

Group Oblivious Message Retrieval

Zeyu Liu¹, Eran Tromer², Yunhao Wang²

¹Yale University

²Columbia University

April 13, 2023

Abstract

Anonymous message delivery, as in private communication and privacy-preserving blockchain applications, ought to protect recipient metadata: a message should not be inadvertently linkable to its destination. But in this case, how can messages be delivered to each recipient, without every recipient scanning all the messages? Recent work constructed Oblivious Message Retrieval (OMR) protocols that outsource this job to untrusted servers in a privacy-preserving manner.

We consider the case of group messaging, where each message may have multiple recipients (e.g., in a group chat or blockchain transaction). A direct use of prior OMR protocols in the group setting increases the servers' work linearly in the group size, rendering it prohibitively costly for large groups.

We thus devise new protocols where the servers' cost grows very slowly with the group size, while recipients' cost is low and independent of the group size. Our approach uses Fully Homomorphic Encryption and other lattice-based techniques, building on and improving on prior work. The efficient handling of groups is attained by encoding multiple recipient-specific clues into a single polynomial or multilinear function that can be efficiently evaluated under FHE, and via preprocessing and amortization techniques.

We formally study several variants of Group Oblivious Message Retrieval (GOMR), and describe corresponding GOMR protocols. Our implementation and benchmarks show, for parameters of interest, cost reductions of orders of magnitude compared to prior schemes. For example, the servers' cost is \sim \$3.36 per million messages scanned, where each message may address up to 15 recipients.

Contents

1	Introduction	4
1.1	Our Contribution	5
1.2	Related Work	6
1.2.1	Private Retrieval	6
1.2.2	Multi-Recipient Encryption and Broadcast Encryption	6
2	Overview	6
2.1	Model Overview	7
2.2	Summary and Comparisons	7
2.3	Main Techniques	9
3	Defining Group Oblivious Message Retrieval	10
3.1	Ad-hoc Group OMR	10
3.2	Fixed Group OMR	12
4	Preliminaries	14
4.1	Notation	14
4.2	PVW Encryption	15
4.3	Fully Homomorphic Encryption	15
5	Optimization to Prior Work	16
5.1	The Original OMR Construction	16
5.2	Optimizing PVUnpack	17
5.3	Optimizing OMRp2	18
6	Ad-hoc Group OMR	19
6.1	Trivial Solution	19
6.2	Improved Ad-hoc Group OMR	21
6.3	Extension to Distinct Payloads	25
7	Fixed Group OMR	26
7.1	Multi-Recipient Encryption	26
7.2	Naive Solutions	28
7.3	Efficient Key-private MRE	28
7.4	Applying MRE to FGOMR	31
7.5	Applying FGOMR to AGOMR	32
8	DoS Resistance	34
8.1	Ad-hoc Group OMR	34
8.2	Fixed Group OMR	36
9	Performance Evaluation	38
9.1	Methodology	38
9.2	Evaluation Results	39
10	Applications	40
10.1	Secure Group Messaging	40
10.2	Private Blockchains	40
	Acknowledgements	42

References	43
A Strengthened Privacy Definition of FGOMR	45
B Sparse Random Linear Coding (SRLC)	45

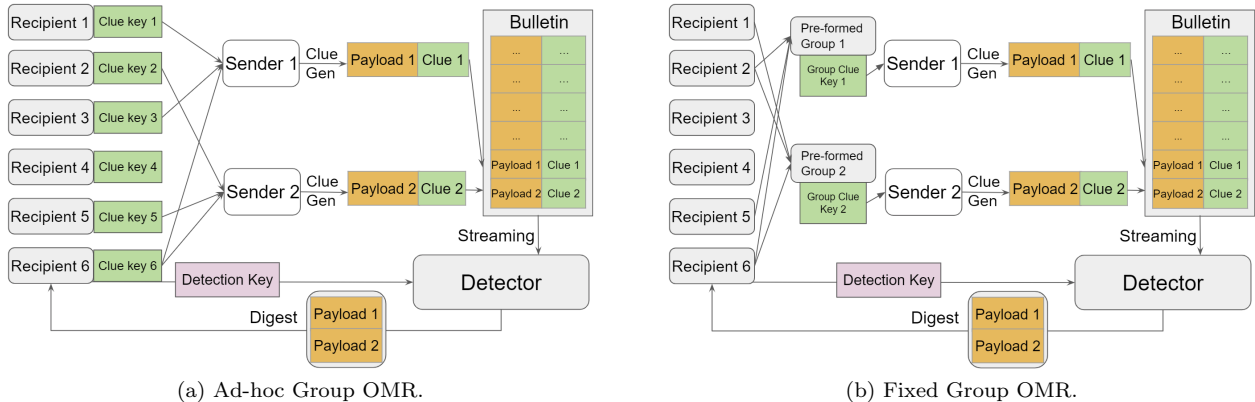


Figure 1: The high-level illustration of AGOMR and FGOMR. For AGOMR, the sender obtains the individual clue keys to form a group message. For FGOMR, individuals form a group clue key that the sender obtains.

1 Introduction

The protection of message contents in messaging applications is well studied, and end-to-end encryption is nowadays widely practiced. Protecting the metadata, such as sender and recipient identity, is likewise crucial in anonymous message delivery systems [WCGFJ12, CGBM15, Lun18, BEM⁺17, BLMG21], especially when messages are posted on a publicly-visible blockchain [Noe15, BSCG⁺14, HBHW, BCG⁺20].

Recipient privacy is especially difficult to achieve at scale, because of seemingly contradictory requirements. On one hand, no one should be able to discern the intended recipient of a message (other than that recipient, and the sender). On the other hand, each recipient wishes to retrieve just the messages addressed to them, without scanning all messages in existence for potentially-pertinent ones. There is thus a need to outsource the detection task to a server (or a *detector*) in a privacy-preserving way.

Fuzzy Message Detection [BLMG21] was the first to address this, proposing a decoy-based approach in which the server sees the set of messages pertinent to the recipient buried among many additional randomly-chosen messages (which is a weak security notion [SPB22]). Two follow-up works [MSS⁺22, LT22] improved this to entirely hide the set of pertinent messages, and in [LT22], furthermore ensure unlinkability of keys and resistance to Denial-of-Service or spamming attacks.

Those works focused on the case that a message is sent to a single recipient. What if the sender wants to send a message to multiple recipients, as in group messaging, mailing lists, or blockchain transactions with multiple parties?

Specifically, we consider the following *Group Oblivious Message Retrieval (GOMR)* setting, generalizing [LT22]: A sender sends a message to (up to) G recipients, attaching a *clue* generated using the recipients’ *clue keys*. The clue indicates to which recipients this message is addressed. The sender then publishes the message (along with the clue) onto a public *bullet board*. A *detector*, which is an untrusted third-party server, is enlisted by a recipient: given the recipient’s *detection key*, the detector processes the bulletin board and derives a *digest* that contains the pertinent messages, in encrypted form, by an oblivious computation that checks clues against the detection key. The detector sends back the digest to the recipient, who can then extract the messages via their *secret key*.

A trivial solution is to utilize prior OMR schemes [MSS⁺22, LT22], and have each message’s sender attach a separate clue for each of the G recipients, inducing a corresponding linear increase in the detector’s work — which is impractical for large groups.

In this paper, we show that much more efficient GOMR schemes exist. We study these in two models, which differ in how groups are formed (see Fig. 1), motivated by different applications:

The first flavor is the *Ad-hoc GOMR (AGOMR)*, which allows the senders to send messages to a group of

recipients chosen arbitrarily. This suits cases such as messaging protocols (e.g., WhatsApp Broadcast Lists) that let a message be addressed to any set of recipients chosen on the fly, or blockchains where transactions may have many recipients chosen arbitrarily.

The second flavor is *Fixed GOMR (FGOMR)*, where groups are pre-formed by their members and then addressed collectively. This suits applications with a notion of persistent groups, such as mailing lists or group chats. It also suits blockchains applications in which transactions need to be visible to a set of parties in addition to the recipient (e.g., auditors or jurisdictional law enforcement). The FGOMR setting is a special case of AGOMR, where having pre-formed groups FGOMR allows for more efficient constructions, and a stronger Denial-of-Service property (two honest recipients cannot be spammed jointly if they did not agree to join the same group).

These two flavors can be further extended to include different payloads to different recipients. For example, in the case of Zcash transactions [HBHW], a recipient of funds transaction could retrieve just the output notes they can spend, instead of the whole transaction (which may include additional notes).

1.1 Our Contribution

Definitions of GOMR. We formally define the two aforementioned models of Group Oblivious Message Retrieval, with fixed groups (FGOMR) and ad-hoc groups (AGOMR).

Our definitions capture natural notions of correctness and privacy for both notions. Moreover, our definitions include the DoS Resistance and detection-key-unlinkability notions from prior work [LT22] and extend them to the group setting.

Main GOMR constructions. We provide constructions of AGOMR and FGOMR that are far more efficient than applying prior works. Both constructions achieve the strongest functionalities, including the DoS-resistance and detection-key-unlinkability properties introduced in [LT22] (adapted to the group setting).

Our schemes are based on (leveled) homomorphic encryption and achieve privacy guarantees proven under a standard lattice hardness assumption. Our approach encodes multiple recipient-specific clues into a single polynomial or multilinear function that can be efficiently evaluated under FHE. We utilize preprocessing and amortization of several computational steps and tune homomorphic operation scheduling.

Additional GOMR constructions. In addition to our two main constructions above, we construct four variants, serving as either a stepping stone to our main constructions or an extension of our main constructions. Compared to our main constructions, most variants sacrifice some properties for efficiency; thus may serve better for other applications. In more detail, if we only accept honestly selected groups instead of arbitrarily formed ones, we can improve the runtime both asymptotically and concretely. Similarly, if we do not require detection-key unlinkability, our schemes can achieve even better detector runtime.

Extension to multi-payload GOMR. We show how to extend our main AGOMR construction to take different payloads each addressed to a different recipient at a small cost. A similar technique can be applied to all other GOMR constructions.

Improving the original OMR scheme. We also improve the single-recipient OMR scheme [LT22] by optimizing procedures and tightening the parameter analysis, resulting in 2.5x detector runtime speedup and 8x shrinkage in digest size.

Lattice-based Key-private Multi-Recipient Encryption scheme. As a component of independent interest, we construct the first lattice-based key-private Multi-Recipient Encryption (KP-MRE) scheme. It also serves as a building block for our main FGOMR construction.

Implementation and evaluation. We implemented our schemes as an open-source C++ library [?] and measured their concrete performance for a variety of parameters in comparison to prior work. With parameters of interest, our schemes can be two to three orders of magnitude faster than the baseline scheme. The FGOMR schemes show better efficiency and scalability than the AGOMR schemes.

Concretely, to retrieve half a million messages each addressing 15 recipients, the total cost of only about \sim \$1.68, each message takes \sim 0.018 second.

1.2 Related Work

1.2.1 Private Retrieval

Oblivious Message Retrieval. OMR [LT22] addresses the recipient-privacy problem for the case of a single recipient. We extend this to the group setting. Our privacy, key-unlinkability, and DoS resistance properties are all adapted and generalized from [LT22]. For the ciphertext compression component, an alternative technique is introduced in [FLS22].

Fuzzy Message Detection. FMD [BLMG21] uses a decoy-based privacy notion for single recipients. This is a weaker privacy guarantee [Lew21, SPB21], and does not provide DoS resistance and key unlinkability.

Private Signaling. PS [MSS⁺22] has the same base privacy guarantee as OMR (i.e., a message cannot be linked to a recipient). However, their constructions rely upon strong environmental assumptions (trusted hardware and two communicating but non-colluding servers, respectively), and their security notion does not provide DoS resistance or key unlinkability. Furthermore, their work does not directly deal with retrieval but just detection. Moreover, as in OMR, their definition is intended to have a message addressed to a single recipient.

PIR. Other related problems are *Private Information Retrieval (PIR)* [CGKS95] and its variant *Keyword PIR* [CGN98]; and in particular, since we retrieve multiple messages, the most related primitive is the variant called *multi-query (keyword) PIR* or *batch (keyword) PIR*. As in OMR, our setting differs in that recipients do not know the indices or labels of messages pertinent to them; rather, the clues are randomized (which is necessary for unlinkability) and require nontrivial computation (rather than simple comparison) to detect. Hence, we utilize general homomorphic encryption, and the resulting costs drive our schemes' design.

Private Stream Search. In Private Stream Search (PSS) [OS05, DD07, BSW09, FR13], a client can search a keyword over a database of documents and retrieve the ones with such a keyword without revealing the keyword to the server. As for Keyword PIR, this cannot be directly used to solve OMR or GOMR. In [LT22], the authors use similar techniques as in PSS works. Since our work is built on [LT22] constructions, our schemes also involve those techniques to perform group OMR. However, since those parts are not the main focus of this work, we refer the readers to the original OMR paper for more details.

1.2.2 Multi-Recipient Encryption and Broadcast Encryption

Multi-Recipient Encryption. Multi-Recipient Encryption (MRE) [Kur02, BBS02, BBKS07, PPS14] focuses on efficiently encrypting a message (or multiple different messages) to different recipients. In these works, they use the technique of randomness reuse, based on schemes like El Gamal, and achieve an efficiency of 2x compared to the naive one-to-one encryption method.

Broadcast Encryption. In Broadcast Encryption (BE) [FN94], a central server controlling a master key is required for distributing keys to clients/recipients. However, in the applications we are interested in, like permissionless blockchain, it is hard to distribute a key from a central server. Moreover, we would also like to protect users' privacy against the central server. Therefore, the BE model is unsuitable for this paper's focus.

[BBKS07] claims that a single-message MRE can also be called broadcast encryption. In this paper, to better distinguish the two different primitives, we only consider the model with a central server as BE, and call the model without a central server MRE.

2 Overview

The following section gives an overview of the model, results, and the main techniques. Section 3 formally defines Ad-hoc Group OMR and Fixed Group OMR. Section 4 covers preliminaries. Section 5 recalls the main protocol of [LT22], and introduces several optimizations later used to build our GOMR constructions. Section 6 and Section 7 introduce our main constructions for AGOMR and FGOMR respectively, and briefly

discuss different trade-offs between properties and efficiency which give birth to some variants of our main constructions. Section 8 proposes stronger definitions of GOMR that take DoS attacks into consideration. Section 9 reports on implementation benchmarks and comparisons for the optimized OMR, AGOMR, and FGOMR protocols.

2.1 Model Overview

We first define the following components for GOMR.

A *bulletin board* (or *board* for short), denoted as BB , contains N messages (e.g., blockchain transactions). Each message is sent from a sender, addressed to up to G recipients, whose identities are supposed to remain private. (By contrast, OMR supports only a single recipient.) A message consists of a pair (x_i, c_i) where x_i is the message *payload* to convey, and c_i is a *clue* string which helps notify the intended recipients (in a privacy-preserving way) that the message is addressed to them.

To generate the clue, the sender uses the individual *clue keys* of the intended recipients, or alternatively, a *group clue key* jointly generated by the intended group of recipients. Clue keys, or group clue keys, are assumed to be published or otherwise communicated by some authenticated channels (whose details are outside our scope).

The whole board BB (i.e., all payloads and clues) is public. Typically, payloads will be end-to-end encrypted. We let $\mathcal{P} = \{0, 1\}^{\tilde{n}}$ denote the payload space for some $\tilde{n} \in \mathbb{Z}^+$, and \mathcal{C} denote the clue space (which depends on the construction).

At any time, potential recipient p may retrieve the messages addressed to the group including p in BB . We denote these messages as *pertinent* (to recipient p), and the rest as *impertinent*.

A server, called a *detector*, helps the recipient *retrieve* the payloads of those pertinent messages. The retrieval is performed obliviously: even a malicious detector learns nothing about which messages are pertinent to the recipient. The recipient gives the detector its *detection key* and a bound \bar{k} on the number of pertinent messages it expects to receive. The detector then aggregates all pertinent messages into a single *digest* string M and sends it to the recipient p . The digest M should be much smaller than the board BB (ideally, proportional to \bar{k}).

Assuming the detector is semi-honest and the number of pertinent messages is less than \bar{k} , the recipient p should be able to recover all pertinent messages from M with high probability. We denote the probability that a pertinent message is not recovered from the digest as *false negative rate* ϵ_n and the probability that the recovery procedure outputs an impertinent message as *false positive rate* ϵ_p . We require both ϵ_n and ϵ_p to be small (e.g., under 10^{-9} for ϵ_n and 10^{-6} for ϵ_p).

The board can be written incrementally, and the retrieval can be done by specifying which portion of the messages is needed.

We discuss two variants of GOMR, in this model: in AGOMR the sender can arbitrarily address any group of (up to G) recipients; whereas in FGOMR, any set of (up to G) recipients can jointly form a group with a corresponding group clue key, and the sender can address any of the already-formed groups. These notions are formally defined Section 3.

2.2 Summary and Comparisons

Table 1 summarizes the asymptotic performance and properties of all the schemes constructed in this paper (see Section 9 for concrete benchmarks). The two main constructions, alongside OMR [LT22] as a baseline, are shown first. additional variants are listed below. For reasonable group sizes, all of our schemes are preferred over the baseline; yet they offer various tradeoffs between achieved properties and computation cost, so the choice is application-dependent.

The detector shoulders the heaviest computation work in the original OMR scheme, and thus our schemes aim to improve the performance of the detector. As seen in the “Detector homomorphic ops” columns of Table 1, the main improvement lies in replacing the expensive homomorphic ciphertext decryption and homomorphic digest coding with relatively cheaper homomorphic matrix multiplications. These multiplications

Scheme	Efficiency (per msg)					Properties		
	Clue Size	Detector homomorphic ops			Sender time	Functionality	Key unlinkability	DoS resistance
		Matrix mul size (cheap)	Decrypt (expensive)	Digest coding (expensive)				
OMRp2 [LT22]	$O(G(n + \ell))$	0	G	G	$O(G(n + \ell)w)$	arbitrary ad-hoc groups	Full	Ad-hoc DoS
AGOMR3 Theorem 6.4	$O(G(n + \ell))$	$O(G^2 \log(P) + G(n + \ell))$	1	1	$O(G^2(G + n + \ell + \log(P)) + G(n + \ell)w)$	arbitrary ad-hoc groups	Detection-key unlinkability	Ad-hoc DoS
FGOMR1 Theorem 7.3	$O(n + G\ell)$	$\tilde{O}(G \log(P)\ell)$	1	1	$O(G^2(G\ell + \log(P)) + w(n + \ell G))$	arbitrary fixed groups and ad-hoc subgroup	Detection-key unlinkability	Fixed DoS
Stepping-stone and variants								
AGOMR1 Remark 6.1	$O(G(n + \ell))$	0	1	1	$O(G^2(G + n + \ell) + G(n + \ell)w)$	arbitrary ad-hoc groups	No	Ad-hoc DoS
AGOMR2 Remark 6.2	$O(G(n + \ell))$	$O(Gn)$	1	1	$O(G^2(G + n + \ell) + G(n + \ell)w)$	ad-hoc groups formed honestly	Detection-key unlinkability	Ad-hoc DoS
FGOMR2 Remark 7.4	$O(n + G\ell)$	$O(G\ell)$	1	1	$O(G^3\ell + w(n + \ell G))$	fixed groups formed honestly, arbitrary ad-hoc subgroup	Detection-key unlinkability	Fixed DoS
FGOMR-based AGOMR §7.5	$O(n + G\ell)$	$O(G\ell)$	1	1	$O(G^3\ell + w(n + \ell G))$	ad-hoc groups formed honestly	Detection-key unlinkability	No

Table 1: The table shows the asymptotic efficiency of our different GOMR constructions in this paper compared to directly using OMR from [LT22] for GOMR. n, ℓ, t, w are PVW encryption parameters. Group size G is the number of recipients in a group. P denotes the total number of recipients in the system. ϵ_p denotes the false positive rate when $G = 1$ (i.e., the false positive rate for plain OMR). “Fixed DoS” is stronger than “Ad-hoc DoS”. See more details in Section 8. “Arbitrary groups” means that a group can be any subset of recipients, while “groups formed honestly” means that honest senders choose groups with only honest recipients.

consist of just a single layer of plaintext-by-ciphertext multiplications (or two layers for AGOMR3). By contrast, the homomorphic decryption and digest coding each require multiple levels of ciphertext-by-ciphertext multiplications (each of which is orders of magnitude more expensive than plaintext-by-ciphertext multiplication): homomorphic decryption of PVW encryption scheme [PVW08] involves a range check, which needs to be done using a high-degree polynomial evaluation; and homomorphic digest coding limits the ability to use the SIMD-like batching of the underlying BFV scheme (see details in [LT22]).¹

Generally, FGOMR schemes have the smallest clue size, and a better sender time than the main AGOMR scheme. The sender time of all our schemes is worse than that of the baseline, since we have the sender proactively performs more computation to alleviate the detector’s work (which is more important, since it grows with the *total* number of messages).

Functionality-wise, our major constructions achieve the strongest functionality defined for AGOMR and FGOMR respectively: the groups can be formed arbitrarily. Some of the variants provide better detector performance assuming groups are formed honestly.

For detection-key-unlinkability, since all our schemes require the recipients to include a unique identifier in their clue keys, all clue keys of the same recipients are linkable. However, we achieve clue-key-to-detection-key and detection-key-to-detection-key unlinkability for all schemes except AGOMR1.

DoS resistance for AGOMR is weaker than the one for FGOMR. In FGOMR, it requires that two honest recipients should not be spammed (except with a small probability) as long as they are not in the same group, which is inherently impossible for AGOMR schemes. Details are in Section 8.

For a single payload (i.e., all recipients receive the same payload), all schemes have a similar digest size of $\tilde{O}_\lambda(\tilde{n} \cdot (\bar{k} + N \cdot \epsilon_p))$ (practically identical digest size for most parameters; see Section 9 for details), where \tilde{n} is the payload size, \bar{k} is the upper bound on the number of pertinent messages, N is the totally number of messages on the board, and ϵ_p is the input false positive rate. .

Moreover, all schemes in the table can be extended to take distinct payloads (see Section 6.3), without

¹Within the additional schemes, AGOMR4 is a direct extension from our baseline scheme, and does not introduce additional matrix multiplication operations, but still requires G expensive decryptions compared to our main schemes. Conversely, AGOMR1 utilizes an evaluation of a degree- G polynomial over some finite field (e.g., $\mathbb{Z}_{2^{127}-1}$) in plaintext, which is very cheap, to avoid homomorphic matrix multiplications; the drawback is that AGOMR1 does not provide any detection unlinkability. The other schemes aims to capture detection-key unlinkability, at the cost of using homomorphic matrix multiplication.

any change in terms of asymptotic behaviors in the table. The digest size of all schemes also remains the same as before, except for the baseline scheme, where the digest would be G times larger. Thus, our schemes are all better off than the baseline scheme in terms of detector-recipient communication when extended to having distinct payloads.

2.3 Main Techniques

We briefly summarize the main techniques used in our schemes.

Two optimizations to prior work. We propose two techniques to improve the overall efficiency of the previous OMR constructions: 1) an optimization to PVUnpack procedure, previously accounting for a large fraction of the detector’s runtime, which reduces the number of homomorphic operations from $O(D \log(D))$ to $O(D)$; 2) a tighter parameter analysis by better utilizing the SIMD property of BFV; 3) an approach of compressing messages by concatenating $G > 1$ payloads together to reduce detector cost.

AGOMR using polynomial encoding. To send a message to G recipients, the sender first generates G clues as in [LT22]. Then, the sender encodes multiple clues into a single polynomial function by interpolating the polynomial at points representing the intended recipients’ identities. The detector evaluates this polynomial at a specific point corresponding to a recipient’s ID, to recover that recipient’s clue. This polynomial evaluation can be done using homomorphic encryption to protect the recipients’ identities. The detector then only needs to perform a single homomorphic decryption and encoding, instead of G such operations for G different clues (as in the baseline method).

Multilinear encoding and clue size compression. For more efficient homomorphic evaluation, we replace the polynomial encoding by multilinear encoding. However, naively interpolating a multilinear function suspects to attacks: an adversarially-formed group may cause the equation system to be unsolvable, breaking completeness. This can be solved by making the ID space sufficiently large, but that increases the clue size. To compress these large IDs, the sender applies a shrinking linear transformation to the IDs, chosen pseudorandomly *after* the group is chosen, and solves the system induced by these compressed IDs.

FGOMR through LWE-based Multi-Recipient Encryption. We show how to construct FGOMR given a key-private Multi-Recipient Encryption (KP-MRE), by instantiating the clues as MRE ciphertexts.

We thus construct such a KP-MRE based on Learning With Error (LWE), allowing for an FHE-friendly instantiation. The key insight is that given G public keys from G recipients, the sender can construct a single ciphertext (\vec{a}, b) such that for all the corresponding secret keys $(\text{sk}_i)_{i \in [G]}$, it holds that $\vec{a} \cdot \text{sk}_i \approx b + m$, for a bit m . To do so we divide \vec{a} and all the secret keys sk_i into two parts, $\vec{a} = \vec{a}' \parallel \alpha$ and $\text{sk}_i = \text{sk}'_i \parallel \text{dk}_i$, where dk is a public auxiliary key randomly sampled by the recipients. \vec{a}' and \vec{b} are shared among all recipients in the group. The recipient-specific component α_i is computed by the sender to bridge the gap between $\langle \vec{a}', \text{sk}'_i \rangle$ and \vec{b} , i.e., $\langle \vec{a}' \parallel \alpha_i, \text{sk}_i \rangle \approx \vec{b} + \vec{m}$.

This construction naturally extends to a multi-bit message \vec{m} and moreover extends to allow each recipient to have their own message (i.e., \vec{m}_i for recipient i).

Generalized snake-eye conjecture. To prove the DoS resistance property of our AGOMR scheme, we generalize the snake-eye conjecture proposed in [LT22]. For FGOMR, we prove it under the same snake-eye conjecture in [LT22]. We also propose another general conjecture and prove it equivalent to the old conjecture in [LT22].

GOMR with distinct payloads. The polynomial encoding technique extends to the case where the recipients in a group each have a different payload. This allows our GOMR schemes to accommodate different payloads for different recipients with a very small overhead.

