Poster: Multi-Client and Quantum-Resilient Searchable Symmetric Encryption

Debadrita Talapatra
IIT Kharagpur, India
debadritat.fg2219@kgpian.iitkgp.ac.in

Arnab Bag
IIT Kharagpur, India
arnabbag@iitkgp.ac.in

Sikhar Patranabis
IBM Research India
sikhar.patranabis@ibm.com

Debdeep Mukhopadhyay
IIT Kharagpur, India
debdeep@cse.iitkgp.ac.in

***Abstract*—Searchable Symmetric Encryption (SSE) enables privacy-preserving keyword searches over encrypted document collections. In this poster abstract, we present new techniques for designing practically efficient and highly scalable SSE schemes that support a rich class of Boolean queries on large real-world databases in the multi-client setting. As an additional contribution, we propose novel SSE schemes that are quantum-resilient based on the hardness of well-studied lattice problems. Our contributions represent significant progress towards making SSE schemes practically efficient, quantum-resilient, and suitable for deployment in multi-user cloud environments.**

I. Introduction

The advent of cloud computing potentially allows individuals and organizations to outsource storage and processing of large volumes of data to third party servers. However, this leads to privacy concerns – clients typically do not trust service providers to respect the confidentiality of their data. Consider a client that offloads an encrypted database of (potentially sensitive) emails to an untrusted server. At a later point of time, the client might want to issue a query of the form "*retrieve all emails received from xyz@foobar.org* or "*retrieve all emails with the keyword "research" in the subject field*". Ideally, the client should be able to perform this task without revealing any sensitive information to the server. Unfortunately, techniques such as fully homomorphic encryption (FHE) or oblivious RAM (ORAM), that potentially allow achieving such an "ideal" notion of privacy, are sometimes unsuitable for practical deployment due to large performance overheads.

Searchable Symmetric Encryption. Searchable symmetric encryption (SSE) [1]–[3] allows a client to execute keyword search queries *directly* over symmetrically encrypted document collections stored at an *untrusted* third-party server. SSE schemes typically achieve fast and efficient query processing over encrypted data while allowing the server to learn some controlled amount of information (called "leakage") during query execution. Taking cognizance of the vast literature on SSE over the past 20 years or so, the following have emerged as some desirable requirements for any SSE scheme: (i) support keyword queries represented as expressive Boolean formulae (in particular, conjunctive keyword searches), (ii) handle dynamic databases (i.e., support updates and keyword queries efficiently and in tandem), (iii) support queries from multiple clients apart from the data owner (referred to as *multi-client* SSE), and (iv) resist quantum attacks (this is particularly relevant for SSE schemes such as [3] that adopt classical public-key cryptographic techniques for query and storage efficiency, and hence are vulnerable to quantum attacks).

Relevant Prior Work. We briefly summarize some relevant prior work by the authors of this poster abstract that made progress towards designing SSE schemes that achieve one or more of the above goals. We subsequently summarize the main technical results that we intend to present in this poster, which build upon these prior works.

Oblivious Dynamic Cross-Tags (ODXT). At NDSS' 21 [4], we introduced ODXT – the first dynamic SSE scheme capable of supporting conjunctive keyword queries with sub-linear query complexity and linear storage requirements, while ensuring *both forward and backward privacy*[1]. ODXT is a dynamic counterpart to the seminal but static OXT scheme from [3]. Unfortunately, both OXT and ODXT are designed for the single-client setting (where the client is also the data owner), and lack support for multi-client search. Trivially extending ODXT to the multi-client setting incurs significant leakage that renders such extension insecure in practice. This leads to the following question: *can we design a conjunctive dynamic multi-client SSE scheme with forward and backward privacy?*

TWINSSE. At PoPETS' 23 [5], we introduced TWINSSE – the first SSE scheme capable of supporting conjunctive, disjunctive and more generally expressive Boolean queries (in the conjunctive and disjunctive normal forms) with sub-linear query complexity and linear storage requirements. TWINSSE builds upon OXT [3], which in turn relies crucially on discrete-log hard groups (typically implemented in practice

[1] Forward and backward privacy are the de facto security requirements for any dynamic SSE scheme. Forward privacy requires that adding a new document should not reveal whether it contains keywords that have been previously queried, while backward privacy requires that searching for a keyword w should reveal no information about files containing w that have already been deleted from the database.

using elliptic-curve-based cryptography) for efficient and secure query processing. As a result, OXT and by extension, TWINSSE are quantum-unsafe, and will be devastatingly broken once scalable quantum computers capable of solving the discrete log problem become feasible.

