Setup-Free Secure Search on Encrypted Data: Faster and Post-Processing Free

Adi Akavia*[1], Craig Gentry[2], Shai Halevi[2] and Max Leibovich*[1]

[1] University of Haifa
[2] IBM Research

Abstract. We present a novel *secure search* protocol on data and queries encrypted with Fully Homomorphic Encryption (FHE). Our protocol enables organizations (client) to (1) securely upload an unsorted data array $x = (x[1], \ldots, x[n])$ to an untrusted honest-but-curious sever, where data may be uploaded over time and from multiple data-sources; and (2) securely issue repeated search queries q for retrieving the first element $(i^*, x[i^*])$ satisfying an agreed matching criterion $i^* = \min\{\, i \in [n] \mid \mathsf{IsMatch}(x[i], q) = 1\}$, as well as fetching the next matching elements with further interaction. For security, the client encrypts the data and queries with FHE prior to uploading, and the server processes the ciphertexts to produce the result ciphertext for the client to decrypt. Our secure search protocol improves over the prior state-of-the-art for secure search on FHE encrypted data (Akavia, Feldman, Shaul (AFS), CCS'2018) in achieving:

- *Post-processing free* protocol where the server produces a ciphertext for the correct search outcome with overwhelming success probability. This is in contrast to returning a list of candidates for the client to post-process, or suffering from a noticeable error probability, in AFS. Our post-processing freeness enables the server to use secure search as a sub-component in a larger computation without interaction with the client.
- *Faster protocol:* (a) Client time and communication bandwidth are improved by a $\log^2 n / \log\log n$ factor. (b) Server evaluates a polynomial of degree linear in $\log n$ (compare to cubic in AFS), and overall number of multiplications improved by up to $\log n$ factor. (c) Employing only GF(2) computations (compare to GF(p) for $p \gg 2$ in AFS) to gain both further speedup and compatibility to all current FHE candidates.
- *Order of magnitude speedup exhibited by extensive benchmarks* we executed on identical hardware for implementations of ours versus AFS's protocols.

Additionally, like other FHE based solutions, out solution is setup-free: to outsource elements from the client to the server, no additional actions are performed on x except for encrypting it element by element (each element bit by bit) and uploading the resulted ciphertexts to the server.

Keywords: Secure search, Fully homomorphic encryption, Randomized algorithms, Razborov-Smolensky, Low degree approximation, Universal hash functions

1 Introduction

Following the rapid advancement and widespread availability of cloud computing it is a common practice to outsource data storage and computations to cloud providers. Placing cleartext (i.e unencrypted) data on the cloud compromises data security. To regain data privacy one could encrypt the data prior to uploading to the cloud. However, if using standard encryption (e.g. AES), this solution nullifies the benefits of cloud computing: when given only ciphertexts the cloud provider cannot process the underlying cleartext data in any meaningful way.

*This work was supported in part by the Center for Cyber Law & Policy at the University of Haifa, and by the BIU Center for Research in Applied Cryptography and Cyber Security. Both in conjunction with the Israel National Cyber Directorate in the Prime Minister's Office.

Fully homomorphic encryption (FHE) [55, 23] is an encryption scheme that allows processing the underlying cleartext data while it still remains in encrypted form, and without giving away the secret key (see Definition 1). With FHE it is possible for the client to securely outsource computations to the server as follows: The client first encrypts its data x with an FHE scheme to obtain the ciphertext $[\![x]\!] \leftarrow \mathsf{Enc}_{\mathsf{pk}}(x)$, and sends $[\![x]\!]$ to the server. The server can now compute any function f on the underlying clear-text data x by evaluating a homomorphic version of f on the ciphertext $[\![x]\!]$. The outcome of this computation is a ciphertext $[\![y]\!] \leftarrow \mathsf{Eval}_{\mathsf{pk}}(f, [\![x]\!])$ that decrypts to the desired output $y = f(x)$. The server can now send the ciphertext $[\![y]\!]$ to the client who would decrypt $y \leftarrow \mathsf{Dec}_{\mathsf{sk}}([\![y]\!])$ to obtain the result.

The homomorphic computations achievable by the known FHE candidates (e.g. [8, 48, 20, 25]) are specified by a polynomial over a finite ring (i.e. by repeated application of homomorphic-addition and homomorphic-multiplication for that ring). For example, for data in binary representation, bitwise operations on plaintext bits (addition and multiplication modulo 2) can be replaced by their homomorphic counterparts on encrypted bits (homomorphic-addition and homomorphic-multiplication).

Key factors influencing the running-time of such homomorphic computations are the degree and overall multiplications of the polynomial. Leading to the main two constraints in designing algorithms that compute on FHE encrypted data: they must be realized by a polynomial of low degree and low amount of overall multiplications.

Note that this FHE approach for securely outsourcing to the server the computation of $y = f(x)$ has the benefits of requiring only a single round of communication, and with low communication bandwidth (communicating only the encrypted input $[\![x]\!]$ and output $[\![y]\!]$). Furthermore, the server in this protocol learns no new information about x or y (assuming the FHE is semantically secure).

Secure search is a fundamental computational problem, useful in numerous data analysis and retrieval tasks. An abundance of proposed solutions were presented to solve it using different cryptographic tools (see Section 1.1 and Tables 1-2). In particular, Gentry [23] proposed using FHE to securely search on encrypted data.

In this work we address the natural and simple formulation for *secure search on FHE encrypted data* (*secure search*) as considered by [3]: Secure search is a two party protocol between a server and a client. The server holds an unsorted array $[\![x]\!] = \left([\![x[1]]\!], \ldots, [\![x[n]]\!]\right)$ of encrypted elements (not necessarily distinct) that were previously encrypted and uploaded by the client, as well as a specification of a predicate $\mathsf{IsMatch}(a, b) \in \{0, 1\}$ specifying the matching condition. The client submits encrypted queries $[\![q]\!]$ to the server in order to retrieve the first matching element. The server returns to the client the encrypted index and element pair $[\![y]\!] = \left([\![i^*]\!], [\![x[i^*]]\!]\right)$ for i^* the index of the first element satisfying the matching condition, $i^* = \min\{i \in [n] \mid \mathsf{IsMatch}(x[i], q) = 1\}$. See detailed definition and extensions in Section 3 and 7.

Restrictions on protocols. Note that the above secure search formulation, as addressed in this work, focuses on protocols that involve (a) a **single server**, and where the client-server interaction is of (b) a **single round**, and (c) **low communication** complexity. Furthermore, (d) **no initial setup** is performed on x except for encrypting it element-by-element (each element encrypted bit-by-bit) and uploading the resulting ciphertexts to the server.

We point out that the latter condition, among other things, prevents speeding up the search by using standard data structures such as search-trees, hash-tables, or sorted arrays (on top of, or instead of, the encrypted unsorted array $[\![x]\!]$). A *linear scan lower bound* is thus implied by the addressed formulation, even if we were to search on clear-text data. This restriction is nonetheless motivated by many use-cases, as discussed next.

Use-cases motivating the aforementioned no-setup restriction arise in settings where, for example:

- Matching criteria are unknown in advance, thus precluding appropriate indexing or sorting at setup;
- High dimensional range queries, where index size is exponential in the number attributes and infeasible to compute or store;
- Streaming data with client discarding each element immediately after encrypting and uploading to the server, thus precluding client's setup or maintenance of the desired data-structures (and where for the server, seeing only ciphertexts, secure maintenance of advanced data structures seems even harder than secure search);

Table 1: Comparing search solutions that **allow setup** in the single-server settings.

Setup Allowed	Security				Sub-Linear Complexity in $\|x\|$?	Single Round
	Hide Query Content	Hide Elements Content	Hide Search Pattern	Hide Access Pattern		
Cleartext	×				✓$_{(iv)}$	✓
SE	✓$_{(i)}$	✓$_{(i)}$	×$_{(ii)}$	×$_{(ii)}$		✓
SE + ORAM	✓$_{(iii)}$	✓$_{(iii)}$	×$_{(iii)}$	×$_{(iii)}$		×
PIR by Keywords	✓	×	✓	✓		✓

Rows correspond to related works (see Section 1.1.2); Columns correspond to properties; Cells contain ✓ if the column's property is attained by the row's work (× if not attained). ***Acronyms and Abbreviations:*** *Comm. – Communication; SE – Searchable Encryption; PIR – Private Information Retrieval; FHE – Fully Homomorphic Encryption; ORAM – Oblivious RAM.* ***Comments:*** *(i) Known methods exploit leakage of SE protocols to obtain query and/or elements content [11, 69, 37, 1, 26, 30, 31, 53]. (ii) SE deliberately leaks information to enable highly efficient search over encrypted data, see additional info in [7]. (iii) Employing ORAM to completely prevent leakage was shown impractical in achieving sub-linear complexity protocols (e.g. [49]). (iv) Sub-linear client, server and communication complexity in $|x|$.*

- Low capacity clients that are too weak to run setup over the entire cleartext array prior to encrypting and uploading it to the server;
- Fragmented data uploaded to the server from multiple distinct client endpoints (data-sources) with no single endpoint that can perform setup over the entire cleartext data.

The single-round and low-communication restrictions are motivated by use-cases in settings where communication is a major bottleneck, e.g. in being intermittent or unreliable, or where communicating is with data-sources that are mostly offline, or have restricted battery capacity as in sensors-networks or some Internet-of-Things (IoT) devices.

The single server restriction is motivated, not only by the simplicity of such architecture, but also by its stronger security guarantee: requiring no non-collusion assumption on servers.

Threat model. We address computationally-bounded semi-honest adversaries that follow the protocol but may try to learn additional information. Our security requirement is that adversaries controlling the server cannot distinguish between two adversarially-chosen equal size queries or data arrays. See Section 3.3.

The leakage of our protocols include only size information (specifically, upper-bounds on array size, elements' sizes, number of queries, and queries' sizes); see detailed leakage discussion in section 3.3.

1.1 Prior Works

We survey related works, focusing primarily on works addressing similar secure-search formulation as addressed in this work: *single-server, single-round, low-communication* and *no-setup*. See Tables 1-2.

1.1.1 Secure Search on FHE Encrypted Data

The most relevant works are those addressing the same secure-search formulation as considered in our work (aka, secure search on FHE encrypted data); See above and Definition 4.

Folklore solutions for secure search on FHE encrypted data suffered from inefficient server runtime due to evaluating degree $\Omega(n)$ polynomials, for n the number of elements; see discussion in [3].

SPiRiT. The prior state-of-the-art for secure-search on FHE encrypted data appeared in a recent work of Akavia, Feldman and Shaul CCS'2018 [3], where the server evaluates a polynomial of logarithmic degree $\log^3 n$ (instead of degree at least linear in the folklore solution). Their work proposes both a deterministic and a randomized variant.

Their deterministic variant (SPiRiT Det.) uses modern data summarization techniques known as sketches [66] alongside multi-ring simultaneous evaluations of their search polynomial to retrieve a

poly-logarithmic short list of candidates for the first matching element. This novel technique essentially reduces the degree of the polynomial evaluated by the server from linear to poly-logarithmic in the number of elements n.

Their randomized variant (SPiRiT Rand.) offers an efficiency improvement by working over a single random ring instead of several different rings. The disadvantage of this randomized variant is that it achieves an error probability that is only polynomially small in n (i.e., noticeable error) rather than a negligible error probability.

1.1.2 Other Related Works

We next discuss other related methods and works.

Setup vs. no-setup – server's efficiency gap. There is a major efficiency gap between works allowing initial setup to works disallowing setup, as in our work. This is because setup allows sub-linear search time (e.g. using indexing, search trees or hash table), whereas disallowing setup necessitates –even on cleartext data– a linear scan of the data. Our work focuses on the no-setup case.

Secure two-party computation (2PC w/o FHE). Seminal works dating back to the 1980s [67, 27] showed that two parties can compute any polynomial-time computable function of their private inputs via an interactive protocol that reveals no information beyond what can be inferred from the function's output. In particular, parties can securely compute the search functionality.

However, secure two-party protocols preceding the constructions of FHE schemes (2PC w/o FHE) suffer from a communication complexity, and hence also the client's time, grows with the complexity of the computed function. This is in contrast to growing only with the input and output sizes $|q| + |(i, x[i])|$ in FHE based solutions.

We note that, while we focus on secure search in the two parties settings (client and a single server), promising results have been shown for settings where the server can be partitioned into several non-colluding entities that secret-share the data; a comprehensive survey of such works is beyond our scope. A few examples include [4, 5, 65], all addressing search on cleartext data held by the multiple servers, namely, protecting the query but not the data against the server.

Searchable Encryption (SE) focuses on inherent efficiency versus security trade-off when searching on encrypted data. Specifically SE focuses on achieving sublinear search time. Main primitives for SE include searchable symmetric encryption (SSE) and public key encryption with keyword search (PEKS) with schemes first introduced by Song et al. [62] and Boneh et al. [6] respectively. See also [7] for a thorough survey.

To achieve sublinear time SE incorporates setup such as sorting, indexing or usage of auxiliary data structures. Furthermore SE deliberately leaks information to enable highly efficient search over encrypted data. This leakage typically includes the access pattern (which elements match a given keyword) and/or search pattern (whether different queries were generated for the same searched keyword).

Starting with the work of [18], research on SE formalized security by defining a *leakage profile* that characterizes the information an adversary may learn. In some cases, an adversary can exploit a schemes leakage to completely reveal the content of elements and queries it comes across. See discussions and example attacks in [11, 69, 37, 1, 26, 30, 31, 47, 53].

In contrast, in FHE based solutions such as our work there is no leakage other than size information; see detailed leakage discussion in Section 3.3.

SE using Oblivious RAM (SE+ORAM). ORAM [28] is a cryptographic primitive that allows a client to conceal its access pattern to a remote storage held by the server via continuous re-shuffling of a dedicated structure and re-encryption of the data as it is being accessed. Typical ORAM constructions (e.g. [57, 63]) achieve poly-logarithmic server run-time by requiring the client to perform a "download-decrypt-compute-encrypt-upload" each round, for at least $\mathcal{O}(\log n)$ rounds.[1]

In the context of SE, it has been believed that there exists a method of employing ORAM to eliminate access and search patterns leakage. However, Naveed [49] show that using ORAM to completely eliminate

[1] An exception is TWORAM [22] that achieves 2 rounds by utilizing garbled circuits.

leakage either renders the communication performance worse than the trivial approach of streaming all of the outsourced data for each query or they do not provide any meaningful reduction in leakage.

Private Information Retrieval (PIR) [17] in a single server scenario [43, 50], allows a client holding the *physical address* $i \in [n]$ to retrieve the ith element $x[i]$ from a data array $x = (x[1], \ldots, x[n])$ held by the server. The above is achieved while ensuring the server learns no new information on i or $x[i]$ and while keeping the communication complexity strictly smaller than $|x|$. The state-of-the-art communication complexity, as achieved by FHE based PIR [23, 9, 19], is proportional only to the size of i, $x[i]$ and the security parameter. We note however that the server's run-time in a single server PIR (whether or not FHE based) is inherently linear in $|x|$.

PIR-by-Keywords via setup on unencrypted data. PIR-by-keywords protocols [16, 12, 56] remove the requirement to know the physical address i of the sought element $x[i]$, and allow instead to retrieve elements using keywords search. These works however address different settings than our work: their server holds *unencrypted data* and performs *setup* on that data to produce auxiliary data-structures used to speedup search (in contrast to the server's holding encrypted data and performing no-setup in our work).

Several disadvantage of this PIR-by-keyword approach make it unsuitable for the use-cases and settings considered in our work: (1) Their server holds unencrypted data, offering no data protection against the server. This is inherent for their setup, as efficient setup on encrypted data is a challenging problem not addressed by these works.[2] (2) They require setup for producing their data-structures, which is unsuitable for use-cases where setup is disallowed or infeasible; See use-cases examples above. (3) Their search-space is restricted by the initial setup choices, e.g., the keywords used in producing their index [12, 56].

In contrast, in our work (1) Data is encrypted by the client to guarantee data protection against server. (2) No initial setup or maintenance of additional data structures is required, neither by the server nor by the client. (3) We enable arbitrary matching criteria to be chosen by the client, on the fly, for each query.

Private Set Intersection (PSI) [21] enables two parties each holding a private set to securely compute the intersection of their sets. PSI protocols were constructed using various cryptographic primitives (e.g. [51, 52, 41], including FHE [13]). PSI protocols can be employed to solve the decision problem of whether the lookup value (size 1 set) appears in the data array (size n set); however, with server holding *unencrypted data*.

Secure Pattern Matching (SPM) on FHE encrypted data [14, 15, 64, 39, 40, 44, 68], given an encrypted lookup value, returns a vector of n ciphertexts $(c_1, \ldots, c_n)$, for n the number of elements, where c_i indicates whether the ith data element is a match to the lookup value (or sometimes returning only a YES/NO answer of whether a match exists). The main drawback of these protocols is that the communication complexity and client's running time are proportional to the number of stored elements $\Omega(n)$.

1.2 Our Contributions

In this work we present a new and improved solution for secure search on FHE encrypted data, analyze its efficiency compared to the prior state-of-the-art and demonstrate its concrete run-time performance by providing an implementation (built on top of HElib C++ library [32]) together with extensive experiments.

Our secure-search protocol is a single-server, single-round, low-communication protocol that requires no initial setup or maintenance of additional data-structures.

Our protocol is compatible with generic matching criteria, such as exact and wild-card matching; similarity search in metric spaces, e.g., with Hamming/Euclidean/Edit distance; Boolean and range queries; and so forth. See Section 7 for specific instantiations examples; efficiency improvements via universal hashing; and employment for fetch-next queries.

[2] In contrast, in standard PIR when given the physical address i, it is irrelevant whether x is encrypted, because the server simply retrieves whatever content is in $x[i]$.

Table 2: Comparing search solutions that **disallow setup** in the single-server, single-round settings.

Setup Disallowed	Security	Sub-Linear Complexity in $\|x\|$?			Allows Multiple Matches	Retrieval of Index and Element	Post Processing Free	Negligible Error Probability	Compatibility with all FHE Schemes
		Server	Client	Comm.					
Cleartext	×	×	✓	✓	✓	✓	✓	✓	N/A
2PC w/o FHE	✓		×	×	✓	✓	✓	✓	N/A
PIR	✓		✓	✓	×	✓	✓	✓	✓
PSI	✓		✓	✓	×	×	✓	✓	✓
SPM	✓		×	×	✓	×	×	✓	✓
Folklore	✓		✓	✓	✓	✓	✓	✓	✓
SPiRiT Det.	✓		✓	✓	✓	✓	×	✓	×
SPiRiT Rand.	✓		✓	✓	✓	✓	✓	×	×
Binary Raffle	✓		✓	✓	✓	✓	✓	✓	✓

Rows correspond to related works (see Section 1.1); Columns correspond to properties (see Section 3.2); Cells contain ✓ if the column's property is attained by the row's work (× if not attained, N/A if not applicable). ***Acronyms and Abbreviations:*** *Comm. – Communication; 2PC – Two Party Computation; PIR – Private Information Retrieval; FHE – Fully Homomorphic Encryption; PSI – Private Set Intersection; SPM – Secure Pattern Matching; Folklore – Natural secure search on FHE encrypted data; SPiRiT Det. and Rand. – deterministic and randomized protocols of AFS [3];* ***Binary Raffle – This Work****.* ***Comments:*** *(i) All the works in this table attain* ***Single Round*** *protocols.*

Table 3: Complexity comparison of first match index (i^*) computation phase between SPiRiT (rows (i)-(ii)) vs. our work (rows (iii)-(v)).

