

SophOMR: Improved Oblivious Message Retrieval from SIMD-Aware Homomorphic Compression

Keewoo Lee*
Ethereum Foundation

Yongdong Yeo*
Seoul National University

Abstract

Privacy-preserving blockchains and private messaging services that ensure receiver-privacy face a significant UX challenge: each client must scan every payload posted on the public bulletin board to avoid missing messages intended for them. Oblivious Message Retrieval (OMR) addresses this issue by securely outsourcing this expensive scanning process to a service provider using Homomorphic Encryption (HE).

In this work, we propose a new OMR scheme that substantially improves upon the previous state-of-the-art, PerfOMR (USENIX Security'24). Our implementation demonstrates reductions of 3.4x in runtime, 2.2x in digest size, and 1.5x in key size, in a scenario with 65536 payloads (each 612 bytes), of which up to 50 are pertinent. At the core of these improvements is a new homomorphic compression mechanism, where ciphertexts of length proportional to the number of total payloads are compressed into a digest whose length is proportional to the upper bound on the number of pertinent payloads. Unlike previous approaches, our scheme fully exploits the native homomorphic SIMD structure of the underlying HE scheme, significantly enhancing efficiency. In the setting described above, our compression scheme achieves 7.5x speedup compared to PerfOMR.

1 Introduction

Consider a classic scenario where a sender wants to deliver a message to a receiver without revealing any information to eavesdroppers. While end-to-end encryption ensures *message-privacy*, it alone does not guarantee *metadata-privacy*. In particular, eavesdroppers may still gather details such as who is communicating, when, and how often, enabling traffic analysis [6, 52] and compromising privacy.¹² As a former NSA gen-

eral counsel said:³ “Metadata absolutely tells you everything about somebody’s life. If you have enough metadata you don’t really need content.” In this regard, mechanisms for metadata-privacy are critical in private messaging systems [4, 24, 61, 63] and privacy-preserving cryptocurrencies [5, 10, 53].

In this work, we specifically focus on *receiver-privacy*, which aims to hide the receiver’s identity. A plausible solution to receiver-privacy is the *public bulletin board* approach. In this framework, the sender encrypts the message using the receiver’s public key and posts it on a public bulletin board (e.g., blockchain). Over time, the bulletin board accumulates messages from various senders to various receivers. To retrieve their messages, a receiver scans the entire bulletin board and attempts to decrypt each message one by one. Messages that decrypt correctly are those intended for the receiver, while unsuccessful decryptions are ignored. This framework ensures that eavesdroppers without access to the decryption key cannot link messages to their receivers.

In fact, many privacy-preserving blockchain projects that aim to provide receiver-privacy are essentially following this approach (e.g., Zcash [10, 35], Aztec [5], Monero [53], ERC-5564 [62] for Stealth Address). However, this approach poses a significant UX challenge: clients must individually scan every payload posted on a public bulletin board to avoid missing messages intended for them. In the context of privacy-preserving blockchains, this means that even light clients are required to scan all transactions.

One approach to addressing this issue is to *securely outsource* the expensive scan to a powerful server. *Homomorphic Encryption* (HE) [27, 57], a cryptosystem that allows computation on encrypted data, is a natural cryptographic primitive for such a secure outsourcing scenario [16, 17]. With HE, we can delegate the scan to an untrusted server without compromising the privacy of receivers. Liu and Tromer [43] formalized this HE-based secure scan outsourcing solution as a primitive called *Oblivious Message Retrieval* (OMR). They also explored various design choices and optimization

*Equal contribution.

¹<https://www.propublica.org/article/how-facebook-undermines-privacy-protections-for-its-2-billion-whatsapp-users>

²<https://www.forbes.com/sites/thomasbrewster/2022/02/23/meet-the-secretive-surveillance-wizards-helping-the-fbi-and-ice-wiretap-facebook-and-google-users/>

³<https://www.wired.com/2015/03/data-and-goliath-nsa-metadata-spying-your-secrets/>




Figure 1: Oblivious Message Retrieval

techniques and provided a proof-of-concept implementation of OMR. However, this initial scheme required over 20 hours to process 2^{19} payloads (155 ms/payload), making it impractical for real-world use. A follow-up work, PerfOMR [46] by Liu-Tromer-Wang, advanced the practicality of OMR through further optimizations, achieving significantly improved efficiency. Specifically, PerfOMR reported $15\times$ faster runtime compared to the original OMR construction [43].

1.1 Our Contribution

In this work, we propose a new OMR scheme, SophOMR, that substantially improves upon PerfOMR [46]. Specifically, it takes less than 3 minutes to process 2^{16} payloads in a single thread (2.5 ms/payload).

- While SophOMR largely follows the high-level framework of prior works [43, 46], our key conceptual contribution is in viewing OMR through the lens of *homomorphic compression* and in identifying compression as the primary bottleneck in prior works. At the heart of our improvement is our new homomorphic compression mechanism based on [21]. The main advantage of our approach lies in its compatibility with the homomorphic SIMD (Single Instruction, Multiple Data) structure of the underlying HE scheme. This SIMD-awareness avoids complexities found in prior works [43, 46], resulting in a more efficient and arguably simpler scheme. Notably, SophOMR requires only one level of multiplicative depth and $O(\sqrt{N})$ homomorphic rotations to compress N payloads, in contrast to PerfOMR, which requires two levels of depth and $O(N)$ rotations. In our primary setting, this yields a $7.5\times$ speedup in homomorphic compression. See also Sec. 1.3 and Tab. 1.
- We also suggest several other optimization techniques to further improve the performance of the scheme: ring-switching (Sec. 4.2), hybrid key-switching (Sec. 4.3), and BSGS-style homomorphic matrix multiplication (Sec. 4.4).
- We implement SophOMR in a C++ library and compare its concrete performance to that of PerfOMR. Our implementation demonstrates reductions of $3.4\times$ in runtime, $2.2\times$ in

digest size, and $1.5\times$ in key size, in a scenario with 65536 payloads (each 612 bytes), of which up to 50 are pertinent. We highlight that the comparison was conducted in a manner that favors PerfOMR, whose parameters achieve 94-bit security and 18-bit completeness, whereas our parameters offer 128-bit security and 30-bit completeness. (See Sec. 5.)

- We also consider OMD (Oblivious Message Detection), where clients retrieve only pertinent indices rather than the payloads (Sec. 4.6). OMD can be combined with PIR (Private Information Retrieval) schemes [23] to achieve OMR-like functionality. Although this approach adds an extra communication round, it provides a generic way to achieve OMR-like functionality while leveraging the specific advantages of the chosen PIR scheme. In the same setting, our OMD variant achieves an overall $4.6\times$ speedup and a specific $479\times$ speedup in homomorphic compression compared to PerfOMR (Sec. 5).

1.2 System Model

More precisely, OMR operates as follows (Fig. 1). For more details, refer to Sec. 2.5 and Sec. 3.

- ① Each user runs KeyGen, publishes their signal key pk_{sig} , and send the detection key pk_{det} to Detector.
- ② Sender, who wants to deliver a payload⁴ x to Receiver, posts the pair (sig, x) on the bulletin board BB, tagging the payload with a signal sig generated using the Receiver’s signal key pk_{sig} . The signal allows only the intended receiver to detect the payload.
- ③ Upon Receiver’s requests, Detector uses the Receiver’s detection key pk_{det} to summarize the bulletin board BB into a digest D and sends it to Receiver.
- ④ Finally, Receiver decodes the digest D to recover its pertinent payloads.

⁴Note that we treat the payload x simply as a bitstring. In practice, these payloads are typically end-to-end encrypted, possibly with accompanying authentication. Since OMR focuses on receiver privacy, we assume message privacy and integrity are handled by the payload format itself and therefore abstract them away as a bitstring.

At a high level, Detector (3) performs *homomorphic detection* of the signals, given homomorphically encrypted Receiver’s secret as a part of the detection key pk_{det} . This yields an encrypted sparse binary vector, where the ones indicate the indices pertinent to Receiver. The encrypted binary vector is then used to mask the payloads on Bulletin Board, preserving only the pertinent payloads while zeroizing the rest. However, notice that the result is as large as the entire Bulletin Board; the Detector cannot simply discard irrelevant payloads since the detection result is encrypted. Therefore, homomorphic detection alone does not yield a *compact* digest.

This is where *homomorphic compression* comes into play. It is a technique that obviously compresses a lengthy, homomorphically encrypted sparse vector into a digest whose size depends only on an *upper bound for the number of non-zero components* in the input ciphertexts. Applied to the output of homomorphic detection, it condenses the sparse vector of pertinent payloads into a compact digest. However, note that this requires the client to know an a priori upper bound on the number of pertinent payloads.

Remark 1 (A Priori Upper Bound). *The upper bound for the number of pertinent payloads can be predicted from application scenarios and past statistics. If the number of relevant items exceeds this bound, it can be easily detected (e.g., through a checksum). In such cases, the query can simply be rerun with a larger bound. We leave the bound as a client-side configuration, which can be easily updated over time. While activity volume can be inferred through the bound, this leakage is unavoidable if the client’s communication cost is to remain independent of the Bulletin Board size.*

Threat Model. The security of OMR (Sec. 2.5) considers a passive adversary attempting to link specific messages to their receivers. The adversary has access to the public bulletin board, public keys, and any communication among the parties in the system. It may also collude with any party except the sender and receiver of the target messages.

1.3 Technical Overview

SIMD-Aware Homomorphic Compression. Our main technical contribution is identifying that the bottleneck of prior works [43, 46] is at homomorphic compression and devising a new mechanism based on the SIMD-aware homomorphic compression scheme of [21]. Prior approaches rely on a *hashing-based* compression scheme (Sec. 3.4) in which each index is *hashed* into *buckets*, making the scheme incompatible with the homomorphic SIMD structure. As a result, before the actual compression, each slot in the detection result must be *unpacked* into separate ciphertexts, using costly subroutines like SlotToCoeff [30, 47] and *oblivious expansion* [3]. The latter step involves $O(N)$ homomorphic rotations, where N is the size of the bulletin board, creating a significant bottleneck in the digesting process.

Table 1: Cost Comparison for Digest Compression

	PerfOMR	SophOMR
#(PMul)	$\left(\frac{kt}{n} + 1 + \frac{vt}{n \log p}\right) N$	$\left(\frac{kt}{n} + \frac{k}{n}\right) N$
#(Rot)	$\left(4 + \frac{2}{\sqrt{n}} - \frac{4}{n}\right) N$	$\approx \sqrt{\frac{4k(t+1)}{n}} \sqrt{N}$
Depth	2	1

* N : # of payloads

* n : BFV ring dimension

* t : # of BFV slots per payload

* k : # of pertinent payloads

* p : BFV plaintext modulus

* 2^{-v} : false negative probability

In contrast, our SIMD-aware homomorphic compression mechanism based on [21] eliminates the need for a complex and expensive unpacking process, enabling a seamless transition between the detection and compression phase (Sec. 4.1). Furthermore, our compression scheme ultimately reduces to a single homomorphic matrix-vector multiplication, a well-studied operation in the HE literature [18, 32, 33, 38, 48]. As a result, our scheme is arguably simpler to understand and implement, while also being significantly more efficient (see Tab. 1): the number of required homomorphic rotation is reduced from $O(N)$ to $O(\sqrt{N})$, and the number of required plaintext-ciphertext multiplications is also substantially reduced under realistic parameter settings ($kt < n$), compared to the previous hashing-based approach of PerfOMR. Furthermore, SophOMR achieves this compression using only a single multiplicative depth, whereas PerfOMR requires two—an additional depth is needed for the *unpacking* step—which directly impacts the overall efficiency of OMR. Another advantage of our compression scheme is that, unlike the previous hashing-based approach, it is *deterministic* and does not introduce any additional failure probability, which substantially simplifies the analysis for parameter setting.

A limitation of the SIMD-aware homomorphic compression of [21] is that it only considers small payloads that fit within a single message *slot* of the underlying HE (e.g., 16-bit). However, in many applications, including OMR, payloads are significantly larger and require multiple slots. To address this, we introduce *precomputation* and *stacking* techniques to accelerate compression of multi-slot payloads, achieving a quadratic reduction in the required number of homomorphic rotations, relative to the number of slots, compared to naively compressing each slot individually.

Optimizations. We also suggest several other optimization techniques to further improve the performance of our scheme.

- **Ring-Switching** (Sec. 4.2): We adopt *ring-switching* [14, 28, 29], which reduces the HE parameter after completing the desired homomorphic computations, to achieve further compression of the digest.

- *Hybrid Key-Switching Optimization* (Sec. 4.3): We optimize the parameter for *hybrid key-switching* [31, 34, 39] in the HE scheme to reduce the key size.
- *BSGS Optimization for Affine Transform* (Sec. 4.4): We apply *BSGS (Baby-Step-Giant-Step)* optimization to speed up homomorphic matrix-vector multiplications.