Amortization and preprocessing. As our schemes highly depend on homomorphic matrix multiplications, we introduce additional techniques that encode the clues and detection keys in special forms (NTT representations) to reduce the cost. The clue encoding and some other homomorphic operations can be performed offline and amortized across multiple recipients.

3 Defining Group Oblivious Message Retrieval

In this section, we formally define the problem of Group Oblivious Message Retrieval, using the model and notation of Section 2.1.

3.1 Ad-hoc Group OMR

Ad-hoc Group OMR addresses the case where senders arbitrarily choose up to G recipients, given the recipients' clue keys. No advance action or cooperation is necessary from the recipients other than publishing their clue keys (as in standard OMR). For privacy, an adversary who corrupted the detector should not tell which group a message is addressed to. Furthermore, even if it has corrupted some of the group members (including maliciously generating their clue keys), the adversary should not learn who the *remaining* recipients are.

This definition generalizes that of OMR [LT22, Section 4.3]; see Remark 3.1 for detail.

Definition 3.1 (Ad-hoc Group Oblivious Message Retrieval (AGOMR)). An AGOMR scheme has the following PPT algorithms:

- $\text{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$: takes a security parameter λ , a false positive rate ϵ_p , a false negative rate ϵ_n , an upper bound of the number of recipients G per message, the total number of recipients P ; outputs public parameters pp .
- $(\text{sk}, \text{pk} = (\text{pk}_{\text{clue}}, \text{pk}_{\text{detect}})) \leftarrow \text{KeyGen}(\text{pp})$: takes a public parameters pp ; outputs a secret key sk and a public key pk consisting of a clue key pk_{clue} and a detection key $\text{pk}_{\text{detect}}$.
- $c \leftarrow \text{GenClue}(\text{pp}, \text{pk}_{\text{clue}_1}, \dots, \text{pk}_{\text{clue}_G}, x)$: takes a public parameter pp , up to G clue keys and a payload $x \in \mathcal{P}$; outputs a clue $c \in \mathcal{C}$.
- $M \leftarrow \text{Retrieve}(\text{pp}, \text{BB}, \text{pk}_{\text{detect}}, \bar{k})$: takes a public parameter pp , a board $\text{BB} = \{(x_1, c_1), \dots, (x_N, c_N)\}$ of size N , a detection key $\text{pk}_{\text{detect}}$ from the recipient, an upper bound \bar{k} on the number of pertinent messages for that recipient; outputs a digest M .
- $\text{PL} \leftarrow \text{Decode}(\text{pp}, M, \text{sk})$: takes a public parameter pp , the digest M , a secret key sk , and outputs either a decoded payload list $\text{PL} \in \mathcal{P}^k$ or an overflow indication $\text{PL} = \text{overflow}$.

An AGOMR scheme should fulfill completeness, soundness, and computational privacy, as defined below. These definitions use the notion of a board generation procedure:

Definition 3.2 (AGOMR board generation). Given the total number of messages N , the total number of recipients P and a public parameter pp : for each recipient $j \in [P]$ ², generate keys $(\text{sk}_j, \text{pk}_j) \leftarrow \text{KeyGen}(\text{pp})$; for $i \in [N]$, arbitrarily choose a set $Y_i \subseteq [P]$, $1 \leq |Y_i| \leq G$, representing the set of recipients for the i -th message. Define sets S_1, \dots, S_P such that $S_j = \{i \mid j \in Y_i\}$ representing the indices of messages addressed to recipient j . Arbitrarily choose unique payloads (x_1, \dots, x_N) . Generate clues $c_i \leftarrow \text{GenClue}(\text{pp}, \{\text{pk}_{\text{clue}_j}\}_{j \in Y_i}, x_i)$ for $i \in [N]$ ³. Then, output the set S_1 , $(\text{sk}_1, \text{pk}_{\text{detect}_1})$, and the board $\text{BB} = \{(x_1, c_1), \dots, (x_N, c_N)\}$ ⁴.

² $[N] := \{1, \dots, N\}$, as defined in Section 4.1.

³Payloads being unique is w.l.o.g. See more details in paragraph "Repeating Payloads" in [LT22, Section 4.3]

⁴ S_1 is the set containing indices of messages pertinent to the recipient holding keys sk_1, pk_1 , which w.l.o.g is the first recipient. All the properties defined below symmetrically applies to other recipients.

Adversary	Challenger
	$\text{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$
Choose $Y_0, Y_1 \subseteq [P]$, where $ Y_0 , Y_1 \leq G$	$(\text{pk}_i, \text{pk}_i = (\text{pk}_{\text{clue}_i}, \cdot)) \leftarrow \text{KeyGen}(\text{pp})$ for $i \in [P]$
Let $Z = Y_0 \cap Y_1$	
Choose $(\text{sk}_i, \text{pk}_i)_{i \in Z}$ overwriting their old values	
Choose a payload x	$b \leftarrow \mathcal{S}\{0, 1\}$
	$c \leftarrow \text{GenClue}(\text{pp}, (\text{pk}_{\text{clue}_i})_{i \in Y_b}, x)$
	(Adversary wins if $b = b'$)

Figure 2: Computational Privacy game for Ad-hoc Group OMR

- (Completeness) Let $\text{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$. For any $N = \text{poly}(\lambda)$, and $0 < \bar{k} \leq N$, let the set S of pertinent messages, the key pair $(\text{sk}, \text{pk}_{\text{detect}})$, and the board BB be generated as in Definition 3.2, for any set choice and payloads therein. Let $M \leftarrow \text{Retrieve}(\text{pp}, \text{BB}, \text{pk}_{\text{detect}}, \bar{k})$, $\text{PL} \leftarrow \text{Decode}(\text{pp}, M, \text{sk})$, and $k = |S|$ (the number of pertinent messages in S), either $k > \bar{k}$ and $\text{PL} = \text{overflow}$, or

$$\Pr[x_j \in \text{PL} \mid j \in S] \geq 1 - \epsilon_n - \text{negl}(\lambda), \text{ for } j \in [N].$$

The randomness is over the coins of KeyGen and GenClue .

- (Soundness) For the same quantifiers as in Completeness:

$$\Pr[(x_j \in \text{PL} \mid j \notin S)] \leq (\epsilon_p + \text{negl}(\lambda)), \text{ for } j \in [N].$$

- (Computational privacy) An AGOMR scheme is computationally private if there does not exist any PPT adversary that can win the game in Fig. 2 with probability $> 1/2 + \text{negl}(\lambda)$, where Z, Y_0, Y_1 represent groups with at most G recipients.

An AGOMR scheme is *compact* if it moreover satisfies the following:

- (Compactness) For $\text{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, $(\text{sk}, \text{pk} = (\text{pk}_{\text{clue}}, \text{pk}_{\text{detect}})) \leftarrow \text{KeyGen}(\text{pp})$, for any board $\text{BB} = \{(x_1, c_1), \dots, (x_N, c_N)\}$, for $M \leftarrow \text{Retrieve}(\text{pp}, \text{BB}, \text{pk}_{\text{detect}}, \bar{k})$, it holds that:

$$|M| = \text{poly}(\lambda, \log N, \log G) \cdot \log \epsilon_p^{-1} \cdot \tilde{O}(\bar{k} + \epsilon_p N)$$

The compactness definition is the same as [LT22] (except for some minor interface changes). In particular, the digest size is independent of the group size G .

$\tilde{O}(\bar{k} + \epsilon_p N)$ (where $\tilde{O}(x) = x \text{polylog}(x)$) accounts for the number of messages detected as pertinent, including false positives; and the remaining factors account for the cost of representing each such message, taking the payload size as constant.

Remark 3.1 (Relation to OMR). When $G = 1$, this AGOMR definition implies OMR ([LT22, Definition 4.1]), with a minor difference in privacy. The adversary in AGOMR definition is given P honestly generated recipient keys and can arbitrarily choose two of them as the challenge, where P represents the total number of recipients in the system passed as a parameter in GenParams . In contrast, the adversary in [LT22] is given two honestly generated keys as the challenge. Other parts are trivially identical.

Detection-key unlinkability. In addition to hiding the results of a retrieval query, we may wish to hide which recipient is even doing the retrieval. Thus, essentially, we require the property that given a detection key of a recipient, one cannot tell who the recipient is.

To satisfy this property, we need the recipient to be able to regenerate its detection key to do the retrieval, while still having the correctness, soundness, privacy, and compactness hold, using the new detection key. Furthermore, the new detection key should not be able to be linkable to the recipient even given the original set of public keys.

We formally capture this property by the following definition, adapted from [LT22, Definition 9.2].

Definition 3.3. (AGOMR detection-key-unlinkability) An AGOMR scheme is said to be detection-key-unlinkable if it further has an interface $(sk', pk'_{\text{detect}}) \leftarrow \text{RegenDetectKey}(pp, sk)$, such that:

1. Correctness, soundness, computational privacy, and compactness hold also after replacing $M \leftarrow \text{Retrieve}(pp, BB, pk_{\text{detect}}, \bar{k})$ with $M \leftarrow \text{Retrieve}(pp, BB, pk'_{\text{detect}}, \bar{k})$, and replacing $PL \leftarrow \text{Decode}(pp, M, sk)$ with $PL \leftarrow \text{Decode}(pp, M, sk')$, where $(sk', pk'_{\text{detect}}) \leftarrow \text{RegenDetectKey}(pp, sk)$.
2. Let $pp \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, $(sk, pk = (pk_{\text{clue}}, pk_{\text{detect}})) \leftarrow \text{KeyGen}(pp)$, $(sk', pk' = (pk'_{\text{clue}}, pk'_{\text{detect}})) \leftarrow \text{KeyGen}(pp)$, then for any $n = \text{poly}(\lambda)$, for all $i \in [n]$, let $(\cdot, pk_{\text{detect}_i}) \leftarrow \text{RegenDetectKey}(pp, sk)$, $(\cdot, pk'_i = (\cdot, pk'_{\text{detect}_i})) \leftarrow \text{KeyGen}(pp)$, it holds that $(pk, pk_{\text{detect}_1}, \dots, pk_{\text{detect}_n}) \approx_c (pk', pk'_{\text{detect}_1}, \dots, pk'_{\text{detect}_n})$.

Remark 3.2. This *detection-key unlinkability* directly implies *detection-to-clue-key-unlinkability* [LT22, Def 9.1], but is weaker than *full-key-unlinkability* [LT22, Def 9.2]. The difference from the latter is that *detection-key unlinkability* does not achieve the clue-key-to-clue-key unlinkability, i.e., the ability for recipients to generate multiple indistinguishable-but-functionally-identical clue keys, which is useful in some applications for creating ephemeral/stealth addresses.

3.2 Fixed Group OMR

In some applications, groups are predetermined by their member recipients (e.g., for mailing lists). A recipient joining multiple groups should be able to use a single secret key to detect all pertinent messages addressed to the groups, which include that recipient as a group member. By sending to a fixed group, we may further improve the efficiency, including the clue size, server computation time, and so on. (For simplicity, a sending to a single recipient will be handled as sending to a fixed group consisting of just that recipient.)

Besides the difference in efficiency, FGOMR has a slightly weaker privacy notion. If an adversary corrupts some members in a fixed group, it may distinguish messages addressed to that group from those addressed to other groups (which in reality is often possible anyway by inspecting the payload's plaintext). Thus, we require for messages addressed to groups *without* corrupted recipients, the groups are indistinguishable.

Furthermore, we allow a sender to select an arbitrary subgroup Y' of some fixed group Y and send a message only pertinent to this subgroup. The excluded recipients $Y \setminus Y'$ should not discover such an exclusion (i.e., the clue generated for the subgroup is indistinguishable from a clue generated for a group that does not include the excluded recipients at all).

We capture the notion of an FGOMR scheme as follows. We replace KeyGen in OMR with PersonalKeyGen , which only generates secret and detection keys. A clue key for an entire group of up to G recipients will be constructed by first invoking GroupKeyGenAux to generate a share with respect to that group, and then invoking GroupKeyGen .

Definition 3.4 (Fixed Group Oblivious Message Retrieval (FGOMR)). An FGOMR scheme has the following PPT algorithms:

- $pp \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$: takes a security parameter λ , a false positive rate ϵ_p , a false negative rate ϵ_n , the upper bound of the number of recipients of each message G , and the total number of recipient P ; outputs public parameters pp .

- $(\text{sk}, \text{pk}_{\text{detect}}) \leftarrow \text{PersonalKeyGen}(\text{pp})$: takes public parameters pp ; outputs a secret key sk and a detection key $\text{pk}_{\text{detect}}$.
- $\text{gPKshare} \leftarrow \text{GroupKeyGenAux}(\text{pp}, \text{sk}, Y)$: takes public parameters pp , a secret key sk , and a group of recipients Y ; outputs a group key share gPKshare for a single recipient.
- $\text{pk}_{\text{clue}} \leftarrow \text{GroupKeyGen}(\text{pp}, \text{gPKshare}_1, \dots, \text{gPKshare}_G)$: takes public parameters pp , up to G public key shares gPKshare_i ; outputs a clue key.
- $c \leftarrow \text{GenClue}(\text{pp}, \text{pk}_{\text{clue}}, Y, x)$: takes public parameters pp , a clue key pk_{clue} , a group of recipients Y , a payload $x \in \mathcal{P}$; outputs a clue $c \in \mathcal{C}$.
- $M \leftarrow \text{Retrieve}(\text{pp}, \text{BB}, \text{pk}_{\text{detect}}, \bar{k})$: takes public parameters pp , a board $\text{BB} = \{(x_1, c_1), \dots, (x_N, c_N)\}$ of size N , a detection key $\text{pk}_{\text{detect}}$, and an upper bound \bar{k} on the number of pertinent messages addressed to that recipient; outputs a digest M .
- $\text{PL} \leftarrow \text{Decode}(\text{pp}, M, \text{sk})$: takes public parameters pp , a digest M and a secret key sk ; outputs either a decoded payload list $\text{PL} \in \mathcal{P}^k$ or an overflow indication $\text{PL} = \text{overflow}$.

A FGOMR scheme should fulfill completeness, soundness, and computational privacy, as defined below. These definitions use a different notion of board generation:

Definition 3.5 (FGOMR board generation). Given the total number of messages N , the total number of recipients P and public parameter pp : for each recipient $j \in [P]$, generate keys $(\text{sk}_j, \text{pk}_{\text{detect}_j}) \leftarrow \text{PersonalKeyGen}(\text{pp})$; for $i \in [N]$, arbitrarily choose $Y_i \subseteq [P], 1 \leq |Y_i| \leq G$, where Y_i represents the group of recipients for the i -th message, and $Y'_i \subseteq Y_i, Y'_i \neq \phi$. Define sets S_1, \dots, S_P such that $S_j = \{i \mid j \in Y'_i\}$, representing the indices of messages addressed to recipient j . For all $i \in [N], j \in Y_i$, generate $\text{gPKshare}_{i,j} \leftarrow \text{GroupKeyGenAux}(\text{pp}, \text{sk}_j, Y_i)$, and $\text{pk}_{\text{clue}_i} \leftarrow \text{GroupKeyGen}(\text{pp}, \{\text{gPKshare}_{i,j}\}_{j \in Y_i})$. Arbitrarily choose unique payloads (x_1, \dots, x_N) ; generate clues $c_i \leftarrow \text{GenClue}(\text{pp}, \text{pk}_{\text{clue}_i}, Y'_i, x_i)$. Finally, output $S_1, (\text{sk}_1, \text{pk}_{\text{detect}_1})$, and the board $\text{BB} = \{(x_1, c_1), \dots, (x_N, c_N)\}$.

- (Completeness and Soundness) Same as AGOMR completeness and soundness in Definition 3.1, replacing Definition 3.2 with Definition 3.5.
- (Computational privacy) An FGOMR scheme is computationally private if there does not exist any PPT adversary that can win the game in Fig. 3 with probability $> 1/2 + \text{negl}(\lambda)$, where Y_j, Z_j represent groups of recipients.

A FGOMR scheme is *compact* if it further satisfies the compactness definition in Theorem 6.1 (w.r.t. the interfaces defined above).

Detection-key unlinkability. Similarly, we can also define detection-key unlinkability for FGOMR.

Definition 3.6. (FGOMR detection-key-unlinkability) An FGOMR scheme is said to be detection-key-unlinkable if it further has an interface $(\text{sk}', \text{pk}'_{\text{detect}}) \leftarrow \text{RegenDetectKey}(\text{pp}, \text{sk})$, such that:

1. Correctness, soundness, computational privacy, and compactness hold also after replacing $M \leftarrow \text{Retrieve}(\text{pp}, \text{BB}, \text{pk}_{\text{detect}}, \bar{k})$ with $M \leftarrow \text{Retrieve}(\text{pp}, \text{BB}, \text{pk}'_{\text{detect}}, \bar{k})$, and replacing $\text{PL} \leftarrow \text{Decode}(\text{pp}, M, \text{sk})$ with $\text{PL} \leftarrow \text{Decode}(\text{pp}, M, \text{sk}')$, where $(\text{sk}', \text{pk}'_{\text{detect}}) \leftarrow \text{RegenDetectKey}(\text{pp}, \text{sk})$.
2. Let $\text{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, $(\text{sk}, \text{pk}_{\text{detect}}) \leftarrow \text{PersonalKeyGen}(\text{pp})$, $(\text{sk}', \text{pk}'_{\text{detect}}) \leftarrow \text{PersonalKeyGen}(\text{pp})$, for all $Y, Y' \subseteq [P]$, let $\text{gPKshare} \leftarrow (\text{pp}, \text{sk}, Y)$, $\text{gPKshare}' \leftarrow (\text{pp}, \text{sk}', Y')$, for any $n = \text{poly}(\lambda)$, for all $i \in [n]$, let $(\cdot, \text{pk}_{\text{detect}_i}) \leftarrow \text{RegenDetectKey}(\text{pp}, \text{sk})$, $(\cdot, \text{pk}'_i = (\cdot, \text{pk}'_{\text{detect}_i})) \leftarrow \text{KeyGen}(\text{pp})$, it holds that $(\text{gPKshare}, \text{pk}_{\text{detect}}, \text{pk}_{\text{detect}_1}, \dots, \text{pk}_{\text{detect}_n}) \approx_c (\text{gPKshare}', \text{pk}'_{\text{detect}}, \text{pk}'_{\text{detect}_1}, \dots, \text{pk}'_{\text{detect}_n})$

Adversary		Challenger
	← PP	pp ← GenParams($1^\lambda, \epsilon_p, \epsilon_n, G, P$)
		$(sk_j, \cdot) \leftarrow \text{PersonalKeyGen}(pp)$ for $j \in [P]$
Query Phase I: for $i = 1, \dots, N$		
Choose recipient set $Y_i \subseteq [P]$, where $ Y_i \leq G$	→ Y_i	
	← $(gPKshare_{j,i})_{j \in Y_i}$	$gPKshare_{j,i} \leftarrow \text{GroupKeyGenAux}(pp, sk_j, Y_i)$, for $j \in Y_i$
Challenge Phase:		
Choose $i_0 \neq i_1 \in [N]$, and for $i \in [i_0, i_1]$:		
Choose $Z_i \subseteq Y_i$		
Choose $(gPKshare'_{j,i})_{j \in Z_i}$,		
and $gPKshare'_{j,i} = gPKshare_{j,i}$, $j \in Y_i \setminus Z_i$		
Choose a payload x	→ i_0, i_1, x, Z_0, Z_1 ← $((gPKshare'_{j,i})_{j \in Y_i})_{i \in [i_0, i_1]}$	$b \leftarrow_{\$} \{0, 1\}$
		$pk_{clue_{i_b}} \leftarrow \text{GroupKeyGen}(pp, (gPKshare'_{j,i_b})_{j \in Y_{i_b}})$
	← c	$c \leftarrow \text{GenClue}(pp, pk_{clue_{i_b}}, Y_{i_b} \setminus Z_{i_b}, x)$
Query Phase II: same as Query Phase I.	→ b'	(Adversary wins if $b = b'$)

Figure 3: Computational Privacy game for Fixed Group OMR

Remark 3.3. A stronger notion of FGOMR computational privacy is possible, where privacy holds even if the keys are chosen semi-maliciously (i.e., with malicious choice of randomness). This is captured by the game of Fig. 10, which allows the adversary to rewrite some secret keys with its own maliciously-chosen randomness.

An even stronger notion of privacy allows the keys to be crafted fully maliciously (as in AGOMR). This reduces to the above semi-malicious notion, by adding a non-interactive zero-knowledge (NIZK) proof that the keys were correctly generated from some randomness.

Neither strengthening is essential for the envisioned application, so for readability we omit these from subsequent discussion. Note, however, that all of our FGOMR constructions satisfy the semi-malicious strengthening, and thus can be made private even for fully-malicious key generation by addition of the NIZK proof to the public keys.

4 Preliminaries

4.1 Notation

All logarithms are to base 2 and rounded up to an integer, unless otherwise specified. The notation \vec{v} denotes row vectors, and $[n]$ denotes the set $\{1, \dots, n\}$. By $\mathcal{D}(x)$ we denote the distribution of a random variable x .

For a variable (e.g., PV) representing a BFV ciphertext encrypting a vector of D elements in \mathbb{Z}_t for some $D \in \mathbb{Z}^+$, let $PV[i]$ denote the i -th element in the vector. Similarly, when the encrypted vector is a flattened matrix of size $D \times T$ for some $D, T \in \mathbb{Z}^+$, we let $PV[i][j]$ denote the (i, j) -th element of the matrix.

We summarize the main symbols throughout this paper in Table 2.

$N \in \mathbb{N}$	Size of bulletin board, $ \mathbb{BB} $
$P \in \mathbb{N}$	Total number of recipients
$k \in \mathbb{N}$	Upper bound of the number of messages each recipient has
$\tilde{k} \in \mathbb{N}$	Upper bound of the number of messages each recipient has, including false positives
$x \in \mathcal{P}$	Payload
$c \in \mathcal{C}$	Clue
M	Digest
$0 < \epsilon_p < 1$	False positive rate
$0 < \epsilon_n < 1$	False negative rate
$n \in \mathbb{N}$	Dimension of PVW secret key
$w \in \mathbb{N}$	Number of PVW instances in a PVW public key
$\ell \in \mathbb{N}$	Number of bits encrypted in a PVW ciphertext
$t \in \mathbb{N}$	Ciphertext modulus of the underlying PVW scheme; Plaintext modulus of the underlying BFV ciphertext
$\sigma \in \mathbb{R}$	Error standard deviation of the underlying PVW scheme
$D \in \mathbb{N}$	Ring dimension of the underlying BFV ciphertext
$G \in \mathbb{N}$	Maximum number of recipients for each message
$Y \subseteq [P]$	A set of recipients, $ Y \leq G$
$I \in \mathbb{N}$	Size of ID for each recipient in AGOMR
$\text{id} \in \mathbb{Z}_t^I$	ID of each recipient
$L \in \mathbb{N}$	Size of auxiliary key for each recipient in FGOMR
$\text{dk} \in \mathbb{Z}_t^L$	Auxiliary key of each recipient

Table 2: A list of common parameters used in all the OMR/GOMR constructions.

4.2 PVW Encryption

Our constructions are based on lattice-based encryption schemes as in [LT22]. We use the Peikert-Vaikuntanathan-Waters (PVW) [PVW08] variant of Regev’s LWE-based encryption [Reg09] defined as follows:

- $\text{pp}_{\text{LWE}} = (n, \ell, w, q, \sigma) \leftarrow \text{PVW.GenParams}(1^\lambda, \ell, q, \sigma)$: Choose a secret key dimension n , and $w = \text{poly}(\lambda, n, \ell, q)$. Set ciphertext modulus q , number of bits of plaintext modulus ℓ , and standard deviation for Gaussian distribution for ciphertext noise generation σ . n, w, q, σ are chosen as in [PVW08]. pp_{LWE} is assumed to be implicitly taken by the following algorithms.
- $(\text{sk}, \text{pk}) \leftarrow \text{PVW.KeyGen}(\text{pp}_{\text{LWE}})$: Choose a secret key $\text{sk} \leftarrow \mathbb{Z}_q^{n \times \ell}$ uniformly at random. Randomly sample $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times w}$ and a error matrix $\mathbf{X} \in \mathbb{Z}_q^{\ell \times w}$ from some Gaussian distribution χ_σ , and compute $\text{pk} = (\mathbf{A}, \mathbf{P} = \text{sk}^\top \mathbf{A} + \mathbf{X})$.
- $\text{ct} = (\vec{a}, \vec{b}) \leftarrow \text{PVW.Enc}(\text{pp}_{\text{LWE}}, \text{pk}, \vec{m})$: To encrypt a vector $\vec{m} \in \mathbb{Z}_2^\ell$, define a vector $\mathbf{t} = \frac{q}{2} \cdot \vec{m} \in \mathbb{Z}_q^\ell$. Choose a random vector $\mathbf{e} \leftarrow \{0, 1\}^w \in \mathbb{Z}_2^w$ uniformly at random. The ciphertext is the pair $(\vec{a}, \vec{b}) = (\mathbf{A}\mathbf{e}, \mathbf{P}\mathbf{e} + \mathbf{t}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$.
- $\vec{m} \leftarrow \text{PVW.Dec}(\text{pp}_{\text{LWE}}, \text{sk}, \text{ct} = (\vec{a}, \vec{b}))$: $\vec{d} = \vec{b} - \text{sk}^\top \vec{a} \in \mathbb{Z}_q^\ell$, $\vec{m} = \lfloor \frac{\vec{d} + q/4}{q/2} \rfloor \in \mathbb{Z}_2^\ell$

The PVW scheme is unconditionally correct and sound. Under the LWE hardness assumption, it also fulfills the standard definitions of semantic security (IND-CPA) and key privacy [Reg09, APS15]. Moreover, it satisfies the property of wrong-key decryption defined in [LT22, Definition 5.1].