	• **Server's Degree,** • **Server's Overall Multiplications** • **Client's Decryptions**
(i) ***SPiRiT* Det.**	• $\log^3(n) \cdot d$ • $k \cdot n \cdot \left(\log^2(n) + \mu\right)$ • $k \cdot \log(n)$
(ii) ***SPiRiT* Rand.**	• $\frac{c}{2\varepsilon} \log^3(n) \cdot d$ • $n \cdot \left(\log\left(\frac{n}{\varepsilon} \cdot \frac{c}{2} \cdot \log n\right) + \mu\right)$ • $\log(n)$
(iii) ***Binary Raffle***	• $\log(n/\varepsilon) \cdot d$ • $n \cdot (\log(n/\varepsilon) + \mu)$ • $\log(n)$
(iv) ***Binary Raffle* + Universal Hash**	• $2 \cdot \log^2(2n/\varepsilon)$ • $n \cdot (3 \cdot \log(2n/\varepsilon))$ • $\log(n)$
(v) ***Binary Raffle* + Client Probability Amplification**	• $\log(3n) \cdot d$ • $\alpha \cdot n \cdot (\log(3n) + \mu)$ • $\alpha \cdot \log(n)$

Notations: n – array size; ε – failure probability; d, μ – degree and overall multiplications of IsMatch*; $k = \log^2 n / \log\log n$; $\alpha = \mathcal{O}(\log(1/\varepsilon))$; c – a constant depending on the density of prime numbers.*

The client's complexity is optimal in the sense of only encrypting the input and decrypting the output. The communicating consists only of the encrypted input and output. The server sees only ciphertexts for both data and queries, encrypted with FHE in a black-box fashion and compatible with all known FHE candidates. The server evaluates a search polynomial over the encrypted data and encrypted query. This polynomial for computing both index i^* and element $x[i^*]$ is of degree $\log(n/\varepsilon) \cdot d$ and overall multiplication $n\,(\log(n/\varepsilon) + \mu + w)$. Here n is the number of data elements, ε the failure probability, w the binary representation length of $x[i^*]$, and d, μ the degree and overall multiplications respectively for the polynomial realizing the matching criterion. See Table 3.

The security guarantee against semi-honest adversaries controlling the server is that our protocol leaks no information on data and queries, except for size information (aka, full-security). Namely, the leakage profile consists solely of upper bounds on the counts and sizes of queries and data elements. In particular, the adversary cannot tell whether two queries are for the same keyword, or whether the client issued a fresh query or a fetch next query, etcetera. See Section 3.3.

Comparison to prior works. We next compare our protocol to prior works on secure search, focusing on single-server, single-round protocols, and discussing both works that allow and disallow setup; See Table 1 and Table 2, respectively.

Our work is incomparable to works allowing setup; See SE, SE+ORAM and PIR-by-Keywords in Section 1.1.2. On the one hand, setup enables attaining search with sub-linear server complexity, which is impossible without setup even on cleartext data and query. On the other hand, we attain a stronger security guarantee of hiding all the following: data content, query content, access pattern, and search pattern; in contrast to leaking at least some of the former in the aforementioned works; See Table 1.

Our work strictly improves over prior secure-search works that disallow setup (see Folklore, SPiRiT, 2PC w/o FHE, PIR, PSI and PSM in Section 1.1), in the following sense. Our secure-search protocol is the first to simultaneously attain all desired properties that follows (Properties 1-9, Section 3.2): full security (i.e., completely hiding all the following: query content, data elements content, search pattern, access pattern); efficient client and communication (i.e., polynomial in input and output size, and not in the time to compute the search functionality), single and efficient server (in the sense of evaluating over encrypted data a polynomial of degree poly-logarithmic in the number of data elements); unrestricted search functionality; retrieval of both index and element; post-processing free; negligible error probability; and compatibility with all current FHE schemes. In contrast, all prior secure search solutions achieve only a strict subset of these properties. See Table 2.

In particular, when comparing to the prior state-of-the-art secure-search on FHE encrypted data (SPiRiT) [3] our protocol offers the following contributions:

Contribution 1. Our protocol simultaneously achieves both the properties of post-processing free client and negligible error probability. In contrast, the protocols of [3] achieve either post-processing free client or negligible error probability, but not both. Simultaneously achieving both properties, as in our work, is highly motivated as it allows the server to employ secure search as a sub-component in a larger computation without interaction with the client.

Contribution 2. Our secure search solution is asymptotically faster than [3], in attaining: (1) Optimal client run-time in the sense of requiring only encrypting the input and decrypting the output. (2) Considerable improvement of the server's run-time: we reduce the degree of the evaluated polynomial from cubic to linear in $\log n$, and from linear to logarithmic in $1/\varepsilon$, and reduce the overall multiplications by up to $\log n$ factor. See Table 3.

Contribution 3. Our secure search solution requires computations solely over $\mathrm{GF}(2)$ (instead of $\mathrm{GF}(p)$ for primes $p > 2$ in [3]). This leads to compatibility with all currently known candidate FHE schemes including GSW [25], unlike [3]. This also allows further run-time speedup when using current FHE schemes implementations, including HElib [32] that implements BGV [8] scheme. The reason for the speedup is that in all current FHE schemes, working over $\mathrm{GF}(p)$ for larger primes $p > 2$ causes a general slowdown of all the homomorphic operations and size inflation of the keys and ciphertexts.

Contribution 4. Our secure search solution is concretely faster than [3] by an order of magnitude. This is demonstrated by our implementation, based on the FHE HElib C++ library [32], and our extensive run-time benchmarks experiments, performed on a mid-range Linux server of 16 CPU cores and 16GB RAM. A few examples of comparing our results on same server and with similar parameters for bits per element, execution time and error probability follow; See more details in Section 6.3.

(i) We securely search on $\approx 3 \times 10^6$ (in contrast to $\approx 0.2 \times 10^6$ for SPiRiT) 16-bit elements in 4.5 hours, with error probability 2^{-80}.

(ii) We securely search on $\approx 3 \times 10^6$ (in contrast to $\approx 0.3 \times 10^6$ for SPiRiT) 16-bit elements in 1 hour, with error probability 1/2.

(iii) We securely search on $\approx 1 \times 10^6$ (in contrast to running out of RAM and being unable to complete the experiment for SPiRiT) 64-bit elements in 1 hour, with error probability 1/2.

(iv) We securely search on $> 10 \times 10^6$ (in contrast $\approx 1.5 \times 10^6$ for SPiRiT) 1-bit elements in 1 hour, with error probability 1/2.

1.3 Our Techniques Highlights

When considering secure search over unsorted FHE encrypted data, approaches like binary-search are rejected immediately.

The approach proposed by Akavia et. al. [3] for solving secure search on FHE encrypted array includes the following steps: (1) Obtaining a binary array of indicators after executing the desired IsMatch predicate between the given query and each array element; (2) Calculating an array of prefix-sums of the array of binary indicators; (3) Transforming the prefix-sums array to a binary step-function array with value 1 at every non-zero prefix-sum, namely, the first 1 bit is in the index of the first match; (4) Transforming the step-function array to a selector array where only the index of the first match contains 1 (and all other indices contain 0); (5) Utilizing this selector array to calculate and return this first match (index and element).

Realizing this approach however is challenging: Step (2) (computing prefix-sums) has high degree if working with binary plaintext space and using standard addition circuits such as full-adders. Step (3) (zero-testing each prefix-sum) has high degree if working over plaintext spaces larger than the array size and utilizing Fermat's Little Theorem for the zero-test.

To address this challenge Akavia et. al. [3] propose combining steps (2)-(3) above to a single probabilistic step that returns the required step-function (albeit, with noticeable error probability). They later show how to eliminate the error using few repetitions over multiple rings $\mathrm{GF}(p)$ for $p > 2$ together with client post-processing for selecting the correct result.

To avoid the aforementioned post-processing we propose an alternative for the probabilistic test combining steps (2)+(3). First, we make the straightforward observation that instead of testing if the sum of binary indicators is zero (as done in [3]), we can compute the logical-OR of these indicator values. However, this would result in high degree, as the logical-OR over n variables has degree n. Next, to reduce the degree, we employ the method of Razborov and Smolensky [54, 61] for low-degree approximation of the logical-OR function. This method yields a polynomial of degree logarithmic in both n and $1/\varepsilon$, for ε the failure probability.

Elaborating on the above, the Razborov-Smolenski method is applicable in $\mathrm{GF}(q)$ for any $q \geq 2$; we apply it with $q = 2$ on all k-th prefix $\big(v[1], \ldots, v[k]\big) \in \{0,1\}^k$ of the aforementioned vector of n binary indicator values. Their low-degree approximation for $\mathrm{OR}\big(v[1], \ldots, v[k]\big) \in \{0,1\}$ is computed as follows. First, for $N(\varepsilon) = \lceil \log_2(n/\varepsilon) \rceil$ uniformly random i.i.d. $r_1, \ldots, r_{N(\varepsilon)} \in \{0,1\}^n$, we compute the parity of the corresponding random subset of entries,

$$\mathsf{p}(r_j) = \sum_{i=1}^{k} r_j[i] \cdot v[i] \mod 2.$$

The parity bit $\mathsf{p}(r_j)$ is always zero when $v = 0^k$ and it is one with probability half when $v \neq 0^k$. Next, we compute the OR of these parity values using the standard degree $N(\varepsilon)$ polynomial for the logical-OR of $N(\varepsilon)$ binary values:

$$\mathrm{OR}\Big(\mathsf{p}\big(r_1\big), \ldots, \mathsf{p}\big(r_{N(\varepsilon)}\big)\Big) = 1 - \prod_{j=1}^{N(\varepsilon)} \big(1 - \mathsf{p}\big(r_j\big)\big) \mod 2.$$

This is equal to $\mathrm{OR}(v[1],\ldots,v[k])$ with probability $1-\frac{\varepsilon}{n}$.

We note that the Razborov's and Smolensky's [54, 61] approximation method has numerous uses in computer science. In particular, in the context of secure search Barkol and Ishai [4], building on [54, 61, 36], gave a generic transformation from constant-depth unbounded fan-in boolean circuits to low-degree polynomials. They employ their technique for secure multi-party computation of common search functionalities; albeit, in settings of multiple-servers holding *unencrypted data* (cf. single-server holding encrypted data in our work).

1.4 Article Road-map

The rest of this paper is organized as follows. Preliminary definitions and notations in Section 2; Problem statement and threat model in Section 3; Our protocol in Section 4; Theoretical results in Section 5; Experimental results in Section 6; Instantiations of $\mathsf{IsMatch}$ demonstrating performance and functionality enhancements and extensions in Sections 7-8; Conclusions in Section 9. We deffer to the appendix missing proof details for our main theorem (Appendix A), and details regarding the natural secure search solution on FHE encrypted data (folklore) (Appendix B).

2 Preliminaries

We state some preliminary notations and definitions.

2.1 Notations

For natural numbers $k<n$, denote $[n]=\{1,\ldots,n\}$, $[k,n]=\{k,\ldots,n\}$, and $(k,n)=\{k+1,\ldots,n-1\}$. For array v denote $v[i]$ the i-th element in v. Similarly, for $x\in\{0,1\}^*$, $x[i]$ denotes its i-th bit. We follow the convention of enumerating array entries starting from entry number 1 (not 0), unless stated otherwise. For matrix M the element in row i and column j will be denoted as $M[i,j]$. For a field $\mathbb{F}$, vectors $v,u\in\mathbb{F}^n$ and $k\in[n]$, denote: $\langle v,u\rangle=\sum_{i=1}^{n}v[i]\cdot u[i] \pmod 2$, $\mathsf{prefix}_k(v)=(v_1,\ldots,v_k)\in\mathbb{F}^k$, $\mathsf{suffix}_k(v)=(v_{k+1},\ldots,v_n)\in\mathbb{F}^{n-k}$, and $|v|$ the size (length, dimension) of v $(=n)$.

For $k,n\in\mathbb{N}$, denote by $r_1,\ldots,r_k\leftarrow_{\$}\{0,1\}^n$ the sampling of k arrays independently at random from the uniform distribution over $\{0,1\}^n$. As standard, PPT denotes *probabilistic polynomial time*; and a function $\nu\colon\mathbb{N}\to\mathbb{R}$ is called *negligible* in κ, denoted $\mathsf{negl}(\kappa)$, if for every constant $c>0$ there exists n_0 such that for all $n>n_0$, $\nu(n)<\kappa^{-c}$.

2.2 Fully Homomorphic Encryption

Definition 1 (FHE)**.** A *leveled homomorphic encryption (FHE)* scheme is defined by a quadruple of PPT algorithms $\mathcal{FHE}=(\mathsf{KGen},\mathsf{Enc},\mathsf{Dec},\mathsf{Eval})$ as follows.

- **Key generation.** $(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{KGen}(1^\kappa,1^L)$ takes a security parameter κ and a circuit depth upper-bound L, and outputs public key pk and secret key sk.

- **Encryption.** $[\![b]\!]\leftarrow\mathsf{Enc}_{\mathsf{pk}}(b)$ takes the public key pk and a message $b\in\{0,1\}$, and outputs a ciphertext $[\![b]\!]$. For $x\in\{0,1\}^n$, we denote its bit-by-bit encryption $[\![x[i]]\!]\leftarrow\mathsf{Enc}_{\mathsf{pk}}(x[i])$ by $[\![x]\!]=\big([\![x[1]]\!],\ldots,[\![x[n]]\!]\big)$.

- **Decryption.** $x'\leftarrow\mathsf{Dec}_{\mathsf{sk}}([\![x]\!])$ takes the secret key sk and a ciphertext $[\![x]\!]$, and outputs a message $x'\in\{0,1\}^*$. When $[\![x]\!]$ is an array of ciphertexts, decryption is ciphertext-by-ciphertext.
  Correctness says that $\mathsf{Dec}_{sk}(\mathsf{Enc}_{pk}(x))=x$.

- **Homomorphic evaluation.** $[\![y]\!]\leftarrow\mathsf{Eval}_{\mathsf{pk}}(f,\,[\![x[1]]\!],\ldots,[\![x[t]]\!])$ takes pk, a function $f:\{0,1\}^t\to\{0,1\}$ represented as an arithmetic circuit over $\mathrm{GF}(2)$ and a set of t ciphertexts $\big([\![x[i]]\!]\big)_{i=1}^{t}$ outputs a ciphertext $[\![y]\!]$ such that $\mathsf{Dec}_{\mathsf{sk}}([\![y]\!])=f(x[1],\ldots,x[t])$. As a shorthand notation we write $f([\![x]\!])$ in place of $\mathsf{Eval}_{\mathsf{pk}}(f,[\![x]\!])$.

We will use the standard equality operator, for $a, b \in \{0,1\}^w$, of degree w and $w-1$ overall multiplications:

$$\mathsf{IsEqual}_w(a,b) = \prod_{i\in[w]} \Big(1 + a[i] + b[i]\Big) \qquad \bmod 2 \tag{1}$$

We will also use the standard "greater than" operator ($a > b$), for $a, b \in \{0,1\}^w$, of degree $w+1$ and $2w$ overall multiplications:

$$\mathsf{IsGrt}_w(a,b) = \sum_{i\in[w-1]} \Big[\big(a[i]\cdot(b[i]+1)\big)\cdot \mathsf{IsEqual}_{w-i}\big(\mathsf{suffix}_i(a), \mathsf{suffix}_i(b)\big)\Big] + \big(a[w]\cdot(b[w]+1)\big) \qquad \bmod 2 \tag{2}$$

2.3 BGV FHE Scheme and HElib Implementation

The HElib [32] implementation of the BGV [8] homomorphic encryption scheme includes Smart-Vercauteren [60] Single Instruction Multiple Data (SIMD) optimization alongside many other optimizations [24, 33, 34]. We now point out several characteristics of BGV and HElib that will play key roles in the run-time efficiency analysis of our implementations. More details on these characteristics and general design of HElib can be obtained from [35, 32].

Plaintext Space This BGV variant in HElib is defined over polynomial rings of the form $\mathbb{A} = \mathbb{Z}[X]/\Phi_m(X)$ where m is a parameter and $\Phi_m(X)$ is the m'th cyclotomic polynomial. Let $\mathbb{A}_q = \mathbb{A}/q\mathbb{A} = \mathbb{Z}[X]/(\Phi_m(X), q)$ for integer q, and identify $\mathbb{A}_q$ as the set of integer polynomials of degree up to $\phi(m)-1$ reduced module q (where $\phi(m)$ is Euler's totient function). The "native" plaintext space for HElib is the ring $\mathbb{A}_2$, although after HElib incorporated additional optimizations, other plaintext spaces, including $\mathbb{A}_p$ for any arbitrary prime p, are also available. Enlarging the plaintext space $\mathbb{A}_p$ to larger prime values leads to a general slowdown in all the homomorphic operations and size inflation of the keys and ciphertexts.

Plaintext Packing and SIMD Smart-Vercauteren optimization allows to "pack" many plaintext elements in a single ciphertext and apply to them operations in a SIMD manner. We refer to the different plaintext values in a single ciphertext as the "plaintext slots" of that ciphertext. This is achievable by factoring the polynomial $\Phi_m(X)$ into s irreducible factors modulo 2, $\Phi_m(X) = F_1(X)\cdot\ldots\cdot F_s(X) \pmod 2$, all of degree $d = \phi(m)/s$. Now we can view a polynomial $a \in \mathbb{A}_2$ as representing a vector $(a \mod F_i)_{i=1}^s$ that holds s encodings of plaintext values. Notice that the amount of plaintext slots s is dependent on the value $\phi(m)$ (the degree of the cyclotomic polynomial $\Phi_m(X)$), thus higher values of $\phi(m)$ will usually enable more plaintext slots. It is important to mention that in addition to enabling more plaintext slots, incrementing $\phi(m)$ will also cause a general slowdown of all the homomorphic operations and size inflation of the keys and ciphertexts.

Leveled FHE As BGV is a leveled FHE scheme, the ciphertext space for this scheme consists of vector over $\mathbb{A}_q$, where q is a large odd modulus that evolves with the homomorphic evaluation. Specifically, the system is parametrized by a "chain" of moduli of decreasing size, $q_0 < q_1 < \ldots < q_L$ and freshly encrypted ciphertexts are defined over $\mathbb{A}_{q_L}$. During homomorphic evaluation, after each multiplication, to handle the increasing "noise", a switching to smaller and smaller moduli is preformed until a ciphertext over $\mathbb{A}_{q_0}$ is obtained, on which further computations are impossible.

When working over plaintext space $\mathbb{A}_p$ for $p > 2$ the above operation of ciphertext "refresh" after each multiplication can "consume" several such levels. This is caused by the rounding operation during modulus switching where additional "noise", proportional to $\mathsf{poly}(p)$, is added to the ciphertext. Therefore, when working with plaintext space with higher value of p, additional levels (and additional primes q_i) are necessary in order to enable successful computation.

Security Parameter We set the security parameter of HElib to 80 bit (same as Akavia et al. [3] and other works [38, 45, 39]). Changing this parameter is possible in the initialization phase of HElib.

Another aspect of the security parameter is that it brings about a linear increase in $\phi(m)$, which in term will cause a longer computation time of homomorphic operations and increase in ciphertexts and keys sizes.