1.4 Related Work

Oblivious Message Retrieval. Besides the two works [43, 46] previously mentioned, OMR has been extended further by several studies. Group OMR [45] extends OMR to group messaging scenarios, offering enhanced efficiency than naive use of OMR, where the detector’s workload scales linearly with the group size. DoS-resistant OMR [42] considers an extended security notion to safeguard OMR schemes against spamming attacks by malicious senders.

Receiver Privacy. Several non-HE-based solutions have been proposed that offer receiver privacy along with OMR/OMD-like functionality. Fuzzy Message Detection (FMD) [9] is a *decoy-based* solution that employs a detection algorithm specifically designed to misidentify impertinent messages at a predetermined false positive rate, effectively hiding the exact set of pertinent messages. However, the decoy-based security notion of FMD is indeed relatively complex and weak, as analyzed in [59]. In contrast to FMD, some works achieve full privacy [36, 37, 51], but they rely on stronger trust assumptions than OMR. HomeRun [37] and Scalable Private Signaling [36] are roughly two-orders-of-magnitude faster than SophOMR, assuming non-colluding servers and trusted hardware, respectively.

SIMD-Aware Homomorphic Compression. Several homomorphic compression methods have been proposed [11, 21, 22, 26, 43]. However, many of these methods [22, 26, 43] are not particularly efficient when implemented, as they are not SIMD-friendly due to their *hashing-based* approach. (See Sec. 3.4.2.) Although Cheon-Lee-Park-Yeo [21] is the first to investigate SIMD-friendliness, some earlier works feature structures that can be implemented in a SIMD-friendly manner. However, these approaches have their own limitations: The FFT-based method by Fleischhacker-Larsen-Simkin [26] has a costly decompression that requires solving small instances of discrete logarithms, and the method by Binstock-Patel-Seo-Yeo [11] assumes that the entity performing the decompression already knows the indices of interest, which is not applicable in the context of OMR. Cheon et al. [21] proposed a SIMD-aware homomorphic compression scheme by reinterpreting and extending PS-COIE of Choi et. al. [22], which only considered compressing a binary vector and disregarded its SIMD-friendly structure. This new scheme achieves

both an asymptotically optimal compression rate and asymptotically good decompression complexity.

2 Preliminaries

2.1 Notations

We use $[n]$ to denote the set of integers $\{0, \dots, n-1\}$, and \log refers to the base-2 logarithm. Throughout the paper, we follow the column vector notation. For a vector \mathbf{v} , we use $\mathbf{v}[i]$ to denote its i -th element, and for a matrix \mathbf{M} , we use $\mathbf{M}[i][j]$ to denote its (i, j) -th element. The Hadamard product (a.k.a. element-wise product) is denoted by \odot , and $\langle \cdot, \cdot \rangle$ denotes the inner product. For $i < n$, we use Rot^i to denote the left rotation of a vector by i positions, i.e., $\text{Rot}^i(v_0, \dots, v_{n-1}) = (v_i, \dots, v_{n-1}, v_0, \dots, v_{i-1})$. The vector $\mathbf{1}$ refers to a vector of ones $(1, \dots, 1)$. We use λ for the security parameter and $\text{negl}(\lambda)$ to denote a negligible function with respect to λ . The notation $\text{ct}(m)$ denotes a ciphertext encrypting a message m , and $\text{ct}(\mathbf{m})$ denotes possibly multiple ciphertexts encrypting the entries of a vector \mathbf{m} .

2.2 Homomorphic Encryption

A homomorphic encryption (HE) scheme allows us to compute an encryption of $f(m)$ when only given a circuit f and a ciphertext encrypting m without decryption. Let \mathcal{M} be the message space and \mathcal{C} be the ciphertext space. An HE scheme is a tuple of probabilistic polynomial time algorithms ($\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval}$) described below:

- $\text{KeyGen}(1^\lambda)$: The key generation algorithm takes a security parameter λ and returns a secret key sk and public key pk .
- $\text{Enc}_{\text{pk}}(m)$: The encryption algorithm takes the public key pk and a message $m \in \mathcal{M}$ as inputs. Then, it returns a ciphertext $\text{ct}(m) \in \mathcal{C}$ encrypting m .
- $\text{Dec}_{\text{sk}}(\text{ct}(m))$: The decryption algorithm takes the secret key sk and a ciphertext $\text{ct}(m) \in \mathcal{C}$ encrypting a message $m \in \mathcal{M}$. Then, it returns the message m .
- $\text{Eval}_{\text{pk}}(f, \text{ct}(\mathbf{m}))$: The evaluation algorithm takes the public key pk , a circuit f , and ciphertext $\text{ct}(\mathbf{m}) \in \mathcal{C}^*$ as inputs. Then, it returns ciphertext $\text{ct}(f(\mathbf{m})) \in \mathcal{C}^*$ that encrypts $f(\mathbf{m}) \in \mathcal{M}^*$ as the message.

Definition 1 (Security of HE). *Let $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)$.*

- **Correctness:** *For any circuit f (in some class \mathcal{F}) and input $\mathbf{m} \in \mathcal{M}^*$, the following holds with overwhelming probability for $\text{ct}(\mathbf{m}) \leftarrow \text{Enc}_{\text{pk}}(\mathbf{m})$:*

$$\text{Dec}_{\text{sk}}(\text{Eval}_{\text{pk}}(f, \text{ct}(\mathbf{m}))) = f(\mathbf{m})$$

- **Semantic Security:** For any probabilistic polynomial time adversary \mathcal{A} and $m_1, m_2 \in \mathcal{M}$, the following holds for $\text{ct}_i \leftarrow \text{Enc}_{\text{pk}}(m_i)$:

$$|\Pr[\mathcal{A}(\text{pk}, \text{ct}_1) = 1] - \Pr[\mathcal{A}(\text{pk}, \text{ct}_2) = 1]| \leq \text{negl}(\lambda)$$

2.2.1 BFV Scheme

As in the previous works [43, 46], we employ the BFV scheme [13, 25] in our SophOMR construction. The BFV scheme supports homomorphic SIMD arithmetic in \mathbb{Z}_p for a prime p , which is critical to our SophOMR construction, especially in our homomorphic compression scheme.

More precisely, we consider the BFV scheme with the ring dimension n that is a power of 2, along with the plaintext modulus p satisfying $p \equiv 1 \pmod{2n}$. In this case, the plaintext space is given by $\mathcal{R}_p := \mathbb{Z}_p[X]/(X^n + 1) \cong \mathbb{Z}_p^{2 \times (n/2)}$, while the ciphertext space is \mathcal{R}_Q^2 , where $\mathcal{R}_Q := \mathbb{Z}_Q[X]/(X^n + 1)$ for a ciphertext modulus Q . For a plaintext $\mathbf{m} \in \mathbb{Z}_p^{2 \times (n/2)}$, we often refer to each plaintext entry as a *slot*.

The BFV scheme supports (1) SIMD arithmetic in \mathbb{Z}_p and (2) homomorphic rotations (horizontal and vertical). That is, we can homomorphically add and multiply a ciphertext encrypting $\mathbf{v} \in \mathbb{Z}_p^{2 \times (n/2)}$ with a ciphertext (or plaintext) of $\mathbf{w} \in \mathbb{Z}_p^{2 \times (n/2)}$ to obtain a ciphertext of $\mathbf{v} + \mathbf{w}$ and $\mathbf{v} \odot \mathbf{w}$, respectively. For homomorphic rotations, to be precise, let $\text{Rot}^r(\cdot)$ denote the rotation operation which maps $[v_1, \dots, v_{n/2}] \mapsto [v_{1+r}, \dots, v_{n/2}, v_1, \dots, v_r]$. Then, the BFV scheme supports homomorphic rotations computing Rot_{row} and Rot_{col} , defined as follows for $\mathbf{m} = [\mathbf{m}_1 | \mathbf{m}_2]^T \in \mathbb{Z}_p^{2 \times (n/2)}$, where $\mathbf{m}_1, \mathbf{m}_2 \in \mathbb{Z}_p^{n/2}$.

$$\begin{aligned} \text{Rot}_{\text{row}}^r \left(\begin{bmatrix} \mathbf{m}_1^T \\ \mathbf{m}_2^T \end{bmatrix} \right) &= \begin{bmatrix} \text{Rot}^r(\mathbf{m}_1^T) \\ \text{Rot}^r(\mathbf{m}_2^T) \end{bmatrix} \\ \text{Rot}_{\text{col}}^1 \left(\begin{bmatrix} \mathbf{m}_1^T \\ \mathbf{m}_2^T \end{bmatrix} \right) &= \begin{bmatrix} \mathbf{m}_2^T \\ \mathbf{m}_1^T \end{bmatrix} \end{aligned}$$

In this work, we will often abuse notation and identify a BFV plaintext $[\mathbf{m}_1 | \mathbf{m}_2]^T \in \mathbb{Z}_p^{2 \times (n/2)}$ with the vector $[\mathbf{m}_1^T | \mathbf{m}_2^T]^T \in \mathbb{Z}_p^n$.

2.2.2 Sparse Packing

When working with a vector \mathbf{v} of length k that is smaller than the HE ring dimension n , we often employ *Sparse Packing* (a.k.a. Repetitive Slot Packing), where multiple copies of \mathbf{v} are packed into a plaintext. More precisely, when $n \equiv 0 \pmod{k}$, we pack \mathbf{v} (n/k) times, resulting in $[\mathbf{v}^T | \dots | \mathbf{v}^T]^T$. The rationale for using sparse packing is that it preserves the rotation structure. For example, $\text{Rot}_{\text{row}}^k([\mathbf{v}^T | \dots | \mathbf{v}^T]^T) = [\mathbf{v}^T | \dots | \mathbf{v}^T]^T$, which would not hold if we did not use sparse packing, such as when padding the remaining slots naively with zeros. Throughout this paper, whenever $n \equiv 0 \pmod{\dim \mathbf{v}}$, $\text{ct}(\mathbf{v})$ should be interpreted as a sparsely packed ciphertext.

2.3 Homomorphic Matrix Multiplication

At the heart of our SophOMR scheme is homomorphic matrix-vector multiplication (MatMul) [18, 32, 33, 38, 48]. The goal of MatMul is to multiply a plaintext matrix \mathbf{M} with an encryption of vector \mathbf{v} and obtain an encryption of $\mathbf{M}\mathbf{v}$. Here, we review an efficient MatMul algorithm that we use in SophOMR.

Diagonal Packing. The idea of *diagonal packing*, first introduced by Halevi and Shoup [32], is fundamental to MatMul algorithms. Given a matrix $\mathbf{M} \in \mathbb{Z}_p^{m \times k}$ and a vector $\mathbf{v} \in \mathbb{Z}_p^k$, let $\text{diag}_i(\mathbf{M})$ denote the i -th diagonal packing of \mathbf{M} (w.r.t. dimension n), which is defined as follows (for $0 \leq j < n$).

$$\text{diag}_i(\mathbf{M})[j] = \mathbf{M}[j \bmod m][(i+j) \bmod k]$$

The rationale behind diagonal packing is that it enables a concise representation of a matrix-vector multiplication with native HE operations, specifically Hadamard products and rotations, as follows (for $n = m = k$).

$$\mathbf{M}\mathbf{v} = \sum_{i=0}^{k-1} \text{diag}_i(\mathbf{M}) \odot \text{Rot}^i(\mathbf{v}) \quad (1)$$

BSGS-Style MatMul. Halevi and Shoup [33] introduced the baby-step-giant-step (BSGS) style optimization to reduce the number of homomorphic rotations required in a MatMul algorithm. The core idea is captured by the following equalities, where $k = \tilde{g} \cdot \tilde{b}$ and $\mathbf{m}_{\tilde{g}\tilde{b}+b}$ denotes $\text{Rot}^{-\tilde{g}\tilde{b}}(\text{diag}_{\tilde{g}\tilde{b}+b}(\mathbf{M}))$. Notice that, by reusing the $\text{Rot}^b(\mathbf{v})$ terms, we can reduce the number of rotations to $\tilde{g} + \tilde{b}$, rather than $k = \tilde{g} \cdot \tilde{b}$ as in Eq. 1.

$$\begin{aligned} \mathbf{M}\mathbf{v} &= \sum_{i=0}^{k-1} \text{diag}_i(\mathbf{M}) \odot \text{Rot}^i(\mathbf{v}) \\ &= \sum_{g=0}^{\tilde{g}-1} \sum_{b=0}^{\tilde{b}-1} \text{diag}_{\tilde{g}\tilde{b}+b}(\mathbf{M}) \odot \text{Rot}^{\tilde{g}\tilde{b}+b}(\mathbf{v}) \\ &= \sum_{g=0}^{\tilde{g}-1} \text{Rot}^{\tilde{g}\tilde{b}} \left(\sum_{b=0}^{\tilde{b}-1} \mathbf{m}_{\tilde{g}\tilde{b}+b} \odot \text{Rot}^b(\mathbf{v}) \right) \end{aligned}$$

Among several follow-ups on [33], we adopt the approach of PEGASUS [48], which exhibits the best performance in the case of non-square matrices. The original MatMul algorithm of PEGASUS is designed for the CKKS scheme [20] and its rotation structure. In this work, we use its straightforward adaptation to our BFV setting [13, 25]. Additionally, we extend the approach to handle cases where the matrix dimension exceeds the underlying HE ring dimension, requiring the encrypted vectors to span multiple ciphertexts.

