $e_{i,j} = K_X[I_i] \cdot \rho_i$ used in GENTOKEN . $\text{xtokenList}[c, i]$ can be expressed as

$\text{xtokenList}[c, i] = (H(w_i))^{z_c \cdot e_{i,j} \cdot \frac{K_X[I(w_i)]}{K_X[I_i]}}$. In summary, the discrete log of this expression is equal to $K_X[I_i]/K_X[I(w_i)]$. K_X is not utilised by G_{11} in any other way for this validation. As K_X is generated randomly from \mathbb{Z}_p^* and allowed to make only polynomial number of checks, the probability of successful checks for $I(w_i) \neq I_i$ is negligible and G_{11} is indistinguishable from G_{10} . \square

Game G_{12} . The game aborts if it receives the following - env , an index i , two counter c, c' and two keywords $w_i \neq w'_i$, such that $\text{xtokenList}[c, i] = (H(w_i))^{K_X[I_i] \cdot z_c \cdot \rho_i}$ and $\text{xtokenList}[c', i] = (H(w'_i))^{K_X[I_i] \cdot z_{c'} \cdot \rho_i}$. The game considers the SEARCH invocations that will collide under these modifications. Following the exponent notation $e_{i,j} = K_X[I_i] \cdot \rho_i$ in the j 'th invocation of GENTOKEN , the xtoken expressions is modified to $\text{xtokenList}[c, i]^{1/z_c} = (H(w_i))^{e_{i,j}}$ and $\text{xtokenList}[c', i]^{1/z_{c'}} = (H(w'_i))^{e_{i,j}}$.

Lemma E.12. *The probability of getting two $(w_i, \text{xtokenList}[c, i], z_c)$ and $(w'_i, \text{xtokenList}[c', i], z_{c'})$ for $w_i \neq w'_i$ is negligible. Thus G_{12} is indistinguishable from G_{11} .*

Proof. Let μ be the probability of getting two such tuples. Similar to G_8 , an OM-GDH challenge (h_1, h_2) is identical G_{12} . However, on each query w by \mathcal{A} to H , it selects (r_1, r_2) in \mathbb{Z}_p^* and outputs $H(w) = (h_1)^{r_1} (h_2)^{r_2}$. It also samples a random number j between $[1, l]$ where l is the maximum number of time GENTOKEN invocation by \mathcal{A} is allowed. During each invocation, it selects i at random from $[2, n]$ and sends the a_i from the GENTOKEN to OM-GDH challenger, and gets back $b_i \leftarrow (a_i)^t$, where t is chosen by the OM-GDH challenge game. \mathcal{A} gets back this b_i in response. In each SEARCH invocation where \mathcal{A} sends env , for each c the game takes $\text{xtoken}[c, i]$ received from \mathcal{A} , and each query w issued by \mathcal{A} to H , the DDH oracle is consulted if $(a_1, b_1, H(w), (\text{xtokenList}[c, i])^{1/z_c})$ is a DDH tuple. If for two instances $(w, \text{xtokenList}[c, i], z_c)$ and $(w', \text{xtokenList}[c', i], z_{c'})$ are the same for $w \neq w'$, the game computes h_1^t, h_2^t as $((\text{xtokenList}[c, i])^{r_2/z_c} (\text{xtokenList}[c', i])^{-r_2/z_{c'}})^{1/(r_1 r_2' - r_1' r_2)}$ and $((\text{xtokenList}[c, i])^{r_1/z_c} (\text{xtokenList}[c', i])^{-r_1/z_{c'}})^{1/(r_1' r_2 - r_1 r_2')}$ where $H(w) = (h_1)^{r_1} (h_2)^{r_2}$ and $H(w') = (h_1)^{r_1'} (h_2)^{r_2'}$. The probability of such an event is $\mu/(l \cdot n)$ where n is the number of w s in the query. This is negligible and G_{12} is indistinguishable from G_{11} . \square