2.4 Universal Family of Hash Functions

Carter and Wegman [10] defined the notion of a universal family of hash functions:

Definition 2 ([10])**.** A family of hash functions $\mathcal{H} : \mathcal{W} \to \mathcal{V}$ is said to be **strongly universal** if for every $h \in \mathcal{H}$ and for all $x \neq y \in \mathcal{W}$,

$$\Pr_{h \in \mathcal{H}} \left[h(x) = h(y) \right] \leq \frac{1}{|\mathcal{V}|}$$

Theorem 2.1. *Given a hash function h sampled randomly from a family of **strongly universal** hash functions $\mathcal{H} : \mathcal{W} \to \mathcal{V}$ and n distinct values $x_1 \ldots, x_n \in \mathcal{W}$ ($\forall i \neq j \in [n] : x_i \neq x_j$), the hashed values $h(x_1), \ldots, h(x_n)$ are all distinct with probability*

$$\Pr_{h \in \mathcal{H}} \left[\forall i \neq j \in [n] : h(x_i) \neq h(x_j) \right] > 1 - \frac{\binom{n}{2}}{|\mathcal{V}|}$$

Proof. There are exactly $\binom{n}{2}$ possible value pairs among $x_1, \ldots, x_n$. From the property of the hash function being strongly universal we get that for each value pair $x_i \neq x_j$ for $i \neq j$ the probability for a collision ($h(x_i) = h(x_j)$) is at most $|\mathcal{V}|^{-1}$. Finally, by applying the union bound over all possible value pairs we get that the total probability for any collision is at most $\binom{n}{2} \cdot |\mathcal{V}|^{-1}$ as required. □

Corollary 2.1. *Given security parameter κ, and suppose that $|\mathcal{V}| = 2^v$, if $v = 2\log_2(n) + \omega(\log_2(\kappa))$ then we get*

$$\Pr_{h \in \mathcal{H}} \left[\forall i \neq j \in [n] : h(x_i) \neq h(x_j) \right] = 1 - \mathsf{negl}(\kappa)$$

Proof. Immediate by using $2^{-\omega(\log_2(\kappa))} = \mathsf{negl}(\kappa)$ and assigning $v = 2\log_2(n) + \omega(\log_2(\kappa))$ in Theorem 2.1. □

Theorem 2.2 ([10])**.** *Let $\mathcal{W} = \{0,1\}^w$ and $\mathcal{V} = \{0,1\}^v$. We now provide the following construction of hash function family: Sample random matrix $A \in \{0,1\}^{w \times v}$ and vector $b \in \{0,1\}^v$, for any $x \in \{0,1\}^w$ the hash value $h(x) \in \{0,1\}^v$ will be*

$$h(x) = Ax + b \pmod 2$$

*This is a **strongly universal** family of hash functions that contains $2^{(w+1)v}$ functions.*

In the above construction, describing the hash function requires $\mathcal{O}(w \cdot v)$ bits. We can reduce this storage overhead with the corollary below.

Definition 3 (Toeplitz Matrix)**.** $A \in \{0,1\}^{w \times v}$ is defined as following: Fill the first row $A_{1,1}, \ldots, A_{1,w}$ and the first column $A_{1,1}, \ldots, A_{v,1}$ with random bits. For every other entry $A_{i,j}$ for $i > 1$ and $j > 1$ define $A_{i,j} = A_{i-1,j-1}$. So all entries in each "northwest-southeast" diagonal in A are the same.

Corollary 2.2 ([42])**.** *Let $\mathcal{W} = \{0,1\}^w$ and $\mathcal{V} = \{0,1\}^v$. The construction of the following hash function family will be identical to the one in Theorem 2.2 with the sole difference that $A \in \{0,1\}^{w \times v}$ will be a random Toeplitz Matrix. This is a **strongly universal** family of hash functions that contains $2^{(w+v-1)+v}$ functions, its description requires only $\mathcal{O}(w + v)$ bits.*

3 Problem Statement

Suppose a client (Alice) wants to use a server (cloud service provider, Bob) for data storage, management and retrieval (Search, Insert, Update, Delete). To protect her privacy Alice uploads only encrypted data to the cloud. She encrypts it using FHE so that Bob has processing capabilities on the data, with single round and low communication protocols that hide Alice's data, queries, returned results and access pattern from Bob.

3.1 Secure-Search on Encrypted Data

In this paper we focus on the problem of setup-free secure-search on FHE encrypted data, following [3]; see Definition 4 below. In this problem given an unsorted and encrypted data array $[\![x]\!] = ([\![x[1]]\!], \ldots, [\![x[n]]\!])$ and encrypted query $[\![q]\!]$ the goal is to find the encrypted index $[\![i^*]\!]$ and element $[\![x[i^*]]\!]$ so that $x[i^*]$ is the first match for query q in array x (formally, $\mathsf{IsMatch}(x[i^*], q) = 1$ and $\forall j < i^* : \mathsf{IsMatch}(x[j], q) = 0$). The predicate IsMatch can be generic (see below).

Definition 4 (Secure search)**.** The server holds an array of encrypted elements (previously encrypted and uploaded by the client to the server, and where the server has no access to the secret decryption key):

$$[\![x]\!] = ([\![x[1]]\!], \ldots, [\![x[n]]\!])$$

The original array $x = (x[1], \ldots, x[n])$ is **unsorted** and its elements are **not necessarily distinct**. The client sends to the server an encrypted query $[\![q]\!]$. The server returns the client an encrypted index $[\![i^*]\!]$ and element $[\![x[i^*]]\!]$ where the index i^* is satisfying the condition that it is the index of the first match for query q in array x: $i^* = \min\{i \in [n] \mid \mathsf{IsMatch}(x[i], q) = 1\}$

Setup-free secure-search protocol employing a solution to the above secure search problem follows (cf. Figure 1):

1. *Keys Generation*: Alice initializes the scheme $\mathcal{FHE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ and generates the keys $(\mathsf{pk}, \mathsf{sk})$. Alice keeps pk, sk and sends pk to Bob.

2. *Array Upload*: Alice gradually, over time, encrypts and uploads elements to Bob. At any given moment, Bob holds an encrypted and unsorted array $[\![x]\!] = ([\![x[1]]\!], \ldots, [\![x[n]]\!])$.

3. *Secure search:* At any time, Alice may issue a search query q by encrypting it and sending $[\![q]\!]$ to Bob. Bob employs the secure search solution to obtain and send to Alice the encrypted search outcome $([\![i^*]\!], [\![x[i^*]]\!])$ for i^* the index of the first match and $x[i^*]$ the corresponding element. Alice then decrypts to obtain i^* and $x[i^*]$.

Usage could be versatile. For example, the client (Alice) may upload additional data over time with search queries interleaved between uploads; See Section 8.2. Furthermore, Alice could be instantiated by multiple parties with distinct roles: a key generation authority in Step 1, and multiple data-sources and search-clients in Steps 2 and 3 respectively, where the key generation authority sends pk to the data-sources and server, and sends pk, sk to the search-clients.

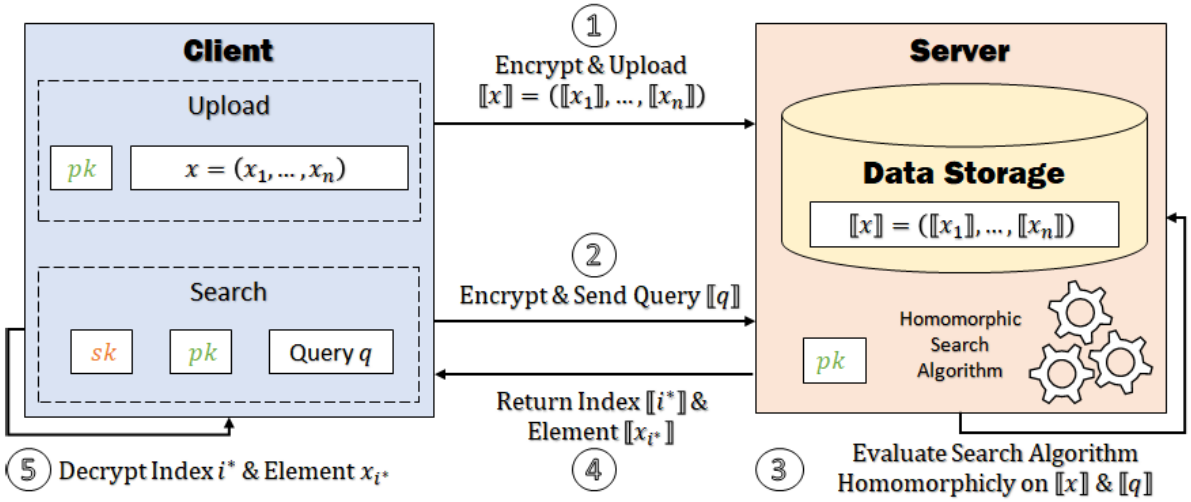

Figure 1: Depiction of *secure search* on FHE encrypted data

Generic IsMatch. In order for secure search to be applicable to versatile settings we emphasize that it can be instantiated with various IsMatch predicates: any predicate that when evaluated on two elements a, b, for a from the space of data elements and b from the space of queries, returns a binary indicator

accepting value $\mathsf{IsMatch}(a,b)=1$ if a,b are considered a match in the context of the given settings (0 otherwise). The server's complexity depends on the complexity of the $\mathsf{IsMatch}$ predicate.

3.2 Desired Properties

We define properties desired from a secure search protocol (following [3]); see Tables 1-2:

1. ***Full security:*** *two (equal size) adversarially-chosen queries and data arrays are computationally-indistinguishable from the search and upload protocols; see formal statement in Definition 5.*

2. ***Efficient client:*** *client's running-time is polynomial in the time to encrypt the query q and decrypt the search outcome ciphertexts $(\llbracket i^* \rrbracket, \llbracket x[i^*] \rrbracket)$.*

3. ***Efficient server:*** *the server evaluates polynomials $f(\llbracket x \rrbracket, \llbracket q \rrbracket)$ of degree polynomial in $\log n$ and the degree of $\mathsf{IsMatch}$, and of size (i.e. the overall number of multiplication and addition operations) polynomial in n and the size of $\mathsf{IsMatch}$.*

4. ***Efficient communication:*** *the protocol has single-round protocol and communication bandwidth polynomial in $|q|$, $|i^*| = \log n$ and $|x[i^*]|$ (for $|z|$ denoting the binary representation length of z).*

5. ***Unrestricted search functionality:*** *no restrictions are placed on the number of array elements that match the query.*

6. ***Retrieval of both index and element:*** *client's output consists of both index and element $(i^*, x[i^*])$.*

7. ***Post-processing free:*** *the server sends the encrypted search outcome $(\llbracket i^* \rrbracket, \llbracket x[i^*] \rrbracket)$ for the client to decrypt, with no client's post-processing needed.*

8. ***Negligible error probability:*** *with overwhelming probability the client's output $(i^*, x[i^*])$ is the correct search outcome, i.e., $i^* = \min\{\, i \in [n] \mid \mathsf{IsMatch}(x_i, q) = 1 \}$.*

9. ***Compatibility with all FHE schemes:*** *the protocol can employ (as a black-box) any FHE scheme.*

3.3 Threat Model

The untrusted party in our scenario is the *honest-but-curious* (also called, semi-honest) and *computationally-bounded* adversary controlling the server (as in the case of hacked cloud servers). Semi-honest means, as standard, that the adversary follow the protocol, but may try to learn sensitive information. Namely, for the upload functionality the server provides the storage facility and is prohibited from modifying or destroying the encrypted array ($\llbracket x \rrbracket$). Likewise, for the search functionality the server receives encrypted queries ($\llbracket q \rrbracket$) and is obligated to follow the protocol and return encrypted search outcomes accordingly ($\llbracket i^* \rrbracket$, $\llbracket x[i^*] \rrbracket$). On the other hand, the adversary can try to derive sensitive information from the stored elements, received queries, data access patterns and search outcomes. Computationally-bounded means, as standard, that the adversary's actions are captured by a probabilistic polynomial time (PPT) algorithm.

We mention that there is no need to consider adversaries controlling the client as it can trivially simulate the entire protocols (Upload, Search) by herself. This is because the server's role is not to provide input or receive output, but rather to take the bulk of computational burden off the client.

Our security requirement is that if the client issues the protocol with one of two adversarially-chosen equal size queries $q^{(0)}$, $q^{(1)}$ (similarly, arrays $x^{(0)}$, $x^{(1)}$), the adversary controlling the server cannot distinguish between them; see the formal attack games below.

Definition 5 (Full security)**.** We say that an upload and search protocol provides *full security* if every PPT semi-honest adversary $\mathcal{A}$ controlling the server has no more than a negligible advantage $\mathbf{Adv}(\mathcal{A}, \mathcal{FHE}, \kappa) = \mathsf{negl}(\kappa)$ in winning the attack games on query or data (as specified below).

Attack on query (respectively, data). The attack games involve the adversary $\mathcal{A}$ and a challenger $\mathcal{C}$, both given the FHE scheme $\mathcal{FHE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ and the security parameter κ, proceeding as follows.

1. $\mathcal{C}$ executes the key generation step (Step I, Figure 3) to obtain $(\mathsf{pk}, \mathsf{sk}) \leftarrow_{\$} \mathsf{KGen}(1^{\kappa})$, and sends pk to $\mathcal{A}$.

2. $\mathcal{A}$ chooses parameters n, w, w' for array size, elements size, and query size, respectively; generates and sends to $\mathcal{C}$ the tuple $(x, q^{(0)}, q^{(1)})$ for x an array of n elements of size w, and $q^{(0)}, q^{(1)}$ queries of size w' (respectively, the tuple $(x^{(0)}, x^{(1)}, q)$ for arrays $x^{(0)}, x^{(1)}$ and query q of sizes as specified above).

3. $\mathcal{C}$ samples $b \leftarrow_{\$} \{0, 1\}$ uniformly at random.

4. $\mathcal{C}$ and $\mathcal{A}$ execute the upload and search protocols (Step II-III, Figure 3) playing the roles of client and server respectively. The client's input is x and $q^{(b)}$ (respectively, $x^{(b)}$ and q); the server has no input.

5. $\mathcal{A}$ sends $b' \in \{0, 1\}$ to $\mathcal{C}$, and wins if $b' = b$.

The advantage of $\mathcal{A}$ in the search attack on query (respectively, data) is defined to be

$$\mathbf{Adv}(\mathcal{A}, \mathcal{FHE}, \kappa) = |\Pr[b' = b] - 1/2|$$

We remark that, since $\mathcal{A}$ holds pk, he can simulate on its own the upload step for additional data entries x' of its choice. Likewise, since $\mathcal{A}$ holds $x, q^{(0)}, q^{(1)}$ (respectively, $x^{(0)}, x^{(1)}, q$), he can simulate on its own the search step –excluding the client's final decryption step– for whatever and as many queries q' as $\mathcal{A}$ wishes, including queries $q^{(0)}, q^{(1)}$ (respectively, q). These upload and search steps can occur both before and after the challenge.

Full security leakage profile discussion. Full security implies that the adversary participating in the protocol does not learn new information on *data, queries, and search outcomes*, other than the following size information: (1) plaintext space; (2) array size upper-bound; (3) element size upper-bound; (4) overall count of executed queries. This holds both for data-at-rest (upload) and data-in-use (search).

In particular, the protocol hides *access-patterns* to prevent, for example, identifying frequently searched data elements; and hides *search-patterns* to prevent inferring from search outcomes whether two searches use related query values. The overall count of executed queries does not reveal any information regarding the content or distribution of stored elements; and the server is unable to distinguish between fresh queries and fetch next queries (cf. Section 7.6).

4 Secure Search

We specify our secure search protocol (see Figures 3-4, Section 4.3.3) that we name: *Binary Raffle Protocol*.

This section is organized as follows. The simple keys generation and data upload steps are in Sections 4.1–4.2; the secure search step, which is the heart of this work, in Section 4.3; and our main theorem in Section 5.1.

4.1 Keys Generation Step

In the keys generation step the client executes the key generation algorithm of the leveled scheme $\mathcal{FHE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ (see Definition 2). The input to the KGen algorithm are the security parameter κ and the level $L = \log_2(d)$ for d the degree of the secure search polynomial (see Section 4.3 below).

In details, the level L depends on the following upper-bounds: (1) error probability ε; (2) array size n; (3) degree of the desired matching polynomial $d_{\mathtt{IsMatch}}$. Specifically it needs to be set to $L = \lceil \log \log(n/\varepsilon) + \log(d_{\mathtt{IsMatch}}) \rceil$.

The output is $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}(1^{\kappa}, 1^{L})$ where $(\mathsf{pk}, \mathsf{sk})$ are kept by the client, and pk is sent to the server.

4.2 Data Upload Step

In the array upload step the client encrypts its data array $x = \big(x[1], \ldots, x[n]\big)$ and sends the ciphertexts $[\![x]\!] = \big([\![x[1]]\!], \ldots, [\![x[n]]\!]\big)$ to the server. The encryption is performed element by element. Each element $x[i]$ is given by its binary representation and contains up to w bits. Prior to encryption, to avoid revealing the number of bits in each element the client pads all elements with leading zeros until all of them are of length w. The encryption of each element is then performed bit by bit.

4.3 Secure Search Step

We specify the secure search step, which is the heart of our protocol (Figure 3, Step III).

The starting point of the secure search protocol is after the client obtained sk, pk, the server obtained pk and the encrypted array $[\![x]\!] = \big([\![x[1]]\!], \ldots, [\![x[n]]\!]\big)$ has been uploaded to the server (see Sections 4.1-4.2).

The secure search step (Figure 3, Step III) proceeds as follows. First the client encrypts her search query $[\![q]\!] \leftarrow \mathsf{Enc}_{\mathsf{pk}}(q)$ and sends it to the server (Step III.(a)). Next, the server evaluates the steps specified below on the stored array $[\![x]\!]$ and the query $[\![q]\!]$ to obtain and send to the client $([\![b^*]\!], [\![x[i^*]]\!])$ for $b^* \in \{0,1\}^{\lceil \log_2 n \rceil + 1}$ the binary representation of the index of the first match $i^* = \min\{i \in [n] \mid \mathsf{IsMatch}(x[i], q) = 1\}$ (Step III.(b)). Finally, the client decrypts to obtain the desired output $(b^*, x[i^*])$ (Step III.(c)).

We elaborate below on the server's computations (Step III.(b)). We start by specifying how to return a selector array $s' \in \{0,1\}^n$ accepting value 1 on entry i^* and 0 otherwise (Section 4.3.1), then elaborate on our key algorithm for achieving the former (Section 4.3.2), and finally specify the additional actions for returning the ciphertext for index and element $(b^*, x[i^*])$ (Section 4.3.3).

4.3.1 Binary Raffle for Computing Selector Array s'

We specify how the server computes (the encryption of) a selector array $s' \in \{0,1\}^n$ accepting value 1 on entry i^* and 0 otherwise (Figure 3, Step III.(b).(1)–(3)).

First (Step III.(b).(1)), the server evaluates the specified pattern matching polynomial $\mathsf{IsMatch}$ on each entry of the stored array $[\![x]\!]$ and the lookup value $[\![q]\!]$. This results in an encrypted array $[\![ind]\!]$ that contains in every index $i \in [n]$ the encrypted boolean result $\mathsf{IsMatch}([\![x[i]]\!], [\![q]\!]) \in \{[\![0]\!], [\![1]\!]\}$:

$$[\![ind]\!] \leftarrow \big(\mathsf{IsMatch}([\![x[1]]\!], [\![q]\!]), \ldots, \mathsf{IsMatch}([\![x[n]]\!], [\![q]\!])\big)$$

Next (Step III.(b).(2)), the heart of the protocol is converting $[\![ind]\!]$ to a step function array of size n

$$[\![s]\!] = \big([\![0]\!], \ldots, [\![0]\!], [\![1]\!], \ldots, [\![1]\!]\big)$$

that contains $[\![0]\!]$ in every index before i^* and $[\![1]\!]$ from index i^* and further on. This $[\![s]\!]$ is computed using our randomized algorithm detailed in Section 4.3.2:

$$[\![s]\!] \leftarrow \mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}([\![ind]\!])$$

With probability $1 - \varepsilon$, the result $[\![s]\!]$ of our randomized algorithm will be the encryption of the step function described above.