To be precise, let n be the ring dimension of the underlying BFV scheme. Let us consider the MatMul of $k \times N$ matrix \mathbf{M} with a length- N vector \mathbf{v} . Here, we focus on the *wide* case,

Table 2: Cost Analysis for Alg. 1

#(Ptxt-Mult)	#(Rotation)	#(Rot. Key)
$s \cdot k$	$s \cdot \tilde{b} + \tilde{g} + \log(n/k)$	$\leq 2 + \log(n/k)$

where $N = s \cdot n$ for some positive integer s and $k < n$. For simplicity, we assume k is a power of two. Alg. 1 describes the BFV-adapted PEGASUS MatMul algorithm for the specified setting. We use the Horner-style [33] variant, which optimizes the number of required rotation keys. The costs for Alg. 1 are summarized in Tab. 2. Note that the number of rotations is minimized when $\tilde{g} \approx \sqrt{s \cdot k}$. For more detailed discussions, refer to PEGASUS [48] and references therein, such as [33].

Algorithm 1 MatMul (Wide)

Input: $\mathbf{M} \in \mathbb{Z}_p^{k \times N}$, $(\text{ct}(\mathbf{v}_i))_{i=0}^{s-1} \triangleright k < n, N = s \cdot n, \mathbf{v}_i \in \mathbb{Z}_p^n$,
 k : power-of-two

Step 0: Parameter Setting

$$[\mathbf{M}_0 | \dots | \mathbf{M}_{s-1}] \leftarrow \mathbf{M} \quad \text{and} \quad k = \tilde{g} \cdot \tilde{b}$$

Step 1: Plaintext Packing

for $(i, j) \in [s] \times [k]$ **do**
 $\mathbf{m}_{i,j} \leftarrow \text{Rot}_{\text{row}}^{-g\tilde{b}}(\text{diag}_j(\mathbf{M}_i))$ **for** $g = \lfloor j/\tilde{b} \rfloor$

Step 2: Baby-Step-Giant-Step (BSGS)

for $i = 0, \dots, s-1$ **do** \triangleright Baby-Step
 $\text{ct}_{i,0} \leftarrow \text{ct}(\mathbf{v}_i)$
for $b = 1, \dots, \tilde{b}-1$ **do**
 $\text{ct}_{i,b} \leftarrow \text{Rot}_{\text{row}}^1(\text{ct}_{i,b-1})$

for $g = \tilde{g}-1, \dots, 0$ **do** \triangleright Giant-Step

$\text{ct}_{\text{sum}} \leftarrow \sum_{i=0}^{s-1} \sum_{b=0}^{\tilde{b}-1} (\mathbf{m}_{i,g\tilde{b}+b} \odot \text{ct}_{i,b})$
if $g = \tilde{g}-1$ **then**
 $\text{ct}_{\text{out}} \leftarrow \text{ct}_{\text{sum}}$
else
 $\text{ct}_{\text{out}} \leftarrow \text{Rot}_{\text{row}}^{\tilde{b}}(\text{ct}_{\text{out}}) + \text{ct}_{\text{sum}}$

Step 3: Block Summation

for $j = 0, \dots, \log(n/k) - 2$ **do**
 $\text{ct}_{\text{out}} \leftarrow \text{ct}_{\text{out}} + \text{Rot}_{\text{row}}^{k \cdot 2^j}(\text{ct}_{\text{out}})$
 $\text{ct}_{\text{out}} \leftarrow \text{ct}_{\text{out}} + \text{Rot}_{\text{col}}^1(\text{ct}_{\text{out}})$

return ct_{out}

2.4 Homomorphic Compression

The goal of homomorphic compression [11, 21, 22, 26, 43] is to homomorphically compress an encryption of a *sparse* vector, which has only a few non-zero components, into an encryption of a shorter vector, which can later be efficiently decompressed after decryption. In this section, we provide a brief overview of the homomorphic compression scheme

from [21], which offers high efficiency due to its compatibility with homomorphic SIMD operations. For a more in-depth discussion, please refer to [21].

2.4.1 Index Compression

We first consider the task of compressing a sparse *binary* vector (a.k.a. index compression) as a stepping stone for compressing an arbitrary sparse vector (a.k.a. payload compression). The index compression of [21] leverages so-called *Newton's identity*, which provides conversions between *power sum polynomials* and *elementary symmetric polynomials*. Although the idea of using Newton's identity for index compression has been previously explored in the literature [22], the novelty of [21] lies in (1) recognizing that the power sum encoding can be interpreted as a matrix-vector multiplication and (2) implementing it with efficient SIMD-aware homomorphic algorithms for MatMul (e.g., Alg. 1).

More precisely, for $p > N$, consider a sparse binary vector $\mathbf{v} \in \mathbb{Z}_p^N$ with at most k number of non-zero entries and define a Vandermonde-like $k \times N$ matrix \mathbf{C} defined as follows. The homomorphic compression scheme can then be concisely described as in Alg. 2, with decompression outlined in Alg. 3.

$$\mathbf{C} := \begin{bmatrix} 1 & 2 & \dots & N \\ 1^2 & 2^2 & \dots & N^2 \\ \vdots & \vdots & \ddots & \vdots \\ 1^k & 2^k & \dots & N^k \end{bmatrix} \pmod{p} \quad (2)$$

Algorithm 2 Complx

Input: $\text{ct}(\mathbf{v}) \triangleright$ ciphertexts of $\mathbf{v} \in \mathbb{Z}_p^N$
 $\text{ct}(\mathbf{w}) \leftarrow \text{MatMul}(\mathbf{C}, \text{ct}(\mathbf{v})) \triangleright$ see Sec. 2.3 and Alg. 1
return $\text{ct}(\mathbf{w})$

Algorithm 3 Decomplx

Input: $\mathbf{w} = (w_1, \dots, w_k) \in \mathbb{Z}_p^k \triangleright$ decryption of $\text{ct}(\mathbf{w})$

Step 1: Build a polynomial from \mathbf{w} using Newton's identity.

$a_0 \leftarrow 1$
for $j = 1, \dots, k$ **do**
 $a_j \leftarrow j^{-1} \cdot \sum_{i=1}^j (-1)^{i-1} a_{j-i} \cdot w_i$
 $f_{\mathbf{v}}^k(X) \leftarrow \sum_{i=0}^k (-1)^i a_i X^{k-i}$

Step 2: Reconstruct non-zero indices from the polynomial.

while X divides $f_{\mathbf{v}}^k(X)$ **do**
 $f_{\mathbf{v}}^k(X) \leftarrow f_{\mathbf{v}}^k(X)/X$

$I_{\mathbf{v}} \leftarrow$ Find zeros of $f_{\mathbf{v}}^k(X) \triangleright$ by Cantor-Zassenhaus

return $I_{\mathbf{v}}$

2.4.2 Payload Compression

With index compression at hand, we can now compress an arbitrary sparse vector. Consider a sparse vector $\mathbf{d} \in \mathbb{Z}_p^N$ with at most k non-zero entries. To compress $\text{ct}(\mathbf{d})$, we first obtain an encryption of the binary vector \mathbf{v} , whose indices of non-zero entries match those in \mathbf{d} . For example, Fermat's little theorem can be employed for this purpose (as outlined in Alg. 4). However, in certain scenarios (e.g., private database query [21]), it may not be necessary to derive $\text{ct}(\mathbf{v})$ from $\text{ct}(\mathbf{d})$, as $\text{ct}(\mathbf{v})$ may be already available from other source. In fact, $\text{ct}(\mathbf{d})$ is often computed from $\text{ct}(\mathbf{v})$. Looking ahead, this is the case also for OMR.

Then, we can compress $\text{ct}(\mathbf{v})$ via Alg. 2 and use \mathbf{v} as a *hint* in the decompression phase. Note that any k distinct columns of the matrix \mathbf{C} are linearly independent. Therefore, we can also compress $\text{ct}(\mathbf{d})$ by performing a MatMul with \mathbf{C} and decompress it by solving a linear equation for a submatrix of \mathbf{C} . The homomorphic compression scheme is described in Alg. 4 and the decompression in Alg. 5.

Algorithm 4 Comp

Input: $\text{ct}(\mathbf{d})$ ▷ ciphertexts of $\mathbf{d} \in \mathbb{Z}_p^N$
 $\text{ct}(\mathbf{v}) \leftarrow \text{Eval}(\text{Power}_{p-1}, \text{ct}(\mathbf{d}))^5$ ▷ Fermat's little theorem
 $\text{ct}(\mathbf{w}) \leftarrow \text{Compldx}(\text{ct}(\mathbf{v}))$ ▷ see Alg. 2
 $\text{ct}(\mathbf{e}) \leftarrow \text{Eval}(\mathbf{C}, \text{ct}(\mathbf{d}))$
return $\text{ct}(\mathbf{e}), \text{ct}(\mathbf{w})$

Algorithm 5 Decomp

Input: $\mathbf{e}, \mathbf{w} \in \mathbb{Z}_p^k$ ▷ decryption of $\text{ct}(\mathbf{e}), \text{ct}(\mathbf{w})$
 $I_v \leftarrow \text{Decompldx}(\mathbf{w})$ ▷ see Alg. 3
 $\hat{\mathbf{C}} \leftarrow [\mathbf{C}_{i_1} | \dots | \mathbf{C}_{i_\ell}]$ ▷ $k \times \ell$ matrix whose j -th column is i_j -th column of \mathbf{C} , where $I_v = \{i_1, \dots, i_\ell\}$
 $\hat{\mathbf{d}} = (\hat{d}_1, \dots, \hat{d}_\ell) \leftarrow \text{Solve } \hat{\mathbf{C}}\mathbf{x} = \mathbf{e}$ for \mathbf{x}
return $\hat{\mathbf{d}}$

2.5 OMR and OMD

In this section, we review the formal definitions of Oblivious Message Retrieval (OMR) and Oblivious Message Detection (OMD) [43]. See also Sec. 1.2 and Fig. 1.

2.5.1 Oblivious Message Retrieval (OMR)

Definition 2 (Oblivious Message Retrieval). *An Oblivious Message Retrieval (OMR) scheme (w.r.t. payload space \mathcal{P})*

⁵ Power_ℓ denotes an arithmetic circuit over \mathbb{Z}_p such that $\text{Power}_\ell(\mathbf{d}) = (d_1^\ell, \dots, d_N^\ell)$, where $\mathbf{d} = (d_1, \dots, d_N) \in \mathbb{Z}_p^N$.

is a tuple of probabilistic algorithms (ParamGen, KeyGen, Signal, Digest, Decode) described below:

- $\text{pp} \leftarrow \text{ParamGen}(1^\lambda)$: The parameter generation algorithm takes a security parameter λ and returns a public parameter pp . The following algorithms implicitly take pp as an input.
- $(\text{sk}, \text{pk}_{\text{sig}}, \text{pk}_{\text{det}}) \leftarrow \text{KeyGen}(\text{pp})$: The key generation algorithm outputs a secret key sk , a signal key pk_{sig} , and a detection key pk_{det} .
- $\text{sig} \leftarrow \text{Signal}(\text{pk}_{\text{sig}})$: The signaling algorithm takes a signal key pk_{sig} and outputs a signal sig .

A bulletin board $\text{BB} = ((x_1, \text{sig}_1), \dots, (x_N, \text{sig}_N))$ is formed by several parties uploading a payload $x_i \in \mathcal{P}$ tagged with a signal sig_i with respect to the intended receiver.

- $\text{D} \leftarrow \text{Digest}(\text{BB}, \text{pk}_{\text{det}}, k)$: The digesting algorithm takes a bulletin board $\text{BB} = ((x_1, \text{sig}_1), \dots, (x_N, \text{sig}_N))$ for some size N , a detection key pk_{det} , and an upper bound k on the number of pertinent payloads addressed to the receiver. Then, it returns a digest D .
- $\text{M} \leftarrow \text{Decode}(\text{D}, \text{sk})$: The decoding algorithm takes a digest D and a secret key sk . Then, it returns a list of decoded payloads $\text{M} \in \mathcal{P}^*$.

Definition 3 (Security of OMR). Let $\text{pp} \leftarrow \text{ParamGen}(1^\lambda)$.