4.3 Fully Homomorphic Encryption

Fully Homomorphic Encryption (FHE), introduced by Rivest et al. [RAD78] and first constructed by Gentry [Gen09], enables evaluation of a circuit on encrypted data, such that the result is the encryption of the

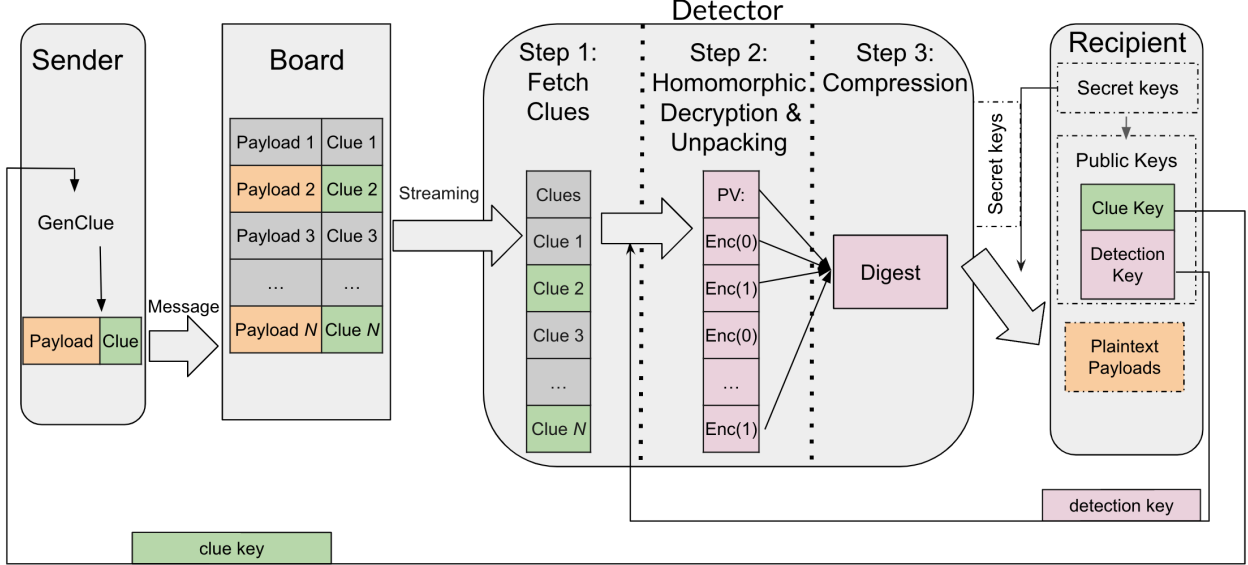


Figure 4: Main components of the original OMR construction. Our main focus in this work is to adapt and optimize the sender side and step 1 on the detector side to the group setting. Step 2 only has very minor changes. Step 3 maintains exactly the same.

corresponding output.

BFV FHE scheme. We use the Brakerski/Fan-Vercauteran (BFV) homomorphic encryption scheme [Bra12, FV12].

The BFV scheme consists of the following PPT algorithms: $\text{GenParams}(1^\lambda)$, $\text{KeyGen}(\text{pp}_{\text{BFV}})$, $\text{Enc}(\text{pp}_{\text{BFV}}, \text{pk}, m)$, $\text{Dec}(\text{pp}_{\text{BFV}}, \text{sk}, c)$ as normal PKE schemes. And an additional algorithm: $\text{Eval}(\text{pp}_{\text{BFV}}, \text{pk}, (\text{ct}_1, \dots, \text{ct}_k), C)$, which takes k ciphertexts encrypting $(m_{i \in [k]})$ and a circuit C , and outputs another ciphertext ct' such that $\text{Dec}(\text{pp}_{\text{BFV}}, \text{sk}, \text{ct}') = C(\text{Dec}(\text{pp}_{\text{BFV}}, \text{sk}, m_{i \in [k]}))$.

BFV is unconditionally correct and sound. Under the *Ring-LWE (RLWE)* [LPR13, Pla18] hardness assumption, it also fulfills the standard definitions of semantic security (IND-CPA) for FHE schemes.

Given a polynomial from the cyclotomic ring $R_t = \mathbb{Z}_t[X]/(X^D + 1)$, the BFV scheme encrypts it into a ciphertext consisting of two polynomials, where each polynomial is from a larger cyclotomic ring $R_q = \mathbb{Z}_q[X]/(X^D + 1)$ for some $q > t$. We refer to t , q , and D as the plaintext modulus, the ciphertext modulus, and the ring dimension, respectively. Each ciphertext can pack D plaintext group elements $(m_1, \dots, m_D) \in \mathbb{Z}_t^D$. We often use “single instruction, multiple data” (SIMD) homomorphic evaluation to BFV ciphertexts.

5 Optimization to Prior Work

This section introduces two optimizations to the original OMR constructions [LT22]. Subsequent sections are based on this optimized OMR construction. We first recall the OMRp2 of [LT22], with its high level algorithms sketched in Fig. 4 and summarized as follows.

5.1 The Original OMR Construction

Setup and sending (OMR.GenClue). A bulletin board BB of size N , $\text{BB}[i] = (x_i, c_i)$, contains messages, each consisting of a payload x_i and clue c_i . Each clue c_i is generated by the sender of that message by

producing a PVW ciphertext of ℓ 1's encrypted under the PVW public key found in the clue key of the intended (secret) recipient.

Retrieval (OMR.Retrieve). To request retrieval, recipient j sends its detection key $\text{pk}_{\text{detect}_j} = (\text{BFV.pk}_j, \text{BFV.Enc}(\text{BFV.pk}_j, \text{PVW.sk}_j))$ to detector. The detector then uses $\text{pk}_{\text{detect}_j}$ to homomorphically process all messages and send back a digest containing payloads addressed to recipient j , via the following stages.

Fetch clues: The first step of the original construction is simply fetching all the clues (and payloads) from the board.

Homomorphic decryption and unpacking: Then, the detector uses the clues on the board to obliviously obtain N encrypted bits, each indicating whether a message is pertinent (0 for no and 1 for yes) (i.e., the left half of Fig. 4), as follows. For each clue c_i , the detector uses the PVW.sk_j encrypted under BFV.pk_j to homomorphically evaluate the decryption of c_i , and then homomorphically computes the pertinency indicator b_i to be 1 if all ℓ decryption results are 1's. Thus, b_i is (with high probability) 1 iff the i -th message is intended for recipient j .

For efficiency, the above is done in large batches of D clues at a time, using BFV's SIMD evaluation. Let PV denote the resulting vector of BFV-encrypted pertinency bits (b_1, \dots, b_D) . The detector then decomposes PV through function PVUnpack into D BFV ciphertexts $\vec{\text{PV}}_{i \in [D]}$, where $\vec{\text{PV}}_i = \text{BFV.Enc}(b_i, \dots, b_i)$.

Compression: The detector compresses the encrypted binary vector of length N , which is sparse (mostly 0), to form a digest of size $o(N)$, using coding techniques as follows.

Assuming at most \bar{k} messages are pertinent, the detector randomly assigns each message to one of $m > \bar{k}$ buckets, each consisting of an accumulator and a counter. Suppose messages $S \subseteq [N]$ are assigned to a bucket. The detector then homomorphically sums up the indices of all the assigned pertinent messages in the same bucket as the value of the accumulator (i.e., $\text{Acc} \leftarrow \sum_{j \in X} \vec{\text{PV}}_j \cdot j$), and records the number of the pertinent messages in each bucket by using the counter (i.e., $\text{Ctr} \leftarrow \sum_{j \in X} \vec{\text{PV}}_j \cdot 1$) (e.g., result = 2 if and only if there are two pertinent messages assigned to the bucket). The compression succeeds only if there is no more than one pertinent message assigned in one bucket (i.e., the counter value is 0 or 1). Otherwise, it is a collision. To amplify the success probability, the detector will repeat this process several times. We denote the number of such repetitions by C . [LT22] further uses partial information gathering (i.e., gather the information from the buckets whose $\text{Ctr} = 1$ in each repetition) to reduce C .

The payloads can also be compressed in the same way as the pertinent indices described above. However, for better efficiency, [LT22] uses Sparse Random Linear Coding (SRLC). See details of the payload processing there (we reuse this component verbatim).

Decryption (OMR.Decode). The recipient, given the digest, decrypts and decodes it and obtains the correct payloads using the PVW secret key PVW.sk .

In this work, we focus on optimizing the detector's cost in the aforementioned stage "Homomorphic decryption and unpacking", which dominates cost — in the case where a message is addressed to a group of recipients. To this end, we also modify the sender's algorithm in the first stage. The later stages are essentially the same as [LT22], except for the optimization introduced in Section 5.3.

5.2 Optimizing PVUnpack

The construction in [LT22] expands a packed PV (i.e., a BFV ciphertext encrypting $\{0, 1\}^D$) into D BFV ciphertexts slot by slot (i.e., extract a single slot into a BFV ciphertext and replicate the result into all the D slots), resulting in $O(D \log(D))$ homomorphic operations. We improve this by batching the expansion step, resulting in $O(D)$ homomorphic operations.

The replication of the bits across all slots takes $\log(D)$ homomorphic rotations per slot, resulting in $D \log(D)$ rotations and D multiplications (see [LT22, Alg. 6]). We observe that these rotations can be amortized over multiple slots, as follows.

Instead of extracting one slot and rotate $\log(D)$ times, we instead extract half of the slots (thus resulting into 2 ciphertexts) and rotate once, and then extract 1/4 of the slots (thus resulting into 4 ciphertexts) and

then rotate them once, and repeat this process totally $\log(D)$ levels. In the end, we obtain D ciphertexts each encrypting a slot as expected. This is essentially a binary tree structure, with $\log(D)$ depth.

However, this also means that the multiplicative depth is $\log(D)$. Thus, instead of having a binary tree, we can have a t -ary tree (i.e., a tree with a branching factor of t instead of 2), resulting in a depth of $\log_t(D)$. Even more generally, we do not need to have the same branching factor at each level. Instead, we can have branching factors $(b_1, \dots, b_L) \in \mathbb{Z}_D^L$, where $b_1 \times \dots \times b_L = D$ for the tree with L levels. For simplicity, we only have branching factors that are power-of-two. This general tree structure thus gives us in total $D \log(x_L) + \frac{D}{x_{L-1}} \log(\frac{x_{L-1}}{x_L}) + \dots + \frac{D}{x_1} \log(\frac{x_1}{x_2})$ rotations, where $x_i = \prod_{j=1}^i b_j$.

We formalize the procedure in Algorithm 1.⁵

Algorithm 1 Optimized PVUnpack

```

1: procedure PVUnpack(pkBFV, ct, L)
2:   Initialize res ← (ct) as a vector
3:   Find power-of-two points (x1, ..., xL) s.t.  $D \log(x_L) + \frac{D}{x_{L-1}} \log(\frac{x_{L-1}}{x_L}) + \dots + \frac{D}{x_1} \log(\frac{x_1}{x_2})$  is minimal
4:   for i ∈ x do
5:     tmp = (1, ..., 0)
6:     ▷ where (1 + k · xi-1, ..., xi + k · xi-1) are 1's for all k ∈ [D/xi-1], for i = 1, k ∈ {0}, x0 = 0
7:     cur_size ←  $\frac{D}{x_{i-1}}$  ▷ x0 = D
8:     Initialize res' as an empty vector
9:     for u = 1 to cur_size do
10:      restmp ← res[u]
11:      for j = 0 to  $\frac{x_{i-1}}{x_i} - 1$  do
12:        tmp = (1, ..., 0)
13:        ▷ where (1 + k · xi-1 + j ·  $\frac{x_{i-1}}{x_i}$ , ..., xi + k · xi-1 + j ·  $\frac{x_{i-1}}{x_i}$ ) are 1's for all k ∈ [D/xi-1]
14:        ▷ for i = 1, k ∈ {0}, x0 = D
15:        restmp1 ← BFV.Eval(pkBFV, restmp, tmp, ×)
16:        for β = log(xi) to log(xi-1) do ▷ x0 = D
17:          restmp1 ← BFV.Rotate(pkBFV, restmp1, 2β)
18:        Append restmp1 to res'
19:   res ← res'
20: return res

```

5.3 Optimizing OMRp2

Another optimization is tightening the parameter choice of m (the aforementioned number of buckets for index retrieval) and C (the aforementioned number of repetitions to reduce failure probability) of OMRp2 in [LT22]. After our optimization, OMRp2 dominates OMRp1 for most practical parameters (including the parameters set by [LT22, Section 10]) unless \bar{k} is huge and N is relatively small.

Originally, OMRp2 chooses a fixed $m = D$ when $\bar{k} \ll D$ ⁶, and then chooses the smallest C such that the error probability is achieved (i.e., $1 - \prod_{i=1}^{\bar{k}-1} (1 - (\frac{i}{m})^C) \leq \epsilon_p/4$). Instead, we let m vary and minimize $m \cdot C$ subject to the same constraint. Since each bucket takes $\log_i(N) + 1$ BFV slots to store the accumulator and the counter, this reduces the number of BFV ciphertexts from $(\log_i(N) + 1) \cdot C \cdot m/D = (\log_i(N) + 1) \cdot C$ to $(\log_i(N) + 1) \cdot m \cdot C/D$ ciphertexts; or concretely: from 15 to 1 for the parameters in [LT22].

Therefore, the optimized OMRp2 performance will dominate OMRp1 for most parameters.

This optimized OMRp2, dubbed OMR-OPT, is given in Algorithm 2.

⁵For simplicity and concrete performance, when invoking our new PVUnpack in our OMR/GOMR constructions, we always invoke it with $L = 2$. Thus, the interface become the same as the old PVUnpack.

⁶When $\bar{k} \ll D$, simply concatenate c multiple BFV ciphertexts together and view them as a single ciphertext, thus having $D \leftarrow c \cdot D$.

Theorem 5.1. The scheme OMR-OPT in Algorithm 2 is an OMR scheme for $N < D \cdot t/2$, assuming security of LWE encryption and security of BFV leveled HE as in [LT22], when instantiated with PRF f and an SRLC scheme SRLC defined in [LT22]. Moreover when instantiated with SRLC1 [LT22] Section 6.3.1, OMR-OPT is also compact.

Proof sketch. Notice that the only difference between OMR-OPT defined in [LT22] and OMR-OPT is the optimized accumulator and counter encoding described in Section 5.3. Thus, we only need to show that the probability of the optimized accumulator/counter overflowing is $\leq \epsilon_n/4$, which will then satisfy the completeness requirement defined in [LT22, Section 4.3]. The proof of soundness, privacy, and compactness remain the same as in [LT22, Thm 7.2].

From the above analysis, we have $d = \lfloor \frac{\bar{d}D}{\lceil \log_t(N) \rceil + 1} \rfloor$ buckets, each expected to be assigned at most N/d messages, as the number of messages that are detected as pertinent is trivially bounded by N . But bucket counters may overflow, i.e., get incremented by more than t assigned messages detected as pertinent (whether true positives or false positives). We bound this overflow probability as follows.

$$\begin{aligned}
\Pr[X \geq t] &< \Pr[X \geq 2N/D] \quad (\text{since } N < Dt/2) \\
&= \Pr[X \geq 2(N/D)] \\
&= \Pr \left[X \geq 2 \left(\frac{N}{d \cdot (\lceil \log(N) \rceil + 1)} \right) \left(\frac{d \cdot (\lceil \log(N) \rceil + 1)}{D} \right) \right] \\
&\leq \exp \left(- \frac{\delta^2}{2 + \delta} \frac{N}{d(\lceil \log(N) \rceil + 1)} \right) \quad (\text{by Chernoff bound, where } \delta = \frac{2d(\lceil \log(N) \rceil + 1)}{D} - 1 = 2\bar{d} - 1) \\
&= \exp \left(- \frac{\delta^2}{2 + \delta} \frac{N}{\bar{d}D} \right) \\
&\leq \exp \left(- \frac{(2\bar{d} - 1)^2}{2\bar{d} + 1} \frac{t/2}{\bar{d}} \right)
\end{aligned}$$

By the union bound, the probability of none of the d buckets overflowing is $d \exp(-\frac{(2\bar{d}-1)^2}{2\bar{d}+1} \frac{t/2}{\bar{d}}) < \epsilon_n/4$, where $d = O(\log(\epsilon_n^{-1}))$. Therefore, for $N < Dt/2$, the condition at line 23 gives us a failure probability $< \epsilon_n/4$.

Therefore, all five conditions together have a failure probability of $\epsilon_n + \text{negl}(\lambda)$ for $k \leq \bar{k}$. \square

Due to these optimizations, our improved OMR-OPT is more efficient than OMRp2 in [LT22] for most applications. Thus, our GOMR construction will be based on OMR-OPT.

6 Ad-hoc Group OMR

We proceed to construct Ad-hoc Group OMR (AGOMR), where a sender can send each message to any arbitrary set of up to G recipients, as defined in Section 3.1.

6.1 Trivial Solution

AGOMR can be directly but inefficiently realized using any OMR scheme (such as those in [LT22]): for each group message, and each of its G recipients, the sender sends a separate single-recipient OMR message with a clue dedicated to that recipient (i.e., a PVW ciphertext in the case of [LT22]). The number of clues in the board then increases by a factor of G , and therefore so does the detector's work.

When the above construction is instantiated with OMR-OPT scheme described in Section 5, the detector first performs G homomorphic PVW decryption to obtain G encrypted bits per message. If the message is pertinent, one of these encrypted bits would be 1, otherwise the bits would be all 0's (with high probability).

After the bits are obtained, the detector performs the remaining procedure G times per message. This is wasteful, since there is at most one 1 in G bits. Ideally, the detector can sum up all G bits and performs the

Algorithm 2 OMR-OPT: Improved Practical Compact Oblivious Message Retrieval

```
1: procedure OMR-OPT.GenParams( $1^\lambda, \epsilon_p, \epsilon_n$ )
2:   Choose  $\text{pp}_{\text{BFV}} = (D, t, \dots)$ ,  $\text{pp}_{\text{PVW}} = (n, w, \ell, q, \sigma)$ , and range  $r$  with one change:
3:   Replace item 3 with  $\ell \cdot (1 - \text{erf}(r/(\sqrt{2}w\sigma))) < \epsilon_n/4$   $\triangleright$  See Setting parameters
4:   return  $\text{pp} = (1^\lambda, \epsilon_n, \epsilon_p, \text{pp}_{\text{BFV}}, \text{pp}_{\text{PVW}}, r)$   $\triangleright$  Provided implicitly below
5: procedure OMR-OPT.KeyGen
6:    $(\text{sk}_{\text{pvw}}, \text{pk}_{\text{pvw}}) \leftarrow \text{LWE.KeyGen}()$ 
7:    $(\text{sk}_{\text{BFV}}, \text{pk}_{\text{BFV}}) \leftarrow \text{BFV.KeyGen}()$ 
8:    $\text{ct}_{\text{pvwSK}} \leftarrow \text{BFV.Enc}(\text{pk}_{\text{BFV}}, \text{sk}_{\text{pvw}})$ 
9:   return  $(\text{sk} = (\text{sk}_{\text{BFV}}), \text{pk} = (\text{pk}_{\text{clue}} = \text{pk}_{\text{pvw}}, \text{pk}_{\text{detect}} = (\text{pk}_{\text{BFV}}, \text{ct}_{\text{pvwSK}})))$ 
10: procedure OMR-OPT.GenClue( $\text{pk}_{\text{clue}}, x$ )
11:    $\vec{m} \leftarrow (0, \dots, 0) \in \mathbb{Z}_t^\ell$ 
12:    $c \leftarrow \text{LWE.Enc}(\text{pk}_{\text{clue}}, \vec{m})$   $\triangleright$  Recall: clue  $c \in \mathbb{Z}_t^{n \times \ell}$ 
13:   return  $c$ 
14: procedure OMR-OPT.Retrieve( $\text{BB}, \text{pk}_{\text{detect}}, \bar{k}$ )  $\triangleright$  Phase 1: Initialization
15:   Draw a random seed  $s = (s_f, s_h)$ 
16:   Parse  $\text{BB} = \{(x_1, c_1), \dots, (x_N, c_N)\}$  and  $\text{pk}_{\text{detect}} = (\text{pk}_{\text{BFV}}, \text{ct}_{\text{pvwSK}})$ 
17:   Let  $C \leftarrow N/(D \cdot \log(t))$ 
18:   Initialize  $(\text{Acc}_i = \text{BFV.Enc}(\text{pk}_{\text{BFV}}, (0, \dots, 0)))_{i \in [C]}$   $\triangleright D$  zeros
19:    $\triangleright$  Phase 2: detection in batches of  $D$  messages
20:   Find the smallest  $d', \bar{d}$  and let  $d = \lceil \frac{\bar{d}D}{\lceil \log(N) \rceil + 1} \rceil$  such that:
21:   (1)  $1 - \prod_{i=1}^{\bar{k}-1} (1 - (\frac{i}{\bar{d}})^{d'}) < \epsilon_n/4$ 
22:   (2)  $d \cdot \exp(-\frac{N(2\bar{d}-1)^2}{(2\bar{d}+1)dD}) \leq \epsilon_n/4$ 
23:   Initialize  $(C_{\text{lhs},z} \leftarrow (\text{BFV.Enc}(\text{pk}_{\text{BFV}}, (0, \dots, 0))_w)_{w \in [d]})$ 
24:    $\triangleright$  encrypted  $d'\bar{d}D$  zeros as a single concatenated BFV ciphertext, to avoid detailed mod calculation
25:   for  $i = 1$  to  $N/D$  do  $\triangleright$  Assume wlog that  $D$  divides  $N$ 
26:     Parse each clue as  $c_{iD+i'} = ((c_{i',\kappa}))_{\kappa \in [n+\ell]} \in \mathbb{Z}_q^{n+\ell}$  for  $i' \in [D]$ 
27:      $\bar{c}_\kappa \leftarrow (c_{i',\kappa})_{i' \in [D]}$  for  $\kappa \in [n+\ell]$ 
28:      $\triangleright \bar{c}_\kappa$  lists the  $\kappa$ -th element of every PVW clue in this batch
29:      $\alpha_1 \leftarrow \text{InnerProd}(\text{pp}_{\text{BFV}}, \text{pk}_{\text{BFV}}, \text{ct}_{\text{pvwSK}}, (\bar{c}_\kappa)_{\kappa \in [n+\ell]})$ 
30:      $\alpha_2 \leftarrow \text{RangeCheck}(\text{pp}_{\text{BFV}}, \text{pk}_{\text{BFV}}, \alpha_1, r)$ 
31:      $\alpha_3 = \prod_{i=0}^{\ell-1} (1 - \alpha_2[i])$ 
32:      $(\text{PV}_{i'})_{i' \in ((i-1)D, iD]} \leftarrow \text{PVUnpack}(\text{pp}_{\text{BFV}}, \text{pk}_{\text{BFV}}, \alpha_3)$ 
33:     for  $i = 1$  to  $N$  do
34:        $j, k \leftarrow f_{s_f}(i)$   $\triangleright j \in [C], k \in [d]$ 
35:        $a \leftarrow (k-1) \cdot \lceil \log(N) \rceil + 1$ 
36:        $b \leftarrow k \cdot \lceil \log(N) \rceil + 1$ 
37:        $\triangleright$  We need to update slot  $[a, b-1]$ ; first  $b-a-2$  slots are for accumulator and last one is for counter
38:        $C_{\text{lhs},j}[a, b-2] = C_{\text{lhs},j}[a, b-2] + i \cdot \text{PV}_i$   $\triangleright$  t-ary addition
39:        $C_{\text{lhs},j}[b-1] = C_{\text{lhs},j}[b-1] + 1 \cdot \text{PV}_i$   $\triangleright$  Slot-wise addition, done homomorphically
40:        $\triangleright$  Phase 3: Finalization
41:        $\hat{k} \leftarrow \bar{k} + N \log(N) \epsilon_p$ 
42:        $(\text{pp}_{\text{SRLC}}, m) \leftarrow \text{SRLC.GenParams}(1^\lambda, \hat{k}, \epsilon_n/4, t)$ 
43:        $\triangleright$  In practice  $(\text{pp}_{\text{SRLC}}, m)$  is preprocessed and tabulated and therefore becomes  $O(1)$ 
44:       Initialize combinations  $\{\text{Cmb} = \text{BFV.Enc}(\text{pk}, 0)\}_{k \in [m]}$ 
45:       for  $i = 1$  to  $N$  do
46:          $S \leftarrow \text{SRLC.GenWeights}(\text{pp}_{\text{SRLC}})$ 
47:         for  $(j, w_j) \in S$  do
48:            $\text{Cmb}_j = \text{Cmb}_j + \text{PV}_i \cdot x_i \cdot w_j$   $\triangleright$  by homomorphic addition and scalar multiplication
49:       return  $M = (s, (\text{Acc}_{\text{lhs},i})_{i \in [C \cdot d/D]}, (\text{Ctr}_{\text{lhs},i})_{i \in [C \cdot d/D]}, (\text{Cmb}_k)_{k \in [m]}, \text{pp}_{\text{SRLC}})$ 
50: procedure OMR-OPT.Decode( $\text{pp}, M, \text{sk}$ )
51:   As in [LT22], details omitted.
```

following procedures once per message. However, with a non-negligible false positive rate and the possibility that the clues may be maliciously formed, there are potentially multiple 1’s, which would cause retrieval failure. This can be solved by additional changes (omitted for brevity), but the cost remains dominated by the aforementioned blowup in the number of homomorphic PVW decryptions.

6.2 Improved Ad-hoc Group OMR

The main detection cost, in the trivial scheme above, lies in the G -fold increase in expensive homomorphic operations: homomorphic decryption of PVW ciphertexts and the digest compression process. In this section, we use a different technique to construct AGOMR scheme so that the decryption process and digest compression only need to be performed once for each message. (We generalize this to the case of multiple payloads in Section 6.3.)

Polynomial interpolation encoding. Let \mathcal{D} be a finite field (to be fixed later). Each recipient randomly draws $\text{id} \leftarrow_{\S} \mathcal{D}$ as its unique public identity and includes the id as part of its pk_{clue} . To send a payload to a group of G recipients, the sender first generates G PVW ciphertexts: $(c_i \in \mathbb{Z}_t^{n+\ell})_{i \in [G]}$ all encrypting 1’s using the corresponding PVW public keys included in the pk_{clue} of those G recipients. But instead of directly sending these G PVW ciphertexts as clues (as in the trivial solution), the sender encodes them into a polynomial to be used as the clue. Specifically, the sender interpolates a degree- $(G - 1)$ polynomial $f : \mathcal{D} \rightarrow \mathbb{Z}_t^{n+\ell}$ such that $f(\text{id}_i) = c_i$ for $i \in [G]$, and publishes the coefficients of this polynomial as the clue. To guarantee all the id ’s are unique, we need \mathcal{D} to be large enough (e.g., $\mathcal{D} = \mathbb{Z}_{2^{127}-1}$).