Oblivious Post-Quantum Secure Cross-Tags (OQXT). At EuroS&P '23 [6], we introduced OQXT – a plausibly quantum-safe counterpart to OXT that relies on the presumed (quantum) hardness of solving certain lattice problems, such as short integer solutions (SIS) and learning with rounding (LWR). Unfortunately, OQXT has two major drawbacks. First of all, like OXT, it only supports conjunctive keyword queries, and not disjunctive or more general Boolean queries. More crucially, OQXT works over integer lattices, and incurs significantly higher storage overheads as compared to OXT, which is not desirable from a practical point of view. This leads to the following question: *can we design a lattice-based plausibly quantum-safe yet practically efficient SSE scheme supporting* ***both*** *conjunctive and disjunctive keyword queries?*

II. Our Contributions

In this poster, we present two contributions that answer the above questions in the affirmative. The first of these will appear at ACM AsiaCCS '24 [7]. The second contribution is an unpublished, in-progress work.

NOMOS. In a recent work (to appear at ACM AsiaCCS '24 [7]), we present the first dynamic multi-client SSE scheme NOMOS supporting efficient conjunctive Boolean queries over an encrypted database. Precisely, NOMOS is a multi-client extension of our ODXT scheme [4] that allows only the data-owner (a trusted entity in the NOMOS framework) to update the encrypted database stored on the adversarial server, but allows multiple clients (distinct from the data owner) to query the encrypted database. NOMOS achieves forward and backward privacy while incurring less leakage than the trivial extension of ODXT to the 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.

Clients in NOMOS obtain search tokens from the trusted data owner, which is allowed to update the encrypted database and holds the keys for token generation. We use an oblivious pseudorandom function (OPRF)-based mechanism to delegate the search token generation process to the data owner while preserving client privacy; thus bypassing the need to share the secret keys for token generation among multiple potentially semi-honest clients. From a security standpoint, NOMOS mitigates a particular leakage in ODXT. To this end, we introduce a de-correlated access pattern based on a novel usage of a variant of Bloom Filters, called the Redundant Bloom Filter. This mitigates the aforementioned leakage while incurring only minimal impact on the query and storage overheads. The token generation process and search phase of NOMOS can work asynchronously. This tokenized search process allows fine-grained user management and enforcing access permissions for each user individually in a multi-client setting, which is a highly desirable capability in multi-user cloud applications where query requests are expected to arrive asynchronously.

NTRU-OQXT. As a second contribution (unpublished, in-progress), we present the first practically efficient SSE scheme with fast conjunctive *and* disjunctive keyword searches, compact storage, and security based on the (plausible) quantum-hardness of well-studied *lattice-based* assumptions. As a stepping stone, we introduce NTRU-OQXT – a highly compact NTRU lattice-based conjunctive SSE scheme that optimizes our proposed OQXT scheme [6] and outperforms *all* existing conjunctive SSE schemes in terms of search latency. We then present an extension of NTRU-OQXT, which we call TWINSSE-NTRU-OQXT, that integrates NTRU-OQXT with our proposed TWINSSE scheme [5] to additionally support disjunctive queries. Technically, both of our proposed schemes rely on a novel oblivious search protocol based on highly optimized Fast-Fourier trapdoor sampling algorithms over NTRU lattices. While such techniques have been used to design other lattice-based cryptographic primitives (such as digital signatures), they have not been applied before in the context of SSE. We present prototype implementations of both schemes, and experimentally validate their practical performance over a large real-world dataset. Our experiments demonstrate that both schemes support extremely fast query processing while scaling smoothly to large datasets. In fact, NTRU-OQXT achieves at least 2× faster conjunctive keyword searches as compared to *all other* conjunctive SSE schemes (including the best quantum-broken conjunctive SSE schemes), and substantially outperforms many of these schemes in terms of storage requirements. These efficiency benefits also translate to TWINSSE-NTRU-OQXT, which is practically competitive with the best quantum-broken SSE schemes capable of supporting both conjunctive and disjunctive queries.