- **Correctness:** Let $(\text{sk}, \text{pk}_{\text{sig}}, \text{pk}_{\text{det}}) \leftarrow \text{KeyGen}(\text{pp})$. Let BB be a bulletin board of length N and $\tilde{\text{M}}$ be the list of pertinent payloads in BB addressed to pk_{sig} . For $k \geq |\tilde{\text{M}}|$, the following holds.

$$\Pr \left[\text{M} \neq \tilde{\text{M}} \mid \begin{array}{l} \text{D} \leftarrow \text{Digest}(\text{BB}, \text{pk}_{\text{det}}, k) \\ \text{M} \leftarrow \text{Decode}(\text{D}, \text{sk}) \end{array} \right] \leq \text{negl}(\lambda)$$

- **Receiver-Privacy:** Signals generated from different signal keys should be computationally indistinguishable to anyone who does not have access to the corresponding secret keys. To be precise, let $(\text{sk}, \text{pk}_{\text{sig}}, \text{pk}_{\text{det}})$ and $(\text{sk}', \text{pk}_{\text{sig}}', \text{pk}_{\text{det}}')$ be two set of keys generated from $\text{KeyGen}(\text{pp})$. Let $\text{sig} \leftarrow \text{Signal}(\text{pk}_{\text{sig}})$ and $\text{sig}' \leftarrow \text{Signal}(\text{pk}_{\text{sig}}')$. For any probabilistic polynomial time adversary \mathcal{A} , the following holds for $\text{aux} = (\text{pp}, \text{pk}_{\text{sig}}, \text{pk}_{\text{det}}, \text{pk}_{\text{sig}}', \text{pk}_{\text{det}}')$.

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

2.5.2 Oblivious Message Detection (OMD)

Oblivious Message Detection (OMD) is a variant of OMR where the receiver only retrieves pertinent *indices*, rather than *payloads*. Note that by combining an OMD scheme with a private information retrieval (PIR) scheme [23] one can obtain an OMR-like functionality at the cost of an additional round of communication: first, retrieve pertinent indices using OMD, and then use PIR to obtain the corresponding payloads. Below are formal definitions for OMD.

Definition 4 (Oblivious Message Detection). An Oblivious Message Detection (OMD) scheme is a tuple of probabilistic algorithms $(\text{ParamGen}, \text{KeyGen}, \text{Signal}, \text{Digest}, \text{Decode})$. Descriptions for ParamGen , KeyGen , Signal , and Digest are identical to those of OMR (Def. 2).

- $M \leftarrow \text{Decode}(D, \text{sk})$: The decoding algorithm takes a digest D and a secret key sk . Then, it returns a list of pertinent indices $M \in \mathbb{N}^*$.

Definition 5 (Security of OMD). Let $\text{pp} \leftarrow \text{ParamGen}(1^\lambda)$. The definition of receiver-privacy is identical to that of OMR (Def. 3).

- **Correctness:** Let $(\text{sk}, \text{pk}_{\text{sig}}, \text{pk}_{\text{det}}) \leftarrow \text{KeyGen}(\text{pp})$. Let BB be a bulletin board of length N and \tilde{M} be the list of pertinent indices addressed to pk_{sig} in BB . For $k \leq |\tilde{M}|$, the following holds.

$$\Pr \left[M \neq \tilde{M} \mid \begin{array}{l} D \leftarrow \text{Digest}(\text{BB}, \text{pk}_{\text{det}}, k) \\ M \leftarrow \text{Decode}(D, \text{sk}) \end{array} \right] \leq \text{negl}(\lambda)$$

3 Review: PerfOMR [46]

In this section, we provide a concise overview of the PerfOMR scheme [46]. First, in Sec. 3.1, we recall a building block for OMR, which we call *private signaling*. Then, we review each component of PerfOMR: ParamGen , KeyGen , Signal (Sec. 3.2), and Digest (Sec. 3.3 and 3.4). Specifically, the Digest algorithm is presented in a modular manner, split into a detection phase (Sec. 3.3) and a compression phase (Sec. 3.4).

3.1 RLWE-based Private Signaling Scheme

A Private Signaling (PS) scheme is a core building block of OMR, enabling a sender to generate a *signal* that only the intended receiver can detect. We begin by presenting a formal definition of private signaling (Sec. 3.1.1), followed by a construction based on the Ring Learning with Errors (RLWE) assumption [49, 56] (Sec. 3.1.2).

3.1.1 Private Signaling

Definition 6 (Private Signaling⁶). A Private Signaling (PS) scheme is a tuple of probabilistic polynomial time algorithms $(\text{ParamGen}, \text{KeyGen}, \text{Signal}, \text{Detect})$ described below:

- $\text{pp} \leftarrow \text{ParamGen}(1^\lambda)$: The parameter generation algorithm takes a security parameter λ and returns a public parameter pp . The following algorithms implicitly take pp as an input.
- $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\text{pp})$: The key generation algorithm returns a secret key sk and a public key pk .

⁶Our definition differs from that of [51] and aims to capture the framework they refer to as *naive approach* based on *key-private* PKE.

- $\text{sig} \leftarrow \text{Signal}(\text{pk})$: The signaling algorithm takes a public key pk and returns a signal sig .
- $0/1 \leftarrow \text{Detect}_{\text{sk}}(\text{sig})$: The detection algorithm takes a secret key sk and a signal sig . Then, it returns 1 if sig was generated from a public key corresponding to sk ; otherwise, it returns 0.

Definition 7 (Security of Private Signaling). Let $\text{pp} \leftarrow \text{ParamGen}(1^\lambda)$.

- **Completeness:** The following holds.

$$\Pr \left[\begin{array}{l} (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\text{pp}); \\ \text{sig} \leftarrow \text{Signal}(\text{pk}); \\ \text{Detect}_{\text{sk}}(\text{sig}) = 1 \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

- **Soundness:** The following holds.

$$\Pr \left[\begin{array}{l} (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\text{pp}); \\ (\text{sk}', \text{pk}') \leftarrow \text{KeyGen}(\text{pp}); \\ \text{sig} \leftarrow \text{Signal}(\text{pk}); \\ \text{Detect}_{\text{sk}'}(\text{sig}) = 1 \end{array} \right] \leq \text{negl}(\lambda)$$

- **Receiver-Privacy:** Signals generated from different public keys should be computationally indistinguishable to anyone who does not have access to the corresponding secret keys. To be precise, let (sk, pk) and (sk', pk') be two set of keys generated from $\text{KeyGen}(\text{pp})$. Let $\text{sig} \leftarrow \text{Signal}(\text{pk})$ and $\text{sig}' \leftarrow \text{Signal}(\text{pk}')$. For any probabilistic polynomial time adversary \mathcal{A} , the following holds for $\text{aux} = (\text{pp}, \text{pk}, \text{pk}')$.

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

3.1.2 Construction

We now present an overview of the private signaling scheme from PerfOMR [46] whose security is based on the hardness of the *Ring Learning with Errors* (RLWE) problem [49, 56]. For a more detailed discussion of the scheme, refer to [46].

Remark 2 (On Choice of LWE-based Scheme). Although we have a generic conversion from any PKE scheme with reasonable properties, private signaling schemes in this line of works [43, 46], are based on (Ring) LWE assumption. The primary reason is that the decryption circuit of LWE-based PKE is better suited to the computational model of HE than other cryptographic assumptions (e.g., group-based assumptions). Specifically, a large part of the decryption process is just a simple linear algebra over a relatively small modulus.

We first introduce a notation to ease the presentation. We denote and define the i -th negacyclic vector of $\mathbf{a} = (a_0, \dots, a_{n'-1}) \in \mathbb{Z}_{q'}^{n'}$ as follows, for $0 \leq i < n'$.

$$\text{Neg}^i(\mathbf{a}) = (a_i, \dots, a_0, -a_{n'-1}, \dots, -a_{i+1})$$

Let $a(X) = \sum_{i=0}^{n-1} a_i X^i$ for \mathbf{a} and similarly for $b(X)$ and $\mathbf{b} = (b_0, \dots, b_{n'-1})$. Notice that $\langle \text{Neg}^i(\mathbf{a}), \mathbf{b} \rangle$ equals to the i -th coefficient of $a(X) \cdot b(X)$ over $\mathcal{R}'_{q'} = \mathbb{Z}_{q'}[X]/(X^{n'} + 1)$. With this in mind, the following construction can be seen, at a high level, as an RLWE-based PKE scheme specifically designed for encrypting zeroes. It can be viewed as a *ring-optimized* variant of the PVW scheme [55], or as a *truncated* version of the LPR scheme [50].

Construction 1. *The private signaling scheme PS = (PS.ParamGen, PS.KeyGen, PS.Signal, PS.Detect) is defined as follows.*

- $\text{pp} \leftarrow \text{PS.ParamGen}(1^\lambda)$: *The parameter generation algorithm returns $\text{pp} = (n', q', \ell, r, \chi_s, \chi_e)$, consisting of polynomial ring dimension n' , modulus q' , repetition parameter ℓ , range parameter r , secret key distribution χ_s over $\mathbb{Z}_{q'}^{n'}$, and error distribution χ_e over $\mathbb{Z}_{q'}$.*
- $(\text{sk}, \text{pk}) \leftarrow \text{PS.KeyGen}(\text{pp})$: *The key generation algorithm first samples a secret key $\text{sk} \leftarrow \chi_s$. Then, it samples $\boldsymbol{\alpha} \leftarrow \mathbb{Z}_{q'}^{n'}$ and compute $\boldsymbol{\beta} = (\langle \text{Neg}^i(\boldsymbol{\alpha}), \text{sk} \rangle_{i=0}^{n'-1} + \mathbf{e} \in \mathbb{Z}_{q'}^{n'}$ with noise $\mathbf{e} \leftarrow \chi_e^{n'}$. Finally, the algorithm returns (sk, pk) , where $\text{pk} = (\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{Z}_{q'}^{n'} \times \mathbb{Z}_{q'}^{n'}$.*
- $\text{sig} \leftarrow \text{PS.Signal}(\text{pk})$: *The signaling algorithm takes a public key $\text{pk} = (\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{Z}_{q'}^{n'} \times \mathbb{Z}_{q'}^{n'}$ and samples noises $\mathbf{e}_0 \leftarrow \chi_s$, $\mathbf{e}_1 \leftarrow \chi_e^{n'}$, and $\mathbf{e}_2 \leftarrow \chi_e^\ell$. Then, it computes $\mathbf{a} = (\langle \text{Neg}^i(\boldsymbol{\alpha}), \mathbf{e}_0 \rangle_{i=0}^{n'-1} + \mathbf{e}_1$ and $\mathbf{b} = (\langle \text{Neg}^i(\boldsymbol{\beta}), \mathbf{e}_0 \rangle_{i=0}^{\ell-1} + \mathbf{e}_2$. Finally, the algorithm returns $\text{sig} = (\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_{q'}^{n'} \times \mathbb{Z}_{q'}^\ell$.*
- $0/1 \leftarrow \text{PS.Detect}_{\text{sk}}(\text{sig})$: *The detection algorithm takes $\text{sig} = (\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_{q'}^{n'} \times \mathbb{Z}_{q'}^\ell$ and computes $\mathbf{d} = \mathbf{b} - (\langle \text{Neg}^i(\mathbf{a}), \text{sk} \rangle_{i=0}^{\ell-1} \in \mathbb{Z}_{q'}^\ell$. Then, it computes c_i 's by checking $\mathbf{d} = (d_1, \dots, d_\ell)$ with the range parameter r , as follows.*

$$c_i = \begin{cases} 0, & \text{if } d_i \in [-r, r] \\ 1, & \text{otherwise} \end{cases}$$

Finally, the algorithm returns $\prod_{i=1}^\ell (1 - c_i) \in \{0, 1\}$.

Remark 3 (Security of Construction 1). *We sketch the security of Construction 1 (Def. 7). For details, refer to [46].*

- **Completeness:** *The range parameter r can be set to guarantee completeness, following standard analysis in LWE-based encryptions to guarantee correctness.*
- **Soundness:** *The repetition parameter ℓ can be set to guarantee soundness as follows: By the hardness of RLWE over $\mathcal{R}'_{q'}$, $\text{sig} \leftarrow \text{PS.Signal}(\text{pk})$ is indistinguishable from a uniform random sample, in the absence of access to the secret key sk . Therefore, the following holds when sk' is sampled independently of sk .*

$$\Pr[\text{PS.Detect}_{\text{sk}'}(\text{sig}) = 1] \leq \left(\frac{2r+1}{q'} \right)^\ell + \text{negl}(\lambda)$$

- **Receiver-Privacy:** *The receiver-privacy directly follows from the hardness of RLWE over $\mathcal{R}'_{q'}$.*

3.2 Setup

Given homomorphic encryption scheme BFV (Sec. 2.2.1) and private signaling scheme PS (Sec. 3.1.2), (ParamGen, KeyGen, Signal) of PerfOMR can be described as follows.