During retrieval, the detector takes the recipient’s id included in the detection key. For each message $(x_i, f_i)_{i \in [N]}$, it evaluates $f_i(\text{id})$ (in plaintext form), to obtain a PVW ciphertext. Thus, if id is one of the points over which f is interpolated, then $f(\text{id})$ is a PVW ciphertext encrypting 1^ℓ . Otherwise, $f(\text{id})$ is an independently and pseudorandomly sampled vector of size $\mathbb{Z}_t^{n+\ell}$ w.r.t to this recipient’s decryption key and will be decrypted to 1^ℓ with probability $\leq \epsilon_p$. Hence, after obtaining the PVW ciphertext (i.e., $f(\text{id})$), we are in the same situation as in the detection procedure of OMR-OPT, and can proceed likewise, performing a *single* homomorphic decryption of a PVW ciphertext .

Since the PVW ciphertexts are computationally indistinguishable from uniformly-drawn vectors in $\mathbb{Z}_t^{n+\ell}$, the interpolated polynomial is thus indistinguishable from a random polynomial, in the absence of recipients’ secret keys. Even to an intended recipient, other intended recipients’ identities are unknown. Computational privacy follows.

Remark 6.1. The above yields an AGOMR scheme, AGOMR1, which achieves computational privacy and in whose detection is particularly efficient — since the extra computation to handle groups is a mere polynomial evaluation in plaintext. The drawback of this construction is that each retrieval query reveals the querying recipient (but not their pertinent messages) by including the identity id in the detection key, so it does not achieve detection key unlinkability (Definition 3.3). The following constructions address this and achieve key-unlinkability.

Detection-key-unlinkability for the poly-interpolation-based method. To achieve detection-key unlinkability, the biggest challenge is that now the detection key includes an id . The other parts of the detection key simply contains the BFV public key, and an encryption of the PVW secret key. Thus, instead of evaluating $f(\text{id})$ in plaintext, we can encrypt id using BFV and evaluate $f(\text{id})$ homomorphically.

Thus, RegenDetectKey is also simple: freshly generate new BFV public keys, encrypt the PVW secret key and the id .

However, in this case, the function f is also evaluated under FHE, so it needs to be FHE-friendly.

FHE-friendly multilinear encoding. As our schemes are based on BFV leveled-HE scheme, we need the multiplicative depth to be small (e.g., ≤ 30 levels) and the plaintext modulus to be reasonable (normally, 16–50 bits, and preferably ≤ 20 bits given our detector circuit). However, homomorphically evaluating the polynomial above requires a multiplicative depth to be $\log(G)$, and requires the FHE plaintext space to be a large finite field (e.g., $\mathbb{Z}_{2^{127}-1}$). While the multiplicative depth can be reduced to 1 by sending $\text{Enc}(\text{id}), \text{Enc}(\text{id}^2), \dots, \text{Enc}(\text{id}^G)$, the large plaintext modulus remains to be an issue. Therefore, we change

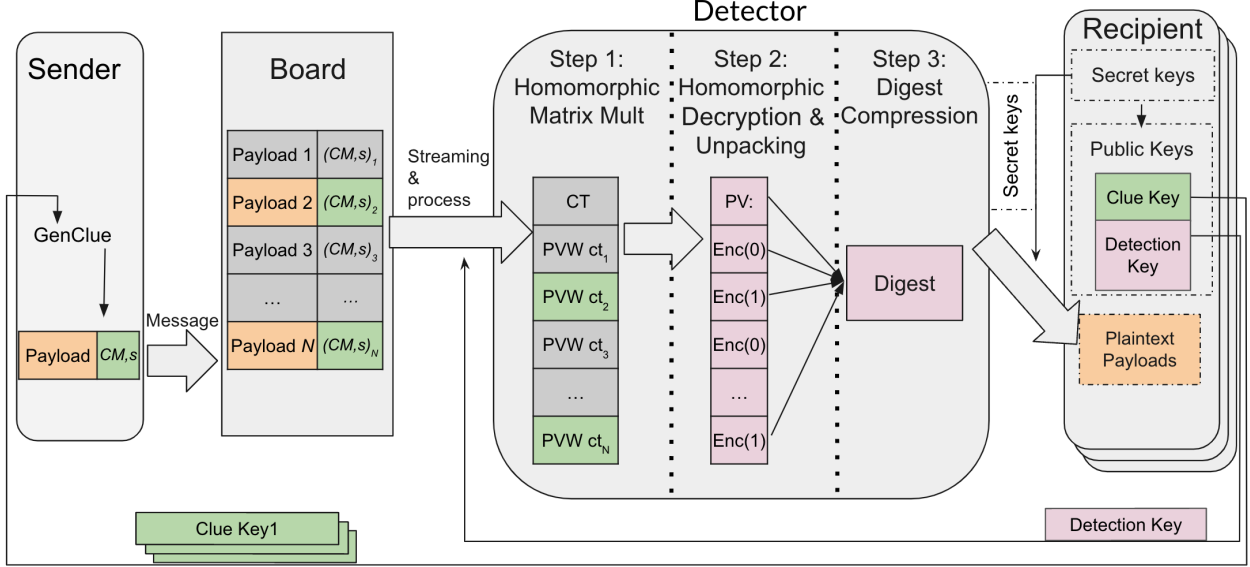


Figure 5: Main components of our Ad-hoc Group OMR construction.

the identities to be vectors $\text{id} \in \mathbb{Z}_t^I$ for some much smaller t (with some large enough I to be fixed later), and then interpolate a *multilinear* function $f : \mathbb{Z}_t^I \rightarrow \mathbb{Z}_t^{n+\ell}$ instead. In this case, the function f can be represented by a matrix $M \in \mathbb{Z}_t^{I \times (n+\ell)}$. To recover the clue, the detector just needs to compute $c_i = \text{id}_i \times M$, which is a homomorphic matrix multiplication of multiplicative depth 1.

Remark 6.2. The above yields an AGOMR scheme, AGOMR2, which is fully secure, efficient, but slightly weaker AGOMR scheme. By encrypting the id vectors, this scheme satisfies the detection-key unlinkability. However, completeness of this scheme requires the groups to be honestly formed; specifically, that the groups are formed independent of the participant id values (details below). Whether assumption is apt depends on the underlying applications, and at worst, a group that cannot form a clue due to maliciously chosen id s can be separated into multiple subgroups to avoid the issue. A better solution is discussed next.

Dealing with a maliciously formed group of ID. For such a matrix M as above to exist, we need all the G id 's used to generate the clue to be linearly independent. Since all id 's are randomly generated, for an honestly formed group, the probability that those randomly chosen G id 's are linearly independent is roughly $1 - t^{-(G-I)}$ for $I > G$ by [SGGC14b, Lemma 1]. However, the groups may be formed maliciously after seeing the id 's. To mitigate this we can increase the length I of the id 's. Let ϵ_{DI} denote the probability that given P randomly drawn id 's, there exists a combination of G of those id 's are linearly dependent. By union bound, $\epsilon_{\text{DI}} = \binom{P}{G} \prod_{i=G+1}^I (1 - 1/t^i)$. Setting $I = O(G \log(P))$ suffices to achieve $\epsilon_{\text{DI}} \leq \text{negl}(\lambda)$, mitigating the issue of maliciously formed groups. Then, to send to G recipients, the sender then solves for $M \in \mathbb{Z}_t^{I \times (n+\ell)}$ as before.

Compressing the clue. With the enlarged I , publishing the whole matrix M as clue results in a $O(G \log(P))$ clue size, which is costly (concretely and asymptotically) when the number P of system participants is large. To reduce its size, we use pseudorandomness derived from a succinct seed, as follows. The sender draws a random seed s , and uses it to generate a pseudorandom matrix $Z \in \mathbb{Z}_t^{I \times G'}$ (where $G' \geq G$ is set below). The sender computes $C \leftarrow W \times Z \in \mathbb{Z}_t^{G \times G'}$, which is indistinguishable from uniform random as shown in Lemma 6.3, where $W = \begin{pmatrix} \text{id}_1 \\ \vdots \\ \text{id}_G \end{pmatrix}$. Let ϵ_{DS} denote the probability that C is not full rank given that

W is full rank. Since C is computationally indistinguishable from random, again, by [SGGC14b, Lemma 1], $\epsilon_{\text{DS}} = \prod_{G'-G+1}^{G'} (1 - 1/t^i)$. Thus, to make $\epsilon_{\text{DS}} = \text{negl}(\lambda)$, we have $G' = G + O(\lambda)$. Then, the sender computes the matrix $M_{\mathbf{s}} \in \mathbb{Z}_t^{G' \times (n+\ell)}$ such that $C \times M_{\mathbf{s}} = (c_1, \dots, c_G)^\top$. The clue is of the form $(M_{\mathbf{s}}, \mathbf{s})$, which is of size $O(G)$.

Lemma 6.3. For any $0 < G \leq \bar{G}$, given any matrix $A \in \mathbb{Z}_t^{G \times \bar{G}}$ with rank G , the distribution $D = \{Y : Y = A \times Z, Z \leftarrow_{\mathfrak{s}} \mathbb{Z}_t^{\bar{G} \times G'}\}$ is equivalent to the uniform distribution $D_U = \{Y : Y \leftarrow_{\mathfrak{s}} \mathbb{Z}_t^{G \times G'}\}$.

Proof sketch. As A has full row rank, $\dim(\text{kernel}(A)) = t^{\bar{G} \times G'} / t^{G \times G'} = t^{G'(\bar{G}-G)}$. Therefore, for any $Y \in \mathbb{Z}_t^{G \times G'}$, $|\{Z \mid Y = A \times Z\}| = \dim(\text{kernel}(A))$. Therefore, if we sample an $Z \in \mathbb{Z}_t^{\bar{G} \times G'}$ uniformly at random, for any $Y \in \mathbb{Z}_t^{G \times G'}$, $\Pr[A \times Z = Y] = \frac{1}{t^{G \times G'}}$. \square

Detector operations. The detector uses \mathbf{s} to recover Z and homomorphically computes $(\text{id} \times Z) \times M_{\mathbf{s}}$ to recover the PVW ciphertext c , where id is encrypted under the recipient's BFV public key included in $\text{pk}_{\text{detect}}$. It then homomorphically decrypts the PVW ciphertext and processes digest encodings, as in OMR-OPT.

Efficiency analysis. With this technique against maliciously chosen groups, the computation cost of recovering the clue grows to $\tilde{O}(G \cdot (G \log(P) + n + \ell))$. However, compared to the baseline scheme OMRp2, which requires G expensive PVW ciphertext decryption, involving expensive range checks, and another G expensive digest compression, those $\tilde{O}(G \cdot (G \log(P) + n + \ell))$ computations are all relatively cheap plaintext-by-ciphertext multiplications. The resulting scheme is thus much lighter.

Combining all of the above, we obtain our main AGOMR construction AGOMR3 given in Algorithm 3. Fig. 5 portrays the high-level components of the resulting scheme. Compared to Fig. 4, the major changes are the clues and step 1. Steps 2 and 3 are unchanged.

Theorem 6.4. The scheme AGOMR3 in Algorithm 3 is an AGOMR scheme for $N < D \cdot t/2$, $P = \text{poly}(\lambda)$, and $G = \text{poly}(\lambda)$, assuming hardness of LWE and security of BFV leveled HE; when instantiated with PRF f and an SRLC scheme SRLC (Definition B.1). Moreover when instantiated with SRLC1 (Algorithm 6), AGOMR3 is also compact.

Above, *SRLC (Sparse Random Linear Code)* refers to schemes for encoding values as linear combinations drawn from certain distributions, as defined in [LT22, Section 6.3]. *SRLC1* refers to a specific such scheme, defined in [LT22, Section 6.3.1]. These are internal components inherited from [LT22], used in parts of *compression* phase of OMR that we inherit unaltered.

Proof sketch. Completeness: Most part of the completeness is the same as the completeness proof of Theorem 5.1. The only additional argument needed is that we now need the clue for the pertinent message can be successfully generated given any subset of recipients, by solving the matrix M with compressed id matrix $\tilde{\text{id}}$ defined at Algorithm 3 line 22.

Based on Lemma 6.3 and the parameter choice at line 4, the matrix $\tilde{\text{id}}$ is of full rank with probability $(1 - \text{negl}(\lambda)) \cdot \prod_{i=G'-G-1}^{G'} (1 - 1/t^i)$ [SGGC14a, Lemma 1]. Thus, by line 5, $\tilde{\text{id}}$ is full rank with probability $1 - \text{negl}$. Therefore, AGOMR3.GenClue can successfully compute M via a solvable linear equation system with $1 - \text{negl}$ probability.

The remainder of the completeness argument is as in [LT22, Theorem 7.2].

Soundness: For some $j \in [N]$, for $1 \notin X_j$ ⁷, given the random id_1 independently generated from M_i , $\text{id}_1 \times M_i$ results in a vector (\vec{a}, \vec{b}) that is indistinguishable from a vector sampled uniformly at random $\in \mathbb{Z}_t^{n+\ell}$. In this case, false positives occur when the inner products of all ℓ parts of this PVW ciphertext with sk_{PVW_1} fall into range $\pm r$. This has probability $((2r + 1)/t)^\ell \leq \epsilon_p$.

Computational privacy: We first argue that the ciphertexts $(c_i)_{i \in Y_0}$ generated in line 18 are indistinguishable from $(c'_i)_{i \in Y_1}$, where Y_0, Y_1 represent groups of recipients. Based on the hardness of LWE, all the honestly generated public keys $\text{pk}_{\text{PVW}} \in \mathbb{Z}_t^{w \times n + \ell}$ are computationally indistinguishable from a matrix

⁷Recall that in Definition 3.2, the board generation process output the corresponding ground truth for recipient 1 and thus our soundness argument is with respect to recipient 1 as well.

Algorithm 3 AGOMR3: Compact Ad-hoc Group Oblivious Message Retrieval

```

1: procedure AGOMR3.GenParams( $1^\lambda, \epsilon_p, \epsilon_n, G, P$ )
2:   Generate  $\text{pp}_{\text{BFV}}, \text{pp}_{\text{PVW}}$  as in OMR-OPT
3:   If  $t \leq G^2$ , choose  $t = G^2$  and  $\text{pp}_{\text{PVW}}$  accordingly.
4:   Choose the smallest  $I$  such that  $\binom{P}{G} \cdot \prod_{j=I-G-1}^I (1 - 1/t^j) \leq 2^{-\lambda}$ .
5:   Choose the smallest  $G' \geq G$  such that  $\prod_{i=G'-G-1}^{G'} (1 - 1/t^i) \leq 2^{-\lambda}$ .
6:   return  $\text{pp} = (1^\lambda, \epsilon_p, \epsilon_n, G, \text{pp}_{\text{BFV}}, \text{pp}_{\text{PVW}}, G', I)$ 
7: procedure AGOMR3.KeyGen( $\text{pp}$ )
8:    $\text{id} \leftarrow_{\S} \mathbb{Z}_t^I$ 
9:    $(\text{sk}_{\text{PVW}}, \text{pk}_{\text{PVW}}) \leftarrow \text{PVW.KeyGen}(\text{pp}_{\text{PVW}})$ 
10:   $(\text{sk}_{\text{BFV}}, \text{pk}_{\text{BFV}}) \leftarrow \text{BFV.KeyGen}(\text{pp}_{\text{BFV}})$ 
11:   $\text{ct}_{\text{SWK}} \leftarrow \text{BFV.Enc}(\text{pk}_{\text{BFV}}, \text{sk}_{\text{PVW}})$ 
12:   $\text{ct}_{\text{ID}} \leftarrow \text{BFV.Enc}(\text{pk}_{\text{BFV}}, \text{id})$ 
13:  return  $(\text{sk} = (\text{sk}_{\text{BFV}}, \text{sk}_{\text{PVW}}, \text{id}), \text{pk} = (\text{pk}_{\text{clue}} = (\text{id}, \text{pk}_{\text{PVW}})), \text{pk}_{\text{detect}} = (\text{pk}_{\text{BFV}}, \text{ct}_{\text{ID}}, \text{ct}_{\text{SWK}}))$ 
14: procedure AGOMR3.GenClue( $\text{pp}, \{\text{pk}_{\text{clue}_i} = (\text{id}_i, \text{pk}_{\text{PVW}_i})\}_{i \in [G]}, x$ )
15:  $\triangleright$  If in the amount of input  $\text{pk}_{\text{clue}}$  is less than  $G$ , run AGOMR3.KeyGen to generate new  $\text{pk}_{\text{clue}}$ 's
    until there are  $G$  of  $\text{pk}_{\text{clue}}$ 's
16:   for  $i = 1 \dots G$  do
17:      $\vec{m} \leftarrow (0, \dots, 0) \in \mathbb{Z}_t^\ell$ 
18:      $c_i \leftarrow \text{LWE.Enc}(\text{pp}_{\text{PVW}}, \text{pk}_{\text{PVW}_i}, \vec{m})$   $\triangleright$  All encryptions are done with distinct randomness
19:      $\triangleright$  Recall:  $\text{clue } c_i \in \mathbb{Z}_t^{n \times \ell}$ 
20:   Draw a random seed  $s$ 
21:   Use  $s$  to sample  $Z$  from  $\mathbb{Z}_t^{I \times G'}$ 
22:    $\tilde{\text{id}} = \begin{pmatrix} \text{id}_1 \\ \vdots \\ \text{id}_G \end{pmatrix} \times Z$   $\triangleright \tilde{\text{id}} \in \mathbb{Z}_t^{G \times G'}$ 
23:   Solve for a matrix  $M \in \mathbb{Z}_t^{G' \times (n+\ell)}$  such that  $\tilde{\text{id}} \times M = \begin{pmatrix} c_1 \\ \vdots \\ c_G \end{pmatrix}$ 
24:   return  $(M, s)$ 
25: procedure AGOMR3.Retrieve( $\text{pp}, \text{BB}, \text{pk}_{\text{detect}}(\text{pk}_{\text{bfv}}, \text{ct}_{\text{ID}}, \text{ct}_{\text{SWK}}), \bar{k}$ )
26:   For all  $i \in [|\text{BB}|]$ , load  $(M_i, s_i)$ 
27:   Use  $s_i$  to sample  $Z_i$  from  $\mathbb{Z}_t^{I \times G'}$ 
28:   Compute  $c_i \leftarrow (\text{ct}_{\text{ID}} \times Z_i) \times M_i$ 
29:   Proceed as OMR-OPT.Retrieve with  $(c_i)_{i \in N}$  from line 17
30: procedure AGOMR3.Decode( $\text{pp}, M, \text{sk}$ )
31:   Same as OMR-OPT.Decode
32: procedure AGOMR3.RegenDetectKey( $\text{pp}, \text{sk} = (\text{sk}_{\text{BFV}}, \text{sk}_{\text{PVW}}, \text{id})$ )
33:    $(\text{sk}'_{\text{BFV}}, \text{pk}'_{\text{BFV}}) \leftarrow \text{BFV.KeyGen}(\text{pp}_{\text{BFV}})$ 
34:    $\text{ct}'_{\text{SWK}} \leftarrow \text{BFV.Enc}(\text{pk}'_{\text{BFV}}, \text{sk}_{\text{PVW}})$ 
35:    $\text{ct}'_{\text{ID}} \leftarrow \text{BFV.Enc}(\text{pk}'_{\text{BFV}}, \text{id})$ 
36:   return  $(\text{sk}' = (\text{sk}'_{\text{BFV}}, \text{sk}_{\text{PVW}}, \text{id}), \text{pk}'_{\text{detect}} = (\text{pk}'_{\text{BFV}}, \text{ct}'_{\text{ID}}, \text{ct}'_{\text{SWK}})$ 

```

$U \leftarrow_{\S} \mathbb{Z}_t^{w \times n + \ell}$ sampled uniformly at random. Thus, by leftover hash lemma, given that all the Enc calls use distinct randomness, $\mathcal{D}((c_i)_{i \in Y_0 \setminus Z}) \approx_c \mathcal{D}((c'_i)_{i \in Y_1 \setminus Z})$. Thus, $\mathcal{D}((c_i)_{i \in Y_0}) \approx_c \mathcal{D}((c'_i)_{i \in Y_1})$

If one can distinguish (M, Z) generated from $(c_i)_{i \in Y_0}$ and (M', Z') generated from $(c'_i)_{i \in Y_1}$, one can distinguish $(c_i)_{i \in Y_0}$ from $(c'_i)_{i \in Y_1}$, as the entire generation process can be simulated by anyone given the public keys, which contradicts to the claim that the distribution of $(c_i)_{i \in Y_0}$ and $(c'_i)_{i \in Y_1}$ are computationally indistinguishable. Hence, computational privacy holds.

Compactness: The digest size is identical to that of OMR-OPT.

Detection-key-unlinkability: pk_{BFV} is simply a fresh BFV public key. $\text{ct}_{\text{ID}}, \text{ct}_{\text{SWK}}$ are both BFV ciphertexts. Thus, by the security of BFV, all these are computationally indistinguishable from other detection keys generated freshly using KeyGen with different sk. \square

Amortization and preprocessing. To further reduce the concrete cost, we take some of the computation offline. A plaintext-by-ciphertext multiplication in BFV contains three components: (1) NTT transformation of the BFV ciphertext (resulting in the “NTT form” of the ciphertext); (2) NTT transformation of the plaintext⁸; (3) multiplication between the two NTT form components. Costs are in descending order. We try to reduce the number the first two operations, and amortize the costs over different recipients in the following way:

- Notice that when performing the linear function evaluation (i.e., matrix multiplication), multiple plaintexts generated based on M_s 's are multiplied with the same ciphertext (i.e., encrypted id). Thus, we cache the NTT form of the encrypted id to reduce the number of the first type of operation.
- The plaintext NTT transformations for all published M's are shared by all recipients; therefore, the runtime of this part can be amortized over all recipients registered under the same detector.

6.3 Extension to Distinct Payloads

In some applications, a sender addresses a group of recipients with different payloads in a single message (e.g., a Zcash transaction may have multiple recipients, each of which cares only about the data describing the output note they can spend). Our AGOMR definition (Definition 3.1) can be extended to fit this multiple payload setting, where the correctness, privacy, and compactness definitions remain exactly the same, and the soundness extends in the most natural way (i.e., payloads not intended for the recipient should be excluded w.h.p.).

First, notice that this can be trivially implemented through our AGOMR (or OMR) scheme by concatenating all the payloads to a giant payload. However, this (1) increases the digest size from PL to $G \cdot \text{PL}$ (which would break the compactness), (2) the soundness guarantee is violated, as impertinent payloads that are concatenated with the pertinent ones.

The following modification to AGOMR3 resolves this issues and achieves compact multi-payload AGOMR.

- Similarly to the multi-linear interpolation performed for different clues, with G distinct payloads $(x_1, \dots, x_G) \in \mathcal{P}^G$ intended for G recipients with id's $(\text{id}_1, \dots, \text{id}_G) \in \mathcal{D}^G$, the sender interpolates a function $h : \mathcal{D} \rightarrow \mathcal{P}$, such that $h(\text{id}_i) = x_i$ for $i \in [G]$. We represent function h as a matrix for the linear transformation from id_i to payload x_i .
- The detector uses the encrypted id in the detection key to evaluate $h(\text{id})$ homomorphically, and obtains a BFV ciphertext encrypting the payload⁹. Then, the rest almost follows exactly what [LT22] has, except that the payload is now a BFV ciphertext, and thus needs to perform ciphertext multiplications when applying Sparse Random Linear Coding (SRLC, [LT22, Section 6.3]). (Alternatively, one can also view the BFV ciphertext as a plaintext payload and proceed exactly as in [LT22].)

⁸BFV first encode messages into a polynomial by embedding the messages $\in \mathbb{Z}_t^D$ in the polynomial coefficients, and the polynomial is called a plaintext

⁹If the payload is larger than the plaintext space of a BFV ciphertext, we simply use multiple BFV ciphertexts, and for simplicity, we omit the details of this case.

Analysis.

- **Correctness and soundness:** As long as the function h can be successfully interpolated (guaranteed by the linear independency of the compressed IDs), the correctness of our scheme is guaranteed by the correctness of BFV, which follows the same as in the proof of Theorem 6.4. Soundness also naturally follows from the same analysis.
- **Privacy:** We now also require that h does not reveal the identity of the recipients. W.l.o.g., we assume all the payloads are drawn from the uniformly random distribution over some payload space \mathcal{P} (if not, we first encrypt the payload through a key-private CPA-secure encryption scheme to make them indistinguishable from random). Privacy is thus guaranteed the same way as the current AGOMR privacy.
- **Detector runtime:** Asymptotically, the runtime does not change much with such a modification. Concretely speaking, the evaluation of h is just a single level and can be done at a very low BFV multiplicative level (as it is only followed by some linear encoding), its cost is very small compared to the rest of the computation. Although the SRLC part will become slower as it requires ciphertext-by-ciphertext multiplication instead of plaintext-by-ciphertext multiplication, it does not contribute much to the overall detector runtime. Thus the total runtime is only slightly affected. (One could also view the BFV encryption of the payload as a BFV plaintext, and proceed exactly as before, at the cost of payload being enlarged.)
- **Digest size:** By homomorphically evaluating $h(\text{id})$, the detector will receive a BFV ciphertext encrypting a single payload. Therefore, the digest size remains the same and is independent of G .
- **Payload size:** The sender publishes the matrix representation of function h as payload, which is of size $\sim (G + \log(\lambda)) \cdot |\mathcal{P}|$ (concretely, for a statistical security parameter of ~ 48 , we have $(G + 4) \cdot |\mathcal{P}|$). This is slightly larger than the naive solution, which is of size $G \cdot |\mathcal{P}|$.

7 Fixed Group OMR

We proceed to construct Fixed Group OMR (FGOMR), where recipients form groups in advance, and the senders can subsequently send messages to these fixed groups (or subsets thereof), as defined in Section 3.2. This allows reduced detection cost compared to AGOMR.