References

[1] D. X. Song, D. A. Wagner, and A. Perrig, "Practical techniques for searches on encrypted data," in *IEEE S&P 2000*, 2000, pp. 44–55.

[2] R. Curtmola, J. A. Garay, S. Kamara, and R. Ostrovsky, "Searchable symmetric encryption: improved definitions and efficient constructions," in *ACM CCS 2006*, 2006, pp. 79–88.

[3] D. Cash, S. Jarecki, C. S. Jutla, H. Krawczyk, M. Rosu, and M. Steiner, "Highly-scalable searchable symmetric encryption with support for boolean queries," in *CRYPTO 2013*, 2013, pp. 353–373.

[4] S. Patranabis and D. Mukhopadhyay, "Forward and backward private conjunctive searchable symmetric encryption," in *NDSS 2021*, 2021.

[5] A. Bag, D. Talapatra, A. Rastogi, S. Patranabis, and D. Mukhopadhyay, "Two-in-one-sse: Fast, scalable and storage-efficient searchable symmetric encryption for conjunctive and disjunctive boolean queries," *Proc. Priv. Enhancing Technol.*, 2023.

[6] D. Talapatra, S. Patranabis, and D. Mukhopadhyay, "Conjunctive searchable symmetric encryption from hard lattices," in *8th IEEE European Symposium on Security and Privacy, EuroS&P*, 2023.

[7] A. Bag, S. Patranabis, and D. Mukhopadhyay, "Tokenised multi-client provisioning for dynamic searchable encryption with forward and backward privacy," in *ACM AsiaCCS (to appear)*, 2024.

Multi-Client and Quantum-Resilient Searchable Symmetric Encryption

Debadrita Talapatra [1] Arnab Bag[1] Sikhar Patranabis [2] Debdeep Mukhopadhyay [1]

[1]Indian Institute of Technology Kharagpur, India [2]IBM Research India

debadritat.fg2219@kgpian.iitkgp.ac.in arnabbag@iitkgp.ac.in sikhar.patranabis@ibm.com debdeep@cse.iitkgp.ac.in

Outsourced Data Storage And Retrieval

Huge volume of data generated by client - offloaded to third-party cloud server

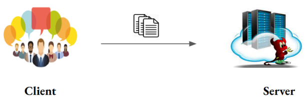

Encrypt the data before offloading?

How do you access encrypted data from the cloud server without decryption?

Privacy Conundrum: To Trust Or Not To Trust The Server?

Security Implications of outsourcing sensitive client-data to third-party cloud server

- Data Confidentiality
- User Revocation
- Scalability and Efficiency
- Collusion between entities

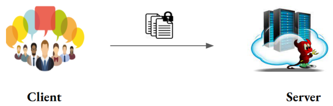

Searchable Symmetric Encryption (SSE)

- Search capability over symmetrically encrypted data offloaded to a third-party cloud server.
- Trades off security for efficiency.
- Encrypted search tokens are used by the server to access encrypted data.
- Server sends the encrypted search result to the client, which decrypts it locally.

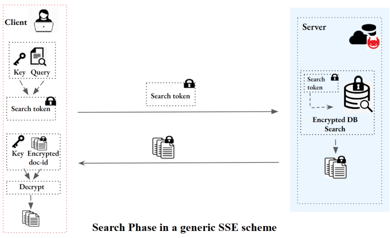

Search Phase in a generic SSE scheme

ODXT [NDSS 2021]

- First forward and backward private dynamic conjunctive SSE scheme.
- Extremely efficient updates (non-interactive + forward private).
- Search overhead scales with the frequency of least frequent conjunct.
- Two-round backward-private searches.