Third (Step III.(b).(3)), we compute the pairwise difference of adjacent indices in $[\![s]\!]$ (i.e. its derivative):

$$\forall i \in [2, n]: \quad [\![s'[i]]\!] \leftarrow [\![s[i]]\!] - [\![s[i-1]]\!] \pmod 2 \quad \text{and}$$

$$[\![s'[1]]\!] \leftarrow [\![s[1]]\!], \quad [\![s'[n+1]]\!] \leftarrow [\![1]\!] - [\![s[n]]\!]$$

The resulting array $[\![s']\!]$ will contain $[\![0]\!]$ in every index except in the index of the first match i^* (or at index $n+1$ if no match exists) where it will be $[\![1]\!]$.

4.3.2 BinaryRaffleStepFunction$_{n,\varepsilon}$ Algorithm

We next describe the BinaryRaffleStepFunction$_{n,\varepsilon}$ algorithm, which is the heart of our secure search protocol (Figure 3, Step III.(b).(2)).

The BinaryRaffleStepFunction$_{n,\varepsilon}$ algorithm transforms any array $v \in \{0,1\}^n$ of binary values into an array $t = (0,\ldots,0,1,\ldots,1) \in \{0,1\}^n$ that contains the step function with value 1 starting from the first index i where $v[i] = 1$. This algorithm is a randomized Monte Carlo algorithm with failure probability ε (see Figure 4). In addition we provide an illustration for the main steps of the algorithm in Figure 2.

For clarity of presentation we present the algorithm as preforming computations on plaintext values. Modifying the algorithm to apply it on FHE encrypted data is straightforward: simply replace each addition/multiplication operation with its homomorphic counterpart.

Random Partial Prefix Sums To determine whether a k-prefix vector $\mathsf{prefix}_k(v) = \big(v[1],\ldots,v[k]\big) \in \{0,1\}^k$ of the binary indicator vector $v = \big(v[1],\ldots,v[n]\big) \in \{0,1\}^n$ is non-zero (i.e. not $0^k = (0,\ldots,0)$ of size k), we do the following. We compute the parity of a random subset for the entries of $\mathsf{prefix}_k(v)$, that is $\mathsf{parity}(r) = \sum_{i=1}^{k} r[i] \cdot v[i]$ for uniformly random $r \in \{0,1\}^n$. The parity bit $\mathsf{parity}(r)$ is always zero when $v = 0^k$ and it is one with probability half when $v \neq 0^k$. By repeating for $N(\varepsilon)$ i.i.d. random variables $r_1,\ldots,r_{N(\varepsilon)} \in \{0,1\}^n$ (for sufficiently large $N(\varepsilon)$) and computing the OR of the resulting bits $\mathsf{parity}\big(r_1\big),\ldots,\mathsf{parity}\big(r_{N(\varepsilon)}\big)$, i.e. computing: $t[k] = \mathrm{OR}\Big(\mathsf{parity}\big(r_1\big),\ldots,\mathsf{parity}\big(r_{N(\varepsilon)}\big)\Big)$ we obtain the desired step-function $t = \big(t[1],\ldots,t[n]\big) \in \{0,1\}^n$ with overwhelming probability.

4.3.3 Returning Index and Element

Finally we specify the additional server's steps for returning the ciphertext for index and element $(b^*, x[i^*])$ that are sent then to the client (Figure 3, Step III.(b).(4)-(5)).

Computing $[\![b^*]\!]$. The server computes $[\![b^*]\!] = B \cdot [\![s']\!]$ for $B \in \{0,1\}^{(\lceil \log_2 n \rceil + 1) \times n}$ the matrix that contains in each column $k \in [n]$ the binary representation of k. The resulting array $[\![b^*]\!]$ will hold the binary representation of the index i^* of the single $[\![1]\!]$ in $[\![s']\!]$ (or 0 if the array contains only $[\![0]\!]$'s). This is because multiplying the matrix B by any array of size $n+1$ that contains a single 1 bit in some index $j \in [n]$ results in a array of size $\lceil \log_2 n \rceil + 1$ that holds the binary representation of j (and a array of zeros if $j = n+1$).

Computing $[\![x[i^*]]\!]$. Since the problem of privately retrieving a uniquely identifiable element from an encrypted array has efficient FHE based solutions, we first focused above (Sections 3-4.3.2) on the task of computing and returning the encrypted index alone. We next explain how to retrieve also the corresponding element.

The most straightforward way to retrieve the element $x[i^*]$ in addition to i^* is to utilize a *Private Information Retrieval* (PIR) protocol (see Section 1.1) on the encrypted array $[\![x]\!]$ and index $[\![i^*]\!]$. This would require no further interaction (as the server already had $[\![i^*]\!]$); However it would increase the degree of our secure search protocol by a factor of $d_{\mathtt{PIR}} = \log n$.

Instead we suggest a more efficient alternative for retrieving the matched element $[\![x[i^*]]\!]$. This is by re-using the array $[\![s']\!]$ that already contains $[\![0]\!]$'s in all indices except in the index of the first match i^*, where it contains $[\![1]\!]$. The additional step would be to calculate for each index $j \in [n]$ and each bit $k \in [w]$: $[\![x[i^*][k]]\!] = \sum_{j=1}^{n} ([\![x[j][k]]\!] \cdot [\![s'[j]]\!])$. This method would increase the degree of our secure search protocol only by 1 since $[\![x]\!] = \big([\![x[1]]\!],\ldots,[\![x[n]]\!]\big)$ are freshly encrypted ciphertexts.

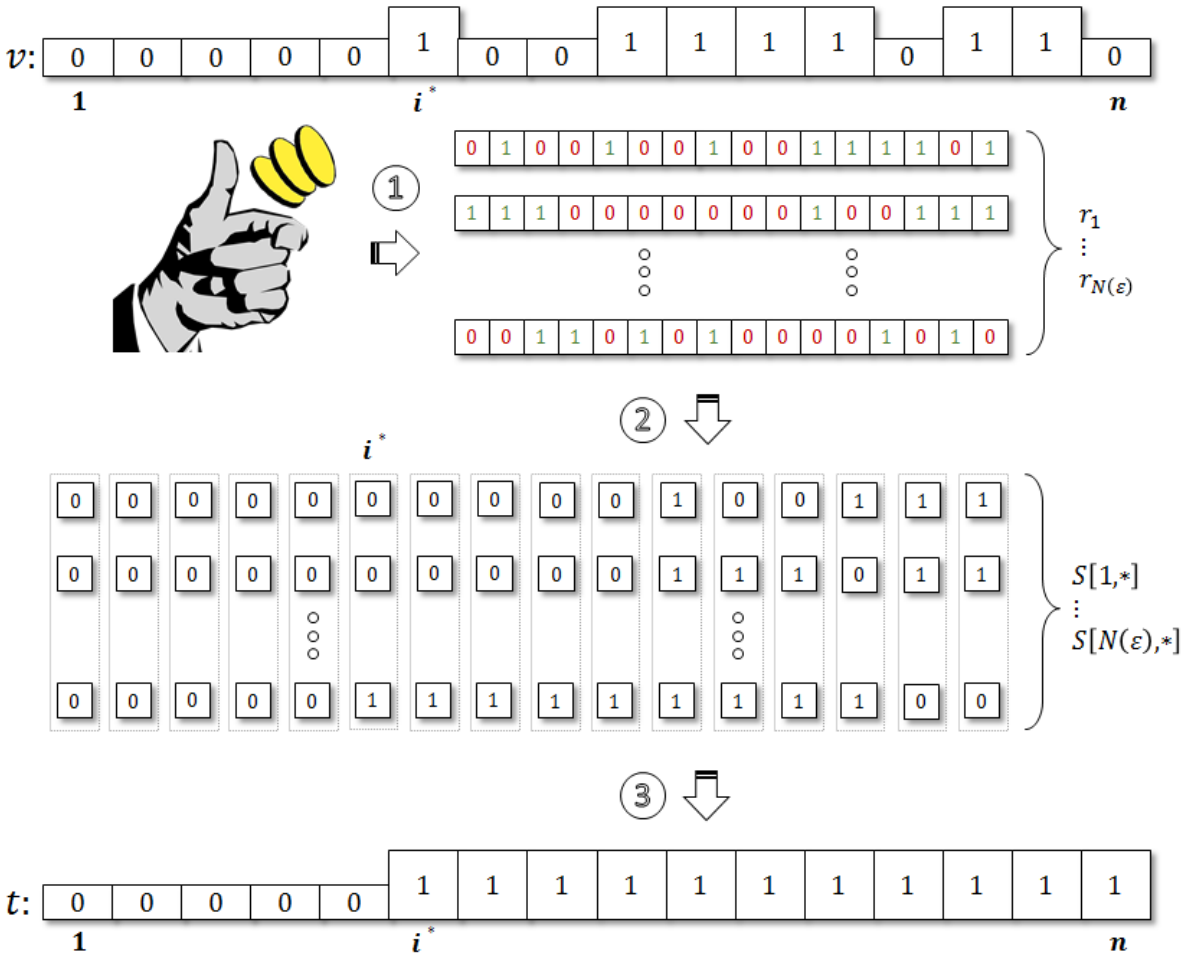

Figure 2: An illustration of the main steps in $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$ algorithm: On input array v the following steps are preformed: (1) Sample $N(\varepsilon)$ random binary arrays $r_1, \ldots, r_{N(\varepsilon)}$ of length n; (2) Matrix S will hold $N(\varepsilon) \cdot n$ random partial prefix sums as follows $\forall j \in [N(\varepsilon)], \forall k \in [n] : \; S[j,k] = \langle \mathsf{prefix}_k(v), \mathsf{prefix}_k(r_j) \rangle$; (3) Each index $k \in [n]$ in t will hold the logical OR between the elements in column k of S. That is, calculate $t[k] = 1 - \left(\prod_{j=1}^{N(\varepsilon)} \left(1 - S[j,k] \right) \right) \pmod 2$. With probability $1 - \varepsilon$, the output is array t that contains a step function with 0s before index i^* and 1s from index i^* and onwards.

Parameters (Shared Input):
- Description of matching condition polynomial $\mathsf{IsMatch}\big(x, y\big) \in \{0,1\}$ with upper bound on the degree d.
- Scheme $\mathcal{FHE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$.
- Security parameter κ.
- Error probability ε (upper-bound).
- Array size n (upper-bound).
- Element bit length w (upper-bound).

Inputs:
The client's inputs are:
- Plaintext array $x = \big(x[1], \ldots, x[n]\big)$, each element $x[i]$ given in binary representation of length w.
- The query/lookup value q.

The server has no input.

Outputs:
With probability $1 - \varepsilon$, the client's output is $(b^*, x[i^*])$ for $i^* = \min\{\, i \in [n] \mid \mathsf{IsMatch}(x[i], q) = 1\}$ and $b^* \in \{0,1\}^{\lceil \log_2 n \rceil + 1}$ the binary representation of the index i^*.
The server has no output.

I **Keys generation.** The client performs the following:
(1) Select the level $L = \lceil \log_2 \log_2(n/\varepsilon) + \log_2(d_{\mathtt{IsMatch}}) \rceil$.
(2) Execute $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}(1^\kappa, 1^L)$.
(3) Keep pk, sk with the client.
(4) Send pk to the server.

II **Data upload.** For each element $x[i]$ the client performs the following:
(1) Pad element $x[i]$ with leading zeros until its size is w bits.
(2) Encrypt the padded element bit by bit to receive $[\![x[i]]\!] = \mathsf{Enc}_{\mathsf{pk}}(x[i])$.
(3) The client sends $[\![x]\!] = \big([\![x[1]]\!], \ldots, [\![x[n]]\!]\big)$ to the server.

III **Secure Search.** The following steps are executed whenever the client issues a search query.

(a) **Client's search query.**
The client encrypts q bit by bit $[\![q]\!] \leftarrow \mathsf{Enc}_{\mathsf{pk}}(q)$ and sends it to the server.

(b) **Server's computation.**
Using the public key pk the server computes the following:
(1) For every $i \in [n]$ evaluate the $\mathsf{IsMatch}$ polynomial on $[\![x[i]]\!]$ and the query value $[\![q]\!]$ to get the indicators array of size n: $[\![ind]\!] \leftarrow \big(\mathsf{IsMatch}([\![x[1]]\!], [\![q]\!]), \ldots, \mathsf{IsMatch}([\![x[n]]\!], [\![q]\!])\big)$
(2) Execute the following sub-routine (see Section 4.3.2) on the $[\![ind]\!]$ array: $[\![s]\!] \leftarrow \mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}\big([\![ind]\!]\big)$
(3) Compute a pairwise difference of adjacent indices in $[\![s]\!]$ (the derivative of $[\![s]\!]$):
$[\![s'[1]]\!] \leftarrow [\![s[1]]\!] \pmod 2$ and $\forall i \in [2, n]: \; [\![s'[i]]\!] \leftarrow [\![s[i]]\!] - [\![s[i-1]]\!] \pmod 2$
(4) Compute $[\![b^*]\!]$, the encrypted binary representation of the location of the single $[\![1]\!]$ in $[\![s']\!]$:
Let $B \in \{0,1\}^{(\lceil \log_2 n \rceil + 1) \times n}$ be the matrix that contains in each column $k \in [n]$ the binary representation of k. Calculate $[\![b^*]\!] = B \cdot [\![s']\!]$.
(5) Compute for each index $j \in [n]$ and each bit $k \in [w]$: $[\![x[i^*][k]]\!] = \sum_{j=1}^{n} ([\![x[j][k]]\!] \cdot [\![s'[j]]\!])$ (see Section 4.3.3)
(6) Send $([\![b^*]\!], [\![x[i^*]]\!])$ to the client.

(c) **Client's decryption.**
The client decrypts $b^* \leftarrow \mathsf{Dec}_{\mathsf{sk}}([\![b^*]\!])$, $x[i^*] \leftarrow \mathsf{Dec}_{\mathsf{sk}}([\![x[i^*]]\!])$ and outputs $(b^*, x[i^*])$.

Figure 3: *Binary Raffle* Secure Search Protocol

Parameters:
• Integer $n \in \mathbb{N}$.
• Failure probability ε.

Input:
• Array $v = \big(v[1], \ldots, v[n]\big) \in \{0,1\}^n$.

Output:
The array is array $t \in \{0,1\}^n$ as follows. If $v \neq (0,\ldots,0)$, then with probability $1-\varepsilon$, $t[0] = \ldots = t[i^*-1] = 0$ and $t[i^*] = \ldots = t[n] = 1$ for $i^* = \min\{i \in [n] \mid v[i] = 1\}$ (step function). Else (if $v = (0,\ldots,0)$), $t = (0,\ldots,0)$ (with probability 1).

Algorithm:
1. Set $N(\varepsilon) = \lceil \log_2(n/\varepsilon) \rceil$ and sample $N(\varepsilon)$ uniformly random arrays $r_1, \ldots, r_{N(\varepsilon)} \leftarrow\!\!\$\ \{0,1\}^n$
2. Recall that for $v \in \{0,1\}^n$ and $k < n$ we denote $\mathsf{prefix}_k(v) = (v_1, \ldots, v_k)$.
   Compute $N(\varepsilon) \cdot n$ random partial prefix sums: $\forall j \in [N(\varepsilon)], \forall k \in [n] :\ S[j,k] = \langle \mathsf{prefix}_k(v), \mathsf{prefix}_k(r_j) \rangle$
3. Compute the binary step function array $t \in \{0,1\}^n$ where each $k \in [n]$, $t[k]$ is the OR of the values in the k-th column of S: $\forall k \in [n] :\ t[k] = 1 - \left(\prod_{j=1}^{N(\varepsilon)} \Big(1 - S[j,k]\Big)\right) \pmod 2$
4. Return array t

Figure 4: $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$ Algorithm

5 Theoretical Results

In this section we present our theoretical contribution of our secure search protocol (Figure 3). Our main theorem appears in Section 5.1; Discussion of key aspects by which we improves over the prior state-of-the-art in Section 5.2.

5.1 Main Theorem

The advantages and properties of our secure search protocol are given in the following theorem.

Theorem 5.1. *There exists a secure search protocol attaining desired properties 1-9 as specified in Section 3.2.*

In particular, the protocol of Figure 3, when executed on shared parameters $(\mathsf{IsMatch}, \mathcal{FHE}, \kappa, \varepsilon, n, w)$ *and client's input data array* $x = (x[1], \ldots, x[n])$ *for* $x[i] \in \{0,1\}^w$ *and query* q*, satisfies the following:*

1. Correctness: *With probability* $1-\varepsilon$*, the client's output is* $(b^*, x[i^*])$ *for*

$$i^* = \min\{\, i \in [n] \mid \mathsf{IsMatch}(x[i], q) = 1\}$$

   and $b^* \in \{0,1\}^{\lceil \log_2 n \rceil + 1}$ *the binary representation of* i^**. The server has no output.*

2. Complexity of the search step (Step III, Figure 3): *The client's running-time is the time compute* $|q|$ *encryptions and* $|b^*| + |x[i^*]|$ *decryptions. The server evaluates a polynomial of degree* $\log(n/\varepsilon) \cdot d$ *and overall multiplications* $n \cdot (\log(n/\varepsilon) + \mu + |x[i^*]|)$ *for* d *and* μ *the degree and overall multiplications of* $\mathsf{IsMatch}$*. The communication is 1-round, consistsing of* $|q|$ *ciphertexts sent from client and* $|b^*| + |x[i^*]|$ *ciphertexts from server. See Table 3.*

3. Security: *The protocol attains full security (see Definition 5), assuming semantic security of* $\mathcal{FHE}$*.*

Proof. Correctness follows from the correctness of $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$ algorithm, which holds with probability $1-\varepsilon$. Complexity analysis follows by inspection. Security easily follows from the semantic security of the underlying FHE scheme. See details in Appendix A.

We next show how properties 1-9 easily follow. Property 1 (full security) follows from the security statement; Properties 2-3 (efficiency of client and server) follow from the complexity statement; Properties 6 and 8 (retrieval of index and element, and with negligible error probability) follow from the correctness statement. Properties 5, 7 and 9 follow from protocol inspection: Property 5 (efficiency communication)

– the communicated consists only of the encrypted query sent by the client, and the encrypted result index and element sent by the server; Property 7 (post-processing free) – the client, upon receiving the encrypted result from the server, only decrypts and outputs the obtained cleartext result without any post-processing (see Step (c), Figure 3); Property 9 (compatibility with all FHE schemes) – the protocol makes no requirements of the FHE scheme beyond using the standard FHE algorithms. □

5.2 Discussion: Comparison to SPiRiT

We highlight key aspect in which our protocol improves over the prior-state-of-the-art secure search solutions on FHE encrypted data with full security, unrestricted search functionality and both index and element retrieval. That is, we compare to Akavia, Feldman, Shaul's (AFS) [3] *SPiRiT Rand.* and *SPiRiT Det.* protocols.