- $\text{pp} \leftarrow \text{ParamGen}(1^\lambda)$: Run $\text{pp}_{\text{PS}} \leftarrow \text{PS.ParamGen}(1^\lambda)$ and return $\text{pp} = (\text{pp}_{\text{PS}}, 1^\lambda)$.
- $(\text{sk}, \text{pk}_{\text{sig}}, \text{pk}_{\text{det}}) \leftarrow \text{KeyGen}(\text{pp})$: First, run $(\text{sk}_{\text{PS}}, \text{pk}_{\text{PS}}) \leftarrow \text{PS.KeyGen}(\text{pp}_{\text{PS}})$ and $(\text{sk}_{\text{BFV}}, \text{pk}_{\text{BFV}}) \leftarrow \text{BFV.KeyGen}(1^\lambda)$. Then, compute $\text{ct}_{\text{sk}} \leftarrow \text{BFV.Enc}(\text{sk}_{\text{PS}})$. Return $\text{sk} = (\text{sk}_{\text{PS}}, \text{sk}_{\text{BFV}})$, $\text{pk}_{\text{sig}} = \text{pk}_{\text{PS}}$, and $\text{pk}_{\text{det}} = (\text{pk}_{\text{BFV}}, \text{ct}_{\text{sk}})$.
- $\text{sig} \leftarrow \text{Signal}(\text{pk}_{\text{sig}})$: Return $\text{sig} \leftarrow \text{PS.Signal}(\text{pk}_{\text{PS}})$.

3.3 Detection

For clarity, we describe the Digest algorithm of PerfOMR in a modular manner, divided into (1) a *detection* phase and (2) a *compression* phase. In this section, we begin by illustrating Detect (Alg. 6) for the detection phase, followed by an explanation of the compression phase in Sec. 3.4.

The goal of Detect is to homomorphically compute the pertinent indices from the bulletin board BB using the detection key pk_{det} . The output of Detect is ciphertexts encrypting a sparse *binary* vector, where the ones indicate the pertinent indices, referred to as the *Pertinency Vector*, PV.⁷ At a high level, Detect can be viewed as a *batched homomorphic detection* of signals, consisting of two main steps: (i) *Affine Transform* and (ii) *Range Check*.

Affine Transform. The goal of Affine Transform (Step 1 of Alg. 6) is to compute *linear* parts of multiple $\text{PS.Detect}_{\text{sk}_{\text{PS}}}(\text{sig}_i)$ in parallel. More precisely, we want to compute $\mathbf{d} = \mathbf{b}_i - (\langle \text{Neg}^j(\mathbf{a}_i), \text{sk} \rangle_{j=0}^{\ell-1} \in \mathbb{Z}_{q'}^\ell$ where $\text{sig}_i = (\mathbf{a}_i, \mathbf{b}_i) \in \mathbb{Z}_{q'}^{n'} \times \mathbb{Z}_{q'}^\ell$ for $1 \leq i \leq N$. Notice that we can stack up the inner products $\langle \text{Neg}^j(\mathbf{a}_i), \text{sk} \rangle$ into a matrix-vector multiplication. In this way, we can leverage homomorphic MatMul algorithms (Sec. 2.3) with respect to ct_{sk} . We will explore this in greater detail in Sec. 4.4.

Range Check. The goal of Range Check (Step 2 of Alg. 6) is to homomorphically compute the remaining *non-linear* parts of multiple $\text{PS.Detect}_{\text{sk}_{\text{PS}}}(\text{sig}_i)$. More precisely, we want to compute $\prod_{i=1}^\ell (1 - c_i) \in \{0, 1\}$ where r is the range parameter and

$$c_i = \begin{cases} 0, & \text{if } d_i \in [-r, r] \\ 1, & \text{otherwise,} \end{cases}$$

⁷Prior works [43, 46] refer to this procedure as *ClueToPackedPV*.

Algorithm 6 Detect

Input: $\text{pk}_{\text{det}}, \text{BB} = (x_i, \text{sig}_i)_{i=1}^N$ $\triangleright N = s \cdot n$
Step 0: Signal Preprocessing $\triangleright \text{sig}_i = (\mathbf{a}_i, \mathbf{b}_i)$
 for $j = 0, \dots, \ell - 1$ **do**
 $\mathbf{A}_j \leftarrow [\text{Neg}^j(\mathbf{a}_1)] \cdots [\text{Neg}^j(\mathbf{a}_N)]^T$ $\triangleright \mathbf{A}_j \in \mathbb{Z}_q^{N \times n'}$
 $\boldsymbol{\beta}_j \leftarrow (b_{1,j}, \dots, b_{N,j})$ $\triangleright \mathbf{b}_i = (b_{i,j})_{j=0}^{\ell-1}$
Step 1: Affine Transform
 for $j = 0, \dots, \ell - 1$ **do**
 $(\text{ct}_{i,j})_{i=0}^{s-1} \leftarrow \boldsymbol{\beta}_j - \text{MatMul}(\mathbf{A}_j, \text{ct}_{\text{sk}})$ \triangleright Sec. 4.4
Step 2: Range Check $\triangleright g(X) = X \cdot \prod_{i=1}^r (X^2 - i^2),$
 $h(X) = X^{p-1}$
 for $i = 0, \dots, s - 1$ **do**
 for $j = 0, \dots, \ell - 1$ **do**
 $\text{ct}_{i,j} \leftarrow \text{BFV.Eval}_{\text{pk}_{\text{BFV}}}(g(X), \text{ct}_{i,j})$
 $\text{ct}_{i,j} \leftarrow \text{BFV.Eval}_{\text{pk}_{\text{BFV}}}(h(X), \text{ct}_{i,j})$
 $\text{ct}_i \leftarrow \prod_{j=0}^{\ell-1} (1 - \text{ct}_{i,j})$
return $(\text{ct}_i)_{i=0}^{s-1}$

for $\mathbf{d} = (d_1, \dots, d_\ell)$ which is the output from Affine Transform. The computation of c_i 's can be done homomorphically by evaluating the following function on d_i 's.

$$f(X) = \begin{cases} 0, & \text{if } -r \leq X \leq r \\ 1, & \text{otherwise} \end{cases} \quad (3)$$

In PerfOMR [46], $f(X)$ is homomorphically evaluated by first computing $g(X) = \prod_{i=0}^r (X^2 - i^2)$ and then $h(X) = X^{q-1}$, rather than directly evaluating the polynomial interpolation of $f(X)$ as done in [43]. Notice that $g(X)$ maps the input to 0 if and only if it lies within the range $[-r, r]$, and $h(X)$ maps non-zero inputs to 1 and zero to 0 by Fermat's little theorem. By decomposing the evaluation of f into g and h , PerfOMR significantly reduced the number of homomorphic multiplications required, while only slightly increasing the circuit depth, resulting in an overall improvement in computational cost.

3.4 Compression

In the compression phase, pertinent payloads are compressed into a digest using the encrypted pertinency vector PV. Roughly speaking, PerfOMR compresses both PV and pertinent payloads; during decoding, PV is decoded first and then used as a *hint* to decode the compressed payloads. Below is a brief overview of PerfOMR's compression phase. For a more detailed discussion, refer to [43, 46].

3.4.1 Unpacking the Pertinency Vector

In PerfOMR, before compressing the pertinency vector PV, each slot of PV is *unpacked* into separate ciphertexts. This

process, referred to as PVUnpack in previous works [43, 46], is necessary for them because they compress PV by *hashing* its components into *buckets*. (For more details, see Sec. 3.4.2.) PVUnpack incorporates so-called SlotToCoeff [30, 47] and *oblivious expansion* [3] as subroutines, both of which are quite costly. As a result, PVUnpack becomes a bottleneck in the entire digesting process.

3.4.2 Hashing-based Index Compression

PerfOMR compresses the unpacked pertinency vector PVUnpack(PV) = $\{\text{ct}_i\}_{i=1}^N$ via *hashing*. Each index $1 \leq i \leq N$ is randomly assigned to a *bucket* Y_j for $1 \leq j \leq m$. The pertinency vector is compressed by computing $\text{Ctr}_j \leftarrow \sum_{i \in Y_j} \text{ct}_i$ and $\text{Acc}_j \leftarrow \sum_{i \in Y_j} i \cdot \text{ct}_i$. Notice that Ctr_j indicates the number of pertinent indices that were hashed into Y_j , revealing whether a collision occurred. If Ctr_j encrypts the value 1, the corresponding value in Acc_j is the pertinent index; otherwise, a collision has occurred, and decoding will fail. To prevent such failures, PerfOMR sets the number of buckets m to be sufficiently larger than k , the upper bound on the number of pertinent indices, and repeats the process multiple times.

3.4.3 Payload Compression with Random Linear Codes

PerfOMR then compresses pertinent payloads using random linear codes. The core idea is to compress the vector of pertinent payloads (i.e., the payloads masked by the pertinency vector) by multiplying it with a random $K \times N$ matrix. By setting K sufficiently larger than k —the upper bound on the number of pertinent indices—it becomes possible to solve the system of equations for the pertinent payloads via Gaussian elimination with high probability, given the pertinent indices.

3.4.4 Bundling: Runtime/Digest-Size Trade-off

To reduce the computational overhead of PVUnpack, PerfOMR processes multiple indices simultaneously via *bundling*. Before applying PVUnpack, PerfOMR combines v indices of the pertinency vector by homomorphically adding them into a single ciphertext. While this reduces the runtime by a factor of v , it introduces a trade-off: now the receiver only knows that a bundle contains a pertinent index, requiring retrieval of all v indices and their corresponding payloads. This approach balances the reduction in PVUnpack runtime against an increase in the final digest size.

3.5 Optimizations

Modulus-Switching. *Modulus-Switching* [14] is a technique of switching the ciphertext modulus of a given HE ciphertext into a smaller one. While modulus-switching is not strictly necessary in the BFV scheme [13, 25] (unlike the BGV scheme [14]), it can be highly beneficial for performance optimizations. In BFV, choosing a large enough

ciphertext modulus is crucial to support the desired computation, as the modulus determines the *computational budget* of a ciphertext. However, a larger modulus increases the ciphertext size and degrades the overall performance of homomorphic operations. Modulus-switching addresses this issue by reducing the modulus as the budget is consumed, preserving the remaining budget. PerfOMR leverages modulus-switching repeatedly to keep the ciphertext modulus as small as possible without compromising the budget, enhancing performance.

Rotation Key Management. The BFV scheme can natively support all types of rotation. However, each requires its own rotation key and storing all these keys can be quite demanding in practice. Therefore, trade-offs between storage and computational time are often advantageous. It is also important to note that the size of rotation keys scales with the size of the ciphertext modulus. However, if the rotation keys will be used only after modulus-switching, generating them in a smaller modulus suffices. For this reason, PerfOMR reduces the number of rotation keys more aggressively for the ones used at higher ciphertext modulus levels. The approach taken by PerfOMR will be further detailed in Sec. 4.4.

4 Our Scheme: SophOMR

In this section, we present our new OMR scheme, SophOMR. While we largely follow the approach of PerfOMR [46], the key distinction lies in the compression phase (Sec. 3.4), where we extend and apply the SIMD-aware homomorphic compression scheme of [21] (Sec. 2.4). A detailed explanation of our homomorphic compression mechanism is provided in Sec. 4.1. Additionally, we propose several optimizations: ring-switching for digest-size reduction (Sec. 4.2), hybrid key-switching optimization for key-size reduction (Sec. 4.3), and BSGS-style Affine Transform step (Sec. 3.3) for run-time reduction (Sec. 4.4).

4.1 SIMD-Aware Digest Compression

As discussed in Sec. 3.4, PerfOMR [46] cannot fully exploit the homomorphic SIMD structure (Sec. 2.2.1) during the compression step due to its hashing-based approach (Sec. 3.4.2). As a result, it requires the costly PVUnpack process (Sec. 3.4.1), which unpacks each slot of the packed PV ciphertext into separate ciphertexts, creating a bottleneck for the entire Digest.

In this regard, we adopt a *SIMD-aware* homomorphic compression scheme from [21] (Sec. 2.4), which removes the need for the heavy PVUnpack step, significantly improving the runtime of the compression step.⁸ However, [21] only considered payloads of the size of a single HE slot, which is

⁸Another advantage of our compression scheme is that, unlike PerfOMR’s, it is *deterministic* and does not introduce any additional failure probability.

unrealistic for many scenarios. For example, in prior works on OMR [43, 46], each payload is assumed to be 612 bytes, the typical size of Zcash transactions [35], meaning that the payload occupies multiple HE slots. Thus, we need to extend the approach of [21] to accommodate these typical OMR settings.