7.1 Multi-Recipient Encryption

We achieve FGOMR via Multi-Recipient Encryption (MRE), defined below. At a high level, MRE enables the sender to encrypt multiple messages to multiple recipients at the same time. FGOMR can leverage the encryption function in MRE to generate a single clue for multiple recipients, and the detector then uses the decryption function in MRE to evaluate all clues on the board.

Our MRE definition is adapted from [BBKS07]. We extend it with a key-privacy property, since the adversary should not learn which group the message is addressed to unless a (fully malicious) adversary corrupts the recipients in the group. We also make a couple of minor relaxations compared to [BBKS07] (see Remark 7.1).

Definition 7.1. (Multi-Recipient Encryption). A Multi-Recipient Encryption scheme has the following PPT algorithms:

- $\text{pp} \leftarrow \text{GenParams}(1^\lambda, G, P)$: takes a security parameter λ , the number of recipients G in a group, and the total number of recipients P in the system; outputs a public parameter pp .
- $\text{sk} \leftarrow \text{SKGen}(\text{pp})$: takes a public parameter pp and outputs a secret key sk .
- $\text{pk} \leftarrow \text{PKGen}(\text{pp}, \text{sk})$: takes a public parameter pp , a secret key sk ; outputs a public key pk .

Adversary	Challenger
	$\text{pp} \leftarrow \text{GenParams}(1^\lambda, G, P)$ $\text{sk}_i \leftarrow \text{SKGen}(\text{pp})$ for $i \in [P]$
Choose groups $Z \subseteq [P], Z < G$	$\text{pk}_i \leftarrow \text{PKGen}(\text{pp}, \text{sk}_i)$ for $i \in [P]$
sample randomness $(r_i)_{i \in Z}$	
Choose $Y \subseteq [P] \setminus Z$, where $1 \leq Y \leq G - Z $	
Choose $m_0, m_1 \in \mathcal{M}^{ Y }, m \in \mathcal{M}^{ Y' }$	$Y' \leftarrow Z \cup Y$
	$\text{sk}_i \leftarrow \text{SKGen}(\text{pp}; r_i), \text{pk}_i \leftarrow \text{PKGen}(\text{pp}, \text{sk}_i)$ for $i \in Z$ overwriting the old ones $b \leftarrow \{0, 1\}$
	$c \leftarrow \text{Enc}(\text{pp}, (\text{pk}_i)_{i \in Y'}, m_b m)$
	(Adversary wins if $b = b'$)

Figure 6: CPA Security game for Multi-Recipient Encryption.

- $\text{ct} \leftarrow \text{Enc}(\text{pp}, \vec{\text{pk}}, \vec{m})$: takes a public parameter pp , a vector of up to G public keys $\vec{\text{pk}}$, a vector of messages $\vec{m}, |\vec{m}| = |\vec{\text{pk}}|$; outputs a ciphertext $\text{ct} \in \mathcal{C}$
- $m \leftarrow \text{Dec}(\text{pp}, \text{sk}, \text{ct})$: takes a public parameter pp , a secret key sk , a ciphertext ct ; outputs a message m

The scheme must satisfy the following properties:

- (Correctness) Let $\text{pp} \leftarrow \text{GenParams}(1^\lambda, G, P)$, for $j \in [P]$, let $\text{sk}_j \leftarrow \text{SKGen}(\text{pp}), \text{pk}_j \leftarrow \text{PKGen}(\text{pp}, \text{sk}_j)$. For any set of recipients $Y \subseteq [P], |Y| \leq G$, and any plaintext vector $(m_j)_{j \in [|Y|]} \in \mathcal{M}^{|Y|}$, for all $i \in [|Y|]$, it holds that:

$$\Pr[\text{Dec}(\text{pp}, \text{sk}_i, \text{Enc}(\text{pp}, (\text{pk}_j)_{j \in Y}, (m_j)_{j \in [|Y|]})) = m_i] \geq 1 - \text{negl}(\lambda)$$

- (CPA security) An MRE scheme is CPA secure if for any PPT adversary \mathcal{A} , it wins the game in Fig. 6 with probability $\leq 1/2 + \text{negl}(\lambda)$, where Y, Y', Z represent groups of recipients.
- (Key Privacy) An MRE scheme is key private if for any PPT adversary \mathcal{A} , it wins the game in Fig. 7 with probability $\leq 1/2 + \text{negl}(\lambda)$, where Y_0, Y_1, Z represent groups of recipients.

Remark 7.1. We adapted this primitive and its CPA security definition from [BBKS07], with the following changes: 1) we introduce key privacy; 2) unlike in [BBKS07], we require G and P to be known when GenParams is invoked (G can be equal to P as in [BBKS07]); 3) we relax correctness probability to $1 - \text{negl}(\lambda)$, rather than 1 as in [BBKS07]; 4) we separate KeyGen into SKGen and PKGen , since they are invoked separately in our case; 5) in [BBKS07], the adversary needs to first send the number of key pairs that it wants to maliciously overwrite before receiving all honestly generated public keys. Regarding the last point, notice that if there exists an adversary that breaks the CPA security in [BBKS07], then it trivially breaks ours. Therefore, our definition implies the one in [BBKS07].

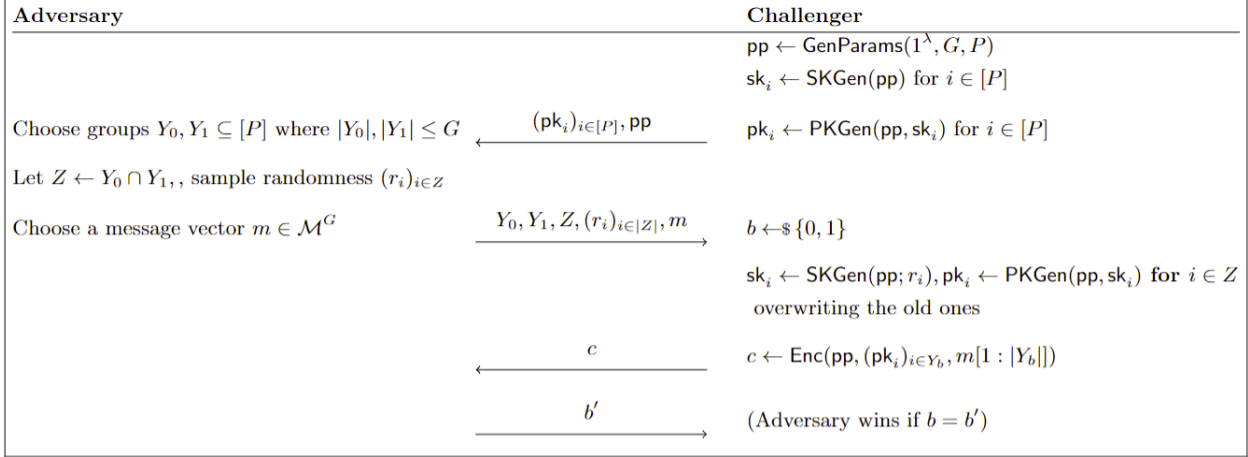


Figure 7: Key Privacy game for Multi-Recipient Encryption.

7.2 Naive Solutions

As a stepping stone, we discuss two naive approaches to constructing MRE schemes based on PVW encryption, and explain why they are inapplicable to our application.

A non-key-private solution. A trivial solution is to use the standard PVW encryption scheme as an MRE scheme. To encrypt to G recipients, simply generate G PVW ciphertexts $(\vec{a}_i, \vec{b}_i)_{i \in [G]}$. Furthermore, all recipients can share the same \vec{a} (i.e., $\vec{a}_i = \vec{a}_j$ for all $i, j \in [G]$), and thus the MRE ciphertext is of the form $(\vec{a}, \vec{b}_1, \dots, \vec{b}_G)$.

However, for correctness, the recipient needs to know which \vec{b}_i ($i \in [G]$) to use during decryption. To guarantee this, one way to do this is that the sender can simply attach a hash function h that hashes a public key pk into a specific \vec{b}_i (the sender can manually avoid collisions). The recipient then computes $h(pk)$ and finds out the correct \vec{b}_i . However, this construction does not satisfy the key privacy requirement, as that the intended recipients do not have collision already leaks information.

An inefficient solution. An alternative way is to let the encryption algorithm encrypts the same message λ times (i.e., the new ciphertext includes λ copies of the ciphertext above without $h(pk)$), where λ is the security parameter. Only if the recipient uses the correct key to decrypt, all the λ copies decrypt into the same message, except with $O(\exp(-\lambda))$ probability, by the wrong-key decryption property of PVW [LT22, Def 5.1]. The drawback is that all costs grow with λ .

Moreover, both solutions are not FHE friendly (the first one needs indexing, and the second one requires an equality check), and thus are hard to incorporate into our application.

If an application does not require key unlinkability for its FGOMR, the hash function can be evaluated directly in plaintext, in which case the first solution above is sufficient and more efficient than the key-private MRE below.

7.3 Efficient Key-private MRE

We now introduce our main Multi-Recipient Encryption construction, which will be used to construct our main FGOMR scheme in Section 7.4.

High-level Idea. Our starting point is PVW encryption as in [PVW08], and we use the same parameters (n, t, w, ℓ, σ) . For simplicity, we start with $\ell = 1$, i.e., given a group of G recipients, the plaintext is represented by $\vec{m} \in \mathbb{Z}_2^G$, one bit per recipient. Our goal is to generate a public key $pk = (A, \vec{\eta})$ such that $Ask_i \approx \vec{\eta}$ for all $i \in [G]$, where sk_i is the secret key for recipient i .

GenParams: The scheme sets up the public parameters \mathbf{pp} according to the PVW security and other probability analysis similar to Section 6.2. The public parameters \mathbf{pp} also include a random seed \mathbf{s} , from which each recipient in the group can generate a matrix $A \in \mathbb{Z}_t^{w \times n}$ and a vector $\vec{b} \in \mathbb{Z}_t^w$.

SKGen: Each recipient generates their secret keys as a random vector $\mathbf{sk} = \mathbf{ck} \parallel \mathbf{dk} \in \mathbb{Z}_t^{n+L}$, where $\mathbf{ck} \leftarrow_{\mathcal{S}} \mathbb{Z}_t^n$ and $\mathbf{dk} \leftarrow_{\mathcal{S}} \mathbb{Z}_t^L$. $L = G + O(\lambda)$ is chosen by a similar procedure as G' in Section 6.2.

PKGen: The recipient then computes $\vec{b}' \leftarrow \mathbf{ck}A^\top + \vec{e}$, where \vec{e} is a Gaussian noise vector, and publishes $\mathbf{pk} = (\vec{b}', \mathbf{dk})$.

Enc: A sender first draws $\vec{e} \leftarrow_{\mathcal{S}} \{0, 1\}^w$ and computes $\vec{a} \leftarrow \vec{e}A \in \mathbb{Z}_t^n$, $\beta \leftarrow \langle \vec{e}, \vec{b} \rangle \in \mathbb{Z}_t$ and $\beta'_i = \langle \vec{e}, \vec{b}'_i \rangle \in \mathbb{Z}_t$, for all $i \in [G]$. The sender then finds $\vec{\gamma} \in \mathbb{Z}_t^L$ such that $\langle \vec{\gamma}, \mathbf{dk}_i \rangle = \beta - \beta'_i$ for all $i \in [G]$. This is done by solving $\vec{\gamma} \mathbf{dk}_{\text{all}}^\top = \vec{\beta}''$, where $\mathbf{dk}_{\text{all}} = \begin{pmatrix} \mathbf{dk}_1 \\ \vdots \\ \mathbf{dk}_G \end{pmatrix} \in \mathbb{Z}_t^{G \times L}$, $\vec{\beta}'' = (\beta - \beta'_1 + m[1] \cdot \lceil t/2 \rceil) \parallel \dots \parallel (\beta - \beta'_G + m[G] \cdot \lceil t/2 \rceil) \in \mathbb{Z}_t^G$.

The equation system is solvable as long as \mathbf{dk}_{all} has full row rank. The ciphertext is $(\vec{a}, \vec{\gamma}, \beta)$.

Dec: The recipient uses its own secret key $\mathbf{sk} = \mathbf{ck} \parallel \mathbf{dk}$ to check $|\beta - \langle \vec{a} \parallel \vec{\gamma}, \mathbf{sk} \rangle| \geq \lceil t/4 \rceil$.

To make the process above resistant to maliciously chosen recipients, we again extend the size of \mathbf{dk} to be long enough. Similarly to Section 6.2, we let ϵ_{DI} denote that the probability that \mathbf{dk}_{all} is not full row rank,

where $\mathbf{dk}_{\text{all}} \leftarrow \begin{pmatrix} \mathbf{dk}_1 \\ \vdots \\ \mathbf{dk}_G \end{pmatrix} \in \mathbb{Z}_t^{G \times L}$, and let ϵ_{DS} denote the probability that $\mathbf{dk}_{\text{all}} \times Z$ is not full row rank given

that \mathbf{dk}_{all} is full row rank, where $Z \leftarrow_{\mathcal{S}'} \mathbb{Z}_t^{L \times G'}$ is generated from a random seed \mathbf{s}' . We can efficiently achieve $\epsilon_{\text{DI}}, \epsilon_{\text{DS}} \leq \text{negl}(\lambda)$ as shown in the lemmas in Section 6.2. As before, the sender then solves for $\vec{a} \in \mathbb{Z}_t^{G'}$, such that $\vec{a}(\mathbf{dk}_{\text{all}} \times Z)^\top = \vec{\beta}''$. The seed \mathbf{s}' is appended to the ciphertext.

To extend this idea to $\ell > 1$, we simply repeat the above procedure with ℓ different secret keys \mathbf{sk} as in the PVW encryption [PVW08]. We use a single \mathbf{dk} instead of having ℓ different \mathbf{dk} for better efficiency (also using the same (\vec{a}, β)). The resulting ciphertext is of form $(\vec{a}, \alpha_1, \dots, \alpha_\ell, \beta, \mathbf{s}') \in \mathbb{Z}_t^n \times \mathbb{Z}_t^{G'} \times \dots \times \mathbb{Z}_t^{G'} \times \mathbb{Z}_t \times \mathbb{Z}_{2^\lambda}$, such that for each intended recipient with $(\mathbf{ck}_j, \mathbf{dk})_{j \in [\ell]}$, $\langle \vec{a} \parallel \alpha_j, \mathbf{ck}_j \parallel (\mathbf{dk} \times Z) \rangle \approx \beta$, where $Z \leftarrow_{\mathcal{S}} \mathbb{Z}_t^{GL \times G'}$, for all $j \in [\ell]$.

Security follows from the following. In the ciphertext, the distribution of $\vec{a}, \beta, \mathbf{s}'$ are all statistically indistinguishable from uniform and independent. Then, the distribution of \vec{b}' is computationally indistinguishable from uniform by the hardness of LWE, and $\vec{\beta}''$ is computationally indistinguishable from uniform by leftover hash lemma. Thus, the distribution of \vec{a} is computationally indistinguishable from uniform.

Algorithm 4 provides the full construction of our MRE construction, including some details omitted above for simplicity.

Theorem 7.2. For any $P = \text{poly}(\lambda), G = \text{poly}(\lambda)$, assuming hardness of LWE, the scheme MRE in Algorithm 4 is a key-private MRE scheme.

Proof sketch. Correctness: Similarly to the proof of correctness in Theorem 6.4, the probability that the equation system is unsolvable has probability $\text{negl}(\lambda)$. Therefore, we have $\text{Ask} \approx \vec{b}$ and the correctness trivially follows.

CPA security: We prove CPA security by showing that given \mathbf{pp} for all $Y \in [P]$ and $|Y| = G$, given $(\mathbf{pk}_i)_{i \in Y}$, for all $\vec{m} \in \mathcal{M}^{|Y|}$, it holds that $\text{ct} = (\vec{a}, \alpha_1, \dots, \alpha_\ell, \beta, \mathbf{s}') \leftarrow \text{Enc}(\mathbf{pp}, (\mathbf{pk}_i)_{i \in Y}, \vec{m})$ is computationally indistinguishable from a uniformly-drawn tuple $u \leftarrow_{\mathcal{S}} \mathbb{Z}_t^n \times \mathbb{Z}_t^{G'} \times \dots \times \mathbb{Z}_t^{G'} \times \mathbb{Z}_t \times \mathbb{Z}_{2^\lambda}$. First, \mathbf{s}' is a random seed, and \vec{a}, β are statistically indistinguishable from random vectors trivially. The only thing to prove is that $(\alpha_i)_{i \in [\ell]}$ is computationally indistinguishable from random vectors.

Notice that for each $\alpha_{i \in [\ell]}$, it is generated by solving equation $\vec{a} \mathbf{X} \mathbf{X}^\top = \beta''$ on line 24 in Algorithm 4. Thus, $(\alpha_i)_{i \in [\ell]}$ is indistinguishable from a random vector following from the fact that β'' is indistinguishable from a random vector by the hardness of LWE assumption.

Since the adversary can only choose the randomness used to generate the malicious keys, the analysis remains the same for those corrupted keys (as the indistinguishability result for other keys remain).

Such indistinguishability implies CPA security, as if there exists a PPT adversary who can break the CPA game, it can be used to break such indistinguishability.

Algorithm 4 MRE : Group Encryption

```

1: procedure MRE.GenParams( $\lambda, G, p$ )
2:   Choose  $(n, w, \ell, t, \sigma)$  according to PVW security and let  $\text{pp}_{\text{PVW}} \leftarrow (n, w, \ell, t, \sigma)$ 
3:   If  $t \leq G^2$ , choose  $t = G^2$  and  $\text{pp}_{\text{PVW}}$  accordingly.
4:   Choose the smallest  $L$  such that  $\binom{P}{G} \cdot \prod_{j=L-G-1}^L (1 - 1/t^j) \leq 2^{-\lambda}$ .
5:   Choose the smallest  $G' \geq G$  such that  $\prod_{i=G'-G-1}^{G'} (1 - 1/t^i) \leq 2^{-\lambda}$ .
6:   Generate a random seed  $s$ .
7:   return  $\text{pp} = (\lambda, \text{pp}_{\text{PVW}}, L, G, G', s)$ 

8: procedure MRE.SKGen( $\text{pp} = (\lambda, \text{pp}_{\text{PVW}} = (n, w, \ell, t, \sigma), L, G, G', s)$ )
9:   Choose a secret key  $\text{ck} = (\text{ck}_1, \dots, \text{ck}_\ell) \leftarrow \mathbb{Z}_t^{n \times \ell}$  uniformly at random.
10:  Choose an auxiliary key  $\text{dk} \leftarrow \mathbb{Z}_t^L$  uniformly at random.
11:  return  $(\text{ck}, \text{dk})$ 

12: procedure MRE.PKGen( $\text{pp} = (\lambda, \text{pp}_{\text{PVW}} = (n, w, \ell, t, \sigma), L, G, G', s), (\text{ck}, \text{dk})$ )
13:  Use  $s$  to randomly generate  $w$  vectors  $A = \{\vec{a}_i \in \mathbb{Z}_t^n\}_{i \in [w]}$ 
14:  Use  $s$  to randomly generate  $w$  vectors  $\vec{b} = \{\vec{b}_i \in \mathbb{Z}_t^\ell\}_{i \in [w]}$ 
15:  for  $i = 1 \rightarrow w$  do
16:    for  $l = 1 \rightarrow \ell$  do
17:       $b'_{i,l} \leftarrow b_{i,l} - \langle a_i, \text{ck}_l \rangle - \vec{e}$  where  $\vec{e}$  is from some Gaussian distribution  $\chi_\sigma$ .
18:  return  $\text{pk} = ((b'_{i,l})_{i \in [w], l \in [\ell]}, \text{dk})$ 

19: procedure EncAux( $\text{pp} = (\lambda, \text{pp}_{\text{PVW}}, L, G, G', s), (\vec{a}, \beta), (\beta'_j, \text{dk}_j)_{j \in [G]}, s'$ )
20:   $\text{dk}_{\text{all}} \leftarrow \begin{pmatrix} \text{dk}_1 \\ \vdots \\ \text{dk}_G \end{pmatrix} \in \mathbb{Z}_t^{G \times L}$ 
21:   $Z \leftarrow_{s'} \mathbb{Z}_t^{L \times G'}$ 
22:   $\text{XK} = \text{dk}_{\text{all}} \times Z, \text{XK} \in \mathbb{Z}_t^{G \times G'}$ 
23:   $\beta'' = (\beta - \beta'_1) \parallel \dots \parallel (\beta - \beta'_G) \in \mathbb{Z}_t^G$ 
24:  Solve the linear equation system  $\alpha \text{XK}^\top = \vec{\beta}''$  and let  $\alpha \leftarrow_{\S} \mathbb{Z}_t^{G'}$  if the equation system is underdetermined
25:  return  $\alpha$ 

26: procedure MRE.Enc( $\text{pp} = (\lambda, \text{pp}_{\text{PVW}}, L, G, G', s), \{\text{pk}_j = ((b'_{i,l,j})_{i \in [w], l \in [\ell]}, \text{dk}_j)\}_{j \in [G]}, \vec{m} \in \mathbb{Z}_t^\ell$ )
27:   $\triangleright$  If the amount of input  $\text{pk}$  is less than  $G$ , run MRE.PKGen to generate new  $\text{pk}$ 's until there are  $G$  of  $\text{pk}$ 's with the same seed  $s$  provided in the input
28:  Use  $s$  to generate  $w$  vectors  $A = \{\vec{a}_i \in \mathbb{Z}_t^n\}_{i \in [w]}$   $\triangleright$  Same  $A$  as generated by each party
29:  Use  $s$  to generate  $w$  vectors  $\vec{b} = \{\vec{b}_i \in \mathbb{Z}_t^\ell\}_{i \in [w]}$   $\triangleright$  Same  $B$  as generated by each party
30:  Define a vector  $\vec{t} = \frac{t}{2} \cdot \vec{m} \in \mathbb{Z}_t^\ell$ 
31:  Sample vector  $\vec{e} \leftarrow \{0, 1\}^w \in \mathbb{Z}_2^w$  uniformly at random.
32:   $(\vec{a}, \beta) \leftarrow (\vec{e}A, \vec{e}\vec{b} + \vec{t}) \in \mathbb{Z}_t^n \times \mathbb{Z}_t^\ell$ 
33:  for  $j = 1$  to  $G$  do
34:     $\beta'_{t,j} \leftarrow \vec{e}(b'_{i,l,j}) \in \mathbb{Z}_t^\ell$ , for  $l \in [\ell]$ 
35:  Generate some randomness  $s'$ 
36:  for  $l = 1$  to  $\ell$  do
37:     $\alpha_l \leftarrow \text{EncAux}(\text{pp}, (\vec{a}, \beta), (\beta'_{t,j}, \text{dk}_j)_{j \in [G]}, s')$ 
38:   $\text{ct} \leftarrow (\vec{a}, \alpha_1, \dots, \alpha_\ell, \beta, s')$ 
39:  return  $\text{ct}$ 

40: procedure MRE.Dec( $\text{pp} = (\lambda, \text{pp}_{\text{PVW}}, L, G, G', s), (\text{ck}, \text{dk}), \text{ct} = (\vec{a}, \alpha_1, \dots, \alpha_\ell, \beta, s')$ )
41:   $Z \leftarrow_{s'} \mathbb{Z}_t^{L \times G'}$   $\triangleright$  Same  $Z$  as generated during encryption
42:   $\text{dk}' \leftarrow \text{dk} \times Z$ 
43:   $\vec{d} = \vec{b} - (\langle \vec{a} | \alpha_1, \text{ck}_1 | \text{dk}' \rangle, \dots, \langle \vec{a} | \alpha_\ell, \text{ck}_\ell | \text{dk}' \rangle) \in \mathbb{Z}_t^\ell$ 
44:   $\vec{m} = \lfloor \frac{\vec{d} + t/4}{t/2} \rfloor \in \mathbb{Z}_2^\ell$ 

```

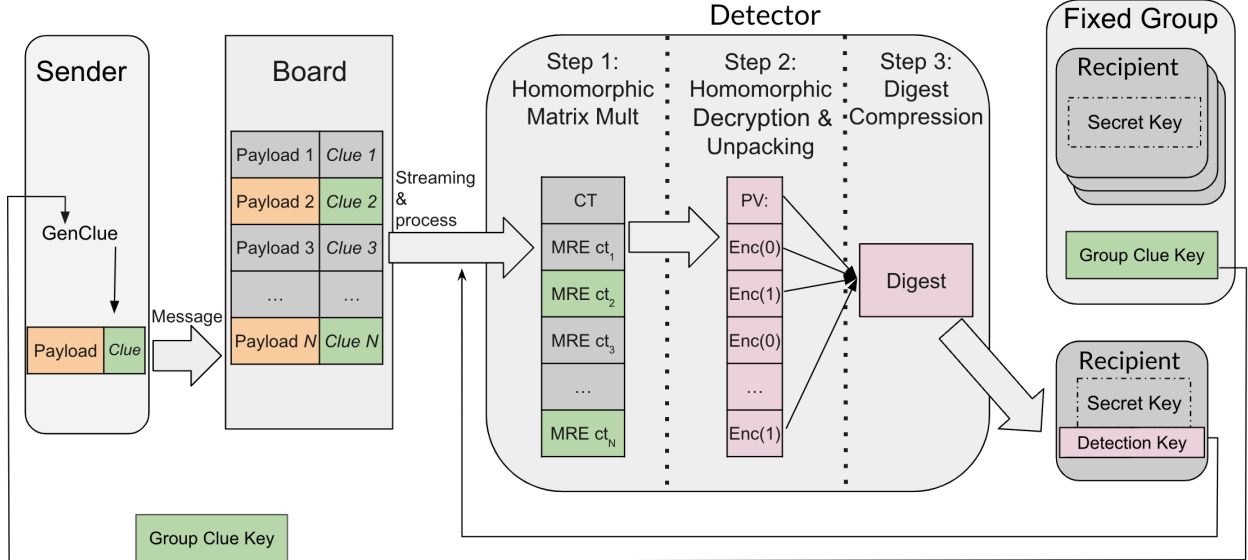


Figure 8: Main components of our Fixed Group OMR construction.