5.2.1 Post-processing Free

In our protocol the server produces and sends to the client a ciphertext for the correct search outcome, with overwhelming success probability. This is in contrast to returning a list of candidates for the client to post-process, or suffering from a noticeable error probability, in AFS. Our post-processing freeness enables the server to use secure search as a sub-component in a larger computation without interaction with the client.

5.2.2 Faster Secure Search

Our protocol is faster than AFS in attaining:

1. Optimal client run-time whose role in search step (Step III, Figure 3) is solely to encrypt the query and decrypt the search outcome. Moreover, our protocol achieves a negligible error probability. In contrast, AFS either require client's post-processing in time $\log^2 n/\log\log n$ (SPiRiT Det.), or exhibit a noticeable error probability (SPiRiT Rand.).

2. Faster server's complexity via reducing the degree of the evaluated polynomial from cubic to linear in $\log n$, and reducing the overall number of multiplications by up to $\log n$ factor.

3. Further speedup due to working over GF(2), in contrast to GF(p) for $p \gg 2$ in AFS; see Section 5.2.3.

See Table 3 and below for detailed comparisons.

Comparing post-processing free solutions (rows (ii) vs. (iii), Table 3): Our solution exhibits a faster server with degree reduced linear in $\log n$ (in contrast to cubic in SPiRiT Rand.) and logarithmic in $1/\varepsilon$ (in contrast to linear in SPiRiT Rand.), and reducing the overall number of multiplications by a factor of $\log\log n$.

Comparing solutions with post-processing in poly$(\log(n))$ time (rows (i) vs. (v), Table 3): Our solution exhibits a faster server with degree linear in $\log n$ (in contrast to cubic in SPiRiT Det.), and overall number of multiplications reduced by a factor of $(k/\alpha)\cdot\log(n) = \mathcal{O}\left(\frac{\log^3 n}{\log\log(n)\cdot\log(1/\varepsilon)}\right)$. We state however that whereas our solution is correct with overwhelming probability, *SPiRiT Det.* is correct with probability 1.

Further optimizations (row (iv) vs. (i)-(iii) in Table 3): For the case that IsMatch is the equality operator, we show how to eliminate the dependence of server's complexity on the the complexity d, μ of the IsMatch by using universal hashing (see row (iv) and Section 7.1). Our server in this case evaluate a polynomial of degree $\mathcal{O}\left(\log^2(n/\varepsilon)\right)$ and $\mathcal{O}(n\cdot\log(n/\varepsilon))$ overall multiplications. This is in contrast to linear dependence on the degree and overall multiplications of IsMatch (d and μ, respectively) in the server's degree and overall multiplications (rows (i)-(iii)). This optimization is especially appealing for scenarios with long binary stored elements and search queries (large files, DNA sequences, etc.), where (i)-(iii) have linear dependence on the elements length w, whereas employing hashing (row (iv)) we completely avoid the dependence on w.

5.2.3 Wider FHE Compatibility and Further Speedup

Our protocol requires computing only over GF(2). This leads to the following two advantages:

First, our solution is compatible with all currently known candidate FHE schemes. Specifically, our solution can use as a black box any FHE scheme that enables homomorphic additions and multiplications over encrypted plaintext values in GF(2) (i.e. $\oplus, \wedge$ over plaintext bits). This is opposed to requiring homomorphic additions and multiplications over rings GF(p) for $p > 2$ as in *SPiRiT*. Thus, for example, our solution is compatible with the GSW [25] FHE scheme, whereas *SPiRiT* is not.

Second, we achieve further run-time speedup. This is because (to the best of our knowledge) in all current FHE schemes implementations, including HElib [32] that implements BGV [8] scheme, homomorphic computations over GF(2) are considerably faster. The reason for that in BGV and HElib is the additional costs induced by ciphertext refresh after each multiplication in plaintext space GF(p) for $p > 2$ (see [35]).

6 Experimental Results

In this section we describe in detail the benchmarks preformed to evaluate our *Binary Raffle* protocol and discuss our results. As a reference point, we executed benchmarks on an implementation of the *SPiRiT* protocol [3], the state-of-the-art secure search solution directly related to our FHE based protocol. We evaluated both the deterministic and randomized variants of *SPiRiT*.

We first present benchmarks with the same matching criteria predicate IsMatch as in [3]: IsEqual (see Equation 1, Section 2). Additionally, is Section 6.3.5, we give selected benchmarks results for several other matching criteria (see more details in Section 7).

6.1 Experimental Setup

We executed the protocols on top of the HElib C++ library [32] that was compiled with NTL [58] running over GMP [29]. We utilized a single Ubuntu Server 16.04.4 LTS Linux machine with Intel Xeon E7-4870 CPU running at 2.40GHz on 16 cores, 30MB Cache and 16GB RAM.

Parallelization and SIMD In all experiments we utilized all available CPU cores by dividing the input array into equally sized segments that were processed by each core. After completing its execution, every core returned the first matched index candidate for its array segment.

We also took advantage of HElib's SIMD [60] capabilities and "packed" multiple plaintext values (at least 500) into each ciphertext. We remark that we did not attempt to optimize the SIMD factor besides setting its minimal required value. By slight modification of HElib's level parameter it is often possible to reach much higher SIMD factors, around several thousands.

To summarize, with each CPU core executed the protocol on an input array of n ciphertexts, the total amount of elements processed in each experiment is given by $n' = n \cdot \text{SIMD} \cdot \text{CORES}$ and the client obtained in the end of the experiment $\text{SIMD} \cdot \text{CORES}$ results.

6.2 Experiments Description

Binary Raffle. Our main focus in *Binary Raffle*'s benchmarks was to evaluate the running-time of the protocol as a parameter of total input array size (n'), word width (i.e. bit length w) and error probability (ε).

For word widths of $w > 1$ we use the equality operator (see Equation 1, Section 2) as the selected IsMatch predicate. Beyond that, we also experiment on elements with single bit ($w = 1$). These experiments on $w = 1$ were meant to filter out the running-time of evaluating the IsMatch polynomial on all elements (step (2.a) in Figure 3) from the remaining steps of protocol. This can be thought as using an IsMatch predicate that is the degenerate identity function that does nothing, in order to evaluate the performance of the rest of the protocol on a binary vector of indicators.

For $w \in \{16, 64\}$ input array sizes ranged up to $n' \approx 3 \cdot 10^6$ elements. For $w = 1$ input array sizes ranged up to $n' \approx 20 \cdot 10^6$ elements.

The failure probabilities we experimented with were $\varepsilon \in \{2^{-80}, 2^{-40}, 2^{-20}, 2^{-10}, 2^{-1}\}$. Regarding above failure probabilities, $\varepsilon \in \{2^{-80}, 2^{-40}\}$ can be viewed as a negligible error probability, and any $\varepsilon > 2^{-1}$ as an error probability that allows standard probability amplification (see Section 8.1).

SPiRiT. As mentioned, we used an implementation of *SPiRiT* as a reference point. On array of sizes $n' = n \cdot \text{SIMD} \cdot \text{CORES}$, the deterministic variant was evaluated using $k = \lceil \log^2 n / \log\log n \rceil$ sequential executions for different primes larger then $\log n$. We would like to mention that we did not parallelize these k executions as we already exhausted all available employed parallelism to partition the input array into segments assigned to each CPU core.

Similarly to *Binary Raffle*, *SPiRiT* was executed on elements with $w = 1$. Additionaly, due to relatively low amounts of RAM (16GB) in our test machine we were unable to execute *SPiRiT* over elements with $w = 64$ for sufficiently large array sizes (n') and had to settle for $w = 16$ only. Given this amount of RAM the maximum array sizes that we were able to process ranged between $n' \approx 0.5 \cdot 10^6$ for $w = 16$ and $n' \approx 2 \cdot 10^6$ for $w = 1$.

The running-time of the randomized variant of *SPiRiT* with error probability $\varepsilon = 2^{-1}$ was obtained by taking the mean and standard deviation over the running-time of $2k = 2 \cdot \lceil \log^2 n / \log\log n \rceil$ executions for different primes larger then $\log n$, as required by the protocol. In some cases, due to RAM restrictions, executions for less than $2k$ (although at least k) primes were performed leading to an outcome of error probability higher than 2^{-1}.

6.3 Experimental Results

Our experimental results are presented below, showing the server's running time for different executions of both *Binary Raffle* and *SPiRiT* protocols. The client's running time for encrypting the query and decrypting the result can be ignored as it is negligible in comparison to the server's operations.

6.3.1 *Binary Raffle* with Negligible Error Probability (vs. *SPiRiT* Deterministic)

First we compare the performance *Binary Raffle* with $\varepsilon = 2^{-80}$ to the deterministic variant of *SPiRiT* for both $w \in \{1, 16\}$ (Figures 6a and 6b).

It can be immediately observed from both graphs that *Binary Raffle* achieves faster execution time in an order of magnitude compared to *SPiRiT*. Also we can observe that for *SPiRiT* with $w = 1$ and $w = 16$ there is an approximate $\times 10$ increase in run time between the first and the second. In comparison, for *Binary Raffle* with $w = 1$ and $w = 16$ the increase in run time is relatively minor.

Notice that the *SPiRiT* curves in both graphs are terminated for smaller array sizes than the ones of *Binary Raffle*. The reason for this is that *SPiRiT* executions for larger array sizes were unable to complete successfully. This occurs due to the increase in required levels for larger primes in HElib and the penalty on RAM that is associated with it.

6.3.2 *Binary Raffle* with Error Probability Half (vs. *SPiRiT* Randomized)

Now we compare the performance *Binary Raffle* with $\varepsilon = 2^{-1}$ to the randomized variant of *SPiRiT* (also with error probability half) for both $w \in \{1, 16\}$ (Figures 6c and 6d).

Again, it can be seen in both graphs that *Binary Raffle* achieves faster execution time in an order of magnitude compared to *SPiRiT*. And again, the *SPiRiT* curves in both graphs are terminated for smaller array sizes than the ones of *Binary Raffle*. This happens, similarly to the described in previous section, because working with large primes in HElib increases RAM consumption.

6.3.3 Impact of Error Probability (ε) on *Binary Raffle*

We executed *Binary Raffle* with different error probabilities $\varepsilon \in \{2^{-80}, 2^{-40}, 2^{-20}, 2^{-10}, 2^{-1}\}$ and observed the effect on run time performance (Figures 6e and 6f).

One can see in the graphs that although we jump from half error probability to a negligible error probability ($\varepsilon = 2^{-80}$) the difference in execution time is around $\times 20 - \times 50$ for words with a single bit and around $\times 2 - \times 3$ for words with 64 bits.

This can be explained by the logarithmic dependence between $1/\varepsilon$ and both degree and overall multiplications of the polynomial executed by the server during the *Binary Raffle* protocol.

6.3.4 Impact of Word Size (w) on *Binary Raffle*

We executed *Binary Raffle* with both $w \in \{1, 64\}$ and observed the effect on run time performance (Figures 6g and 6h).

As opposed to the previous section, in this section the change in execution time is more noticeable as word size w increases. When going from a single bit to 64 bit words the difference in execution is around $\times 20 - \times 40$ for error probability half and around $\times 2$ for negligible error probability ($\varepsilon = 2^{-80}$).

This can be explained by the linear dependence between word size w and both degree and overall multiplications of the polynomial executed by the server during the *Binary Raffle* protocol.

This observation brings into being our improvement to the *Binary Raffle* protocol specified in section 7.1.

6.3.5 Other Matching Criteria

The benchmarks below were performed with the same setup described in Section 6.1 with the exception of using a server with different CPU (less cores yet each core more powerful): Intel Core i7-4790 CPU running at 3.60GHz on 8 cores, 8MB Cache and 16GB RAM.

Experiment 1 was the conjunction between two equality queries over an array with two $w = 64$ bit sub-fields (see Section 7.2). Experiment 2 was a range query, with lower and upper values restriction, over an array with $w = 64$ bit unsigned integers (see Section 7.4). Experiment 3 was of equality queries over an array with $w = 64$ bit with 32 publicly known wildcard positions (see Section 7.3). The results are presented in Table 4.

In the results one can clearly see the relative overhead of the additional multiplications performed during the execution of $\mathsf{IsGreater}_{64}$ (Equation 2, Section 2) in the second experiment compared to $\mathsf{IsEqual}_{64}$ (Equation 1, Section 2) in the first experiment. Similarly, the third experiment, in which the matching criteria includes execution of solely $\mathsf{IsEqual}_{32}$, is even faster than the two previous experiments.

Table 4: Binary Raffle benchmarks for various matching criteria.

	(i) Conjunction		(ii) Range		(iii) Wild-cards	
Error Prob.	Array Size	Time (min.)	Array Size	Time (min.)	Array Size	Time (min.)
2^{-1}	38K	5	44K	23	38K	1
	77K	10	87K	50	77K	3
	154K	19	175K	98	154K	5
	307K	37	349K	195	307K	10
2^{-40}	44K	7	47K	50	44K	3
	87K	15	92K	99	87K	6
	175K	30	184K	199	175K	11
	349K	60	369K	396	349K	23
2^{-80}	44K	8	46K	55	44K	4
	87K	17	92K	111	87K	5
	175K	36	184K	223	175K	16
	349K	68	369K	443	349K	32

The matching criteria include:
(i) Conjunction of two $w = 64$ bits equality-test;
(ii) Range queries over $w = 64$ bit unsigned integers;
(iii) Wild-card queries over $w = 64$ bits with 32 wild-card positions.
Array sizes are in thousands of elements (denoted, K).

(a) *Binary Raffle* With Negligible Error Probability ($\varepsilon = 2^{-80}$) Versus *SPiRiT* Deterministic for Word Width $w = 1$

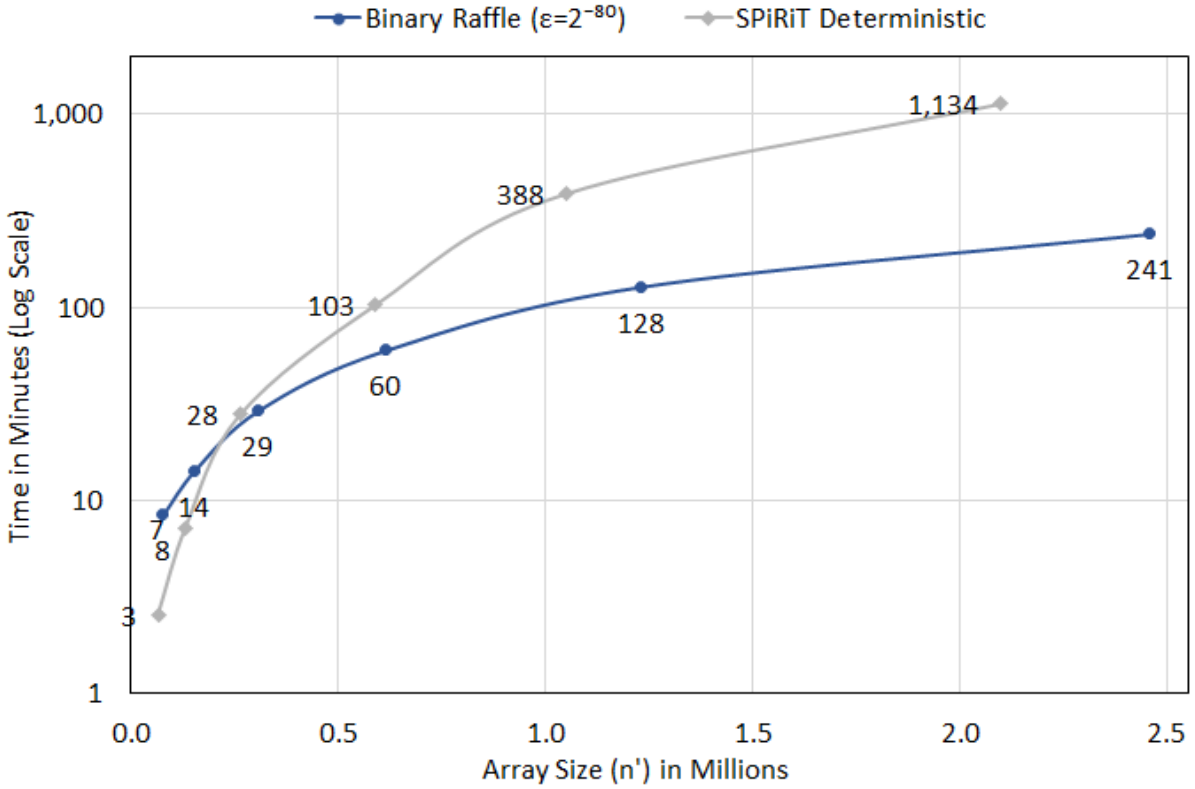

(b) *Binary Raffle* With Negligible Error Probability ($\varepsilon = 2^{-80}$) Versus *SPiRiT* Deterministic for Word Width $w = 16$

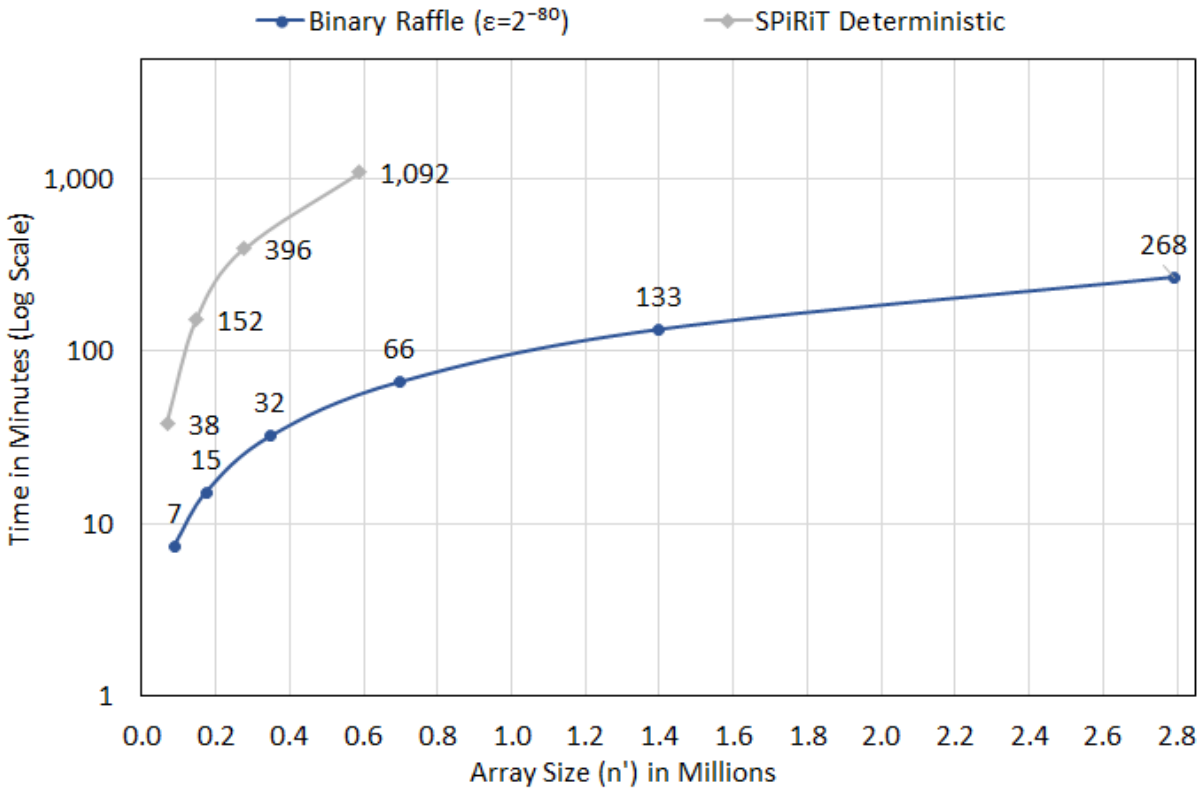

Figure 5: Server's execution time for different experiments. Y axis – minutes in logarithmic scale; X axis – array size in millions.