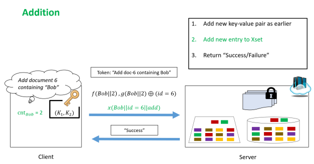
Addition
Add document 6 containing "Bob"
Token: "Add doc-6 containing Bob"
"Success"
Client
Server

1. Add new key-value pair as earlier
2. Add new entry to Xset
3. Return "Success/Failure"

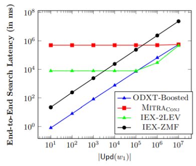

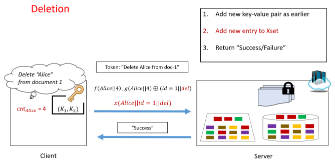
Deletion
Delete "Alice" from document 1
Token: "Delete Alice from doc-1"
"Success"
Client
Server

1. Add new key-value pair as earlier
2. Add new entry to Xset
3. Return "Success/Failure"

TWo-IN-one-SSE or TWINSSE [PoPETS 2023]

- Generic black-box transformation, efficiently supports **conjunctive, disjunctive, and more general Boolean queries** (in both CNF and DNF), with *linear* storage overheads.
- Meta-keywords: Disjunction of certain carefully chosen keywords of the form $\mathtt{mkw_i} = (\mathtt{w_{i_1}} \vee \mathtt{w_{i_2}} \vee \ldots \vee \mathtt{w_{i_\ell}})$.

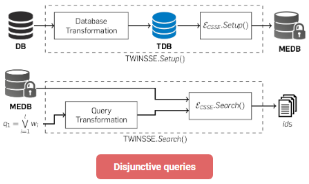

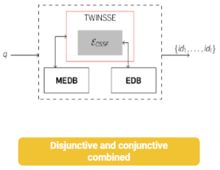

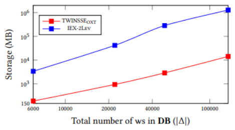

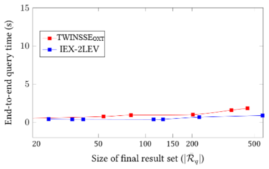

OQXT [IEEE Euro S&P 2023]

PQ-Secure SSE

- **First lattice-based SSE scheme, supports conjunctive queries** over large encrypted databases.
- Elements computed as *Learning With Rounding* samples. *Lattice-trapdoors* used for fast oblivious computations during search operations.

OQXT

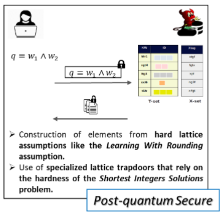

- Construction of elements from **hard lattice assumptions like the *Learning With Rounding* assumption.**
- Use of **specialized lattice trapdoors that rely on the hardness of the *Shortest Integers Solutions* problem.**

Post-quantum Secure

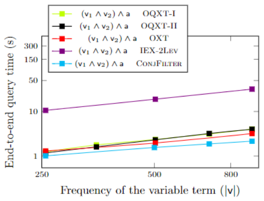

Efficiency Guarantees

- *Sublinear* **conjunctive query complexity** that grows with the frequency of the least frequent keyword in the conjunctive query.

NTRU-OQXT (In Progress)

- **Highly compact NTRU lattice-based conjunctive SSE scheme that optimizes OQXT**.
- TWINSSE-NTRU-OQXT integrates NTRU-OQXT with TWINSSE **to support disjunctive queries**.
- **NTRU-OQXT is at least 2× faster for conjunctive queries and requires significantly less storage** as compared to all other conjunctive SSE schemes.
- These optimized performance benefits of NTRU-OQXT also translates directly to TWINSSE-NTRU-OQXT.