Key Privacy: Key privacy is guaranteed as long as the ciphertexts generated by two keys are from (computationally) indistinguishable distributions. This follows the same way argument above for CPA security. \square

7.4 Applying MRE to FGOMR

Building FGOMR using our MRE construction MRE is straightforward: given the OMR scheme in Section 5, we replace the PVW secret keys with the keys generated by `MRE.SKGen`. The group clue key can simply be a vector of all the members' public keys generated by `MRE.PKGen` ($r \leftarrow h(Y)$ is used as the shared randomness, where h is a random oracle and Y is the set of all the members). During the clue generation, instead of using PVW encryption, we use the MRE scheme to encrypt ℓ ones and proceed with the rest as the original OMR scheme. The detector now homomorphically decrypts a MRE ciphertext, instead of a PVW ciphertext, using (sk, dk) sent in the detection key.

Detection-key-unlinkability. Similarly to the id-based AGOMR construction in Section 6.2, we need to encrypt `dk` using BFV. `RegenDetectKey` is again very easy: simply generate fresh BFV public keys and encrypt `sk` and `dk` accordingly.

The full algorithm is given in Algorithm 5. A high-level overview is given in Fig. 8. Compared to Fig. 4, the clues remain to be encryptions of zeros, but instead of PVW encryption, we have MRE encryption. Thus, step 2 also has small changes.

Theorem 7.3. The scheme FGOMR1 in Algorithm 5 is a FGOMR scheme for $N < D \cdot t/2$, $P = \text{poly}(\lambda)$, $T = \text{poly}(\lambda)$ assuming security of MRE encryption and security of BFV leveled HE and that h is a random oracle, when instantiated with PRF f and an SRLC scheme SRLC (Definition B.1). Moreover when instantiated with SRLC1 (Algorithm 6), FGOMR1 is also compact.

Proof sketch. Completeness: Follows from the correctness of OMR-OPT and the correctness of the underlying MRE scheme.

Soundness: Similarly to Theorem 6.4 soundness proof.

Computational privacy: Flows from the key-privacy property of the underlying MRE scheme.

Compactness: Exactly the same as OMR-OPT digest size.

Detection-key-unlinkability: Same argument as the detection-key-unlinkability argument in the proof of Theorem 6.4. □

Alternating the order of the matrix multiplications. For better performance, we introduce another optimization. Since the entire decryption is under FHE, we would have a homomorphic matrix multiplication to calculate $\text{dk} \times Z$ at Algorithm 4 line 42, which is of size $O(G^2 \log(P))$. Instead, we first perform $\vec{\alpha}'_i \leftarrow \vec{\alpha}_i \times Z^\top$ for all $i \in [\ell]$, and then compute $\langle \vec{a} | \vec{\alpha}'_i, \text{ck}_i | \text{dk} \rangle$ under FHE. The latter operation is an inner product evaluated homomorphically, which is only of size $O(G \log(P) \ell)$. Since ℓ is asymptotically $\sim \log(\lambda)$ and practically with a very small constant (e.g., 4 as in [LT22, Section 10]), the runtime will be both asymptotically and concretely faster.¹⁰

Remark 7.4. As in remark 6.2, by restricting the groups to be honestly formed, we do not need to consider the union bound in ϵ_{DI} , which is the probability of existing a combination of G id's among P recipients to be linearly dependent, but instead, only evaluate the probability of a random group of size G to be linearly independent. Thus, the size of L will be reduced from $O(G \log P)$ to $G + O(\lambda)$. The number of homomorphic operations in the resulting scheme is then reduced from $O(G \log(P) \ell)$ to $O(G \ell)$, which gives a better asymptotic detector runtime. We call this variant FGOMR2.

7.5 Applying FGOMR to AGOMR

An interesting observation is that we can extend our FGOMR scheme to AGOMR by forcing all recipients to share the same A, \vec{b} in their public key and treat this pk as the gPKshare in FGOMR Definition 3.4. Then, everything follows: senders can arbitrarily choose up to G recipients as a subgroup of all P recipients to send a message.

Since we treat all P recipients as a whole group, the recipients do not need to generate different gPKshare for each different group, and thus the efficiency will match our FGOMR scheme without the requirement. However, this comes with two caveats: (1) DoS resistance is almost fully broken, and (2) the privacy guarantee becomes weaker.

Broken DoS resistance. We provide a high-level idea here before formally defining DoS resistance in Section 8: δ -DoS-resistance means that a clue should not be detected as pertinent by more than $\delta - 1$ honest recipients.

When generalizing FGOMR to AGOMR, because all of the honest recipients are sharing the same A , a malicious sender is able to send a single spam message to a large number of recipients $\gg G$ with high probability. In other words, recipient i is publishing some A, B_i such that $A \cdot \text{sk}_i \approx B_i$. Therefore, by crafting some vector e , the adversary can generate eA and check eB_i for all recipients. To spam more than $\delta - 1$ recipients, the adversaries need to find some $\vec{b} \approx B_i e$ for more than δ recipients. This can be done easily by using reject sampling when δ is $O(1)$, which becomes a huge issue for some applications, as discussed in [LT22].

Weaker privacy guarantee than the AGOMR definition. On the other hand, the privacy guarantee of the original AGOMR definition cannot be satisfied. Recall that in Section 3, in AGOMR, for a message sent to G recipients, even if some of the intended recipients have maliciously crafted keys, the identities of other recipients still remain private; while in FGOMR, privacy requires that the intended group does not contain any malicious recipient. Otherwise, that malicious recipient might be able to recover who are the rest of the group. This weaker privacy definition does make sense in the fixed group setting. For example, if a recipient has only joined a single group, it knows that the message must be addressed to that group even though the FGOMR scheme leaks no extra information. Besides, for a recipient joining different groups, it

¹⁰A similar technique can be applied for our AGOMR construction. However, that only changes the matrix multiplication size from $O(G^2 \log(P) + G(n + \ell))$ to $O(G \log(P)(n + \ell))$, which is not necessarily an improvement both asymptotically and concretely

Algorithm 5 FGOMR1: Compact Fixed Group Oblivious Message Retrieval

```
1: procedure FGOMR1.GenParams( $1^\lambda, \epsilon_p, \epsilon_n, G, P$ )
2:   Find some secure MRE parameter  $\text{pp}_{\text{MRE}}$  accordingly to MRE.GenParams in Algorithm 4.
3:   Find some secure BFV parameter  $\text{pp}_{\text{BFV}}$  as in [LT22]
4:   return  $\text{pp} = (1^\lambda, \epsilon_p, \epsilon_n, G, \text{pp}_{\text{MRE}}, \text{pp}_{\text{BFV}})$ 
5: procedure FGOMR1.PersonalKeyGen( $\text{pp} = (1^\lambda, \epsilon_p, \epsilon_n, G, \text{pp}_{\text{MRE}}, \text{pp}_{\text{BFV}})$ )
6:    $(\text{ck}_{\text{MRE}}, \text{dk}) \leftarrow \text{MRE.SKGen}(\text{pp}_{\text{MRE}})$ 
7:    $(\text{sk}_{\text{BFV}}, \text{pk}_{\text{BFV}}) \leftarrow \text{BFV.KeyGen}(\text{pp}_{\text{BFV}})$ 
8:    $\text{ct}_{\text{sk}} \leftarrow \text{BFV.Enc}(\text{pk}_{\text{BFV}}, \text{ck}_{\text{MRE}})$ 
9:    $\text{ct}_{\text{dk}} \leftarrow \text{BFV.Enc}(\text{pk}_{\text{BFV}}, \text{dk})$ 
10:  return  $(\text{sk} = (\text{sk}_{\text{BFV}}, \text{ck}_{\text{MRE}}, \text{dk}), \text{pk}_{\text{detect}} = (\text{pk}_{\text{BFV}}, \text{ct}_{\text{sk}}, \text{ct}_{\text{dk}}))$ 
11: procedure FGOMR1.GroupKeyGenAux( $\text{pp} = (1^\lambda, \epsilon_p, \epsilon_n, G, \text{pp}_{\text{MRE}}, \text{pp}_{\text{BFV}}), \text{sk}, Y$ )
12:    $r \leftarrow h(Y)$  ▷ Hashing  $Y$  to get some random seed
13:   Replace the  $r$  in  $\text{pp}_{\text{PVW}}$  with this new  $s$ 
14:    $((b'_i)_{i \in [w]}, \text{dk}) \leftarrow \text{MRE.PKGen}(\text{pp}_{\text{MRE}}, \text{sk})$ 
15:    $\text{gPKshare} = (r, (b'_i)_{i \in [w]}, \text{dk})$ 
16:   return  $\text{gPKshare}$ 
17: procedure FGOMR1.GroupKeyGen( $\text{pp}, (\text{gPKshare}_j = (r, (b'_{i,j})_{i \in [w]}, \text{dk}_j))_{j \in [G]}$ )
18:   ▷ If the amount of input  $\text{gPKshare}$  is less than  $G$ , run FGOMR1.PersonalKeyGen and
   FGOMR1.GroupKeyGenAux to generate new  $\text{gPKshare}$ 's until there are  $G$  of  $\text{gPKshare}$ 's
19:   return  $\text{pk}_{\text{clue}} = (r, (b'_{i,j})_{i \in [w]}, \text{dk}_j)_{j \in [G]}$ 
20: procedure FGOMR1.GenClue( $\text{pp} = (1^\lambda, \epsilon_p, \epsilon_n, G, \text{pp}_{\text{MRE}}, \text{pp}_{\text{BFV}}), \text{pk}_{\text{clue}} = (r, (b'_{i,j})_{i \in [w]}, \text{dk}_j)_{j \in [G]}, Y', x$ )
21:    $\text{pk}'_{\text{clue}} \leftarrow (r, (b'_{i,j})_{i \in [w]}, \text{dk}_j)_{j \in Y'}$ 
22:    $c \leftarrow \text{MRE.Enc}(\text{pp}_{\text{MRE}}, \text{pk}'_{\text{clue}}, 1^\ell)$ 
23:   return  $(x, c)$ 
24: procedure FGOMR1.Retrieve( $\text{pp}, \text{BB}, \text{pk}_{\text{detect}} = (\text{pk}_{\text{BFV}}, \text{ct}_{\text{sk}}, \text{ct}_{\text{dk}}), \bar{k}$ )
25:   for  $i \in [|\text{BB}|]$  do
26:     Homomorphically perform  $\text{MRE.Dec}(\text{pp}_{\text{MRE}}, (\text{ck}_{\text{MRE}}, \text{dk}), \text{BB}.c_i)$  using BFV
27:     ▷ The decryption process is almost the same, so the techniques in OMR-OPT can be reused.
28:   Proceed as line 35 in OMR-OPT.
29: procedure FGOMR1.Decode( $\text{pp}, M, \text{sk}$ )
30:   Same as OMR-OPT.Decode
31: procedure FGOMR1.RegenDetectKey( $\text{pp}, \text{sk} = (\text{sk}_{\text{BFV}}, \text{ck}_{\text{MRE}}, \text{dk})$ )
32:    $(\text{sk}'_{\text{BFV}}, \text{pk}'_{\text{BFV}}) \leftarrow \text{BFV.KeyGen}(\text{pp}_{\text{BFV}})$ 
33:    $\text{ct}'_{\text{sk}} \leftarrow \text{BFV.Enc}(\text{pk}'_{\text{BFV}}, \text{ck}_{\text{MRE}})$ 
34:    $\text{ct}'_{\text{dk}} \leftarrow \text{BFV.Enc}(\text{pk}'_{\text{BFV}}, \text{dk})$ 
35:   return  $(\text{sk}' = (\text{sk}'_{\text{BFV}}, \text{ck}_{\text{MRE}}, \text{dk}), \text{pk}'_{\text{detect}} = (\text{pk}'_{\text{BFV}}, \text{ct}'_{\text{sk}}, \text{ct}'_{\text{dk}})$ 
```

might even become an essential functionality for the recipient to know which group (or, for example, emailing list) this message is addressed to. However, in the ad-hoc setting, this is impractical.

The stronger definition cannot be satisfied because the PVW ciphertexts in our FGOMR construction share the same randomness. By controlling some of the keys, the adversary might be able to recover some or all of the randomness used and thus makes it insecure. However, this has a minor effect as G is usually small, and thus the randomness the adversary can recover is limited. Therefore, it is still possible that with certain parameters, our FGOMR construction can achieve the same level of privacy as our AGOMR construction. We leave the formal argument to future work. A simpler way to resolve this, as we have mentioned in Remark 3.3, is to use a zk-proof to show that the keys are indeed generated from `PersonalKeyGen`, as the keys that are semi-maliciously generated cannot break the privacy.

In general, for applications that only require a weaker privacy guarantee or accept the zk-based solution, and do not need DoS resistance, applying the FGOMR scheme to AGOMR is a great way to boost the overall performance.

8 DoS Resistance

Motivation. In a messaging system, an attacker may conduct a Denial of Service (DoS) attack on recipients by simply sending numerous of messages to them, increasing the cost of retrieval and processing. Furthermore, in the case where the number of pertinent messages retrieved is bounded to hide traffic patterns, an excessive number of new messages eventually causes overflows and prevents access to other messages (as in [MSS⁺22, LT22]) unless the bound \bar{k} is increased.

Inevitably, an attacker may spam individual recipients with messages, and pay a corresponding linear cost. However, something worse can happen: perhaps the attacker can send a message that is seen as pertinent by more than one recipient? In this case, the attacker can simply cause overflows in many recipients by sending just a small amount of messages. This problem is studied and mitigated in [LT22], for the case of (single-recipient) OMR.

In the group setting, the attacker may inevitable spams *groups* of up to G recipients, but we want to prevent attacks that are worse than that (e.g., spamming all users with “wildcard clues”). In this section, we adapt and generalize the DoS resistance definition from [LT22] to define this stronger security notion for GOMR. Intuitively, for a message with its corresponding clue maliciously crafted, no more than δ honest recipients should detect it as pertinent. Since by definition, GOMR allows G recipients to be addressed by a single message, by default, $\delta \geq G$, and we want δ to be as close as possible to G . For FGOMR, we define a stronger, saying that (honest) recipients inside different groups cannot be spammed together.

We then show how to achieve these notions of DoS-resistant AGOMR and FGOMR, under suitable assumptions.

8.1 Ad-hoc Group OMR

To formalize these properties, as in [LT22], we define an *indicator* predicate $\mathcal{I}(\text{pp}, x, c, \text{pk}_{\text{clue}}, \text{sk})$ as a ground truth for whether a given message (x, c) is pertinent to a given user specified by one’s keys. This predicate, which may not be efficiently computable, should give the natural answer for honestly generated clues. The indicator may answer arbitrarily for otherwise-generated clues, under the restriction that it should claim no more than δ honest recipients as the intended recipients except with a small probability. We define *collision resistance* property for the *indicator* such that it should be difficult for δ recipients to detect the same message as pertinent. In AGOMR, this property prevents a malicious sender from crafting a message to spam more than δ recipients.

Soundness and completeness are then redefined w.r.t the indicator \mathcal{I} . As in [LT22], to facilitate tight analysis, the completeness bound ϵ_n (false negative rate) in the definition is broken up into two components: the rate ϵ_{in} at which the indicator fails to detect truly pertinent messages (which may be non-negligible because an indicator with high thresholds may help achieve collision resistance), and the rate $\epsilon_n - \epsilon_{\text{in}}$ at

which the scheme fails to retrieve messages flagged by the indicator (which may be non-negligible because of error sources in the concrete scheme).

Definition 8.1 (DoS-resistant AGOMR). Let AGOMR be an Ad-hoc GOMR scheme with error rates ϵ_n, ϵ_p (as in Definition 3.1) and group size upper bound G . An *indicator* for AGOMR, with an *indicator false negative rate* ϵ_{in} (where $\epsilon_{in} \leq \epsilon_n$) and *collision resistance level* δ , is a function $b \leftarrow \mathcal{I}(\mathbf{pp}, x, c, \mathbf{pk}_{clue}, \mathbf{sk})$ on public parameter \mathbf{pp} , message (x, c) , clue key \mathbf{pk}_{clue} , and the corresponding secret key \mathbf{sk} that outputs $b \in \{0, 1\}$, such that:

- (*Indicator completeness*) For $\mathbf{pp} \leftarrow \text{AGOMR.GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, for any $\tau \in [G]$, honestly-generated key pairs: $(\mathbf{sk}_i, \mathbf{pk}_i = (\mathbf{pk}_{clue_i}, \cdot)) \leftarrow \text{AGOMR.KeyGen}(\mathbf{pp})$, $i \in [\tau]$, for any payload x , and honestly-generated clue $c \leftarrow \text{AGOMR.GenClue}(\mathbf{pp}, \mathbf{pk}_{clue_1}, \dots, \mathbf{pk}_{clue_\tau}, x)$, it holds that for $i \in [\tau]$:

$$\Pr[\mathcal{I}(\mathbf{pp}, x, c, \mathbf{pk}_{clue_i}, \mathbf{sk}_i) = 1] \geq 1 - \epsilon_{in} - \text{negl}(\lambda) .$$

- (δ -*Collision resistance*) For any PPT adversary \mathcal{A} , let $\mathbf{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, for any δ honestly-generated key pairs: $(\mathbf{sk}_i, \mathbf{pk}_i = (\mathbf{pk}_{clue_i}, \cdot)) \leftarrow \text{AGOMR.KeyGen}(\mathbf{pp})$, $i \in [\delta]$, and adversarially-generated $(x, c) \leftarrow \mathcal{A}(\mathbf{pk}_1, \dots, \mathbf{pk}_\delta)$, for $b_i \leftarrow \mathcal{I}(\mathbf{pp}, x, c, \mathbf{pk}_{clue_i}, \mathbf{sk}_i)$, it holds that for $i \in [\delta]$:

$$\Pr[b_1 = 1 \wedge \dots \wedge b_\delta = 1] \leq \epsilon_p + \text{negl}(\lambda) .$$

An Ad-hoc GOMR scheme AGOMR is δ -*DoS-resistant* for ϵ_n, ϵ_p , if there exists an indicator \mathcal{I} with indicator false negative rate ϵ_{in} , DoS resistance level δ , such that for any $G = \text{poly}(\lambda), P = \text{poly}(\lambda)$ and for any PPT adversary \mathcal{A} , for $\mathbf{pp} \leftarrow \text{AGOMR.GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, $(\mathbf{sk}, \mathbf{pk} = (\mathbf{pk}_{clue}, \mathbf{pk}_{detect})) \leftarrow \text{AGOMR.KeyGen}(\mathbf{pp})$, and adversarially-generated board $\mathbf{BB} \leftarrow \mathcal{A}(\mathbf{pp}, \mathbf{pk})$ where $\mathbf{BB} = ((x_1, c_1), \dots, (x_N, c_N))$ and $(x_i)_{i \in [N]}$ are unique, for any $0 < \bar{k} \leq N$, let $M \leftarrow \text{AGOMR.Retrieve}(\mathbf{BB}, \mathbf{pk}_{detect}, \bar{k})$, $\text{PL} \leftarrow \text{AGOMR.Decode}(M, \mathbf{sk})$:

- (*DoS-completeness*) Let $k = \sum_{j=1}^N \mathcal{I}(\mathbf{pp}, x_j, c_j, \mathbf{pk}_{clue}, \mathbf{sk})$. Then either $k > \bar{k}$ and $\text{PL} = \text{overflow}$, or $\Pr[x_j \in \text{PL} \mid \mathcal{I}(\mathbf{pp}, x_j, c_j, \mathbf{pk}_{clue}, \mathbf{sk}) = 1] \geq 1 - (\epsilon_n - \epsilon_{in}) - \text{negl}(\lambda)$ for all $j \in [N]$.
- (*DoS-soundness*) $\Pr[x_j \in \text{PL} \mid \mathcal{I}(\mathbf{pp}, x_j, c_j, \mathbf{pk}_{clue}, \mathbf{sk}) = 0] \leq \text{negl}(\lambda)$ for all $j \in [N]$.

Note that this implies the Completeness and Soundness definition in Definition 3.1 analogously to the discussion in [LT22, Section 8.3] and [LT22, Lemma 8.2].

To show our AGOMR scheme is DoS resistant, we need to rely on a generalization of the snake-eye conjecture proposed in [LT22] as formalized below.

Conjecture 8.1 (Regev05 is generalized-snake-eye resistant). For any PPT algorithm \mathcal{A} , for Regev05 encryption with modulus q and remaining parameters for which semantic security holds, for any $\delta = \text{poly}(\lambda)$, for any $1 \leq r < q/4$, for key pairs $(\mathbf{sk}_i, \mathbf{pk}_i) \leftarrow \text{KeyGen}(1^\lambda)$, $i \in [\delta + 1]$, for ciphertexts $(\vec{a}_i, b_i)_{i \in [\delta + 1]} \leftarrow \mathcal{A}((\mathbf{pk})_{i \in [\delta + 1]}, r)$, for some coefficients $(c_1, \dots, c_\delta) \in \mathbb{Z}_q^\delta$, it holds that:

$$\Pr \left[\begin{array}{l} |\langle \vec{a}_1, \mathbf{sk}_1 \rangle + b_1| \leq r \\ \wedge \dots \\ \wedge |\langle \vec{a}_{\delta+1}, \mathbf{sk}_{\delta+1} \rangle + b_{\delta+1}| \leq r \end{array} \wedge \exists i \in [\delta + 1], c_i \vec{a}_i \neq \vec{0} \wedge (\vec{a}_{\delta+1}, b_{\delta+1}) = \sum_{i=1}^{\delta} c_i (\vec{a}_i, b_i) \right] \leq \frac{2r+1}{q} + \text{negl}(\lambda) .$$

This generalized snake-eye conjecture is a stronger assumption than the original snake-eye conjecture [LT22, Conj 8.4].

This conjecture can be further generalized from Regev05 encryption to PVW encryption using reject sampling, analogous to [LT22, Lemma 8.5], reducing the snake-eye probability from $\frac{2r+1}{q}$ above to $(\frac{2r+1}{q})^\ell$ for $\ell = \log(\lambda)$ as shown in Lemma 8.2.

Lemma 8.2 (PVW is generalized-snake-eye resistant). Under Conjecture 8.1, for any PPT adversary \mathcal{A} , for PVW encryption with modulus q and plaintext space \mathbb{Z}_2^ℓ , and r such that $(\frac{2r+1}{q})^\ell = \text{poly}(\lambda)$ and remaining parameters for which semantic security hold; for any $\delta = \text{poly}(\lambda)$, $i \in [\delta + 1]$, key pairs $(\text{sk}_i, \text{pk}_i) \leftarrow \text{KeyGen}(1^\lambda)$, and ciphertexts $(\vec{a}_i, \vec{b}_i)_{i \in [\delta+1]} \leftarrow \mathcal{A}((\text{pk})_{i \in [\delta+1]}, r)$, let $\vec{m}_j \leftarrow \text{sk}_j^\top A + \vec{b}_j$, for some coefficients $(c_1, \dots, c_\delta) \in \mathbb{Z}_q^\delta$, it holds that:

$$\Pr \left[\begin{array}{l} \forall k \in [\ell] : |m_j[k]| \leq r \\ \wedge (\exists i \in [\delta + 1], \vec{a}_i \neq \vec{0}) \end{array} \wedge (\vec{a}_{\delta+1}, \vec{b}_{\delta+1}) = \sum_{i=1}^{\delta} c_i (\vec{a}_i, \vec{b}_i) \right] \leq \left(\frac{2r+1}{q} \right)^\ell + \text{negl}(\lambda).$$

Patched AGOMR3. Analogously to [LT22], to prevent trivial DoS attacks, we want to avoid $M_s \times (\overline{\text{id}} \times Z) = (\vec{0}, \vec{b})$ for more than $\delta - 1$ id's, where (M_s, s) is the clue in Algorithm 3, and the pseudorandom matrix Z is sampled using randomness s . Thus, the detector rejects the clue when $M = 0$.

Theorem 8.3. For any $\epsilon_p > 0, \epsilon_n > 0, P = \text{poly}(\lambda), G = \text{poly}(\lambda)$, AGOMR3 in Algorithm 3 (patched as above) is a $(G' + 1)$ -DoS-resistant Ad-hoc Group Oblivious Message Retrieval scheme where G' is as defined in line 5, when instantiated with any PRF f , assuming the hardness of Ring-LWE and Conjecture 8.1.

Proof sketch. The indicator completeness, DoS-completeness, and DoS-soundness are all straightforward since we perform the homomorphic PVW decryption process as in [LT22]. The only exception is when $M = 0$ for honestly generated clues (M, Z) . Notice that $M = 0$ only if $Z \times [\text{id}_1 \parallel \dots \parallel \text{id}_i^G] \times M = \text{ct}_i = \vec{0}$ for all $i \in [G]$. Thus, the probability of rejecting an honestly generate clue is $\Pr[M = 0] \leq \Pr[\forall i \in [G], \text{ct}_i = 0] = t^{-nG} = \text{negl}(\lambda)$.

We then prove for $(G' + 1)$ -collision resistance based on Conjecture 8.1.

Let $G'' = G' + 1$. Recall that a clue is of the form $(M, Z) \in \mathbb{Z}_t^{G' \times (n+\ell)} \times \mathbb{Z}_t^{I \times G'}$. If a clue is detected as pertinent by G'' honest recipients, then there exist G'' id's $(\text{id}_1, \dots, \text{id}_{G''})$, such that for all $i \in [G'']$, $(\vec{a}_i, \vec{b}_i) \leftarrow (\text{id}_i \times Z) \times M$, it holds that $\Pr[\forall i \in [G''], \text{PVW.Dec}(\text{sk}_i, (\vec{a}_i, \vec{b}_i)) = 1^\ell] \geq (\frac{2r+1}{q})^\ell + p(\lambda)$ for some non-negligible function $p(\lambda)$. Furthermore, this also means that there exists $(c_1, \dots, c_{G'}) \in \mathbb{Z}_q^\delta$ such that $(\vec{a}_{G''}, \vec{b}_{G''}) = \sum_{i=1}^{G'} c_i (\vec{a}_i, \vec{b}_i)$. At least one of the \vec{a}_i 's is non-zero, since otherwise $M = 0$, which, due to the aforementioned patch, would make Algorithm 3 reject this clue. Therefore, it breaks Lemma 8.2 and thus Algorithm 3 is $(G' + 1)$ -collision resistance. \square

Remark 8.4. Conjecture 8.1 is actually stronger than what we need, as in our AGOMR scheme, the adversary cannot output $\delta + 1$ clues that are linearly dependent with arbitrary coefficients. Instead, the linear dependency is predetermined by all the ids and the random matrix Z . Although Z is randomly chosen, it is a seed, and thus it works like a random oracle that the adversary needs to call. Therefore, a weaker conjecture would claim that given certain id's and a random oracle determining the linear dependency, the adversary cannot generate pertinent clues according to that certain linear dependency with $\Pr > \frac{2r+1}{q} + \text{negl}(\lambda)$. Although there are exponentially many possible linear dependencies there (i.e., P choose G possibilities), a PPT adversary may only be able to go through a polynomial amount of them. This weaker conjecture is sufficient for us to prove the collision resistance for Algorithm 3. However, for the readability and simplicity of the paper, we use this stronger conjecture. A possible clean change of underlying conjecture is left to future work.

