Anamorphism Beyond One-To-One Messaging: Public Key with Anamorphic Broadcast Mode

Xuan Thanh Do[1], Giuseppe Persiano[2,3], Duong Hieu Phan[4], and Moti Yung[3,5]

[1] Institute of Cryptography Science and Technology, Vietnam
[2] Università di Salerno, Italy
[3] Google LLC, USA
[4] Telecom Paris, Institut Polytechnique de Paris, France
[5] Columbia University, USA

Abstract. To date, Anamorphic Cryptography [EC22] has been developed to support the addition of an *anamorphic* plaintext message to a ciphertext carrying a primary message. The anamorphic message remains hidden even in the presence of a strong adversary that possesses the receiver's key and/or determined the sent primary message. In this paper, we expand one-to-one encrypted anamorphic communication to one-to-many anamorphism, naturally assuming communication over a broadcast channel. What we show is that using a previously designed public-key encryption scheme, two things can happen: First, the receiver of an added hidden message may be a party different from the actual receiver (i.e., a shadow party) who has initially collaborated with the sender. Secondly, and perhaps more surprisingly, the receiving party need not be a singleton, and can be a number of different shadow (i.e., anonymous) groups, each receiving a different anamorphic message, where all these messages are extracted from one single-receiver ciphertext. The idea of having multiple hidden channels to different shadow groups is highly handy if, for example, the anamorphic messages are warnings with operational instructions, sent to the groups, and will be received by (at least one member of) a group even if the adversary is able to temporarily cut off all but one member of a channel.
More specifically,

- First, we motivate and formalize the notion of *Public-Key Encryption with an Anamorphic Broadcast Mode*.
- We then present, as an initial result of an independent interest, the first lattice-based construction of *Anonymous Multi-Channel Broadcast Encryption*. It is important to note here that all Multi-Channel Broadcast schemes to date are in the pairing-based setting (and are, thus, insecure against quantum adversaries).
- Finally, we show how to transform a strong form of anonymity (where the ciphertext also hides the number of channels) into a system with anamorphism in the multi-channel broadcast setting for the well-known Dual Regev Public-Key Encryption scheme. Specifically, we show that, given the single-receiver public key pk for the Dual Regev encryption scheme, and a sequence of ℓ messages for the ℓ channels of the broadcast scheme, it is possible to create a ciphertext that will carry the ℓ messages and is also a legitimate ciphertext for pk.

Table of Contents

1 Introduction **3**
- 1.1 One-to-many anamorphic channels: motivation & challenges . . . 4
- 1.2 Reviewing Our Contributions . . . 5
- 1.3 Related works . . . 7
- 1.4 Paper organization . . . 8

2 Preliminaries **8**
- 2.1 Notations . . . 8
- 2.2 Public Key Encryption . . . 8
- 2.3 Dual Regev Public Key Encryption . . . 9
- 2.4 Lattice and k-LWE problem . . . 9

3 Anonymous Multi-Channel Broadcast Encryption **11**
- 3.1 Broadcast systems . . . 11
- 3.2 IND−CPA security . . . 12
- 3.3 Anonymity . . . 12
- 3.4 Construction . . . 13
- 3.5 Security Proof . . . 14

4 Public Key Encryption with an Anamorphic Broadcast Mode **16**
- 4.1 Formal Definition . . . 16
- 4.2 Construction . . . 18

1 Introduction

We deal with encrypted communication over a broadcast medium where all users have access to the messages. The fundamental goal of encryption in this environment is to ensure that the plaintext carried by a ciphertext is accessible only to the sender (obviously) and to the intended recipient(s). The well established *Kerckhoffs' principle* regarding encryption states that only the decryption key and the message need to be kept secret and all other information about the cryptosystem, including ciphertexts, should be assumed public. This is a natural and minimal requirement because, first, if the message is not secret anymore, there is no need to encrypt it; second, if the decryption key is compromised, anyone holding that key can decrypt and recover the message. Hence, normal encryption relies on two primary assumptions (which are very natural, and hence typically implicit): the *sender-privacy*[6] and the *receiver-privacy* assumption. The sender-privacy assumption states that the encrypted message is freely and privately chosen by the sender, whereas the receiver-privacy assumption says that among all parties, only the recipients have access to the secret key for decryption. The assumptions are the norms in generally free societies, or more generally, when no one attempts to infringe the rights of users to employ cryptography.

In this paper, we consider the anamorphic setting [PPY22] of achieving private communication in the presence of a powerful adversary D that has the power to request all secret information whose existence cannot be denied (given a transmitted ciphertext) in a broadcast communication environment; in addition, the adversary may block (at least temporarily) a bounded number of users from the broadcast medium.

Specifically, the adversary's power can be described by the following three rules (in the given broadcast environment):

The ciphertext rule. The adversary D can request to associate any ciphertext `ct` seen on the network with an encryption key according to which `ct` is a valid ciphertext.

The key rule. The adversary D can ask for the secret key associated with any public key that has been made public by its owner or has been associated with a ciphertext according to the previous rule.

The blocking rule. The adversary D can block communication accessible to any party, making the party unable to receive any communication. (Obviously, the number of blocked parties at any time has to be bounded for the infrastructure to actually fulfill its goal.)