Naive Approach. To be precise, let each payload occupy t number of HE slots (e.g., $t \approx 300$ for typical HE parameters in the example of Zcash transactions). Let i -th slot payload vector PL_i denote the length- N vector consisting of the i -th slot of each payload. Directly applying the homomorphic compression of [21] (Sec. 2.4) to the compression step of OMR (Sec. 3.4), one would first compress the encrypted pertinency vector PV obtained from the detection step (Sec. 3.3). This involves homomorphically computing $\mathbf{C} \cdot PV$ via MatMul (Alg. 1), where \mathbf{C} is defined in Eq. 2. Then, one would compress *masked* payloads $PL_i \odot PV$ by computing $\mathbf{C} \cdot (PL_i \odot PV)$ again with MatMul, for $1 \leq i \leq t$. This naive approach would require a total $(t + 1)$ number of calls to MatMul with respect to the $k \times N$ matrix \mathbf{C} . While this method is already fairly efficient, it can be further optimized as described below.

Stacking Matrices. First, we note a simple yet powerful fact that enables our optimization. Let $\text{diag}(PL_i)$ denote the $N \times N$ diagonal matrix whose diagonal is PL_i and let \mathbf{D}_i denote the $k \times N$ matrix $\mathbf{C} \cdot \text{diag}(PL_i)$. Then, the following holds.

$$\mathbf{C} \cdot (PL_i \odot PV) = \mathbf{D}_i \cdot PV$$

This allows the detector to precompute $\mathbf{D}_i = \mathbf{C} \cdot \text{diag}(PL_i)$ and then compute $\mathbf{D}_i \cdot PV$ homomorphically. Note that this preprocessing is almost free, as both \mathbf{C} and $\text{diag}(PL_i)$ are *plaintexts* with regard to the underlying HE scheme. Furthermore, this approach consumes less *computational budget* (Sec. 3.5) as the homomorphic computation of $PL_i \odot PV$ is no longer performed before MatMul.

But much more importantly, this precomputation allows us to *stack* multiple MatMul into a single MatMul as follows.

$$[\mathbf{C}^T | \mathbf{D}_1^T | \dots | \mathbf{D}_t^T]^T \cdot PV$$

To see why this is useful at all, recall that the number of homomorphic rotations required in $k \times N$ MatMul is $O(\sqrt{k})$. (See Tab. 2.) The stacking method significantly reduces the total number of rotations required for compression: the naive approach, which involves $t + 1$ calls to $k \times N$ MatMul, requires $O((t + 1) \cdot \sqrt{k})$ rotations. In contrast, the stacking approach, which involves a single call to $((t + 1) \cdot k) \times N$ MatMul, only requires $O(\sqrt{(t + 1) \cdot k})$ rotations. This results in roughly a $\sqrt{t + 1}$ reduction in the total number of rotations.

Implementing MatMul. We use MatMul described in Alg. 1 for digest compression. In particular, we employ the Horner-style [33] MatMul, which significantly reduces the

number of required rotation keys. Although this method is incompatible with the so-called *hoisting* technique [33], which optimizes runtime, we prioritize Horner-style MatMul. We believe that in scenarios where a detector handles multiple clients, minimizing the key size per client is much more important than the runtime improvements offered by hoisting.

Remark 4 (Cost Comparison for Digest Compression). *A detailed comparison of the computational cost enabled by the simpler digest compression in SophOMR over the prior PerfOMR can be found in Tab.1 and Section 1.3.*

Remark 5. *Recall that our homomorphic compression scheme requires $p > N$ to ensure the correctness of the index compression step (Sec. 2.4.1). Although this condition may appear restrictive, the same assumption is implicit in prior OMR schemes [43, 46], which also encode indices up to N into the plaintext modulus p . Earlier works address this via p -ary index decomposition. While SophOMR is compatible with this approach, we do not need it in our parameters.*

4.2 Further Compression with Ring-Switching

Similar to modulus-switching (Sec. 3.5), there is a technique called *ring-switching* [14, 28, 29], which switches the ring degree n of a given HE ciphertext into a smaller one. This can be done while preserving the same security level after some modulus-switching.⁹ In our case, ring-switching is useful for further reducing the digest size. Specifically, the ciphertext size scales linearly with the ring dimension n , which corresponds to the number of slots in the BFV scheme (Sec. 2.2.1). This implies that if the final result occupies fewer slots than n , the unused slots are being wasted. For instance, after homomorphic compression (Sec. 4.1), only $(t + 1) \cdot k$ slots are occupied, which is typically smaller than n under standard OMR parameters. Ring-switching addresses this inefficiency by reducing the ring dimension and making the digest more compact.¹⁰ In SophOMR, ring-switching is employed as the final step of Digest to further compress the digest. We note that ring-switching is a relatively inexpensive procedure, whose cost is dominated by a single key-switching operation. For further details, please refer to [29].

4.3 Key Size Reduction through Hybrid Key-Switching Optimization

The *key-switching* technique is a fundamental element of homomorphic multiplications and rotations. There are two main approaches to key-switching. The first, known as the

⁹For the same level of security, a larger ring dimension is required when using a larger ciphertext modulus.

¹⁰Ring-switching is especially effective in the context of OMD (Sec. 2.5.2), where the digest contains only the pertinent indices, excluding payloads, requiring significantly fewer slots. See Sec. 5.

Table 3: Cost Analysis for Affine Transform

	#(Ptxt-Mul)	#(Rot)	#(Rot. Key)
PerfOMR	$s \cdot k$	k	1
Our Approach	$s \cdot k$	$\tilde{b} + s \cdot \tilde{g}$	2

bit-decomposition approach [8, 15] optimizes the *computational budget* (Sec. 3.5). In contrast, the *special-modulus* approach [19, 31] prioritize key-size optimization. In practice, a hybrid approach that combines these two methods is often most effective [31, 34, 39]. PerfOMR suffers from a large key size due to relying on one extreme of this hybrid design. In contrast, SophOMR balances the two extremes, significantly reducing the key size.

4.4 Affine Transform with BSGS-MatMul

We optimize the runtime of Affine Transform step in the Detect procedure (Sec. 3.3) by adopting the BSGS-style MatMul algorithm, which uses two rotation keys. Thanks to hybrid key-switching (Sec. 4.3), this optimization is achieved while maintaining a smaller key-size compared to PerfOMR’s approach of using a single rotation key. (See Sec. 5.)

Recall that the Affine Transform step involves homomorphic matrix multiplications between *tall* plaintext matrices \mathbf{A}_j ’s and the ciphertext vector ct_{sk} (Sec. 3.3, Alg. 6). For clarity, in this section, we consider the goal of homomorphically multiplying a tall $N \times k$ plaintext matrix \mathbf{M} (representing \mathbf{A}_j) with a ciphertext $\text{ct}(\mathbf{v})$ (representing ct_{sk}) that encrypts a vector $\mathbf{v} \in \mathbb{Z}_p^k$ via sparse packing (Sec. 2.2.2).

4.4.1 PerfOMR’s Approach: MatMul with Single Key

Note that the Affine Transform (Sec. 3.3) is the very first step of digesting and occurs before any modulus-switching (Sec. 3.5). As a result, the rotation keys used in this step must be generated at the largest modulus, making them relatively large. In this regard, PerfOMR [46] aggressively minimizes the number of rotation keys required during the Affine Transform, ultimately using only one. The approach skips BSGS-style optimization (Sec. 2.3) and applies Eq. 1 directly. The cost of PerfOMR is summarized in Tab. 3.

4.4.2 Our Approach: MatMul with Two Keys

In SophOMR, thanks to hybrid key-switching (Sec. 4.3), the size of the rotation key generated at the largest modulus is substantially smaller than in PerfOMR. As a result, even with the BSGS-style optimization and the use of two keys, SophOMR’s total key-size remains smaller than PerfOMR’s. The adoption of BSGS reduces the number of homomorphic rotations from $O(k)$ to $O(\sqrt{k})$, enhancing runtime performance. The BSGS-MatMul algorithm used for Affine trans-

form in SophOMR is simply the tall matrix version of Alg. 1. The costs for our approach are summarized in Tab. 3. Note that the number of rotations is minimized when $\tilde{b} \approx \sqrt{s \cdot k}$.

4.5 Security of SophOMR

We sketch the security of SophOMR with respect to Def. 3.

- **Correctness:** The correctness directly follows directly from the correctness of the underlying BFV scheme (Sec. 2.2.1, Def. 1) and the completeness/soundness of the underlying Private Signaling scheme (Construction 1, Def. 7). It is worth noting that our digest compression (Sec. 4.1) is deterministic and guarantees correct output.
- **Receiver-Privacy:** The receiver-privacy directly follows from the semantic security of the underlying BFV scheme (Sec. 2.2.1, Def. 1) and the receiver-privacy of the underlying Private Signaling scheme (Construction 1, Def. 7). More precisely, let $pp \leftarrow \text{ParamGen}(1^\lambda)$ and sample two set of keys (sk, pk_{sig}, pk_{det}) and $(sk', pk_{sig}', pk_{det}')$ from $\text{KeyGen}(pp)$ (see Sec. 3.2). Two signals $sig \leftarrow \text{Signal}(pk_{sig})$ and $sig' \leftarrow \text{Signal}(pk_{sig}')$ are indistinguishable, even given the auxiliary information $aux = (pp, pk_{sig}, pk_{det}, pk_{sig}', pk_{det}')$, due to the receiver-privacy of the Private Signaling scheme. This holds because $pk_{det} = (pk_{BFV}, \text{BFV.Enc}(sk_{PS}))$, which is the only component dependent on the secret key sk_{PS} of the Private Signaling scheme, reveals no information about sk_{PS} by the semantic security of the BFV scheme.

4.6 OMD Variant

Our SophOMR scheme can be readily adapted into an Oblivious Message Detection (OMD) scheme (Sec. 2.5.2). Since OMD does not require payload compression, it is sufficient to compress only the encrypted pertinency vector PV. This is done by homomorphically computing $C \cdot PV$ via MatMul (Alg. 1), where C is defined in Eq. 2, without stacking the matrices as described in Sec. 4.1. This leads to a significant speedup in the compression phase (Sec. 4.1). See Sec. 5 for concrete performance results.

5 Evaluation

Methodology. To evaluate the concrete performance of our SophOMR scheme, we implement it in a C++ library¹¹ using the OpenFHE library [7], with the NTL library [60] to implement Decode. We then compare its performance against the implementation [44] of PerfOMR [46], which uses PALISADE library [54] for the private signaling scheme (Sec. 3.1), the SEAL library [58] for the BFV scheme (Sec. 2.2.1). Both implementations leverage the HEXL library [12] to accelerate

HE libraries. All experiments for SophOMR and PerfOMR were conducted single-threaded on a Google Compute Cloud `n2-standard-8` instance (Intel Ice Lake). Each experiment was repeated five times, and we report the average.

Parameters. We choose parameters that uphold the security claims of PerfOMR’s parameter sets. All bit security is measured using the latest version of the Lattice Estimator¹² [2] available at the time of writing (Commit ID: 5c25455). Detailed parameters are provided in Appendix A. Note that SophOMR uses a larger BFV ring dimension of $n = 2^{16}$, compared to $n = 2^{15}$ of PerfOMR, to meet 128-bit security (Tab. 5). We also employ a larger parameter for the underlying private signaling scheme (Sec. 3.1) that offers 128-bit security and 30-bit completeness to uphold the security claim of PerfOMR’s parameters. The parameters provided in PerfOMR only achieve 94-bit security and 18-bit completeness, leading to easily observable errors (Tab. 7).

Experiments. Experiments are conducted with a total number of payloads $N = 2^{16}$ and 2^{19} . Following PerfOMR, the number of pertinent payloads is set to $k = 50$, and each payload is assumed to be 612 bytes, which corresponds to the typical size of Zcash transactions [35]. For $N = 2^{16}$, we present the results in Tab. 4. Additionally, Fig. 2 illustrate the breakdown of each component contributing to the Digest runtime: Preprocessing (Rmk. 8), Affine Transform (Sec. 4.4), Range Check (Sec. 3.3), and Compress (Sec. 4.1). Refer to Appendix B, for experimental results for $N = 2^{19}$.