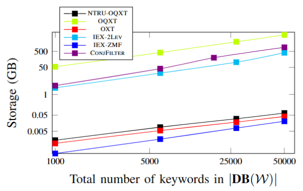

Optimized Storage Overhead

- Storage significantly reduced to **megabytes of memory usage** (from **several gigabytes** required by OQXT).
- Scales **linearly with DB size, competes with classically secure OXT.**
- **Significantly reduces the quadratic storage overheads** incurred by IEX-2LEV and CONJFILTER.

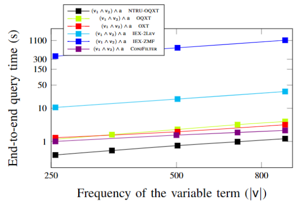

Search Latency

- Search time required for a conjunctive query of NTRU-OQXT is **reduced by** $3-4\times$ from that of OQXT.
- **Reduced by half when compared to OXT and CONJFILTER.**
- **Recuded by** $20\times$ **compared with IEX-2LEV and around** $600\times$ **compared to IEX-ZMF.**

NOMOS [AsiaCCS 2024]

- **First multi-client dynamically updatable** SSE, applicable for practical cloud application.
- **Tokenised Search Functionality:** Fine-grained user management.
- **Redundant BF:** Mitigates cross-term based leakages (otherwise present in state-of-the-art SSE).

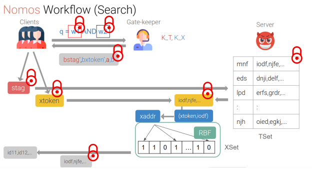

mnf	iodf,njfe,...
eds	dnji,delf,...
lpd	erfs,grdr,...
:	:
njh	oied,egkj,...

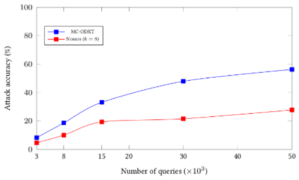

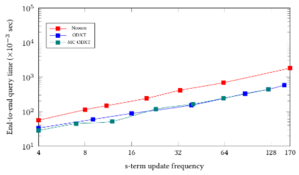

Nomos has a significantly lower attack accuracy compared to MC-ODXT.

Conjunctive search overhead is sub-linear in terms of the total database size.

Discussion And Future Directions

- **NTRU-OQXT:** Lattice-based, quantum-safe SSE scheme supporting both conjunctive and disjunctive keyword queries. Constructed over highly compact and structured NTRU lattices, NTRU-OQXT is highly optimized regarding storage overhead and search latency, and circumvents the inherent limitations of lattice-based constructions.
- **Extension to Dynamic Databases:** Extending both NTRU-OQXT a TWINSSE-NTRU-OQXT to dynamic databases that support updates, to construct a post-quantum secure dynamic SSE scheme supporting conjunctive, disjunctive, and more general Boolean queries.
- **NOMOS:** A conjunctive dynamic multi-client SSE scheme with forward and backward privacy guarantees.
- **NOMOS for MRMW Setting:** Extending the NOMOS framework to a Multi-Reader Multi-Writer setting.

Related Publications

[1] Arnab Bag, Sikhar Patranabis, and Debdeep Mukhopadhyay. Tokenised multi-client provisioning for dynamic searchable encryption with forward and backward privacy. In *ACM AsiaCCS (to appear)*, 2024.

[2] Arnab Bag, Debadrita Talapatra, Ayushi Rastogi, Sikhar Patranabis, and Debdeep Mukhopadhyay. Two-in-one-sse: Fast, scalable and storage-efficient searchable symmetric encryption for conjunctive and disjunctive boolean queries. *Proc. Priv. Enhancing Technol.*, 2023.

[3] Sikhar Patranabis and Debdeep Mukhopadhyay. Forward and backward private conjunctive searchable symmetric encryption. In *28th Annual Network and Distributed System Security Symposium, NDSS 2021*, 2021.

[4] Debadrita Talapatra, Sikhar Patranabis, and Debdeep Mukhopadhyay. Conjunctive searchable symmetric encryption from hard lattices. In *8th IEEE European Symposium on Security and Privacy, EuroS&P 2023, Delft, Netherlands*, 2023.

45th IEEE Symposium on Security and Privacy, 2024