Game G_{13} . Game G_{13} modifies G_{12} as follows. It does not refer to TSet or XSet (it skips the creation process) and samples the keys K_S, K_T, K_Y, K_M . The GENTOKEN instances are invoked the same way as in G_{12} . In SEARCH, for a counter c , $\text{xtokenList}[c, i] = (\text{OPRF}(K_X, w_i))^{z_c \cdot \rho_i}$, $\text{id} \in \text{DB}(w_i)$, and $I(w_i) = I_i$ for $i = 2, \dots, n$, it constructs a query q from w_1, \dots, w_n and send to $\text{SSE}_{\mathcal{L}_c}$. As $\text{av} \in \mathcal{P}$, it returns $(\text{DB}(q), |\text{DB}(w_1)|)$. The game samples $|\text{DB}(q)|$ random indices from $[1, |\text{DB}(w_1)|]$ and assigns to each entry of $\text{DB}(q)$. The game maintains TList such that TList(w_1) stores (c, id, e) entries. This (c, id, e) entries are constructed such that c in $\text{T}[w_1]$ was assigned an id during SEARCH and the associated ciphertext e . The game starts with empty TList and populates using the $\text{id} \in \text{DB}(q)$. G_9 can check that an $\text{id} \in \text{DB}(q)$ belongs to TList(w_1) or not. For any id not in TList, G_9 assigns a random values of c in $[1, \text{TSet}L]$ not used so far. It computes (c, id) for a ciphertext e using K_e and K_S from w_1 of G_{10} . This view of G_{13} is essentially the same as of G_{12} and matches the combined description of the simulators, as detailed in Algorithm 18, 19, and 20, and therefore, G_{13} is indistinguishable from G_{12} .

E.3 Security against Adversarial Server

We present proof of Theorem E.2 that defines the security against the adversarial server. The proof mainly follows from ODXT with modifications following NOMOS and leakage function \mathcal{L}_S . The proof essentially follows the proof outlined in [PM21] with modifications for NOMOS.

The proof follows a series of games, similar to the approach of security against a semi-honest client as the adversarial entity, where we start with the $\text{Real}_{\mathcal{A}}^{\Pi}$ execution (game G_0 is identical to the real execution), and the last game is identical to $\text{Ideal}_{\mathcal{A}, \text{SIM}}^{\Pi}$ (identical to game G_8). We denote the probability of a game G_i to output $b = 1$ by p_i .

Game G_1 . This game replaces the OPRF instances of the form $\text{OPRF}(K_T, \cdot)$ with a randomly sampled value from \mathbb{G} during transcript generation in UPDATE and SEARCH.

Lemma E.13. G_1 and G_0 are indistinguishable.

Proof. By indistinguishability property of OPRF from OMGDH assumption, G_1 is identical to G_0 . \square

Game G_2 . In this game, the OPRF instances of the expression form of $\text{OPRF}(K_X, \cdot)$ are replaced by randomly sampled values from \mathbb{G} during transcript generation in UPDATE and SEARCH.

Lemma E.14. Game G_2 and G_1 are indistinguishable.

Proof. By the indistinguishability property of OPRF from OMGDH assumption, G_2 is indistinguishable from G_1 . \square

Game G_3 . This game replaces the PRF instances of the form $F_p(K_Y, \cdot)$ with a random function in the range \mathbb{Z}_p^* .

Lemma E.15. Game G_3 and G_2 are indistinguishable.

Proof. By the indistinguishability property of PRF, G_3 is indistinguishable from G_2 . \square

Game G_4 . In this game, the PRFs of the form $F_p(K_Z, \cdot)$ are replaced by a random sample from \mathbb{Z}_p^* .

Lemma E.16. *Game G_4 is indistinguishable from G_3 .*

Proof. G_4 is indistinguishable from G_3 from the indistinguishability property of PRF. \square

Game G_5 . This game modifies the way xtoken s are generated. For a conjunctive query $q = w_1 \wedge \dots \wedge w_n$, the game looks up the history updates to retrieve the set of update operations $(\text{op}_j, (w_1, \text{id}_j))$ for the s -term w_1 . Furthermore, for each x -term and the update operation $(\text{op}_j, (w_1, w_j))$, the game computes the dynamic blinding factors $\alpha_{i,j}$ and the $\text{xtag}_{i,j}$ following ODXT construction and generates $\text{xtoken}_{i,j} = \text{xtag}_{i,j}^{\rho_i / \alpha_{i,j}}$.