Remark 6 (On Choice of N). *While PerfOMR conducted experiments with $N = 2^{19}, 2^{21}, 2^{23}$, we focus on $N = 2^{16}, 2^{19}$, considering $N = 2^{16}$ as the primary setting. Our reasoning is that dividing N into smaller sets can be advantageous for latency, provided it maintains similar throughput and keeps communication costs reasonable. For example, PerfOMR uses $N = 2^{19}$ to model Bitcoin-scale applications,¹³ but a client may find it more useful to receive digests for 2^{16} payloads every 3 hours than a single digest for 2^{19} every 24 hours. This issue is even more relevant in private messaging scenarios.*

Remark 7 (Scaling with N). *Digest runtime for both PerfOMR and SophOMR scales approximately linearly with the number of payloads N , as shown in experiments with $N = 2^{16}$ (Fig. 2, Tab. 4) and $N = 2^{19}$ (Fig. 4, Tab. 8). By contrast, the Decode runtime and the resulting digest size are independent of N , depending only on the number of pertinent messages k .*

Remark 8 (Preprocessing). *In Fig. 2 and 4, the blue mesh represents the time that PerfOMR deliberately excluded from*


¹¹<https://github.com/keewoolee/SophOMR>

¹²<https://github.com/malb/lattice-estimator>

¹³There are roughly 2^{19} transactions per day on Bitcoin: https://ycharts.com/indicators/bitcoin_transactions_per_day

Table 4: Performance of OMR/OMD Schemes ($N = 2^{16}$, $k = 50$, single-thread)

	Digest (s)	Decode (ms)	Detection Key (MB) ¹⁴	Digest (KB) ¹⁵
PerfOMR ($v = 2$)	541 (std: 2.9)	14.6 (std: 0.5)	171	567
SophOMR	157 (std: 0.4)	14.0 (std: 0.1)	114	263
PerfOMD ($v = 2$)	493 (std: 1.4)	6.8 (std: 0.4)	171	284
SophOMD	107 (std: 0.3)	5.5 (std: 0.1)	142	132


 Figure 2: Runtime Breakdown for OMR/OMD Schemes ($N = 2^{16}$, $k = 50$)

their end-to-end timing results to account for the preprocessing potential of a significant portion of SlotToCoeff (Sec. 3.4.1). However, this preprocessing requires a substantial amount of RAM. If such an amount of RAM is available, SophOMR can also be accelerated by preprocessing MatMul in Compress (Sec. 4.1). For a comparison in favor of PerfOMR, preprocessing time is excluded from Tab. 4 and 8.

Experimental Results for $N = 2^{16}$. As shown in Tab. 4, SophOMR provides multiple improvements over PerfOMR. SophOMR achieves a $3.4\times$ speedup in Digest runtime while maintaining a reasonably small Decode runtime. The detection key size is reduced by $1.5\times$ through hybrid key-switching optimization (Sec. 4.3). Furthermore, the digest size is reduced by $2.2\times$, thanks to our compact homomorphic compression scheme (Sec. 4.1) and ring-switching (Sec. 4.2) using Param II in Tab. 6; see also Rmk. 10. These improvements are achieved despite SophOMR using larger parameters than PerfOMR to meet security levels that PerfOMR does not satisfy (See Appendix A).

As illustrated in Fig. 2, the improvement in Digest runtime primarily stems from our SIMD-aware homomorphic compression scheme (Sec. 4.1). Specifically, our compression achieves $7.5\times$ speedup compared to PerfOMR, even when excluding the runtime for preprocessing (Rmk. 8). The Affine Transform step also achieves considerable speedup thanks to BSGS-style MatMul (Sec. 4.4). On the other hand, the Range Check step (Sec. 3.3) shows a slight increase in runtime, attributed to our use of larger parameters to meet security levels that PerfOMR does not satisfy (See Appendix A).

Remark 9 (On Choice of v). The bundling parameter v in PerfOMR represents how many payloads are concatenated

and processed simultaneously (See Sec. 3.4.4). We compare only with PerfOMR using $v = 2$ for $N = 2^{16}$ because (i) setting $v = 2$ incurs nearly no additional runtime or digest size costs compared to $v = 1$ due to PerfOMR’s specific utilization BFV slots, and (ii) setting $v > 2$ does not improve runtime further due to the fixed number of BFV slots ($n = 2^{15}$).

Remark 10 (Digest Size). PerfOMR optimizes the digest size by setting the BFV ciphertext modulus at the lowest level to be as small as possible in the SEAL library (Sec. 3.5). However, this optimization is not applied in our SophOMR implementation due to limitations in the OpenFHE library: (i) it does not officially allow the base-level ciphertext modulus to be set independently of the RNS limbs, and (ii) its ciphertext serialization does not further reduce when the ciphertext modulus is below 64-bit. If we apply the same optimization as in PerfOMR, SophOMR’s digest size can be reduced by approximately an additional factor of $2\times$.

Remark 11 (Multiplicative Depth). Recall that Digest consists of three phases: Affine Transform, Range Check, and Compress (Sec. 3.3 and 3.4). PerfOMR consumes multiplicative depths of 1, 22, and 2 for these phases, respectively, while SophOMR uses 1, 27, and 1. Although the Range Check phase consumes substantial multiplicative depth, its circuit requires only a few key-switching operations per level compared to the Affine Transform and Compress phases. As a result, it is not the bottleneck in either scheme.

OMD Variants. We also present the performance of SophOMD, the OMD variant of SophOMR (Sec. 2.5.2, Sec. 4.6). As there is no existing OMD implementation corresponding to PerfOMR, we modified the PerfOMR implementation to run only the relevant components and refer to this version as PerfOMD. The OMD experiments use the same

¹⁴We account for the compression of randomness into the hash seed.

¹⁵Our digest size can be further reduced, as explained in Rmk. 10.




Figure 3: Runtime of SophOMR/SophOMD under varying parameters ($N = 2^{16}$, single-thread)

BFV parameters as those in the OMR experiments, except that the smaller ring-switching parameter set Param I from Tab. 6 is used.

As shown in Tab. 4, SophOMD also provides multiple improvements over PerfOMD, similar to the OMR setting. Specifically, SophOMD achieves a $4.6\times$ speedup in Digest runtime. As illustrated in Fig. 2, this improvement primarily stems from our SIMD-aware homomorphic compression scheme (Sec. 4.1), allowing compression in SophOMD to complete in only one second – a $479\times$ speedup compared to PerfOMD. Notice that the performance gain is even greater than in the OMR setting, as we only compress the indices and not the payloads. On the other hand, the detector key size has slightly increased (though still smaller than PerfOMR) due to the use of additional rotation keys from the reduced number of rows in MatMul (Tab. 2). The bundling technique (Sec. 3.4.4) can still be applied to PerfOMD; however, achieving OMR-like functionality (Sec. 2.5.2) requires performing PIR (Private Information Retrieval) on the *bundles*, causing PIR performance to depend on the bundling parameter v .

Additional Experiments with Varying k and t . To evaluate their impact on performance, we conduct additional experiments that vary the number of pertinent payloads k (Fig. 3a and 3b) and the number of slots per payload t (Fig. 3c). We use different ring-switching parameters, as presented in Tab. 6 in Appendix A, depending on the required number of slots for the digest. Specifically, SophOMR uses Param III for $(k, t) = (128, 250)$ and $(50, 512)$, Param II for $(k, t) = (64, 250)$ and $(50, 256)$, while SophOMD and all remaining SophOMR experiments use Param I.

Regarding Digest, we observe that the runtime of Compress scales linearly with both k and t . This indicates that the performance bottleneck of SophOMR’s Digest Compression lies at plaintext-ciphertext multiplications rather than homomorphic multiplications (See Tab. 1). We note that the runtimes for RangeCheck and Affine are independent of k and t . In contrast, the Decode runtime is affected by multiple

factors. Specifically, Decode consists of a decryption step that is (quasi-)linear in the digest size, primarily determined by the ring-switching dimension, followed by a decompression step that scales linearly with t and super-linearly with k .

Acknowledgments

Yongdong Yeo was financially supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (RS-2025-02217656). We would like to thank Zeyu Liu for the discussion on PerfOMR and its implementation.

Ethical Considerations

This research focuses on improving the efficiency and scalability of Oblivious Message Retrieval (OMR), with the goal of enhancing the user experience in anonymous messaging services. As such, OMR inherently involves ethical considerations related to anonymity and its potential societal impact.

Anonymity can protect users—particularly in surveillance-heavy or oppressive environments—but it may also enable harmful or illicit communication. From a stakeholder perspective, potential risks arise for platform operators, who may face challenges in moderating malicious content, as well as for individuals or communities that may be harmed by those who exploit anonymity. At the same time, enhanced privacy and anonymity can benefit other stakeholders, including platforms seeking to serve end users without collecting or retaining unnecessary personal information. In a broader context, prior work [1, 40, 41] highlights the need to evaluate privacy-preserving technologies, such as anonymous messaging systems and privacy-oriented blockchains, within frameworks of regulatory compliance and responsible governance, emphasizing the importance of balancing user anonymity with legitimate oversight.

Our efficiency improvements lower the barrier to deploying

OMR in practice, which amplifies the importance of these ethical considerations. To support responsible adoption, we provide a transparent and reproducible system design and evaluation methodology. All experiments were conducted in controlled environments and did not involve any live systems.

Open Science

In line with open science principles, we are committed to making our research fully transparent and accessible. To support reproducibility and independent verification, we include comprehensive details of our experimental setup and results in the paper, along with the accompanying code available at: <https://doi.org/10.5281/zenodo.17958466>.

References

- [1] Oluwafemi Akanfe, Diane Lawong, and H. Raghav Rao. Blockchain technology and privacy regulation: Reviewing frictions and synthesizing opportunities. *International Journal of Information Management*, 76:102753, 2024.
- [2] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *J. Math. Cryptol.*, 9(3):169–203, 2015.
- [3] Sebastian Angel, Hao Chen, Kim Laine, and Srinath T. V. Setty. PIR with compressed queries and amortized query processing. In *2018 IEEE Symposium on Security and Privacy*, pages 962–979. IEEE Computer Society Press, May 2018.
- [4] Sebastian Angel and Srinath T. V. Setty. Unobservable communication over fully untrusted infrastructure. In Kimberly Keeton and Timothy Roscoe, editors, *USENIX OSDI 2016*. USENIX Association, November 2016.
- [5] Aztec. <https://aztec.network/>.
- [6] Adam Back, Ulf Möller, and Anton Stiglic. Traffic analysis attacks and trade-offs in anonymity providing systems. In Ira S. Moskowitz, editor, *Information Hiding, 4th International Workshop, IHW 2001, Proceedings*, volume 2137 of *Lecture Notes in Computer Science*, pages 245–257. Springer, April 2001.
- [7] Ahmad Al Badawi, Andreea Alexandru, Jack Bates, Flavio Bergamaschi, David Bruce Cousins, Saroja Erabelli, Nicholas Genise, Shai Halevi, Hamish Hunt, Andrey Kim, Yongwoo Lee, Zeyu Liu, Daniele Micciancio, Carlo Pascoe, Yuriy Polyakov, Ian Quah, Saraswathy R. V., Kurt Rohloff, Jonathan Saylor, Dmitriy Suponitsky, Matthew Triplett, Vinod Vaikuntanathan, and Vincent Zucca. OpenFHE: Open-source fully homomorphic encryption library. *Cryptology ePrint Archive*, Paper 2022/915, 2022. <https://eprint.iacr.org/2022/915>.
- [8] Jean-Claude Bajard, Julien Eynard, M. Anwar Hasan, and Vincent Zucca. A full RNS variant of FV like somewhat homomorphic encryption schemes. In Roberto Avanzi and Howard M. Heys, editors, *SAC 2016*, volume 10532 of *LNCS*, pages 423–442. Springer, Cham, August 2016.
- [9] Gabrielle Beck, Julia Len, Ian Miers, and Matthew Green. Fuzzy message detection. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1507–1528. ACM Press, November 2021.
- [10] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014.
- [11] Alexander Bienstock, Sarvar Patel, Joon Young Seo, and Kevin Yeo. Batch PIR and labeled PSI with oblivious ciphertext compression. In Davide Balzarotti and Wenyan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.
- [12] Fabian Boemer, Sejun Kim, Gelila Seifu, Fillipe DM de Souza, Vinodh Gopal, et al. Intel HEXL (release 1.2). <https://github.com/intel/hexl>, 2021.
- [13] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 868–886. Springer, Berlin, Heidelberg, August 2012.
- [14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012.
- [15] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 505–524. Springer, Berlin, Heidelberg, August 2011.
- [16] Vitalik Buterin. Exploring fully homomorphic encryption. <https://vitalik.eth.limo/general/2020/07/20/homomorphic.html>, 2020.
- [17] Vitalik Buterin. Open problem: improving stealth addresses. <https://ethresear.ch/t/open-problem-improving-stealth-addresses/7438>, 2020.