(c) *Binary Raffle* Versus *SPiRiT* Randomized Both With Error Probability $\varepsilon = 2^{-1}$ for Word Width $w = 1$

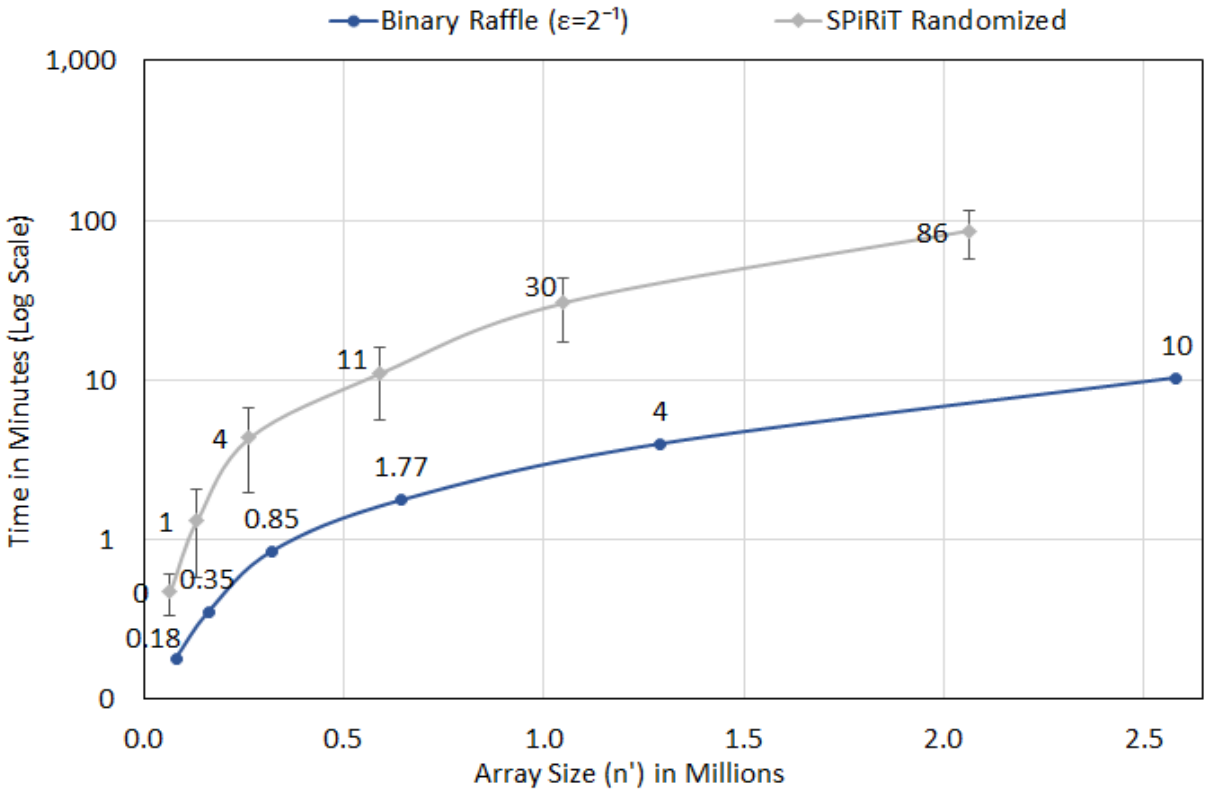

(d) *Binary Raffle* Versus *SPiRiT* Randomized Both With Error Probability $\varepsilon = 2^{-1}$ for Word Width $w = 16$

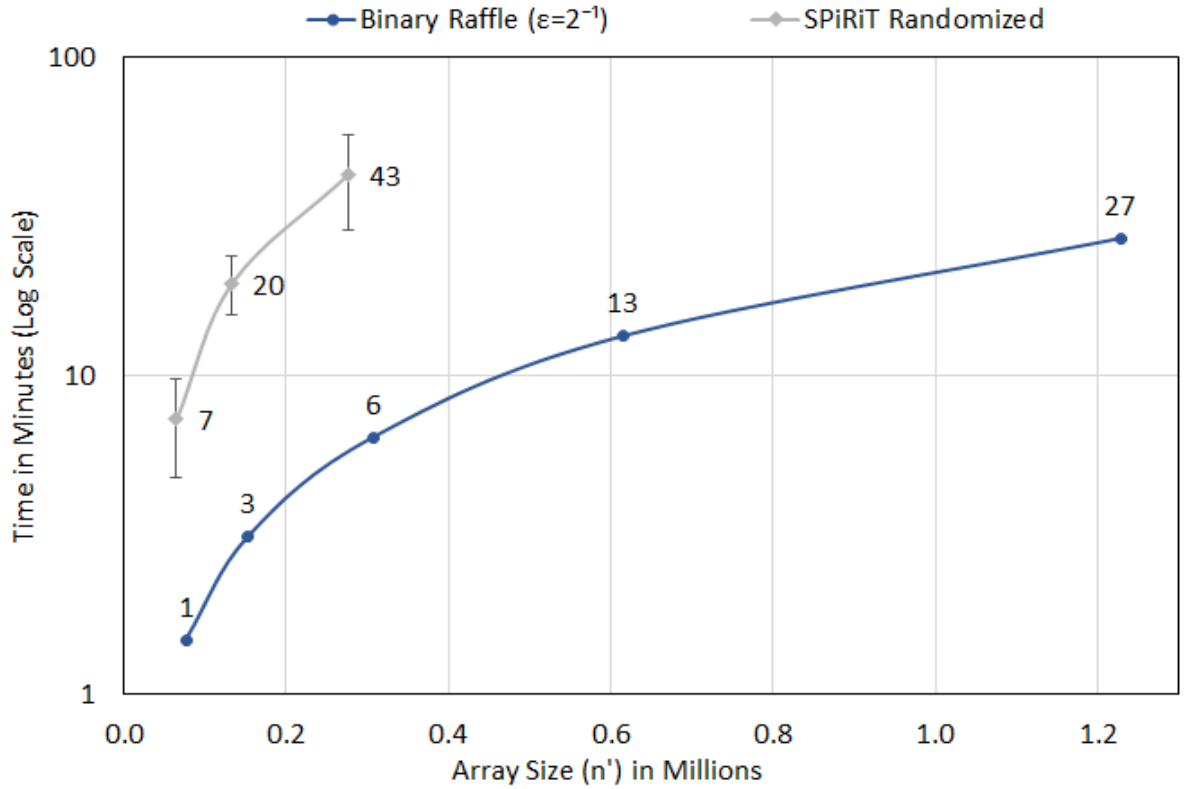

(e) *Binary Raffle* Comparing Different Failure Probabilities (ε) for Word Width $w = 1$

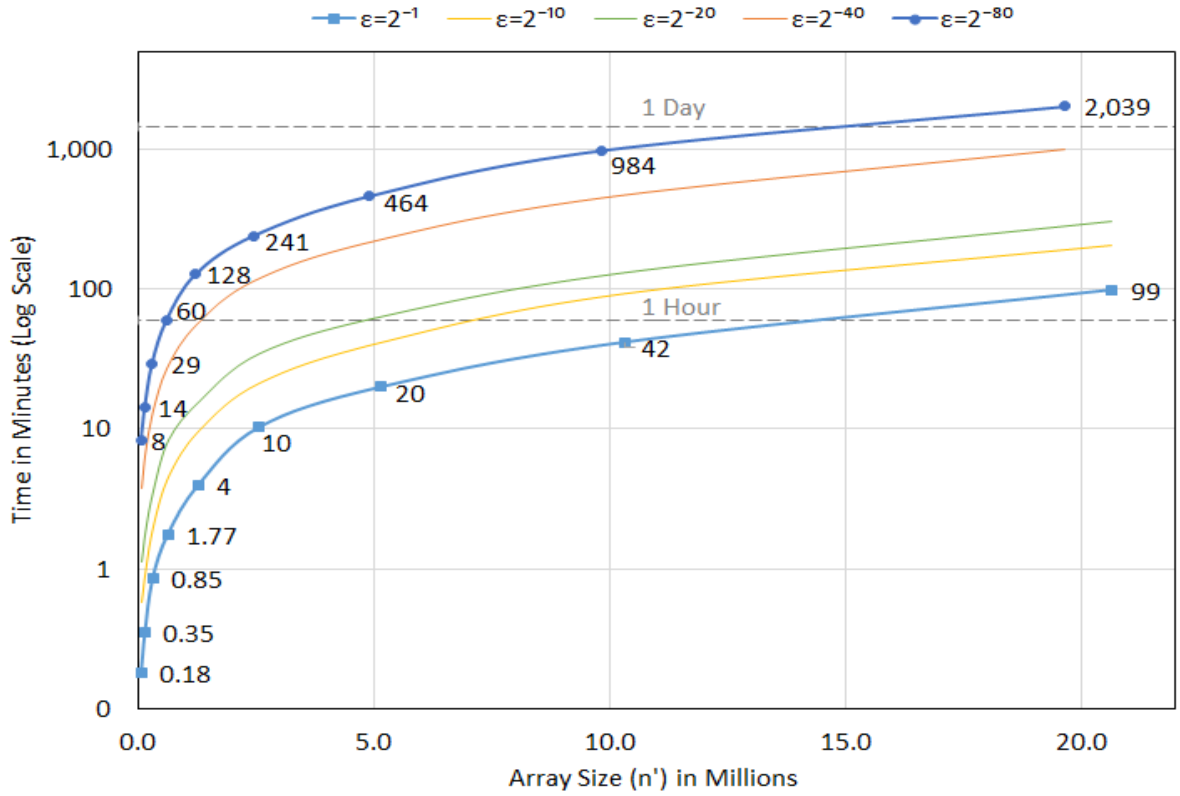

(f) *Binary Raffle* Comparing Different Failure Probabilities (ε) for Word Width $w = 64$

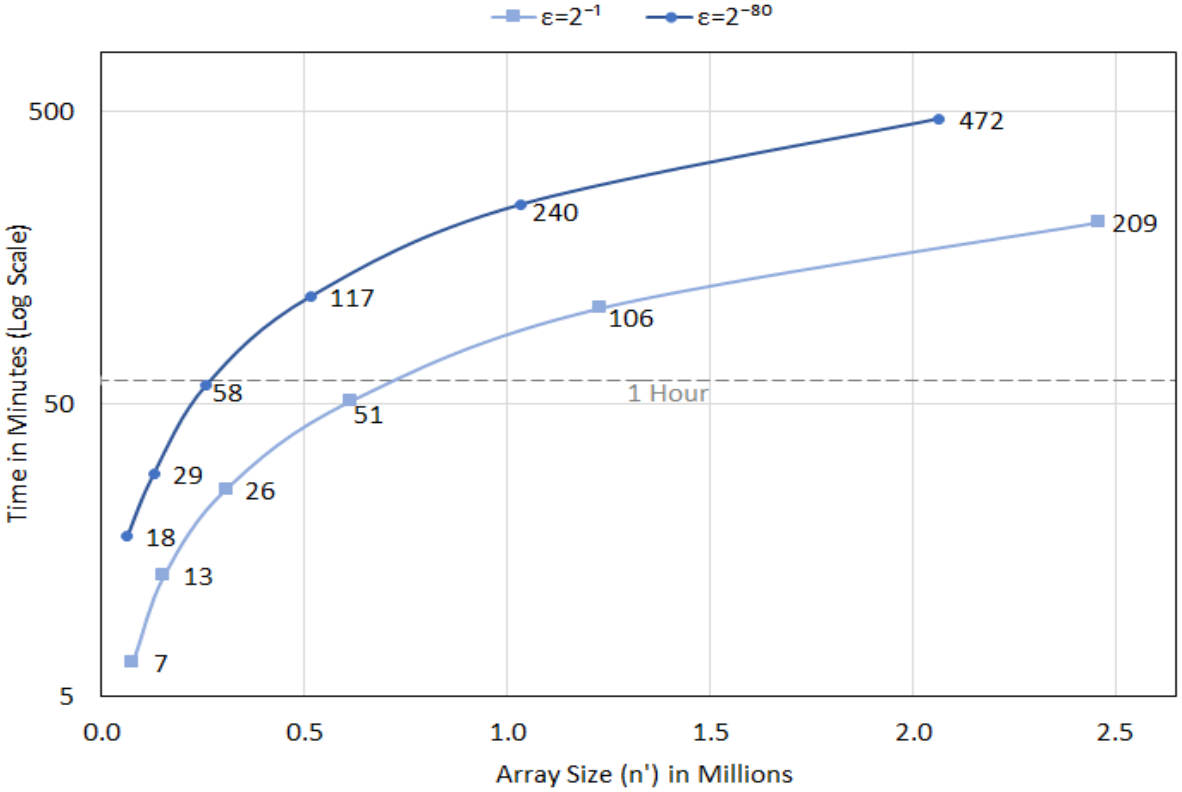

(g) *Binary Raffle* Comparing Different Word Sizes (w) for Error Probability $\varepsilon = 2^{-1}$

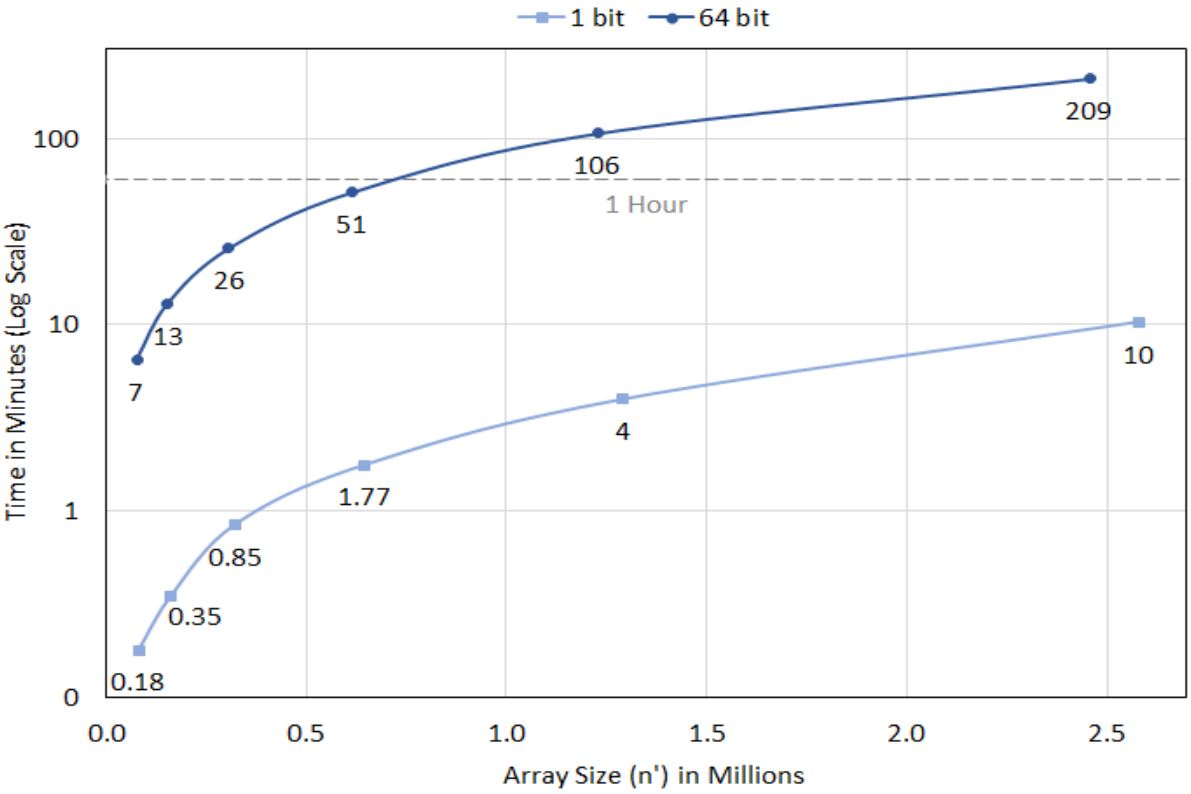

(h) *Binary Raffle* Comparing Different Word Sizes (w) for Error Probability $\varepsilon = 2^{-80}$

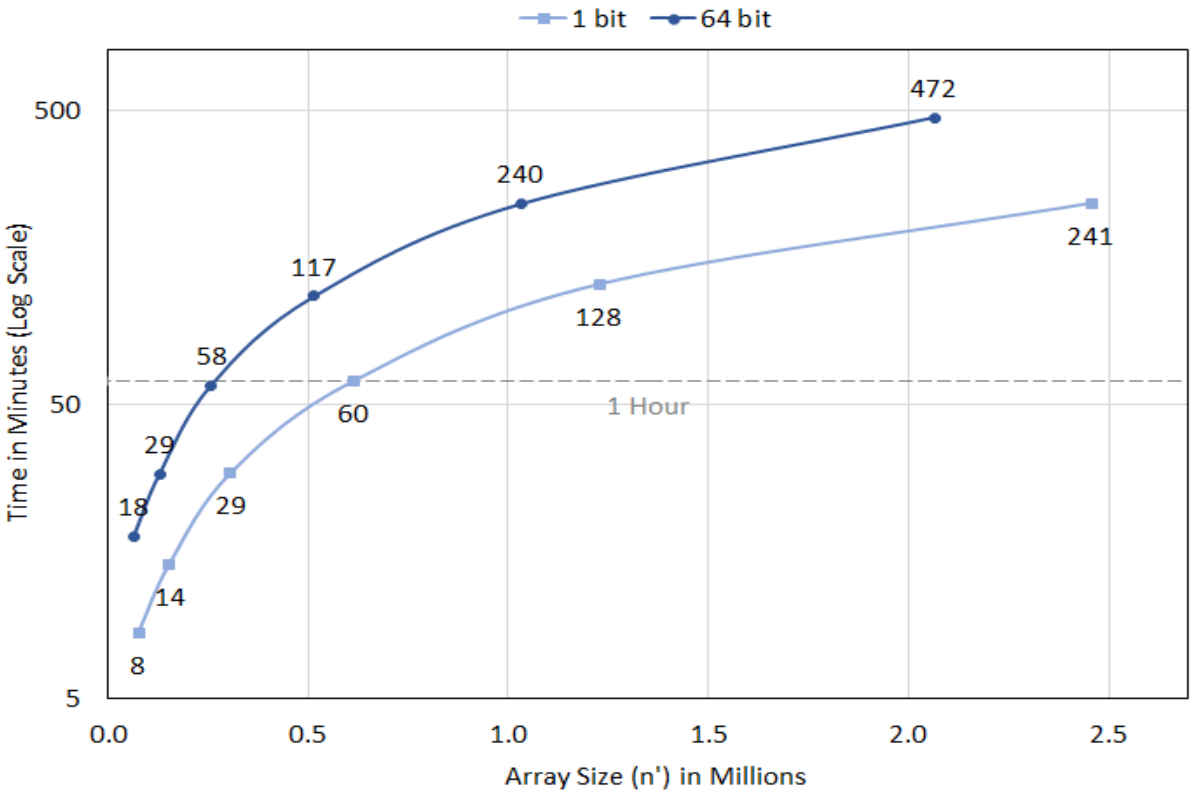

7 Make IsMatch Great Again!

To demonstrate the strength of our compatibility with generic matching criteria we present versatile matching criteria that can be integrated into our protocol, exhibiting advantageous properties for both performance (Section 7.1) and functionality (Sections 7.2-7.6). We stress that security holds for all examples to follow, due to the server obliviously performing homomorphically operations on ciphertexts; See Theorem 5.1, Section 5.1.