Lemma E.17. *Game G_5 is indistinguishable from G_4 .*

Proof. The xtoken values in G_5 and G_4 are identically distributed. For a tuple of the form $(\alpha_{i,j}, \rho_i, \text{xtag}_{i,j}, \text{xtoken}_{i,j})$ in G_4 , $\text{xtoken}_{i,j} = \text{xtag}_{i,j}^{1/\alpha_{i,j}}$, that is same as of the G_5 . Therefore, G_5 and G_4 are indistinguishable. \square

Game G_6 . The game modifies the α values by using randomly sampled values from \mathbb{Z}_p^* .

Lemma E.18. *Game G_6 is indistinguishable from G_5 .*

Proof. In G_4 the PRFs are replaced with random samples from \mathbb{Z}_p^* . This operation is done only once for each UPDATE invocation for xtoken generation, and never repeated for any queries. Also, it is not evaluated on two same values for two different UPDATE invocations. The α values are generated by multiplying a randomly sampled \mathbb{Z}_p^* element in place of PRF with the inverse of a value in \mathbb{Z}_p obtained in the same way. Hence, the distribution of α in G_6 is indistinguishable from G_5 . \square

Game G_7 . This game modifies the way xtags are generated during transcript generation of UPDATE. It samples a $\gamma \leftarrow \mathbb{Z}_p^*$ and computes $\text{xtag} = g^\gamma$, where g is a generator of \mathbb{G} .

Lemma E.19. *Game G_7 is indistinguishable from G_6 .*

To prove Lemma E.19, we construct an alternative version below.

Lemma E.20. *Game G_7 and G_6 are indistinguishable following the polynomial equivalence of the DDH assumption and the extended DDH assumption over any group \mathbb{G} .*

Proof. The xtag values for an update operation $(\text{op}, (\text{id}_j, w_i))$ are computed in the following way in G_7 $\text{xtag}_{i,j,\text{op}} = g^{G_X(w_i) \cdot G_Y(\text{id}_j \parallel \text{op})}$ where $G_X(\cdot)$ and $G_Y(\cdot)$ are random functions uniformly sampled from the set of all functions mapping λ -bit strings to \mathbb{Z}_p^* replacing $\text{OPRF}(K_X, \cdot)$ and $F_p(K_Y, \cdot)$, and g is the generator of group \mathbb{G} . We rewrite the expression as $\text{xtag}_{i,j,\text{op}} = g^{\alpha_i \cdot \beta_{j,\text{op}}}$ where $\alpha_i = G_X(w_i)$ and $\beta_{j,\text{op}} = (G_Y(\text{id}_j \parallel \text{op}))$. Whereas, in previous game we had $\text{xtag}_{i,j,\text{op}} = g^{\gamma_{i,j,\text{op}}}$ where $\gamma_{i,j,\text{op}} \leftarrow \mathbb{Z}_p^*$. The distribution of xtags are indistinguishable from G_6 . Therefore, G_7 and G_6 are indistinguishable. \square

Game G_8 . Game G_8 is identical to G_7 with the following modifications. During transcript generation for UPDATE operation, each $\text{OPRF}(K_T, w \parallel \text{cnt} \parallel b)$, $b \in \{0, 1\}$ instance is replaced with a random function of the form $G_T(t)$ where t is the timestamp of the particular update operation.

Lemma E.21. *Game G_8 is indistinguishable from G_7 .*

Proof. Note that $G_T(t)$ is never evaluated twice on the same input as the real instance has an increasing counter appended to the input, and $G_T(\cdot)$ is uniformly randomly sampled from the set of all λ -bit functions mapping to λ -bit values. Therefore, G_8 is indistinguishable from G_7 . \square

Game G_9 . In this game, simulator SIM' replaces the challenger. SIM' does not have access to the actual queries by \mathcal{A} . It uses the the following leakages for each update or conjunctive queries issued by \mathcal{A} .