8.2 Fixed Group OMR

As mentioned, for FGOMR, we capture an even stronger security notion. Essentially, we require that as long as two honest recipients are not in the same group, they should not be spammed at the same time except with a small probability. Since there are at most G honest recipients in each group, this property is strictly stronger than the property defined in Definition 8.1 even when $\delta = G + 1$ (which is the best security property that can be achieved for AGOMR).

Definition 8.2 (DoS-resistant FGOMR). Let FGOMR be a Fixed GOMR scheme with error rates ϵ_n, ϵ_p and group size upper bound G (as in Definition 3.4). An *indicator* with an *indicator false negative rate* $\epsilon_{in} \leq \epsilon_n$ for FGOMR is a function $b \leftarrow \mathcal{I}(\text{pp}, x, c, \{\text{pk}_{\text{clue}}\}, \text{sk})$ on public parameter pp , message (x, c) , the recipient's secret key sk , clue keys $\{\text{pk}_{\text{clue}}\}$ for all groups that include that recipient, that outputs $b \in \{0, 1\}$, such that:

- (*Indicator completeness*) For $\text{pp} \leftarrow \text{FGOMR.GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, any $g \leq G$ and any set $Y \subseteq [P]$, $|Y| = g$, and any subset $Y' \subseteq Y$; for honestly generated keys $(\text{sk}_i, \cdot) \leftarrow \text{FGOMR.PersonalKeyGen}(\text{pp})$, $\text{gPKshare}_i \leftarrow \text{FGOMR.GroupKeyGenAux}(\text{pp}, \text{sk}_i, Y)$, $i \in Y$, $\text{pk}_{\text{clue}} \leftarrow \text{FGOMR.GroupKeyGen}(\text{pp}, \{\text{gPKshare}_i\}_{i \in Y})$, for any payload x , and honestly generated clue $c \leftarrow \text{FGOMR.GenClue}(\text{pp}, \text{pk}_{\text{clue}}, Y', x)$, it holds that for $i \in Y'$:

$$\Pr[\mathcal{I}(\text{pp}, x, c, \{\text{pk}_{\text{clue}}\}, \text{sk}_i) = 1] \geq 1 - \epsilon_{in} - \text{negl}(\lambda) .$$

- (*Collision resistance*) For any PPT adversary \mathcal{A} , let $\text{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, $P = \text{poly}(\lambda)$; with $g, g' \in [G]$, and any two subsets $Y, Y' \subseteq [P]$ with $|Y| = g$, $|Y'| = g'$; for honestly generated keys $(\text{sk}_i, \cdot) \leftarrow \text{FGOMR.PersonalKeyGen}(\text{pp}), (\text{sk}'_j, \cdot) \leftarrow \text{FGOMR.PersonalKeyGen}(\text{pp})$, $\text{gPKshare}_i \leftarrow \text{FGOMR.GroupKeyGenAux}(\text{pp}, \text{sk}_i, Y)$ ($i \in [g]$), $\text{gPKshare}'_j \leftarrow \text{FGOMR.GroupKeyGenAux}(\text{pp}, \text{sk}'_j, Y')$ ($j \in [g']$), $\text{pk}_{\text{clue}} \leftarrow \text{FGOMR.GroupKeyGen}(\text{pp}, \text{gPKshare}_1, \dots, \text{gPKshare}_g)$, $\text{pk}'_{\text{clue}} \leftarrow \text{FGOMR.GroupKeyGen}(\text{pp}, \text{gPKshare}'_1, \dots, \text{gPKshare}'_{g'})$, and for adversarially-generated $(x, c) \leftarrow \mathcal{A}(\text{pp}, \{\text{gPKshare}_i\}_{i \in [g]}, \{\text{gPKshare}'_j\}_{j \in [g']}, \text{pk}_{\text{clue}}, \text{pk}'_{\text{clue}})$; for any $i \in [g], j \in [g']$, let $b \leftarrow \mathcal{I}(\text{pp}, x, c, \text{pk}_{\text{clue}_i}, \text{sk}_i)$ and $b' \leftarrow \mathcal{I}(\text{pp}, x, c, \text{pk}'_{\text{clue}_j}, \text{sk}'_j)$, it holds that: $\Pr[b = 1 \wedge b' = 1] \leq \epsilon_p + \text{negl}(\lambda)$.

A Fixed GOMR scheme FGOMR is *DoS-resistant* for ϵ_n, ϵ_p , and G if there exists an indicator \mathcal{I} with an indicator false negative rate ϵ_{in} for FGOMR such that for any $P = \text{poly}(\lambda)$, and any PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, for $\text{pp} \leftarrow \text{FGOMR.GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$, $(\text{sk}_i, \text{pk}_{\text{detect}_i}) \leftarrow \text{FGOMR.PersonalKeyGen}(\text{pp})$ for $i \in [P]$; adversarially generate state and groups $(\text{st}, \{Y_j\}_{j \in [W]}) \leftarrow \mathcal{A}_1(\text{pp})$ where $Y_j \subseteq [P]$, $|Y_j| \leq G$, for any $W = \text{poly}(\lambda)$; for all $j \in [W], i \in Y_j$, $\text{gPKshare}_{i,j} \leftarrow \text{FGOMR.GroupKeyGenAux}(\text{pp}, \text{sk}_i, Y_j)$, $\text{pk}_{\text{clue}_j} \leftarrow \text{FGOMR.GroupKeyGen}(\text{pp}, \{\text{gPKshare}_{i,j}\}_{i \in Y_j})$, and adversarially-generated board $\text{BB} \leftarrow \mathcal{A}_2(\text{pp}, \{\text{gPKshare}_{i,j}\}_{j \in [W], i \in Y_j}, \{\text{pk}_{\text{clue}_j}\}_{j \in [W]}, \text{st})$ where $\text{BB} = ((x_1, c_1), \dots, (x_N, c_N))$ with unique x_i , for any $0 < \bar{k} \leq N$, for $i \in [P]$, let $M \leftarrow \text{Retrieve}(D, \text{pk}_{\text{detect}_u}, \bar{k})$, $\text{PL} \leftarrow \text{Decode}(M, \text{sk})$:

- (*DoS-completeness*) Let $k = \sum_{j=0}^N \mathcal{I}(\text{pp}, x_j, c_j, \{\text{pk}_{\text{clue}_j}\}_{j \in [W], i \in Y_j}, \text{sk}_i)$. Then either $k > \bar{k}$ and $\text{PL} = \text{overflow}$, or $\Pr[x_j \in \text{PL} \mid \mathcal{I}(\text{pp}, x_j, c_j, \{\text{pk}_{\text{clue}_j}\}_{j \in [W], i \in Y_j}, \text{sk}_i) = 1] \geq 1 - (\epsilon_n - \epsilon_{in}) - \text{negl}(\lambda)$ for all $j \in [N]$.
- (*DoS-soundness*) $\Pr[x_j \in \text{PL} \mid \mathcal{I}(\text{pp}, x_j, c_j, \{\text{pk}_{\text{clue}_j}\}_{j \in [W], i \in Y_j}, \text{sk}_i) = 0] \leq \text{negl}(\lambda)$ for all $j \in [N]$.

Patched FGOMR1. Analogously to AGOMR3, to prevent trivial DoS attacks, we modify our FGOMR1.Retrieve in Algorithm 5 to reject clues $c = (\vec{a} \parallel \alpha_1 \parallel \dots \parallel \alpha_\ell \parallel \vec{b}, s)$ where $\vec{a} = \vec{0}$.

To prove the DoS-resistance of our FGOMR scheme, we first prove the following lemma:

Lemma 8.5. For any PPT adversary \mathcal{A} , for Regev05 encryption with modulus q and plaintext space \mathbb{Z}_2 , and any $1 \leq r \leq q/4$ and remaining parameters for which the semantic security holds, for $i \in [2]$, $(\text{sk}_i, \text{pk}_i) \leftarrow \text{KeyGen}(1^\lambda)$, for ciphertexts $(\vec{a}, \vec{b}_i)_{i \in [2]} \leftarrow \mathcal{A}(\text{pk}_1, \text{pk}_2, r)$, it holds that:

$$\Pr \left[\begin{array}{l} |\langle \vec{a}, \text{sk}_1 \rangle + b_1| \leq r \\ \wedge |\langle \vec{a}, \text{sk}_2 \rangle + b_2| \leq r \\ \wedge \vec{a} \neq \vec{0} \end{array} \right] \leq \frac{2r+1}{q} + \text{negl}(\lambda)$$

assuming [LT22, Conj 8.4] holds when replacing n there with $n - 1$.

Proof sketch. Suppose that we have an adversary \mathcal{A} that breaks this lemma, we construct an adversary \mathcal{A}' that breaks [LT22, Conj 8.4] as follows:

Given $\mathbf{pk} = (A, \vec{b}) \in \mathbb{Z}_q^{w \times (n-1)} \times \mathbb{Z}_q^{w \times 1}$ and $\mathbf{pk}' = (A', \vec{b}') \in \mathbb{Z}_q^{w \times (n-1)} \times \mathbb{Z}_q^{w \times 1}$ and r , \mathcal{A}' generates $\alpha_1, \alpha_2 \leftarrow_{\$} \mathbb{Z}_q^{w \times 1}$, $s_1, s_2 \leftarrow_{\$} \mathbb{Z}_q$, and computes $\vec{b} \leftarrow \vec{b} + \alpha_1 \times s_1$ and $\vec{b}' \leftarrow \vec{b}' + \alpha_2 \times s_2$. \mathcal{A}' then randomly selects $i \in [n-1]$, and inserts $\alpha_{j \in \{1,2\}}$ to the i -th column, i.e., $A \leftarrow A[0:i] \parallel \alpha_1 \parallel A[i:n-1]$ and $A' \leftarrow A'[0:i] \parallel \alpha_2 \parallel A'[i:n-1]$.

Then, \mathcal{A}' sends $\mathbf{pk}_1 = (A, \vec{b}), \mathbf{pk}_2 = (A', \vec{b}'), r$ to \mathcal{A} . After getting (\vec{a}, b_1, b_2) back, check whether $b_1 + \vec{a}[i]s_1 = b_2 + \vec{a}[i]s_2$. If so, let $\vec{a}' = \vec{a}[0:i-1] \parallel \vec{a}[i+1:n]$, output $(\vec{a}, b_1 + \vec{a}[i]s_1)$. Otherwise, output $\text{Enc}(\mathbf{pk}, 0)$.

Note that trivially, $\mathbf{pk}_1, \mathbf{pk}_2$ are indistinguishable from two honestly generated PVW public keys, and it holds that $|\langle \vec{a}', \mathbf{sk}_1 \rangle + b_1 + \vec{a}[i]s_1| \leq r$ and $|\langle \vec{a}', \mathbf{sk}_1 \rangle + b_2 + \vec{a}[i]s_2| \leq r$ with probability $\geq \frac{2r+1}{q} + f(\lambda)$ for some non-negligible function f . Since s_1, s_2 are all chosen randomly and independently and are masked by α_1, α_2 before given to \mathcal{A} to calculate \vec{a} , we have $\Pr[b_1 + \vec{a}[i]s_1 = b_2 + \vec{a}[i]s_2] \geq 1/q - \text{negl}(\lambda)$.

For the case where $b_1 + \vec{a}[i]s_1 \neq b_2 + \vec{a}[i]s_2$ and $(\vec{a}, b) \leftarrow \text{Enc}(\mathbf{pk}, 0)$, the probability that $|\langle \vec{a}, \mathbf{sk} \rangle + b| \leq r$ and $|\langle \vec{a}, \mathbf{sk} \rangle + b| \leq r$ is $\geq \frac{2r+1}{q} - \text{negl}(\lambda)$. Therefore, we have that this \mathcal{A}' breaks the conjecture with probability $\geq \frac{2r+1}{q} + f(\lambda)/q - \text{negl}(\lambda)$, where $f(\lambda)/q$ is non-negligible. \square

This lemma can be generalized from Regev05 encryption to PVW encryption using reject sampling, analogous to [LT22, lemma 8.5] as follows, which is then used to prove Theorem 8.7.

Lemma 8.6. For any PPT adversary \mathcal{A} , for PVW encryption with modulus q and plaintext space \mathbb{Z}_2^ℓ , and r such that $(\frac{2r+1}{q})^\ell = \text{poly}(\lambda)$ and remaining parameters for which the semantic security holds, for $i \in [2]$, $(\mathbf{sk}_i, \mathbf{pk}_i) \leftarrow \text{KeyGen}(1^\lambda)$, for ciphertext $(\vec{a}, (\vec{b}_i)_{i \in [2]}) \leftarrow \mathcal{A}(\mathbf{pk}_1, \mathbf{pk}_2, r)$, let $\vec{m}_{i \in [2]} \leftarrow \mathbf{sk}_i^\top \vec{a} + \vec{b}_i$, it holds that

$$\Pr \left[(\forall k \in [\ell] : |m_1[k]| \leq r \wedge |m_2[k]| \leq r) \wedge \vec{a} \neq \vec{0} \right] \leq \left(\frac{2r+1}{q} \right)^\ell + \text{negl}(\lambda) \quad (1)$$

assuming [LT22, Conj 8.4] holds when replacing n there with $n-1$.

Theorem 8.7. For any $\epsilon_p > 0, \epsilon_n > 0, P = \text{poly}(\lambda), G = \text{poly}(\lambda)$, FGOMR1 in Algorithm 5 (patched as above) is a DoS-resistant Fixed Group Oblivious Message Retrieval scheme, when instantiated with any PRF f and random oracle h , assuming the hardness of Ring-LWE and [LT22, Conj 8.4]. Moreover when instantiated with SRLC1 defined in [LT22, Section 6.3.1], FGOMR1 is also compact.

Proof sketch. The completeness, soundness, and compactness all follow in a straightforward way. We prove the DoS resistant by contradiction.

Notice that since a group only contains up to G recipients, if the adversary wants the clue $(\vec{a}, \alpha_1, \dots, \alpha_\ell, \beta, s)$ to be detected as pertinent to $G+1$ honest recipients, it should be detected as pertinent by recipients in two different groups. Therefore, for two different honest recipients with auxiliary keys $\mathbf{dk}, \mathbf{dk}'$ in two different groups, let $Z \leftarrow_{\$} \mathbb{Z}_t^{L \times G'}$, we can generate ciphertext $(\vec{a}, \vec{b}_1 \leftarrow \vec{b} - \langle \mathbf{dk} \times Z, \alpha_1 \rangle \parallel \dots \parallel \langle \mathbf{dk} \times Z, \alpha_\ell \rangle, \vec{b}_2 \leftarrow \vec{b} - \langle \mathbf{dk}' \times Z, \alpha_1 \rangle \parallel \dots \parallel \langle \mathbf{dk}' \times Z, \alpha_\ell \rangle)$. This breaks Lemma 8.6 and thus FGOMR1 is DoS resistant. \square

9 Performance Evaluation

9.1 Methodology

We implemented our optimized OMR scheme OMR-OPT in Section 5.3, our AGOMR schemes AGOMR1, AGOMR2, AGOMR3 in Section 6, and our FGOMR schemes FGOMR1, FGOMR2 in Section 7, in a C++ library (released as open source). We used the OMR library [LT22] as our base implementation, the PALISADE library [PAL21] for PVW encryption, and the SEAL library [Mic20] with Intel-HEXL acceleration [BKS⁺21] for the BFV scheme. We benchmarked these schemes on several parameter settings on a Google Compute Cloud e2-standard-2 instance type, 8GB RAM (unless otherwise noted).

Parameters. For comparison to prior works, we reuse all the parameters from [LT22] (except for N): $\bar{k} = 50$, $\epsilon_p = 2^{-21}$, $\epsilon_n = 2^{-30}$, $\text{PVW}.(n, w, \ell, q, \sigma) = (450, 16000, 4, 65537, 1.3)$ and $\text{BFV}.(D, Q, t) = (2^{15}, \sim 2^{850})$,

	Ad-hoc GOMR					Fixed GOMR	
	OMRp2 §6.1 + [LT22]	OMR-OPT §6.1 + Thm 5.1	AGOMR1 Remark 6.1	AGOMR2 Remark 6.2	AGOMR3 Thm 6.4	FGOMR1 Thm 7.3	FGOMR2 Remark 7.4
Detector preprocessing time (sec/msg)	N/A	N/A	N/A	0.0066	0.0232	0.0041	N/A
Detector computation time (sec/msg)	5.205	1.844	0.117	0.168	0.181	0.123	0.119
Sender computation time (sec/msg)	0.131	0.131	0.136	0.132	0.136	0.020	0.020
Clue size (bytes/msg)	11472	11472	11472	15312	15312	1108	1092
Clue key size (bytes)	133K per recipient	133K per recipient	133K per recipient	133K per recipient	133K per recipient	1.56M per group	1.56M per group
Detection key size (bytes)	137M	139M	139M	143M	143M	142M	139M
Digest size (bytes/msg)	35						
Recipient time (sec)	0.026						

Table 3: Comparison of cost metrics, given $N = 2^{15}$, $k = \bar{k} = 50$, $G = 15$, $G' = 19$, $I = L = 81$. Costs are per message, per recipient, except for the preprocessing time, which can be shared by ≥ 1 recipients. For the functionality differences, see Table 1 for details.

65537)¹¹; payload size is 612 bytes. We change N from 500000 to 32768 as for any $N > 32768$, all the benchmarked constructions (including the prior work) simply repeat the scheme $\lceil \frac{N}{32768} \rceil$ times.

Application parameters. We fix $P = 2^{60}$, $G' = G + 4$, ¹² $\epsilon_{\text{DI}} \leq 2^{-128}$ (defined in Section 6.2 ¹³). For AGOMR3, AGOMR2, FGOMR1, FGOMR2, we set $I = L = \max(G', 8)$, such that id and dk are unique with overwhelming probability. For AGOMR1, we set $\mathcal{D} = \mathbb{Z}_{2^{127}-1}$ for the same reason and for that $2^{127} - 1$ is a prime number (such that an inverse exists for any field element).

9.2 Evaluation Results

Representative costs. Table 3 summarizes the main cost metrics of all our schemes and the baseline, for the parameters above for functionality and asymptotic costs, see Table 1. .

We see that for all our GOMR schemes, the detector online time is at least an order of magnitude faster than the baseline scheme (OMRp2[LT22] + §6.1). FGOMR schemes have better sender time and clue size, though the clue key size is larger (since it represents the entire group rather than a single recipient). AGOMR has a comparable sender time and clue size to the baseline. The detection key size, digest size, and recipient runtime are all relatively similar across all schemes.

Furthermore, our OMR-OPT shows an improvement of $\sim 2.5x$ faster detector time standalone as an OMR scheme.

Costs vs. group size G . Fig. 9 shows how the detector runtime, sender time, and clue size scale with the number of recipients per group. (Other metrics are independent of G , and thus given by Table 3.) OMR-OPT has roughly the same behavior as OMRp2, so we leave it out of the comparison for better visualization. For readability, we merge the schemes that have very minor differences ($< 5\%$).

We benchmarked group sizes $\{2, 4, 6, 8, 10, 12, 15, 25, 45, 65, 85, 105, 125, 150, 175, 200, 225, 250, 275, 300, 325, 350, 375, 400\}$. The corresponding id sizes for AGOMR, or the dk size for FGOMR, are $\{19, 28, 38, 47, 57, 66, 81, 128, 223, 318, 413, 508, 603, 722, 841, 959, 1078, 1197, 1316, 1434, 1553, 1672, 1791, 1909\}$. Due to memory constraints (simply because every clue is already too large to store in memory for a large group size),

¹¹We set $\text{BFV}.Q$ according to the homomorphic multiplicative depth in each scheme, for 128-bit security: $\log(Q) \approx 789$ in OMRp2, $\log(Q) \approx 810$ in OMR-OPT, AGOMR4, AGOMR1, FGOMR1, FGOMR2, $\log(Q) \approx 838$ in AGOMR2, and $\log(Q) \approx 868$ in AGOMR3.

¹²This G' guarantees that for any G , we have $\epsilon_{\text{DS}} \approx 2^{-48}$ (defined in Section 6.2 which is the probability that the matrix calculated by multiplying a full-rank matrix and a random matrix is not full rank). In the case that all the extended id's (or dk's) are linearly independent, the probability that our GOMR schemes fail to send a message is $\epsilon_{\text{DS}} \approx 2^{-48}$. However, this can be fixed by performing reject sampling (i.e., resample a Z until a full rank Z is found), at the cost of a minor leakage: given Z , there might exist some combination of recipients that the message cannot be sent.

¹³ ϵ_{DI} is the probability for there to exist any combination of G id's such that their extended version forms a matrix with rank $< G$.

OMR, AGOMR2, and AGOMR3 with group size ≥ 45 use a GCP `e8-highmem-64` instance, 64GB RAM (with a 128GB balanced disk). Other schemes and other group sizes still use the `e2-standard-2` instance type, 8GB RAM. Note that the runtime of the instance `e8-highmem-64` is roughly the same as `e2-standard-2` (about 2.5% faster) for the same schemes with the same group sizes, and thus it is still very continuous in the plot and hard to spot the difference.

As BFV uses power-of-two cyclotomic rings, the underlying ring dimension is a power-of-two. Thus, when doing matrix multiplication between matrices of sizes $(a \times b)$ and $(b \times c)$, for better efficiency, we first round up a, b, c to the nearest power of two. The round-ups in our schemes cause the jumps in Fig. 9.

All schemes have two to three orders of magnitude faster detector runtime compared to the baseline. The detector time of our main schemes AGOMR3, AGOMR2 increases most rapidly among all, because of the large amount of matrix multiplications they need to perform as shown in Table 1. Note that the runtime of our GOMR schemes are calculated as $T_p/10 + T_o$ where T_p is the detector preprocessing time, and T_o is the detector online computation time listed in Table 3. In other words, the runtime is computed based on ten recipients sharing the preprocessing time. If the preprocessing time is not amortized over multiple recipients (as in Table 3), for most of our schemes and parameters, the preprocessing will only take $< 5\%$ of the total time, and in the worst case (AGOMR, $G = 400$) the total runtime of the detector will be only less than two times larger. On the other hand, to further reduce the amortized cost, preprocessing can of course be shared by more recipients.

The clue size for our FGOMR scheme is smaller and grows more gently compared to AGOMR and the baseline. The sender time, in our schemes, grows super-linearly with the group size (due to solving a linear system of equations), compared to linear growth for the baseline; but is still less than a second.

10 Applications

10.1 Secure Group Messaging

The most direct motivation for GOMR is secure group messaging. As discussed in Section 1, Ad-hoc GOMR suits groups of arbitrary recipients chosen on the fly, whereas Fixed GOMR suits applications with well-defined groups such as group chats and mailing lists.

Concrete costs. We use the same parameters as benchmarked in Section 9 (in the absence of detailed data about deployed systems). For a GCP `e2-standard-2` instance (with a 32GB balanced disk), the average cost is $\sim \$1.89 \cdot 10^{-5}/\text{sec}$ ¹⁴. Thus, for 32768 messages, and group size $G = 15$, our GOMR schemes cost range from \$0.07 to \$0.11, while if the baseline OMR-OPT from [LT22] costs $\sim \$3.2$.

Varying group size. Group size varies between messages. The group size parameter G can be adaptively changed between messages (at a cost of leaking the group size), or groups can be padded to some standard sizes or randomized sizes (to minimize information leakage but (at a cost in efficiency)). This tradeoff is information-theoretically inherent.

10.2 Private Blockchains

Transaction Batching. In many blockchain protocols, it is cost-effective to aggregate multiple transactions into a single one with multiple recipients. This is a natural application for AGOMR, especially with the multi-payload extension of Section 6.3 that allows each recipient to retrieve just the pertinent portion.

Private Bitcoin. Concretely, in Bitcoin, transactions with 2 recipients consist of $\sim 74\%$ of all transactions, and transactions with more than 2 recipients take $\sim 10\%$ of all transactions. If Bitcoin became privacy-preserving, using our AGOMR scheme AGOMR4 instead of directly using OMR scheme instantiated with our optimized OMR scheme OMR-OPT results in a $\sim 1.5x$ speedup overall ($\sim 4x$ speedup if instantiated with OMRp2 in [LT22]), at the cost of clue size grown by $\sim 3x$. Note that in the estimation, we do not hide the group size (i.e., group size of each transaction is still public and vary by transactions).

¹⁴Estimate from <https://cloud.google.com/products/calculator>. Communication cost is negligible: $< \$10^{-9}/\text{msg egress}$.