7.1 Faster Exact Match via Hashing

To speedup performance in exact match search we propose applying Universal Hashing [10, 42] (See also Section 2.4)to reduce the degree and overall multiplications for computing $\mathsf{IsMatch}$ from $\mathcal{O}(w)$ to $\mathcal{O}(\log(n/\varepsilon))$, where w is the elements' size, n the number of elements, and ε the probability of error. This is particularly appealing in use-cases with large elements, specifically, when $w > 2\log(2n/\varepsilon)$. See Table 3, row (iv). In details, we propose that the server first chooses a uniformly random Toeplitz Matrix $A \in \{0,1\}^{v\times w}$ and a random vector $b \in \{0,1\}^v$ for $v = 2\log(2n/\varepsilon)$ to specify a hash function $h_A(x) = Ax+b \mod 2$ mapping $\{0,1\}^w$ to $\{0,1\}^v$. The server then homomorphically applies on encrypted values the following equality operator on hashed values:

$$\mathsf{IsMatch}(x[i], q) := \mathsf{IsEqual}_v\left(h_A(x[i]), h_A(q)\right)$$

for $\mathsf{IsEqual}_v$ as specified in Section 2.

To analyze the complexity of the above $\mathsf{IsMatch}$ polynomial note that this hashing requires solely homomorphic additions operations. So the degree and overall multiplications is $d = v$ and $\mu = v - 1$ respectively. The failure probability due to collision is $\varepsilon/2$.

7.2 Boolean Logic Queries

The matching criterion in our protocol can express any Boolean logic, such as conjunction, disjunction, negation or their combination. This logic can be applied, for example, on elements' sub-fields (e.g. first and last names in personal records) or characters.

The expression of Boolean logic as a polynomial over GF(2) is via the standard arithmetization techniques: expressing negation by $\mathsf{not}(a) = 1 - a$, conjunction by $\mathsf{conj}(a_1, \ldots, a_t) = \prod_{j=1}^{t} a_j$ and disjunction by $\mathsf{disj}(a, b) = a \oplus b \oplus a \cdot b$.

The degree d (respectively, overall multiplications μ) is the maximum composition length (respectively, overall number) of disjunction and conjunction operations. See concrete examples in the following sections.

7.3 Wild-Card Queries

Wild-card queries are specified by $q \in \{0, 1, *\}^w$. "Wild-card positions" are the entries j where q accepts $*$. Wild-card match returns is true when $x[i]$ and q agree on all the non wild-card positions.

In case the client is willing to leak the wild-card positions, we simply apply the equality test (see Section 2) on the substrings of $x[i]$ and q corresponding to the non wild-card positions. Complexity is only improved by this (compared to exact match on entire strings).

In case the client wishes to hide the wild-card positions, she can augment the query with the (encrypted) indicator vector $I \in \{0,1\}^w$ accepting 1 on all non wild-card positions, and 0 otherwise. The matching polynomial (to be homomorphically evaluated on encrypted values) is:

$$\mathsf{IsMatch}'\left(x[i], (q, I)\right) = \mathsf{IsMatch}\left(x[i] \cdot I, q \cdot I\right)$$

for $\cdot$ the entry-wise product.

Correctness follows as on wild-card entries both $x[i]$ and q are turned to 0 to guarantee equality, and they keep their original values on the non wild-card position.

Complexity overhead for the client is $|w|$ additional encryptions. The server evaluates a polynomial with degree increased by 1, and overall number of multiplications increased by an additive term of $2w$. Saving a factor of w in the overall multiplication is easy: by replacing each $*$ value in q with 0, and compute $\mathsf{IsMatch}\left(x[i] \cdot I, q\right)$.

7.4 Range Queries

Range queries specify lower and upper boundaries (l, u respectively) to retrieve elements in the range (l, u). I.e., retrieving $(i, x[i])$ for $i = \min\{\, i \in [n] \mid l < x[i] < u\}$. Boundaries and data are encrypted by the client.

Range queries are easily implemented as the conjunction of two boundary-tests. Specifically, for elements $x[i]$ in $\{0,1\}^w$, the matching polynomial (to be homomorphically evaluated on encrypted values) is:

$$\mathsf{IsMatch}\left(x[i], (l, u)\right) = \mathsf{conj}\left(\mathsf{IsGrt}_w\left(x[i], l\right), \mathsf{IsGrt}_w\left(u, x[i]\right)\right)$$

for IsGrt_w and conj operators as in Sections 2 and 7.2.

The complexity of the client is dominated by encrypting $|l| + |u| = 2w$ bits for specifying the query, and by decrypted the outcome. The server's complexity grows with the matching degree and overall multiplications: $d = 2(w+1)$ and $\mu = 4w + 1$, respectively.

7.5 Search In Sub-Array

In sub-array search for $\mathsf{IsMatch}$, the client specifies boundaries $l, u \in [n]$ together with the query q in order to retrieve from the server the first match for q in the sub-array $(x[l+1], \ldots, x[u-1])$. I.e. retrieving $(i, x[i])$ for $i = \min\{\, j \in (l, u) \mid \mathsf{IsMatch}(x[j], q) = 1\}$. Boundaries, query and data are all encrypted by the client.

A sub-array search is easily computed as the conjunction of three requirements: $\mathsf{IsMatch}(x[i], q) = 1$, $i > l$, and $u > i$. Specifically, for indices specified by length $m = \log n$ binary representation, the matching polynomial (to be homomorphically evaluated on encrypted values) is:

$$\begin{aligned}&\mathsf{IsMatch\_InSubArray}_n\left(x[i], (q, l, u)\right) = \\ &\mathsf{conj}\left(\mathsf{IsMatch}\left(x[i], q\right), \mathsf{IsGrt}_m\left(i, l\right), \mathsf{IsGrt}_m\left(u, i\right)\right)\end{aligned}$$

for IsGrt_w and conj operators as in Sections 2 and 7.2.

The complexity overhead compared to the underlying $\mathsf{IsMatch}$ (cf. Table 3) is as follows. The client's overhead is encrypting $|l| + |u| = 2\log(n)$ additional bits. The server evaluates a matching polynomial with degree and overall multiplications $d' = d + 2(m+1)$ and $\mu' = \mu + 4m$ respectively, for d, μ the degree and overall multiplications for the underlying matching criterion $\mathsf{IsMatch}$. Namely, an additive overhead of $\mathcal{O}(\log n)$.

We note that to search on a suffix $[x[l+1], \ldots, x[n]]$ of the array it suffices for the client to specify only the lower boundary l and for the server to compute the conjunction of only two conditions: $\mathsf{IsMatch}(x[i], q) = 1$ and $i > l$. Analogously, for prefix search. This reduces the complexity overhead by a factor of 2.

7.6 Sequential Retrieval ("Fetch-Next")

We extend our secure search functionality to return, not only the first match, but also the next matching element (Fetch-Next), with further interaction. I.e. given a match $(i, x[i])$ retrieving the next match $(i', x[i'])$ for $i' = \min\{\, j \in [i+1, n] \mid \mathsf{IsMatch}(x[j], q) = 1\}$. For this purposes we initiate the protocol with an augmented matching criteria that enables searching in an array suffix $[x[l+1], \ldots, x[n]]$ for boundary l specified by the client; see Section 7.5. Boundary, query and data are all encrypted by the client.

To issue a Fetch-Next query for q, after the client has already retrieved a match $(i, x[i])$, the client simply sets l to be i. This causes the search to be performed on indices $[i+1, \ldots, n]$ (without revealing the sub-array to the server). This routine can go on until the client receives the response that indicates that there are no more matches (possibly padding with dummy queries). To issue a "fresh" query, the client will simply set $l = 0$.

The server cannot distinguish between a "fresh" query and a "fetch-next" query as in both cases the index l and query q are encrypted. The amount of elements that match a query q, out of the total number of queries, is not leaked to the sever.

8 Extensions

We overview extensions to our secure search protocol.

8.1 Client-Side Amplification

To reduce the server's complexity load due to degree's growth in inverse-error $1/\varepsilon$, we can employ standard client-side amplification; See [59] Lemma 10.5. Specifically, by setting the protocols error parameter to $0 < \varepsilon_0 < 1/2$, repeatting the protocol in parallel $\frac{-\log_2(1/\varepsilon)}{\log_2(4\varepsilon_0(1-\varepsilon_0))}$ time, and letting the client select the most frequent result, we get error ε but degree growth only with ε_0; See Table 3, row (v).

8.2 Dynamic Data Management

The client can have the benefits of dynamic data management: Insert, Update and Delete (see below). She can also execute search and the above commands multiple times and in any order that she wants.

Insert: Insertion of additional elements requires the server to append another ciphertext to the end of the encrypted array (here, and throughout this work, we assume that the size of the array n is known to the server, see Section 3.3).

Update: To update a specific element $x[i]$ the client first retrieves its index i using our *Secure Search*. Afterwards, to change the value of $x[i]$, denoted *old*, to a different value *new* the client submits the following tuple the server: $\big(\text{UPDATE}, [\![i]\!], [\![diff]\!] = [\![new - old]\!]\big)$. The server now homomorphically adds to each elements $[\![x[j]]\!]$ of the stored array the value $\mathsf{IsEqual}([\![i]\!], j) \cdot [\![diff]\!]$. This results in a new encrypted array $[\![x']\!]$ satisfying $x'[i] = new$ and $\forall j \neq i,\ x'[j] = x[j]$.

Delete: Deletion of elements can be achieved by updating them to a reserved "Deleted" symbol. Another option is switching the value of the element we wish to delete to that of the last element in the array and reducing the number of elements n by 1 (for cases when the dynamic value of n is either maintained by the client, or is not a secret and can be kept with the server).

9 Conclusions

In this work we presented a new and improved solution for secure search on FHE encrypted data. Our solution improves over the prior state-of-the-art of setup-free searching on FHE encrypted data in being: (1) post-processing free and with negligible error probability, (2) faster for both client and server, and (3) compatible with all FHE candidates. We implemented our secure search protocol and performed extensive benchmarks showing concrete run-time speedup by an order of magnitude.

Acknowledgment

The first author is grateful to Adam Sealfon for pointing out the relevance of the Razborov-Smolenski approximation method.

References

[1] Mohamed Ahmed Abdelraheem, Tobias Andersson, and Christian Gehrmann. "Inference and Record-Injection Attacks on Searchable Encrypted Relational Databases." In: *IACR Cryptology ePrint Archive* 2017 (2017), p. 24.

[2] Adi Akavia, Dan Feldman, and Hayim Shaul. *Personal Communication*. 2018.

[3] Adi Akavia, Dan Feldman, and Hayim Shaul. "Secure Search via Multi-Ring Sketch for Fully Homomorphic Encryption". In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM. 2018.

[4] Omer Barkol and Yuval Ishai. "Secure computation of constant-depth circuits with applications to database search problems". In: *Annual International Cryptology Conference*. Springer. 2005, pp. 395–411.

[5] Dan Boneh et al. “Private database queries using somewhat homomorphic encryption”. In: *International Conference on Applied Cryptography and Network Security*. Springer. 2013, pp. 102–118.

[6] Dan Boneh et al. “Public key encryption with keyword search”. In: *International conference on the theory and applications of cryptographic techniques*. Springer. 2004, pp. 506–522.

[7] Christoph Bösch et al. “A survey of provably secure searchable encryption”. In: *ACM Computing Surveys (CSUR)* 47.2 (2015), p. 18.

[8] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. “(Leveled) fully homomorphic encryption without bootstrapping”. In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ACM. 2012, pp. 309–325.

[9] Zvika Brakerski and Vinod Vaikuntanathan. “Efficient Fully Homomorphic Encryption from (Standard) LWE”. In: *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. IEEE Computer Society. 2011, pp. 97–106.

[10] J Lawrence Carter and Mark N Wegman. “Universal classes of hash functions”. In: *Proceedings of the ninth annual ACM symposium on Theory of computing*. ACM. 1977, pp. 106–112.

[11] David Cash et al. “Leakage-abuse attacks against searchable encryption”. In: *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. ACM. 2015, pp. 668–679.

[12] Gizem S Çetin et al. “Blind Web Search: How far are we from a privacy preserving search engine?” In: *IACR Cryptology ePrint Archive* 2016 (2016), p. 801.

[13] Hao Chen, Kim Laine, and Peter Rindal. “Fast private set intersection from homomorphic encryption”. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM. 2017, pp. 1243–1255.

[14] Jung Hee Cheon, Miran Kim, and Myungsun Kim. “Optimized search-and-compute circuits and their application to query evaluation on encrypted data”. In: *IEEE Transactions on Information Forensics and Security* 11.1 (2016), pp. 188–199.

[15] Jung Hee Cheon, Miran Kim, and Kristin Lauter. “Homomorphic computation of edit distance”. In: *International Conference on Financial Cryptography and Data Security*. Springer. 2015, pp. 194–212.

[16] Benny Chor, Niv Gilboa, and Moni Naor. *Private information retrieval by keywords*. Citeseer, 1997.

[17] Benny Chor et al. “Private information retrieval”. In: *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*. IEEE. 1995, pp. 41–50.

[18] Reza Curtmola et al. “Searchable symmetric encryption: improved definitions and efficient constructions”. In: *Journal of Computer Security* 19.5 (2011), pp. 895–934.

[19] Yarkın Doröz, Berk Sunar, and Ghaith Hammouri. “Bandwidth efficient PIR from NTRU”. In: *International Conference on Financial Cryptography and Data Security*. Springer. 2014, pp. 195–207.

[20] Junfeng Fan and Frederik Vercauteren. “Somewhat Practical Fully Homomorphic Encryption.” In: *IACR Cryptology ePrint Archive* 2012 (2012), p. 144.

[21] Michael J Freedman, Kobbi Nissim, and Benny Pinkas. “Efficient private matching and set intersection”. In: *International conference on the theory and applications of cryptographic techniques*. Springer. 2004, pp. 1–19.

[22] Sanjam Garg, Payman Mohassel, and Charalampos Papamanthou. “TWORAM: Efficient oblivious RAM in two rounds with applications to searchable encryption”. In: *Annual Cryptology Conference*. Springer. 2016, pp. 563–592.

[23] Craig Gentry. *A fully homomorphic encryption scheme*. Stanford University, 2009.

[24] Craig Gentry, Shai Halevi, and Nigel P Smart. “Homomorphic evaluation of the AES circuit”. In: *Advances in cryptology–crypto 2012*. Springer, 2012, pp. 850–867.

[25] Craig Gentry, Amit Sahai, and Brent Waters. “Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based”. In: *Advances in Cryptology–CRYPTO 2013*. Springer, 2013, pp. 75–92.

[26] Matthieu Giraud et al. “Practical passive leakage-abuse attacks against symmetric searchable encryption”. In: *14th International Conference on Security and Cryptography SECRYPT 2017*. SCITEPRESS-Science and Technology Publications. 2017.

[27] Oded Goldreich, Silvio Micali, and Avi Wigderson. “How to play any mental game”. In: *Proceedings of the nineteenth annual ACM symposium on Theory of computing*. ACM. 1987, pp. 218–229.

[28] Oded Goldreich and Rafail Ostrovsky. “Software protection and simulation on oblivious RAMs”. In: *Journal of the ACM (JACM)* 43.3 (1996), pp. 431–473.

[29] Torbjrn Granlund et al. *GNU MP 6.1.2 Multiple precision arithmetic library*. Samurai Media Limited, 2016.

[30] Paul Grubbs et al. “Breaking web applications built on top of encrypted data”. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM. 2016, pp. 1353–1364.

[31] Paul Grubbs et al. “Leakage-abuse attacks against order-revealing encryption”. In: *Security and Privacy (SP), 2017 IEEE Symposium on*. IEEE. 2017, pp. 655–672.

[32] S Halevi and V Shoup. *The HElib library*. 2015.

[33] Shai Halevi and Victor Shoup. “Algorithms in helib”. In: *International cryptology conference*. Springer. 2014, pp. 554–571.

[34] Shai Halevi and Victor Shoup. “Bootstrapping for helib”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2015, pp. 641–670.

[35] Shai Halevi and Victor Shoup. “Design and implementation of a homomorphic-encryption library”. In: *IBM Research (Manuscript)* 6 (2013), pp. 12–15.

[36] Yuval Ishai and Eyal Kushilevitz. “Randomizing polynomials: A new representation with applications to round-efficient secure computation”. In: *focs*. IEEE. 2000, p. 294.

[37] Mohammad Saiful Islam, Mehmet Kuzu, and Murat Kantarcioglu. “Access Pattern disclosure on Searchable Encryption: Ramification, Attack and Mitigation.” In: *Ndss*. Vol. 20. 2012, p. 12.

[38] Alhassan Khedr, Glenn Gulak, and Vinod Vaikuntanathan. “SHIELD: scalable homomorphic implementation of encrypted data-classifiers”. In: *IEEE Transactions on Computers* 65.9 (2016), pp. 2848–2858.

[39] Myungsun Kim et al. “Better Security for Queries on Encrypted Databases.” In: *IACR Cryptology ePrint Archive* 2016 (2016), p. 470.

[40] Myungsun Kim et al. “Private Compound Wildcard Queries using Fully Homomorphic Encryption”. In: *IEEE Transactions on Dependable and Secure Computing* (2017).

[41] Ágnes Kiss et al. “Private set intersection for unequal set sizes with mobile applications”. In: *Proceedings on Privacy Enhancing Technologies* 2017.4 (2017), pp. 177–197.

[42] Hugo Krawczyk. “LFSR-based hashing and authentication”. In: *Annual International Cryptology Conference*. Springer. 1994, pp. 129–139.

[43] Eyal Kushilevitz and Rafail Ostrovsky. “Replication is not needed: Single database, computationally-private information retrieval”. In: *Foundations of Computer Science, 1997. Proceedings., 38th Annual Symposium on*. IEEE. 1997, pp. 364–373.

[44] Kristin Lauter, Adriana López-Alt, and Michael Naehrig. “Private computation on encrypted genomic data”. In: *International Conference on Cryptology and Information Security in Latin America*. Springer. 2014, pp. 3–27.

[45] Kristin E Lauter. “Practical applications of homomorphic encryption”. In: *Proceedings of the 2012 ACM Workshop on Cloud computing security workshop*. ACM. 2012, pp. 57–58.

[46] Yehuda Lindell and Jonathan Katz. *Introduction to modern cryptography*. Chapman and Hall/CRC, 2014.

[47] Chang Liu et al. “Search pattern leakage in searchable encryption: Attacks and new construction”. In: *Information Sciences* 265 (2014), pp. 176–188.

[48] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. "On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption". In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM. 2012, pp. 1219–1234.

[49] Muhammad Naveed. "The Fallacy of Composition of Oblivious RAM and Searchable Encryption." In: *IACR Cryptology ePrint Archive* 2015 (2015), p. 668.

[50] Rafail Ostrovsky and William E Skeith. "A survey of single-database private information retrieval: Techniques and applications". In: *International Workshop on Public Key Cryptography*. Springer. 2007, pp. 393–411.

[51] Benny Pinkas, Thomas Schneider, and Michael Zohner. "Faster Private Set Intersection Based on OT Extension." In: *USENIX Security Symposium*. Vol. 14. 2014, pp. 797–812.