SIM' receives empty leakage from `UPDATE`. It uses the timestamp information to generate the `TSet` entries similar to the procedure in G_8 . It samples uniformly random blinding factors α , and $\gamma \leftarrow \mathbb{Z}_p^*$ and computes $\text{xtag} = g^\gamma$.

SIM' learns the number of updates involving w_1 (the s-term) and the timestamp of these updates. It uses this information to simulate the `xtoken` computation similar to G_8 .

Similarly, SIM' uses the aforementioned information learned to compute $\text{xtoken}_{i,j} = \text{xtag}_{i,j}^{1/\alpha_{i,j}}$. Also, SIM' learns whether two queries have the same s-term from the equality pattern and generates `stoken` values across multiple queries accordingly (consistently). Apart from these, SIM learns the set of `ids` in the final result.

Lemma E.22. *Game G_9 is indistinguishable from G_8 .*

Proof. The transcripts generated by SIM' is identical to the one generated by \mathcal{A} for each update and conjunctive queries. Therefore, game G_9 is identical to G_8 . \square

Game G_{10} . In this game, SIM' is replaced with SIM that has access to the leakage function \mathcal{L}_S as below.

$$\mathcal{L}_S = \{\mathcal{L}_S^{\text{SETUP}}, \mathcal{L}_S^{\text{UPDATE}}, \mathcal{L}_S^{\text{UPDATE}}\}$$

where $\mathcal{L}_S^{\text{SETUP}} = \perp$, $\mathcal{L}_S^{\text{UPDATE}}(\text{op}, (w, \text{id})) = \perp$, and $\mathcal{L}_S^{\text{UPDATE}}(q) = \{\text{TimeDB}(q), \text{UPD}(q)\}$, as defined in Section E.1.

Lemma E.23. *Game G_{10} is indistinguishable from G_9 .*

We construct the following Lemma to prove Lemma E.23.

Lemma E.24. *SIM can efficiently execute SIM' from G_9 as a sub-routine.*

Proof. The proof essentially involves demonstrating that the leakage profile of SIM covers the leakage profile of SIM' .

SIM has access to the same empty update leakage of SIM' . Further, $\text{SIM Updates}(q)$ leakage is covered by `Updates`(w_1) (the s-term). The equality pattern leakage in G_9 is covered by `Updates`(q_1) and `Updates`(q_2) for two different queries q_1 and q_2 . Hence, G_{10} is indistinguishable from G_9 and is identical to the ideal execution of the protocol. \square

Alternative proof approach. The security analysis of `NOMOS` relies on a two-way approach - one considering the leakage to the adversarial client and the other considering the leakage to the adversarial server. Considering that the existing literature lacks well-established proof approaches for multi-client constructions, alternative proof approaches, such as the framework proposed by Jaeger et al. [JT20], can be an interesting direction to assess the semantic security of multi-client constructions. We leave this analysis an interesting future work.

F Forward and Backward Privacy of Nomos

We first state the formal generic forward and backward privacy notions introduced by Bost et al. [BMO17a] in the context of dynamic SSE constructions including `OXT` and `NOMOS`. We informally outline the forward and backward privacy notions below.

Dynamic SSE constructions have the additional `UPDATE` routine to update the content in the encrypted database offloaded to the remote cloud server. This update capability brings in the notions of forward and backward privacy in dynamic schemes, which are necessary to prevent certain kind of leakage-abuse attacks [ZKP16, CJJ⁺14]. Informally, forward privacy states that a previous search operation can not be linked to a subsequent update operation involving the same keyword. Whereas, backward privacy states that if a record involving a keyword is added to the database and later deleted, a subsequent search operation involving the same keyword should not link these update operations. Note that, a search operation involving the same keyword after

the first update would trivially leak the update. We outline the formal definition of forward and backward privacy notions as introduced by Bost et al. [BMO17b] below. In this notions, we assume an update is restricted to $\text{op} = \{\text{ADD}, \text{DEL}\}$ operations only.