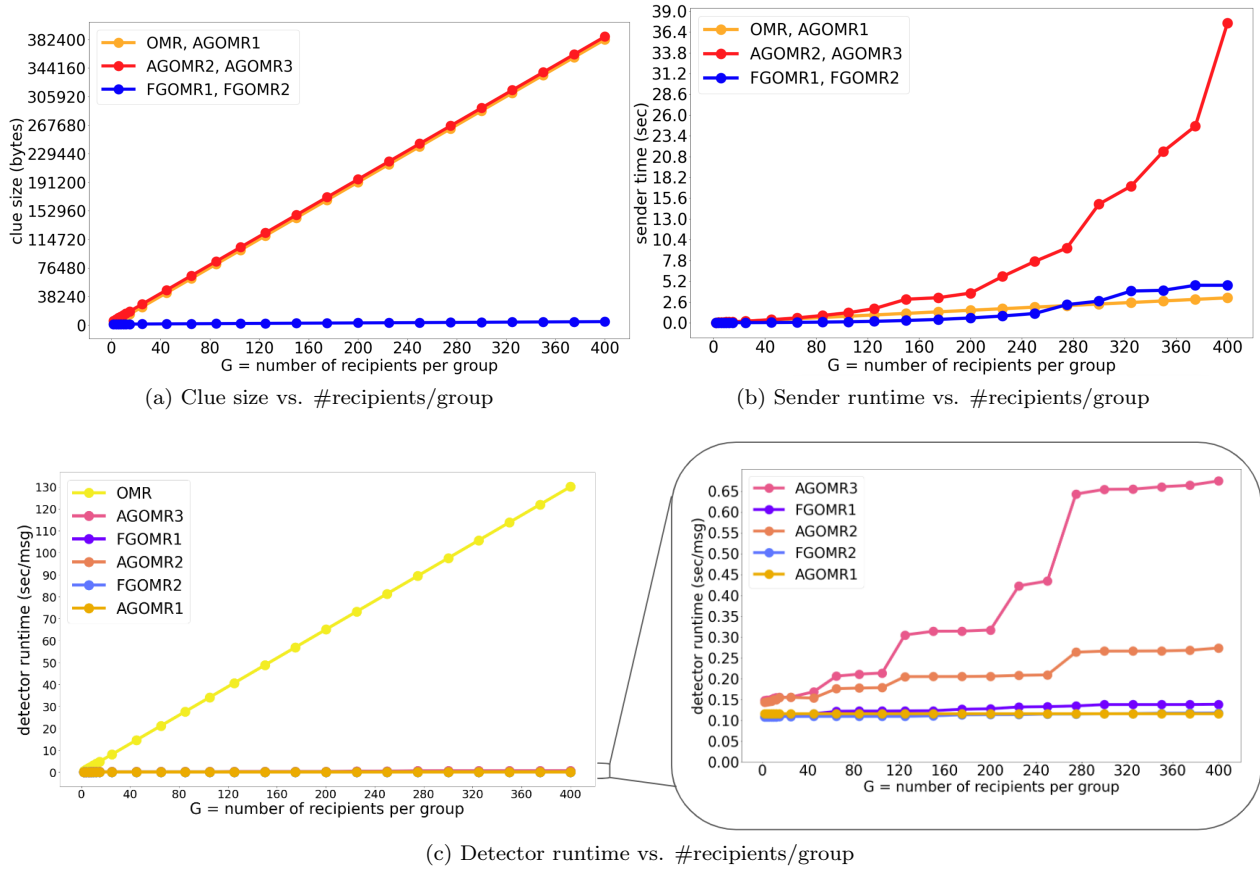


Figure 9: Figures show how clue size, sender time, and detector time scale with G the number of recipients in a group. Other parameters are as in Table 3.

Blockchain Auditing. In privacy-preserving cryptocurrencies (like Zcash or Monero), a user may want to share the transaction information with service providers such as accountants or auditors, or with pertinent authorities. Currently, privacy-preserving cryptocurrencies use a viewing key to allow other users to view the transactions. This entails extra detection costs: every viewing-key holder needs to separately look for transactions in every account they can view. For example, using OMR, an accountant would need to perform a separate retrieval for each of their hundreds of clients.

This can be improved using FGOMR, with groups corresponding to account owners plus their authorized viewers. Each such group is represented by a single FGOMR group clue key, and senders are instructed (at the protocol level) to use this clue key when transacting with this account. Subsequently, each party needs to perform retrievals only for their own key, yet will receive all messages addressed to all the groups (e.g., accounting clients) of which they are members.

A limitation of this approach is that to grant or revoke view keys, the group clue key needs to be regenerated and republished.

Acknowledgements

We are grateful to Xifan Yu for discussing the probability bound on id length for maliciously crafted groups in Section 6.

This material is based upon work supported by Algorand Centres of Excellence programme managed by Algorand Foundation; DARPA under Contract No. HR001120C0085; the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research under award number DE-SC-0001234; JPMorgan Chase & Co; LexisNexis Risk Solutions; and Mastercard Center for Inclusive Growth. Any opinions, views, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Algorand Foundation, the United States Government, DARPA, DOE, JPMorgan Chase & Co. or its affiliates, or other sponsors.

References

- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, pages 169–203, 2015.
- [BKKS07] Mihir Bellare, Alexandra Boldyreva, Kaoru Kurosawa, and Jessica Staddon. Multirecipient encryption schemes: How to save on bandwidth and computation without sacrificing security. *IEEE Transactions on Information Theory*, 53(11):3927–3943, 2007. doi:10.1109/TIT.2007.907471.
- [BBS02] Mihir Bellare, Alexandra Boldyreva, and Jessica Staddon. Randomness re-use in multi-recipient encryption schemes. In Yvo G. Desmedt, editor, *Public Key Cryptography — PKC 2003*, pages 85–99, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [BCG⁺20] Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. Zexe: Enabling decentralized private computation. In *2020 IEEE S&P (SP)*, pages 947–964, 2020.
- [BEM⁺17] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *SOSP*, pages 441–459, 2017.
- [BJT18] Suzie Brown, Oliver Johnson, and Andrea Tassi. Reliability of broadcast communications under sparse random linear network coding. *IEEE Transactions on Vehicular Technology*, 67(5):4677–4682, 2018.
- [BKS⁺21] Fabian Boemer, Sejun Kim, Gelila Seifu, Fillipe DM de Souza, Vinodh Gopal, et al. Intel HEXL (release 1.2). <https://github.com/intel/hexl>, September 2021.
- [BLMG21] Gabrielle Beck, Julia Len, Ian Miers, and Matthew Green. Fuzzy message detection. The ACM Conference on Computer and Communications Security (CCS) 2021, 2021.
- [Bra12] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In *CRYPTO 2012*, LNCS. Springer, August 19–23, 2012.
- [BSCG⁺14] 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 S&P*, pages 459–474, 2014.
- [BSW09] John Bethencourt, Dawn Xiaodong Song, and Brent Waters. New techniques for private stream searching. *ACM Trans. Inf. Syst. Secur.*, 12:16:1–16:32, 2009.
- [CGBM15] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *2015 IEEE S&P*, pages 321–338, 2015.
- [CGKS95] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *36th FOCS*, pages 41–50, Milwaukee, Wisconsin, October 23–25, 1995. IEEE Computer Society Press. doi:10.1109/SFCS.1995.492461.
- [CGN98] Benny Chor, Niv Gilboa, and Moni Naor. Private information retrieval by keywords, 1998. Appeared in the THEORY OF CRYPTOGRAPHY LIBRARY and has been included in the ePrint Archive. benny@cs.technion.ac.il 10500 received February 3rd, 1998. URL: <http://eprint.iacr.org/1998/003>.
- [DD07] George Danezis and Claudia Diaz. Space-efficient private search with applications to rateless codes. In *FC’07*, page 148–162. Springer, 2007.

- [FLS22] Nils Fleischhacker, Kasper Green Larsen, and Mark Simkin. How to compress encrypted data. Cryptology ePrint Archive, Paper 2022/1413, 2022. <https://eprint.iacr.org/2022/1413>. URL: <https://eprint.iacr.org/2022/1413>.
- [FN94] Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *Advances in Cryptology — CRYPTO’ 93*, pages 480–491, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.
- [FR13] Matthieu Finiasz and Kannan Ramchandran. Private stream search at almost the same communication cost as a regular search. In Lars R. Knudsen and Huapeng Wu, editors, *Selected Areas in Cryptography*, pages 372–389, Berlin, Heidelberg, 2013. Springer.
- [FV12] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012. <https://ia.cr/2012/144>.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *ACM Symposium on Theory of Computing*, STOC ’09, page 169–178. ACM, 2009.
- [HBHW] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash Protocol Specification Version 2021.2.14. <https://github.com/zcash/zips/blob/master/protocol/protocol.pdf>.
- [KC16] Amjad Saeed Khan and Ioannis Chatzigeorgiou. Improved bounds on the decoding failure probability of network coding over multi-source multi-relay networks. *IEEE Communications Letters*, 20(10):2035–2038, 2016.
- [KS07] Tali Kaufman and Madhu Sudan. Sparse random linear codes are locally decodable and testable. In *FOCS’07*, 2007.
- [Kur02] Kaoru Kurosawa. Multi-recipient public-key encryption with shortened ciphertext. In David Naccache and Pascal Paillier, editors, *Public Key Cryptography*, pages 48–63, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [Lew21] Sarah Jamie Lewis. Discreet log #1: Anonymity, bandwidth and Fuzzytags, Feb 2021. URL: <https://openprivacy.ca/discreet-log/01-anonymity-bandwidth-and-fuzzytags/>.
- [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 2013.
- [LT22] Zeyu Liu and Eran Tromer. Oblivious message retrieval. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 753–783, Cham, 2022. Springer Nature Switzerland. Full version: Cryptology ePrint Archive <https://ia.cr/2021/1256>, 2021; internal citations follow the latter’s numbering.
- [Lun18] Joshua Lund. Technology preview: Sealed sender for signal. <https://signal.org/blog/sealed-sender/>, Oct. 2018.
- [Mic20] Microsoft SEAL (release 3.6). <https://github.com/Microsoft/SEAL>, November 2020. Microsoft Research, Redmond, WA.
- [MSS⁺22] Varun Madathil, Alessandra Scafuro, István András Seres, Omer Shlomovits, and Denis Varlakov. Private signaling. *USENIX Security 2022*, 2022.
- [Noe15] Shen Noether. Ring signature confidential transactions for monero. *IACR Cryptology ePrint Archive*, 2015:1098, 2015.
- [OS05] Rafail Ostrovsky and William E. Skeith. Private searching on streaming data. In *CRYPTO*, 2005.

- [PAL21] PALISADE lattice cryptography library (release 11.2). <https://palisade-crypto.org/>, June 2021.
- [Pla18] Rachel Player. *Parameter selection in lattice-based cryptography*. PhD thesis, Royal Holloway, University of London, 2018.
- [PPS14] Alexandre Pinto, Bertram Poettering, and Jacob C.N. Schuldt. Multi-recipient encryption, revisited. In *Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '14*, page 229–238, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2590296.2590329.
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO 2008*, pages 554–571. Springer, 2008.
- [RAD78] R L Rivest, L Adleman, and M L Dertouzos. On data banks and privacy homomorphisms. *Foundations of Secure Computation*, pages 169–179, 1978.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, September 2009.
- [SGGC14a] Daniel Salmond, Alex Grant, Ian Grivell, and Terence Chan. On the rank of random matrices over finite fields. *arXiv preprint arXiv:1404.3250*, 2014.
- [SGGC14b] Daniel Salmond, Alex J. Grant, Ian Grivell, and Terence Chan. On the rank of random matrices over finite fields. *CoRR*, 2014. URL: <http://arxiv.org/abs/1404.3250>, arXiv:1404.3250.
- [SPB21] István András Seres, Balázs Pejő, and Péter Burcsi. The effect of false positives: Why fuzzy message detection leads to fuzzy privacy guarantees? Cryptology ePrint Archive, Report 2021/1180, 2021. <https://ia.cr/2021/1180>.
- [SPB22] István András Seres, Balázs Pejő, and Péter Burcsi. The effect of false positives: Why fuzzy message detection leads to fuzzy privacy guarantees? In Ittay Eyal and Juan Garay, editors, *Financial Cryptography and Data Security*, pages 123–148, Cham, 2022. Springer International Publishing.
- [WCGFJ12] David Isaac Wolinsky, Henry Corrigan-Gibbs, Bryan Ford, and Aaron Johnson. Dissent in numbers: Making strong anonymity scale. In *OSDI 12*, pages 179–182. USENIX, October 2012.

A Strengthened Privacy Definition of FGOMR

We give the strengthened privacy definition of FGOMR in Fig. 10 as discussed in Theorem 3.3.

B Sparse Random Linear Coding (SRLC)

Everything below is from [LT22, Section 6.3.1]. We copy it here for easier reference in our construction.

Tightly analyzing SRLC parameters is a longstanding, much-studied open problem [KS07, KC16, BJT18]. We proceed to define a notion of SRLC scheme, which can be used as a black box in our construction, and then specify two concrete schemes. SRLC1 has clean analysis using known bounds, and suffices for our asymptotic results. SRLC2 is simpler, faster, and smaller, but relies on empirical estimation for completeness; it is only used in the compression step on the detector side in all (G)OMR schemes.

Definition B.1. An SRLC scheme consists of two algorithms

Adversary	Challenger
	$\xleftarrow{\text{pp}}$
	$\text{pp} \leftarrow \text{GenParams}(1^\lambda, \epsilon_p, \epsilon_n, G, P)$ $(\text{sk}_j, \cdot) \leftarrow \text{PersonalKeyGen}(\text{pp})$ for $j \in [P]$
Query Phase I: for $i = 1, \dots, N$	
Choose recipient set $Y_i \subseteq [P]$, where $ Y_i \leq G$	$\xrightarrow{Y_i}$
	$\xleftarrow{(\text{gPKshare}_{j,i})_{j \in Y_i}}$
	$\text{gPKshare}_{j,i} \leftarrow \text{GroupKeyGenAux}(\text{pp}, \text{sk}_j, Y_i)$, for $j \in Y_i$
Challenge Phase:	
Choose $i_0 \neq i_1 \in [N]$, let $Y' = Y_{i_0} \cap Y_{i_1}$	
Choose $Z \subseteq Y'$ and coins $(r_j)_{j \in Z}$	
$(\text{sk}'_j, \cdot) \leftarrow \text{PersonalKeyGen}(\text{pp}; r_j)$ for $j \in Z$	
For $i \in [i_0, i_1]$:	
$\text{gPKshare}'_{j,i} \leftarrow \text{GroupKeyGen}(\text{pp}, \text{sk}'_j, Y_i)$, $j \in Z$	
Choose $Z_i \subseteq Y_i$, and $(\text{gPKshare}'_{j,i})_{j \in Z_i}$	
Let $\text{gPKshare}'_{j,i} = \text{gPKshare}_{j,i}$, $j \in Y_i \setminus (Z_i \cup Z)$	
Choose a payload x	$\xrightarrow{i_0, i_1, x, Z_0, Z_1}$ $((\text{gPKshare}'_{j,i})_{j \in Y_i})_{i \in [i_0, i_1]}$
	$b \leftarrow \$\{0, 1\}$ $\text{pk}_{\text{clue}_b} \leftarrow \text{GroupKeyGen}(\text{pp}, (\text{gPKshare}'_{j,i_b})_{j \in Y_{i_b}}, Y_{i_b})$
	\xleftarrow{c}
	$c \leftarrow \text{GenClue}(\text{pp}, \text{pk}_{\text{clue}_b}, Y_{i_b} \setminus Z_{i_b}, x)$
Query Phase II: same as Query Phase I.	$\xrightarrow{b'}$
	(Adversary wins if $b = b'$)

Figure 10: Computational Privacy game for Fixed Group OMR with semi-maliciously chosen secret keys

- $(\text{pp}_{\text{SRLC}}, m) \leftarrow \text{GenParams}(1^\lambda, \kappa, \epsilon_F, t)$: takes as input a security parameter λ , κ (number of columns), ϵ_F (defective rate), and a prime number t , and outputs an SRLC public parameter pp_{SRLC} .
- $\{(j, w_j)\} \leftarrow \text{GenWeights}(\text{pp}_{\text{SRLC}})$: takes as input an SRLC public parameter pp_{SRLC} , and outputs a set of indices and weights $\{(j, w_j)\}$ where $j \in [m]$, $w_j \in \mathbb{Z}_t \setminus \{0\}$, representing a sparse vector of length m .

that satisfy the following:

- (Completeness) For any $\kappa \in \mathbb{Z}^+$, and $0 < \epsilon_F < 1$, let $(\text{pp}_{\text{SRLC}}, m) \leftarrow \text{GenParams}(1^\lambda, \kappa, \epsilon_F, t)$, and $(S_i \leftarrow \text{GenWeights}(\text{pp}_{\text{SRLC}}))_{i \in [\kappa]}$. Then the matrix $A \in \mathbb{Z}_t^{m \times \kappa}$ defined by

$$A_{i,j} = \begin{cases} w_j & \text{if } (j, w_j) \in S_i \\ 0 & \text{otherwise} \end{cases}$$

fulfills (over the randomness of the algorithms):

$$\Pr[\text{rank}(A) = \kappa] \geq 1 - \epsilon_F - \text{negl}(\lambda) .$$

We construct two different SRLC algorithms as follows. The first SRLC algorithm **SRLC1**, given in Algorithm 6, is meant for ease of analysis. **GenWeights** outputs set of indices and weights such that the resulting matrix has independently-drawn entries, with a density of nonzeros set by **GenParams**.

Lemma B.1. **SRLC1** in Algorithm 6 is an SRLC scheme for any κ and ϵ_F .

Algorithm 6 SRLC1: Analytically-Bounded SRLC

```
1: procedure SRLC1.GenParams( $1^\lambda, \kappa, \epsilon_F, t$ )
2:   Find the smallest  $m$  such that  $1 - \prod_{i=1}^{\kappa} (1 - (1 - \frac{3}{\kappa})^{m-i+1}) \leq \epsilon_F$  and  $m \geq 6$ 
3:   return ( $m, \text{pp}_{\text{SRLC}} = (\gamma = 3m/\kappa, \kappa, \epsilon_F, t, \lambda)$ )  $\triangleright \gamma$  is the expected number of nonzeros
4: procedure SRLC1.GenWeights( $\text{pp}_{\text{SRLC}} = (\gamma = 3m/\kappa, \kappa, \epsilon_F, t, \lambda)$ )
5:    $\gamma' \leftarrow \mathbf{B}(m, \gamma/m)$   $\triangleright$  binomially-distributed number of nonzeros
6:    $S \leftarrow \{\}, J \leftarrow \{\}$ 
7:   for  $i = 1$  to  $\gamma'$  do
8:      $j \leftarrow_{\S} [m] \setminus J$   $\triangleright$  draw a new index
9:      $w_j \leftarrow_{\S} \mathbb{Z}_t \setminus \{0\}$   $\triangleright$  draw a nonzero weight
10:     $S \leftarrow S \cup \{(j, w_j)\}$ 
11:   return  $S$ 
```

Algorithm 7 SRLC2: Empirically Bounded SRLC

```
1: procedure TestRank( $\gamma, m, \lambda, \epsilon_F, \kappa, t$ )
2:   Ctr  $\leftarrow 0$ 
3:   for  $i = 1$  to  $\log(\lambda) \log \log(\lambda) \epsilon_F^{-1}$  do
4:     Sample  $(S_i)_{i \in [\kappa]}$  where  $S_i \leftarrow \text{SRLC2.GenWeights}(\gamma, \text{pp}_{\text{SRLC}} = (\gamma, m, \epsilon_F, t, \lambda))$ 
5:     Let  $A \in \mathbb{Z}_t^{m \times \kappa}$  be
```

$$A_{i,j} = \begin{cases} w_j & \text{if } (j, w_j) \in S_i \\ 0 & \text{otherwise} \end{cases}$$

```
6:     If  $\text{rank}(A) < \kappa$ : return False
7:   return True
8: procedure SRLC2.GenParams( $1^\lambda, \kappa, \epsilon_F, t$ )
9:    $\gamma \leftarrow 3$ 
10:  while True do
11:    for  $m = \kappa, \dots, \kappa \cdot \gamma/2$  do
12:      if TestRank( $\gamma, m, \lambda, \epsilon_F, \kappa, t$ ) = True then
13:        return ( $m, \text{pp}_{\text{SRLC}} = (\gamma, \kappa, \epsilon_F, t, \lambda)$ )
14:     $\gamma \leftarrow \gamma + 1$ 
15: procedure SRLC2.GenWeights( $\text{pp}_{\text{SRLC}} = (\gamma, \kappa, \epsilon_F, t, \lambda)$ )
16:    $S \leftarrow \{\}, J \leftarrow \{\}$ 
17:   for  $i = 1$  to  $\gamma$  do
18:      $j \leftarrow_{\S} [m] \setminus J$   $\triangleright$  draw a new index
19:      $w_j \leftarrow_{\S} \mathbb{Z}_t \setminus \{0\}$   $\triangleright$  draw a nonzero weight
20:      $S \leftarrow S \cup \{(j, w_j)\}$ 
21:   return  $S$ 
```

Proof. A matrix $A \in \mathbb{Z}_t^{m \times \kappa}$, constructed as in Definition B.1 using SRLC1, has i.i.d. entries that are nonzero with probability γ/m , and uniform when nonzero. The probability that such A has less than full rank κ is upper bounded by $1 - \prod_{i=1}^{\kappa} (1 - \beta^{m-i+1})$ as shown in [KC16, Lemma 1], where $\beta = \max(1 - \gamma/m, \gamma/(m(t-1)))$, and since we fix $\gamma/m = 3/\kappa$, β is always $1 - \gamma/m$ for $m \geq 6$ because $t \geq 2$. Line 2 thus bounds the failure probability by ϵ_F . Therefore, SRLC1 satisfies the completeness requirement of SRLC. \square

Complexity. SRLC1 generates sparse vectors of length $m = O(\kappa \log^2 \kappa \log(\epsilon_F))$, as shown by the following lemma.

Lemma B.2. The value m in SRLC1 chosen by line 2 in Algorithm 6 is $m = O(\kappa \log^2(\kappa) \log(\epsilon_F^{-1}))$, and the value γ chosen by line 3 in Algorithm 6 is $m = O(\log^2(\kappa) \log(\epsilon_F^{-1}))$

Proof. The expression used to choose m can be bounded as follows (using $0 < \epsilon_F < 1$ and $1 < \kappa$):

$$\begin{aligned}
& 1 - \prod_{i=1}^{\kappa} \left(1 - \left(1 - \frac{3}{\kappa}\right)^{m-i+1}\right) \\
& \leq 1 - \prod_{i=1}^{\kappa} \left(1 - \left(1 - \frac{3}{\kappa}\right)^{m-\kappa+1}\right) \\
& \leq 1 - \left(1 - \left(1 - \frac{3}{\kappa}\right)^{m-\kappa+1}\right)^{\kappa} \\
& \leq 1 - \left(1 - \kappa \left(1 - \frac{3}{\kappa}\right)^{m-\kappa+1}\right) \\
& \leq \kappa \left(1 - \frac{3}{\kappa}\right)^{m-\kappa}
\end{aligned}$$

Therefore, $\kappa \left(1 - \frac{3}{\kappa}\right)^{m-\kappa} \leq \epsilon$ suffices for the expression in 2 to be fulfilled. Furthermore:

$$\begin{aligned}
& \kappa \left(1 - \frac{3}{\kappa}\right)^{m-\kappa} \leq \epsilon_F \\
& \Leftrightarrow \log\left(\kappa \left(1 - \frac{3}{\kappa}\right)^{m-\kappa}\right) \leq \log(\epsilon_F) \\
& \Leftrightarrow \log(\kappa) + (m - \kappa) \log\left(1 - \frac{3}{\kappa}\right) \leq \log(\epsilon_F) \\
& \Leftrightarrow (m - \kappa) \log\left(1 - \frac{3}{\kappa}\right) \leq \log(\epsilon_F/\kappa) \\
& \Leftrightarrow m - \kappa \geq \frac{\log(\epsilon_F/\kappa)}{\log\left(1 - \frac{3}{\kappa}\right)} \\
& \Leftrightarrow m \geq \frac{\log(\epsilon_F/\kappa)}{\log\left(1 - \frac{3}{\kappa}\right)} + \kappa = \frac{\log(\epsilon^{-1}\kappa)}{-\log\left(1 - \frac{3}{\kappa}\right)} + \kappa
\end{aligned}$$

Since $\frac{\log(\epsilon_F^{-1}\kappa)}{-\log\left(1 - \frac{3}{\kappa}\right)} = O(\kappa \log^2(\kappa) \log(\epsilon_F^{-1}))$ and we choose the smallest m in line 2, we indeed get $m = O(\kappa \log^2(\kappa) \log(\epsilon_F^{-1}))$. By line 3, $\gamma = 3m/\kappa = O(\log^2(\kappa) \log(\epsilon_F^{-1}))$. \square

The second SRLC algorithm SRLC2, given in Algorithm 7, has simpler **GenWeights** which just outputs a *fixed* number γ of nonzero elements (with uniformly-random nonzero weight each). However, its completeness relies γ and m being chosen by an empirical estimate (encapsulated within **GenParams**), since adequately tight analytical bounds are not known.

Lemma B.3. SRLC2 in Algorithm 7 is an SRLC scheme for any κ and ϵ_F .

Proof. Let $p_{\gamma,m}$ denote the probability that a matrix generated as in Definition B.1 is not full-rank, for given γ, m . For SRLC completeness, it suffices to prove that if $p_{\gamma,m} > \epsilon_F$, then with probability $1 - \text{negl}(\lambda)$, **TestRank** will output False and thus **GenParams** will not use output these γ, m .

The probability of passing **TestRank** is the probability that $\epsilon_F \log(\lambda) \log \log(\lambda)$ independent $p_{\gamma,m}$ -biased Bernoulli tests are 0. Thus, if $p_{\gamma,m} > \epsilon_F$ then the probability of passing is negligible: $(1 - p_{\gamma,m})^n < (1 - \epsilon_F)^n = ((1 - \epsilon_F)^{\epsilon_F^{-1}})^{\log(\lambda) \log \log(\lambda)} = O(e^{-\log(\lambda) \log \log(\lambda)}) = O(\lambda^{-\log \log(\lambda)}) = \text{negl}(\lambda)$.

It follows that the output of **GenParams** fulfills $p_{\gamma,m} \leq \epsilon_F + \text{negl}(\lambda)$.¹⁵ \square

¹⁵Assuming **GenParams** runs in time $\text{poly}(\lambda)$. We do not prove this directly, but **GenParams.SRLC2** will be empirically executed honestly, and the security model assumes that it is infeasible even for the adversary to run for longer than $\text{poly}(\lambda)$.