[52] Benny Pinkas et al. "Phasing: Private Set Intersection Using Permutation-based Hashing." In: *USENIX Security Symposium*. Vol. 15. 2015, pp. 515–530.

[53] David Pouliot and Charles V Wright. "The shadow nemesis: Inference attacks on efficiently deployable, efficiently searchable encryption". In: *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. ACM. 2016, pp. 1341–1352.

[54] Alexander A Razborov. "Lower bounds on the size of bounded depth circuits over a complete basis with logical addition". In: *Mathematical Notes of the Academy of Sciences of the USSR* 41.4 (1987), pp. 333–338.

[55] Ronald L Rivest, Len Adleman, and Michael L Dertouzos. "On data banks and privacy homomorphisms". In: *Foundations of secure computation* 4.11 (1978), pp. 169–180.

[56] Sujoy Sinha Roy et al. "Hardware assisted fully homomorphic function evaluation and encrypted search". In: *IEEE Transactions on Computers* 66.9 (2017), pp. 1562–1572.

[57] Elaine Shi et al. "Oblivious RAM with O ((logN) 3) worst-case cost". In: *International Conference on The Theory and Application of Cryptology and Information Security*. Springer. 2011, pp. 197–214.

[58] Victor Shoup. "NTL: A library for doing number theory, 10.5.0". In: *http://www. shoup. net/ntl/* (2017).

[59] Michael Sipser. *Introduction to the Theory of Computation*. Vol. 2. Thomson Course Technology Boston, 2006.

[60] Nigel P Smart and Frederik Vercauteren. "Fully homomorphic SIMD operations". In: *Designs, codes and cryptography* 71.1 (2014), pp. 57–81.

[61] Roman Smolensky. "Algebraic methods in the theory of lower bounds for Boolean circuit complexity". In: *Proceedings of the nineteenth annual ACM symposium on Theory of computing*. ACM. 1987, pp. 77–82.

[62] Dawn Xiaoding Song, David Wagner, and Adrian Perrig. "Practical techniques for searches on encrypted data". In: *Security and Privacy, 2000. S&P 2000. Proceedings. 2000 IEEE Symposium on*. IEEE. 2000, pp. 44–55.

[63] Emil Stefanov et al. "Path ORAM: an extremely simple oblivious RAM protocol". In: *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM. 2013, pp. 299–310.

[64] Haixu Tang et al. "Protecting genomic data analytics in the cloud: state of the art and opportunities". In: *BMC medical genomics* 9.1 (2016), p. 63.

[65] Frank Wang et al. "Splinter: Practical Private Queries on Public Data." In: *NSDI*. 2017, pp. 299–313.

[66] David P Woodruff et al. "Sketching as a tool for numerical linear algebra". In: *Foundations and Trends® in Theoretical Computer Science* 10.1–2 (2014), pp. 1–157.

[67] Andrew Chi-Chih Yao. "How to generate and exchange secrets". In: *Foundations of Computer Science, 1986., 27th Annual Symposium on*. IEEE. 1986, pp. 162–167.

[68] Masaya Yasuda et al. "Secure pattern matching using somewhat homomorphic encryption". In: *Proceedings of the 2013 ACM workshop on Cloud computing security workshop*. ACM. 2013, pp. 65–76.

[69] 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, pp. 707–720.

A Proof of Main Theorem 5.1

In this section we give the missing proof details for Theorem 5.1.

A.1 Proof of Theorem 5.1, Correctness

We prove correctness of the search-step (Step III) in the BinaryRaffle protocol (Figure 3), showing that, with probability $1-\varepsilon$, the client's output in is $(b^*, x[i^*])$ for $i^* = \min\{\, i \in [n] \mid \mathsf{IsMatch}(x[i], q) = 1\}$ and $b^* \in \{0,1\}^{\lceil \log_2 n \rceil + 1}$ the binary representation of i^*. The server has no output.

The search is initialized by the client who encrypts her query q and sends $[\![q]\!]$ to the server (Step III-(a)). The server then does the following.

First, the server homomorphically evaluates the matching criterion $\mathsf{IsMatch}$ on each pair of array element $[\![x[i]]\!]$ and query $[\![q]\!]$, $i \in [n]$ (Step (III-(b)-(1)). This results in an array $[\![ind]\!]$ of encrypted binary values $ind = \mathsf{IsMatch}(x[i], q)$, indicating for each element whether it is a match. Namely, ind is a binary vector with first 1 in index i^* (all zeros, if no match exists).

Second, the server computes homomorphically evaluates the $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$ algorithm on $[\![ind]\!]$ to obtain an array of ciphertext $[\![s]\!]$. (Step III-(b)-(2)). By Theorem A.1 with probability $1-\varepsilon$, the outcome is the encryption of a step-function $s[1] = \ldots = s[i^*-1] = 0$ and $s[i^*] = \ldots = s[n] = 1$ (all zeros, if no match exists).

Third, the server homomorphically evaluates the derivative $[\![s']\!]$ for $s'[i] = s[i] - s[i-1]$ and $s'[1] = s[1]$ (Step (III-(b)-(3)). By the above $s' \in \{0,1\}^n$ accept 1 value only on entry i^* (all zeros, if no match exists).

Fourth, the server homomorphically evaluates the product $b^* = Bs'$ on cleartext matrix B and encrypted vector $[\![s']\!]$ to obtain the encrypted vector $[\![b^*]\!]$ (Step III-(b)-(4)). Here $B \in \{0,1\}^{(\lceil \log_2 n \rceil + 1) \times n}$ contains is each column $k \in [n]$ the binary representation of k. Since s' has a single non-zero entry, b^* is equal to the i^*-th column of B. By construction of B this is the binary representation of i^*.

Fifth, the server homomorphically evaluates the sum $\sum_{j=1}^{n} x[j]s'[j]$ to obtain the encrypted output value $[\![x[i^*]]\!]$ (Step III)-(b)-(5)).

The server then sends $([\![b^*]\!], [\![x[i^*]]\!])$ to the client (Step III-(b)-(6)) who decrypts and obtains the desired output: $(b^*, x[i^*])$. □

A.2 Proof of Theorem 5.1, Complexity

We analyze the complexity of the search-step (Step III) in our BinaryRaffle protocol (Figure 3), to show the following. The client's running-time is the time compute $|q|$ encryptions and $|b^*| + |x[i^*]|$ decryptions. The server evaluates a search polynomial over the encrypted data and encrypted query. This polynomial for computing both index i^* and element $x[i^*]$ is of degree $\log(n/\varepsilon) \cdot d$ and overall multiplication $n\,(\log(n/\varepsilon) + \mu + w)$. Here n is the number of data elements, ε the failure probability, w the binary representation length of $x[i^*]$, and d, μ the degree and overall multiplications respectively for the polynomial realizing the matching criterion $\mathsf{IsMatch}$. The communication is 1-round, consisting of $|q|$ ciphertexts sent from client and $|b^*| + |x[i^*]|$ ciphertexts from server.

The only operations performed by the client are encrypting the query q and decrypting the result b^*.

The server's evaluated polynomial includes n executions of $\mathsf{IsMatch}$ (III)-(b)-(1), a single execution of $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$ (III)-(b)-(2), another $n+1$ homomorphic additions for the pairwise differences (III)-(b)-(3), and at most $n\log(n)$ homomorphic additions for the calculation of the binary representation of the result index (III)-(b)-(4).

All the homomorphic multiplications come from the n executions of $\mathsf{IsMatch}$ and the single execution of $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$. Regarding degree, the first introduces a degree of $d_{\mathtt{IsMatch}}$ and the second introduces a degree of $\log(n/\varepsilon)$ (see Theorem A.2). Regarding multiplications, the first introduces $n \cdot \mu_{\mathtt{IsMatch}}$ multiplications (see Theorem A.2) and the second introduces additional $n \cdot \log(n/\varepsilon)$ multiplications. The total degree is the product of both degrees above. The overall multiplications are the sum of all multiplications above. □

A.3 Proof of Theorem 5.1, Security

We prove the that our BinaryRaffle protocol (Figure 3) attains full security (see Definition 5), assuming semantic security of $\mathcal{FHE}$. The security proof follows immediately from the semantic security of the underlying FHE; details follow.

Consider first the query attack game (See Section 3.3). We show below how to construct, given a PPT algorithm $\mathcal{A}$ for the query attack game, a PPT algorithm $\mathcal{A}'$ for the (single message) IND-CPA game for $\mathcal{FHE}$ (see Definition 11.2 and Proposition 11.3 in [46]). We show that the advantage of $\mathcal{A}$ in the former game is equal to the advantage of $\mathcal{A}'$ in the latter. The semantic security of $\mathcal{FHE}$ implies that no PPT adversary $\mathcal{A}'$ has non-negligible advantage in the latter. We conclude that no PPT adversary $\mathcal{A}$ has non-negligible advantage in the former.

We construct the adversary $\mathcal{A}'$ (playing the role of the adversary in the IND-CPA game for $\mathcal{FHE}$, and the challenger in the query attack game): (1) Upon receiving $(\mathsf{pk},\mathsf{sk}) \leftarrow_{\$} \mathsf{KGen}(1^{\kappa})$ from the challenger in the IND-CPA game for $\mathcal{FHE}$ ($\mathcal{C}_{\mathcal{FHE}}$), $\mathcal{A}'$ invokes the query attack game with $\mathcal{A}'$ playing the challenger's role and $\mathcal{A}$ the adversary, by sending $(\mathsf{pk},\mathsf{sk})$ to $\mathcal{A}$. (2) Upon receiving from $\mathcal{A}$ the array x and two queries $q^{(0)}, q^{(1)}$, $\mathcal{A}'$ sends the two queries to $\mathcal{C}_{\mathcal{FHE}}$. (3) Upon receiving from $\mathcal{C}_{\mathcal{FHE}}$ the challenge ciphertext $c = [\![q^{(b)}]\!]$, $\mathcal{A}'$ encrypts (element-by-element and bit-by-bit) to produce the ciphertexts array $[\![x]\!]$, and sends $[\![x]\!], c$ to $\mathcal{A}$. (4) Upon receiving a guess b' from $\mathcal{A}$, $\mathcal{A}'$ sends b' to $\mathcal{C}_{\mathcal{FHE}}$.

To show that the advantage of $\mathcal{A}'$ in the (single message) IND-CPA game for $\mathcal{FHE}$ is equal to the advantage of $\mathcal{A}$ in the query attack game, observe that $\mathcal{A}'$ perfectly simulates the challenger in the query attack game and wins the (single message) IND-CPA game for $\mathcal{FHE}$ if-and-only-if $\mathcal{A}$ wins the query attack game.

Consider next the data attack game (See Section 3.3). The proof that no PPT adversary has non-negligible advantage in this game is analogous to the first case above. The main difference is that the reduction is to the multi-message IND-CPA game for $\mathcal{FHE}$ (see Definition 11.5 and Theorem 11.6 in [46]), where the challenge is the array of ciphertexts for the data $x^{(b)}$, (for $b \leftarrow_{\$} \{0,1\}$ chosen by the challenger). Details omitted. □

A.4 BinaryRaffleStepFunction Analysis

The main properties of BinaryRaffleStepFunction Algorithm are stated below.

Lemma A.1 (Correctness). *Let $v \in \{0,1\}^n$ be binary vector and $t = \big(t[1],\ldots,t[n]\big)$ be the vector returned after executing* $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}(v)$. *Then the following holds:*

1. *If $\forall i \in [n] : v[i] = 0$, then with probability 1 it holds that $\forall j \in [n] : t[j] = 0$.*
2. *If $\exists i \in [n] : v[i] = 1$, then with probability $1-\varepsilon$ it holds that $t[1] = \ldots = t[i^*-1] = 0$ and $t[i^*] = \ldots = t[n] = 1$ for $i^* = \min\{i \in [n] \mid v[i] = 1\}$.*

Proof. The first case is trivial: if $v = (0,\ldots,0)$ then all the random partial prefix sums are zero (i.e., $\forall j \in [N(\varepsilon)], \forall k \in [n] : S[j,k] = 0$). This hold for all samples $r_1,\ldots,r_{N(\varepsilon)}$, so the resulting binary step function vector is $t[1] = \ldots = t[n] = 0$.

Next we analyze the second case, namely, when $\exists i \in [n] : v[i] = 1$, we denote $i^* = \min\{i \in [n] \mid v[i] = 1\}$. We prove the following:

1. For any $\eta \in [N(\varepsilon)]$, the probability of a random partial prefix sum in index $\ell \in [i^*, n]$ (the element $S[\eta,\ell]$ in the matrix) to be either zero or one is exactly half.
2. For any index $\ell \in [i^*, n]$, the probability that $t[\ell]$ equals to zero is exactly $2^{-N(\varepsilon)}$.
3. The probability that there exists an index $\ell \in [i^*, n]$ so that $t[\ell] = 0$ is at most $n \cdot 2^{-N(\varepsilon)}$.

Assigning $N(\varepsilon) = \log_2(n/\varepsilon)$ in (3) above concludes the proof. □

Lemma A.2 (Complexity). $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$ *algorithm is realized by a polynomial of degree $\log_2(n/\varepsilon)$ which performs $n \cdot \log_2(n/\varepsilon)$ overall multiplications.*

Proof. The only multiplications performed in the algorithm are during the computation of OR between the elements in each column of the matrix S (step 3). Thus, the degree of the polynomial equals to on $N(\varepsilon) = \log_2(n/\varepsilon)$ (the length of each column in S), and the overall number of multiplications is $n \cdot N(\varepsilon)$ ($N(\varepsilon)$ multiplications for each of the cells of t). □

Corollary A.1 (Negligible Error Probability)**.** *Given security parameter κ, by choosing $N(\varepsilon) = \log_2(n) + \omega(\log_2(\kappa))$ we get that the failure probability of the* $\mathsf{BinaryRaffleStepFunction}_{n,\varepsilon}$ *algorithm is* $\mathsf{negl}(\kappa)$.

B Folklore

In this section we describe the natural secure search solution on FHE encrypted data (folklore) [2], it's correctness, complexity and experimental results comparing it to Binary Raffle.

B.1 Description

Given array $x = (x[1], \ldots, x[n]))$ and lookup value q, the folklore solution operates as follows: First apply the $\mathsf{IsMatch}$ polynomial to produce an vector of indicators $ind[i] = \mathsf{IsMatch}(x[i], q)$ same as in Step (III)-(b)-(1) in Figure 3. Next, to retrieve the index i of the first match to q in x, or equivalently, the first non-zero entry in $ind = (v[1], \ldots, v[n]) \in \{0,1\}^n$ evaluate the polynomial:

$$\mathsf{Folklore}(ind) = \sum_{i=1}^{n} v[i] \cdot \prod_{j=1}^{i-1} (1 - v[j]) \cdot i \mod p,$$

where for $i = 1$ we define $\prod_{j=1}^{0} (1 - v[j]) = 1$, and $p > n$.

B.2 Correctness

Correctness follows by observing that at most a single summand of the computed in $\mathsf{Folklore}(ind)$ is non-zero. This is because in each summand, the expression $v[i] \cdot \prod_{j=1}^{i-1} (1 - v[j])$ is a test of whether all preceding entries are zero and the current entry is 1. The test outcome true (value 1) only on the first positive entry of ind. In this case, the summand adds the value i to the sum, which is the total outcome.

B.3 Complexity and Optimizations

The degree of $\mathsf{Folklore}(ind)$ is n. The overall number of multiplications is $n^2/2$ with a naive implementation. The number of multiplications can be reduced to n with an optimized implementation that computes re-using the product computed for the previous summand. The ring $\mathbb{Z}_p$ can be reduced to binary by working over the binary representations of i; in this case the overall number of multiplications is increased by $\log n$ as we repeat the computation for producing each bit in the binary representation of i.

B.4 Experimental Results

We executed benchmarks (with the same setup as described in Section 6.1) on the implementation of the folklore protocol that includes all the aforementioned optimizations on binary representations of i and word size $w = 1$ (Figure 6). For values larger than $2 \cdot 10^6$ the folklore experiments were unable to complete due to increased RAM consumption for higher HElib levels. In comparison, executions of Binary Raffle for word size $w = 1$ and error probabilities $\varepsilon \in \{2^{-80}, 2^{-1}\}$ are also present in the graph and show the execution time of array sizes up to $\approx 20 \cdot 10^6$.

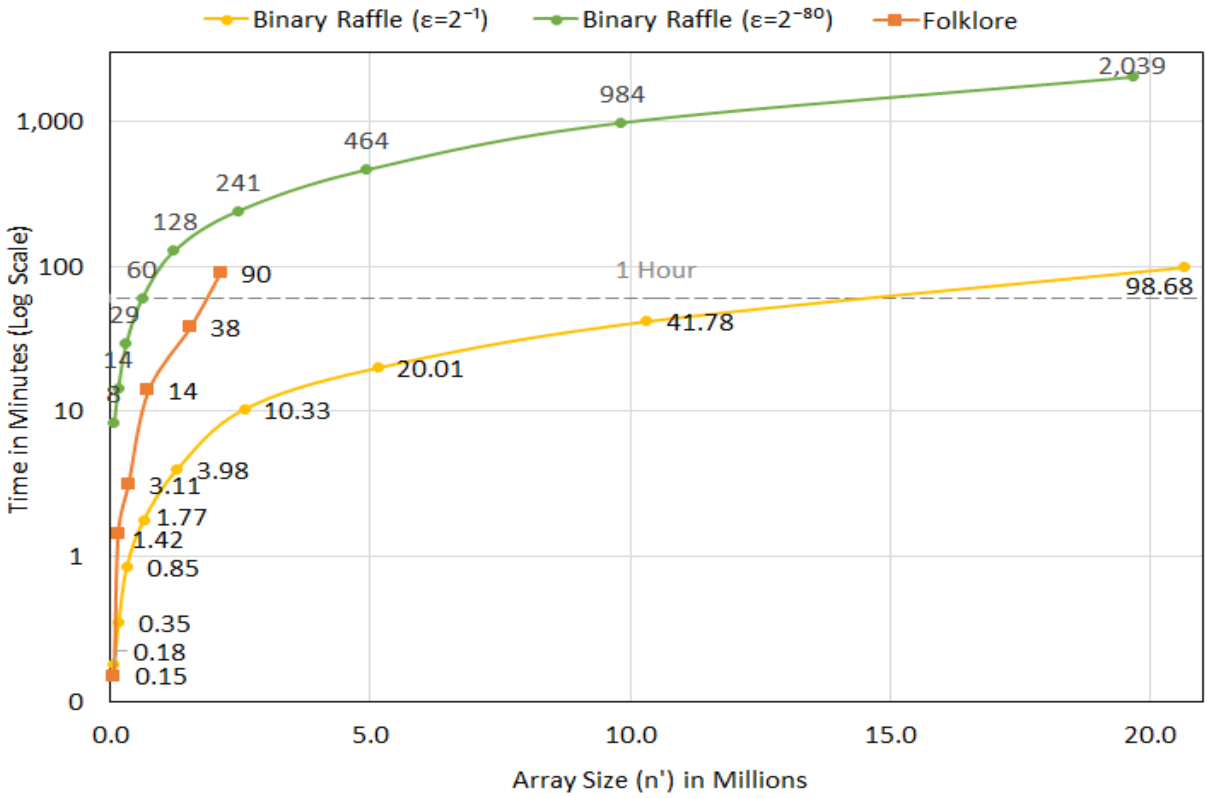

Figure 6: Folklore over binary representation