F.1 Forward Privacy of a Dynamic SSE Scheme

We denote the leakage profile of an adaptively secure dynamic conjunctive SSE as below.

$$\mathcal{L} = (\mathcal{L}^{\text{SETUP}}, \mathcal{L}^{\text{UPDATE}}, \mathcal{L}^{\text{SEARCH}})$$

Adaptive forward privacy states that the update leakage can be expressed as

$$\mathcal{L}^{\text{UPDATE}}(\text{op}, (w, \text{id})) = \mathcal{L}'(\text{op}, \text{id})$$

where \mathcal{L}' is a stateless function and $(\text{op}, (w, \text{id}))$ is a triplet. This essentially implies that the update operation hides the w being updated and therefore can not be linked to any search query containing w by a polynomially bounded \mathcal{A} .

F.2 Backward Privacy of a Dynamic SSE Scheme

Denote the leakage profile of an adaptively secure dynamic conjunctive SSE construction as given below.

$$\mathcal{L} = (\mathcal{L}^{\text{SETUP}}, \mathcal{L}^{\text{UPDATE}}, \mathcal{L}^{\text{SEARCH}})$$

The following three classes of information about updates are necessary for the exposition of backward privacy notions, that captures specific leakages from update operations.

- **TimeDB(w):** Given a query list Q , TimeDB(w) stores the records of matching ids, excluding the deleted ones, for w along with the timestamp of insertion.

$$\begin{aligned} \text{TimeDB}(w) = \{ & (t, \text{id}) \mid (t, \text{ADD}, (w, \text{id})) \in Q, \\ & \text{and } \forall t', (t', \text{DEL}, (w, \text{id})) \notin Q \} \end{aligned}$$

- **Updates(w):** Updates(w) stores the timestamps of the update history of a w .

$$\text{Updates}(w) = \{ t \mid (t, \text{ADD}, (w, \text{id})) \text{ or } (t, \text{DEL}, (w, \text{id})) \in Q \}$$

- **DelHist(w):** DelHist(w) stores the timestamps of update operations, specially which DEL operation cancels which ADD operation for a w .

$$\begin{aligned} \text{DelHist}(w) = \{ & (t^{\text{ADD}}, t^{\text{DEL}}) \mid \exists \text{id s.t. } (t^{\text{DEL}}, \text{DEL}, (w, \text{id})) \in Q \\ & \text{and } (t^{\text{ADD}}, \text{ADD}, (w, \text{id})) \in Q \} \end{aligned}$$

With these captured information defined, the three levels of backward privacy as presented in [BMO17b] are defined as below.

- *Type-I* An adaptively secure dynamic SSE scheme is type-I backward private iff the update and search leakages can be expressed as below.

$$\begin{aligned} \mathcal{L}^{\text{UPDATE}}(\text{op}, (w, \text{id})) &= \mathcal{L}'(\text{op}) \\ \mathcal{L}^{\text{SEARCH}}(w) &= \mathcal{L}''(\text{TimeDB}(w), a_w) \end{aligned}$$

Here, a_w is the update count of w , and $\mathcal{L}', \mathcal{L}''$ are stateless.

- *Type-II* An adaptively secure dynamic SSE scheme is type-II backward private iff the update and search leakages can be expressed as below.

$$\begin{aligned} \mathcal{L}^{\text{UPDATE}}(\text{op}, (w, \text{id})) &= \mathcal{L}'(\text{op}, w) \\ \mathcal{L}^{\text{SEARCH}}(w) &= \mathcal{L}''(\text{TimeDB}(w), \text{Updates}(w)) \end{aligned}$$

Here, $\mathcal{L}', \mathcal{L}''$ are stateless.

- *Type-III* An adaptively secure dynamic SSE scheme is type-III backward private iff the update and search leakages can be expressed as below.

$$\begin{aligned}\mathcal{L}^{\text{UPDATE}}(\text{op}, (\mathbf{w}, \text{id})) &= \mathcal{L}'(\text{op}, \mathbf{w}) \\ \mathcal{L}^{\text{SEARCH}}(\mathbf{w}) &= \mathcal{L}''(\text{TimeDB}(\mathbf{w}), \text{DelHist}(\mathbf{w}))\end{aligned}$$

Here, \mathcal{L}' , \mathcal{L}'' are stateless.