- [18] Hao Chen, Wei Dai, Miran Kim, and Yongsoo Song. Efficient multi-key homomorphic encryption with packed ciphertexts with application to oblivious neural network inference. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 395–412. ACM Press, November 2019.
- [19] Jung Hee Cheon, Kyoohyung Han, Andrey Kim, Miran Kim, and Yongsoo Song. A full RNS variant of approximate homomorphic encryption. In Carlos Cid and Michael J. Jacobson, Jr., editors, *SAC 2018*, volume 11349 of *LNCS*, pages 347–368. Springer, Cham, August 2019.
- [20] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. Homomorphic encryption for arithmetic of approximate numbers. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 409–437. Springer, Cham, December 2017.
- [21] Jung Hee Cheon, Keewoo Lee, Jai Hyun Park, and Yongdong Yeo. SIMD-aware homomorphic compression and application to private database query, 2024. <https://arxiv.org/abs/2408.17063>.
- [22] Seung Geol Choi, Dana Dachman-Soled, S. Dov Gordon, Linsheng Liu, and Arkady Yerukhimovich. Compressed oblivious encoding for homomorphically encrypted search. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2277–2291. ACM Press, November 2021.
- [23] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *36th FOCS*, pages 41–50. IEEE Computer Society Press, October 1995.
- [24] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *2015 IEEE Symposium on Security and Privacy*, pages 321–338. IEEE Computer Society Press, May 2015.
- [25] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012.
- [26] Nils Fleischhacker, Kasper Green Larsen, and Mark Simkin. How to compress encrypted data. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 551–577. Springer, Cham, April 2023.
- [27] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- [28] Craig Gentry, Shai Halevi, Chris Peikert, and Nigel P. Smart. Ring switching in BGV-style homomorphic encryption. In Ivan Visconti and Roberto De Prisco, editors, *SCN 12*, volume 7485 of *LNCS*, pages 19–37. Springer, Berlin, Heidelberg, September 2012.
- [29] Craig Gentry, Shai Halevi, Chris Peikert, and Nigel P. Smart. Field switching in bgv-style homomorphic encryption. *J. Comput. Secur.*, 21(5):663–684, 2013.
- [30] Craig Gentry, Shai Halevi, and Nigel P. Smart. Better bootstrapping in fully homomorphic encryption. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 1–16. Springer, Berlin, Heidelberg, May 2012.
- [31] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the AES circuit. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 850–867. Springer, Berlin, Heidelberg, August 2012.
- [32] Shai Halevi and Victor Shoup. Algorithms in HELib. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 554–571. Springer, Berlin, Heidelberg, August 2014.
- [33] Shai Halevi and Victor Shoup. Faster homomorphic linear transformations in HELib. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 93–120. Springer, Cham, August 2018.
- [34] Kyoohyung Han and Dohyeong Ki. Better bootstrapping for approximate homomorphic encryption. In Stanislaw Jarecki, editor, *CT-RSA 2020*, volume 12006 of *LNCS*, pages 364–390. Springer, Cham, February 2020.
- [35] Daira-Emma Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash Protocol Specification. <https://zips.z.cash/protocol/protocol.pdf>.
- [36] Sashidhar Jakkamsetti, Zeyu Liu, and Varun Madathil. Scalable private signaling. In *38th IEEE Computer Security Foundations Symposium, CSF 2025*, pages 143–158. IEEE, June 2025.
- [37] Yanxue Jia, Varun Madathil, and Aniket Kate. HomeRun: High-efficiency oblivious message retrieval, unrestricted. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024*, pages 2012–2026. ACM Press, October 2024.
- [38] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. GAZELLE: A low latency framework for secure neural network inference. In William Enck and Adrienne Porter Felt, editors, *USENIX Security 2018*, pages 1651–1669. USENIX Association, August 2018.

- [39] Andrey Kim, Yuriy Polyakov, and Vincent Zucca. Revisiting homomorphic encryption schemes for finite fields. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 608–639. Springer, Cham, December 2021.
- [40] Rob Kling, Ya-Ching Lee, Al Teich, and Mark S. Frankel. Assessing anonymous communication on the internet: Policy deliberations. *The Information Society*, 15(2):79–90, 1999.
- [41] Yannan Li, Willy Susilo, Guomin Yang, Yong Yu, Xiaojiang Du, Dongxi Liu, and Nadra Guizani. Toward privacy and regulation in blockchain-based cryptocurrencies. *IEEE Network*, 33(5):111–117, 2019.
- [42] Zeyu Liu, Katerina Sotiraki, Eran Tromer, and Yunhao Wang. Snake-eye resistant PKE from LWE for oblivious message retrieval and robust encryption. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 126–156. Springer, Cham, May 2025.
- [43] Zeyu Liu and Eran Tromer. Oblivious message retrieval. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 753–783. Springer, Cham, August 2022.
- [44] Zeyu Liu, Eran Tromer, and Yunhao Wang. PerfOMR: proof of concept C++ implementation for OMR (Oblivious Message Retrieval) with Reduced Communication and Computation. <https://github.com/ObliviousMessageRetrieval/ObliviousMessageRetrieval/tree/perfomr>.
- [45] Zeyu Liu, Eran Tromer, and Yunhao Wang. Group oblivious message retrieval. In *2024 IEEE Symposium on Security and Privacy*, pages 4367–4385. IEEE Computer Society Press, May 2024.
- [46] Zeyu Liu, Eran Tromer, and Yunhao Wang. PerfOMR: Oblivious message retrieval with reduced communication and computation. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.
- [47] Zeyu Liu and Yunhao Wang. Amortized functional bootstrapping in less than 7 ms, with $\tilde{O}(1)$ polynomial multiplications. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VI*, volume 14443 of *LNCS*, pages 101–132. Springer, Singapore, December 2023.
- [48] Wen-jie Lu, Zhicong Huang, Cheng Hong, Yiping Ma, and Hunter Qu. PEGASUS: Bridging polynomial and non-polynomial evaluations in homomorphic encryption. In *2021 IEEE Symposium on Security and Privacy*, pages 1057–1073. IEEE Computer Society Press, May 2021.
- [49] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010.
- [50] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013.
- [51] Varun Madathil, Alessandra Scafuro, István András Seres, Omer Shlomovits, and Denis Varlakov. Private signaling. In Kevin R. B. Butler and Kurt Thomas, editors, *USENIX Security 2022*, pages 3309–3326. USENIX Association, August 2022.
- [52] Nick Mathewson and Roger Dingledine. Practical traffic analysis: Extending and resisting statistical disclosure. In David M. Martin Jr. and Andrei Serjantov, editors, *Privacy Enhancing Technologies, 4th International Workshop, PET 2004, Revised Selected Papers*, volume 3424 of *Lecture Notes in Computer Science*, pages 17–34. Springer, May 2004.
- [53] Shen Noether. Ring signature confidential transactions for monero. Cryptology ePrint Archive, Report 2015/1098, 2015.
- [54] PALISADE Lattice Cryptography Library (release 1.11.3). <https://palisade-crypto.org/>, May 2021.
- [55] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571. Springer, Berlin, Heidelberg, August 2008.
- [56] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
- [57] Ronald L. Rivest, Len Adleman, and Michael L. Demtousos. On data banks and privacy homomorphisms. *Foundations of Secure Computation, Academia Press*, pages 169–179, 1978.
- [58] Microsoft SEAL (release 4.1). <https://github.com/Microsoft/SEAL>, January 2023. Microsoft Research, Redmond, WA.
- [59] István András Seres, Balázs Pejó, and Péter Burcsi. The effect of false positives: Why fuzzy message detection leads to fuzzy privacy guarantees? In Ittay Eyal and Juan A. Garay, editors, *FC 2022*, volume 13411 of *LNCS*, pages 123–148. Springer, Cham, May 2022.

Table 5: Parameters for BFV

	n	p	$\log Q$	$\log PQ$	λ
PerfOMR	2^{15}	65537	857	917	120.9
SophOMR	2^{16}	786433	1140	1740	128.6

- [60] Victor Shoup. NTL: A Library for doing Number Theory. <https://libnt1.org/>.
- [61] Jelle van den Hooff, David Lazar, Matei Zaharia, and Nickolai Zeldovich. Vuvuzela: scalable private messaging resistant to traffic analysis. In Ethan L. Miller and Steven Hand, editors, *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP 2015*, pages 137–152. ACM, October 2015.
- [62] Toni Wahrstätter, Matt Solomon, Ben DiFrancesco, and Vitalik Buterin. ERC-5564: Stealth Addresses. Ethereum Improvement Proposals, no. 5564, August 2022. Available: <https://eips.ethereum.org/EIPs/eip-5564>.
- [63] David Isaac Wolinsky, Henry Corrigan-Gibbs, Bryan Ford, and Aaron Johnson. Dissent in numbers: Making strong anonymity scale. In Chandu Thekkath and Amin Vahdat, editors, *USENIX OSDI 2012*. USENIX Association, October 2012.

A Parameters

The BFV parameters used in our experiments are presented in Tab. 5. For SophOMR and SophOMD, we also present the BFV parameters for ring-switching (Sec. 4.2) in Tab. 6. The private signaling parameters are presented in Tab. 7.

BFV Parameters. In Tab. 5 and Tab. 6, $\log PQ$ denotes the largest modulus used in the scheme, including *special modulus* used in the hybrid key-switching (Sec. 4.3), which impacts the security level. For all BFV parameter sets, the error distribution is set to a discrete Gaussian with standard deviation $\sigma = 3.19$. In Tab. 5, notice that SophOMR uses a larger BFV ring dimension of $n = 2^{16}$, compared to $n = 2^{15}$ of PerfOMR, to meet 128-bit security. In this setup, a larger plaintext modulus p is required to satisfy $p = 1 \pmod{2n}$ to fully leverage the homomorphic SIMD structure.

Private Signaling Parameters. In Tab. 7, ρ and ν denote the false-positive and false-negative security parameters, which are associated with soundness and completeness, respectively (Def. 7). Note that our key and signal sizes are slightly larger than PerfOMR due to our larger BFV parameter, which is to meet 128-bit security.

Table 6: BFV Parameters for Ring-Switching

	n	$\log Q$	$\log PQ$	λ
Param I	2^{13}	120	180	154.4
Param II	2^{14}	120	240	244.6
Param III	2^{15}	120	240	536.1

B Experimental Results for $N = 2^{19}$

In this section, we present our experimental results for $N = 2^{19}$. We compare the performance of SophOMR/SophOMD with PerfOMR/PerfOMD for $\nu = 2$ and $\nu = 16$, representing the two extremes in the trade-off between digest size and runtime. (See Sec. 3.4.4 and Rmk. 9.) As shown in Tab. 8 and Fig. 4, SophOMR/SophOMD provides multiple improvements over PerfOMR/PerfOMD for $N = 2^{19}$.

OMR Schemes. Compared to PerfOMR with $\nu = 2$, SophOMR achieves a $3.4\times$ speedup in Digest runtime and a $2.2\times$ reduction in digest size. In contrast, compared to PerfOMR with $\nu = 16$, SophOMR achieves $1.2\times$ speedup in Digest runtime and a $10.8\times$ reduction in digest size. Additionally, note that the Decode runtime for PerfOMR increases when $\nu = 16$.

OMD Schemes. For the OMD variants, SophOMD achieves a $4.5\times$ and $1.4\times$ speedup in Digest runtime compared to PerfOMD with $\nu = 2$ and $\nu = 16$, respectively. The other trends are observed similarly as in the $N = 2^{16}$ setting. Specifically, our SIMD-aware homomorphic compression scheme (Sec. 4.1) enables SophOMD to complete compression for $N = 2^{19}$ in just 2.6 seconds—yielding a $1033\times$ speedup over PerfOMD with $\nu = 2$ and a $145\times$ speedup with $\nu = 16$.

¹⁶We note that the range parameter r is chosen to be as small as possible while satisfying the false-negative bound ν . More precisely, r is set to satisfy $1 - \text{erf}(r/(\sigma\sqrt{2} \cdot (2h+1))) \leq 2^{-\nu}/\ell$. For details, we refer to Sec. 5.1 of [46].

¹⁷We account for the compression of randomness into the hash seed.

Table 7: Parameters for PS

	n'	q'	σ	h	ℓ	r^{16}	λ	ρ	v	PS Key (KB) ¹⁷	Signal (KB)
PerfOMR	1024	65537	0.5	32	2	19	94.0	21.4	17.6	2.2	2.2
SophOMR	1024	786433	0.5	80	2	40	128.4	26.5	30.7	2.6	2.6

Table 8: Performance of OMR/OMD Schemes ($N = 2^{19}$, $k = 50$)

	Digest (s)	Decode (ms)	Detection Key (MB)	Digest (KB)
PerfOMR ($v = 2$)	3704 (std: 40.0)	15.0 (std: 0.0)	171	568
PerfOMR ($v = 16$)	1368 (std: 3.5)	72.0 (std: 0.7)	171	2840
SophOMR	1098 (std: 1.7)	12.0 (std: 4.7)	114	263
PerfOMD ($v = 2$)	3328 (std: 7.4)	7.0 (std: 0.0)	171	284
PerfOMD ($v = 16$)	1028 (std: 7.1)	10.0 (std: 0.0)	171	568
SophOMD	745 (std: 6.8)	5.5 (std: 0.2)	142	132




Figure 4: Runtime Breakdown for OMR/OMD Schemes ($N = 2^{19}$, $k = 50$, single-thread)