Based on these, we state the formal forward and backward privacy notions for NOMOS below. Note that, we apply the single data owner notion to multi client setting of MRSW as in this setting only the gate-keeper is allowed to update the database. For an individual client, the view of an adversarial server remains the same as of an SRSW forward and backward construction, in this case ODXT. The forward and backward privacy notions can be independently applied to this setting as well which elaborate below.

F.3 Forward Privacy of Nomos

Given the leakage profile of an adaptively secure dynamic conjunctive SSE

$$\mathcal{L}_S = (\mathcal{L}_S^{\text{SETUP}}, \mathcal{L}_S^{\text{UPDATE}}, \mathcal{L}_S^{\text{SEARCH}})$$

adaptive forward privacy states that the update leakage can be expressed as

$$\mathcal{L}_S^{\text{UPDATE}}(\text{op}, (\mathbf{w}, \text{id})) = \mathcal{L}'(\text{op}, \text{id})$$

where \mathcal{L}' is a stateless function and $(\text{op}, (\mathbf{w}, \text{id}))$ is an arbitrary triplet. This essentially implies that the update operation hides the \mathbf{w} being updated and therefore can not be linked to any search query containing \mathbf{w} by a polynomially bounded \mathcal{A} .

We note that NOMOS UPDATE leakage is

$$\mathcal{L}_S^{\text{UPDATE}}(\text{op}, (\mathbf{w}, \text{id})) = \perp$$

as discussed in Section E.1. Therefore, NOMOS hides the \mathbf{w} as well as the id from the $(\text{op}, (\mathbf{w}, \text{id}))$ input involved in the update process. The following corollary is straightforward from Theorem E.1 and E.2 of the main manuscript.

Corollary F.1. *Provided NOMOS is instantiated with DH OPRF and DH assumption holds in \mathbb{G} , F_p and F are secure PRFs, and $(\text{AE.Enc}, \text{AE.Dec})$ is an IND-CPA and strongly UF-CMA-secure AE scheme, and all hash functions are modelled as random oracles, NOMOS is adaptively forward private.*

F.4 Backward Privacy of Nomos

Given the leakage profile of an adaptively secure dynamic conjunctive SSE

$$\mathcal{L}_S = (\mathcal{L}_S^{\text{SETUP}}, \mathcal{L}_S^{\text{UPDATE}}, \mathcal{L}_S^{\text{SEARCH}})$$

adaptive type-II backward privacy states that the update and search leakages can be expressed as

$$\begin{aligned}\mathcal{L}_S^{\text{UPDATE}}(\text{op}, (\mathbf{w}, \text{id})) &= \mathcal{L}''(\text{op}, \text{id}) \\ \mathcal{L}_S^{\text{SEARCH}}(\mathbf{w}) &= \mathcal{L}''(\text{TimeDB}(\mathbf{w}), \text{UPD}(\mathbf{w}))\end{aligned}$$

For NOMOS, the update and search leakage has the following profile.

$$\begin{aligned}\mathcal{L}_S^{\text{UPDATE}}(\text{op}, (\mathbf{w}, \text{id})) &= \perp \\ \mathcal{L}_S^{\text{SEARCH}}(q) &= (\text{TimeDB}(q), \text{UPD}(q))\end{aligned}$$

for a conjunctive query q , which is extended to the conjunctive keyword setting from single keyword setting. Therefore, the following corollary is immediate from Theorem E.1 and E.2 of the main manuscript.

Corollary F.2. *Provided NOMOS is instantiated with DH OPRF and DH assumption holds in \mathbb{G} , F_p and F are secure PRFs, and $(AE.Enc, AE.Dec)$ is an IND-CPA and strongly UF-CMA-secure AE scheme, and all hash functions are modelled as random oracles, NOMOS is adaptively backward private.*