To deal with the extended power of the adversary above, we allow parties wishing to communicate covertly to initially and secretly exchange keys. Note further that the covert receiver which can be called the shadow receiver should remain unknown (which is possible since reading a message off, say, a public bulletin board is an oblivious operation by any party). This is a minimal requirement as well, since otherwise the powerful adversary can simply block the shadow receiver. Then we ask:

> *Is private communication for shadow receivers (other than the intended receiver of the primary message) possible in the presence of such an adversary?*

Note that the postulated adversary is in blatant violation of the Kerckhoffs' principle as it has access to secret key material and, in addition, the adversary monitors (i.e., reads) and controls (i.e., allows or blocks) arbitrary communication over the broadcast media.

The setting of interest and a partial solution. Consider a simple scenario in which A wants to privately send a message `am` to a shadow receiver C. Due to the key rule, we cannot use any asymmetric key that C has published but rather C will initially share a *hidden* key with A. Still, this is not sufficient because, even if C's key is hidden from D, the moment A injects a ciphertext `ct` into the network, by the ciphertext rule, D will ask to link `ct` to a key. We therefore suggest the following setting. Let B be a third party with a published key. A produces a ciphertext `ct` that has the following property: `ct` is a valid ciphertext for B (in the sense that decryption of `ct` with the secret key related to B's published key will be successful); moreover, if `ct` is decrypted with C's secret key, then message `am` will result.

[6] This assumption was termed *sender-freedom* in the original paper on anamorphic encryption [PPY22].

This scenario admits a solution based on receiver anonymous cryptosystems and, in particular, a very simple implementation with the ElGamal encryption scheme. Specifically, B and C both have a pair of public and secret keys for ElGamal over the same group, $(\mathsf{pk}_B, \mathsf{sk}_B)$ and $(\mathsf{pk}_C, \mathsf{sk}_C)$, respectively. B publishes the public key pk_B, while C privately shares pk_C with A. Whenever A wishes to send a message $\mathtt{am}$ to C, they can simply compute $\mathtt{ct}$ by encrypting $\mathtt{am}$ with pk_C. Note that $\mathtt{ct}$, despite being computed on input pk_C, is a valid ciphertext (though of a different message) with respect to pk_B.

In other words, the answer to our previous question is partially positive.

> *Yes. Private (covert) communication is possible, provided that the adversary is non-blocking.*

Indeed, we observe that even though A's anamorphic message $\mathtt{am}$ is not exposed, D might have other reasons to be suspicious of C and couls block C from receiving any message. Note that this is within the ability of D as stated by the blocking rule. Also, the adversarial setting covers the case where C might simply not be online to receive A's message and, in the case that this message is a distress message, this will have consequences. Then what does it mean that A wants to communicate with C, if C is cut off the network, voluntarily or not?

In fact, we are interested in the case in which A wants to send a distress message and A wishes that the message be read and acted upon more reliably by any member of a group $C_1, \ldots, C_\ell$ of trusted users. Clearly, A could repeat the above solution by using ℓ different ciphertexts, one for each of $C_1, C_2, \ldots, C_\ell$, hoping that the adversary has not blocked all of them and that at least one of them is online. However, this approach imposes an unreasonable burden on A and on B who will receive ℓ ciphertexts from A for no apparent reason. Since in a broadcast medium an ℓ-times increase in messaging is noticed by everyone, this "flooding" will be realized and will not remain hidden (by the usual simple message statistics). We thus reformulate our research question as follows:

> *Is* efficient flooding-free *private (covert) communication possible in the presence of a blocking adversary?*

This state of things calls for a *one-to-many* approach to our problem in which A can use a single ciphertext that can actually be read by multiple clients $C_1, C_2, \ldots, C_\ell$ who share a hidden secret key with A. Furthermore, our solution will allow more than that: it employs a single ciphertext that include different shadow messages to different shadow groups, without flooding the broadcast medium. Next, we present our approach.

1.1 One-to-many anamorphic channels: motivation & challenges

We note the following: previous anamorphic schemes, and sender-anamorphic ones, in particular, all facilitate one-to-one anamorphic communication (from a single sender to a single receiver who is the receiver of the regular message). Therefore, in previous works, if the intended single receiver is unavailable or blocked by the adversary, anamorphic communication fails [PPY22, KPP+23b, KPP+23a, WCHY23, BGH+24, CGM24a]. In a global broadcast medium setting (a bulletin board, a blockchain, etc.) the presence of a blocking adversary has to be considered, since in asynchronous communication environments users may be offline at times, and a malicious adversary may even block users purposely. Thus, blocking is a natural and significant challenge, and overcoming this challenge lies in establishing a one-to-many anamorphic communication method. This method attempts to reach a group of different shadow receivers, ensuring that the anamorphic message will be successfully received even when only one of the receivers is available at the broadcast medium to receive the single ciphertext with an additional anamorphic message/key broadcast.

One-to-many transmission would effectively prevent adversarial control over the transmission process (in case the intended receiver is taken off the communication medium maliciously) and help the sender broadcast the anamorphic message more efficiently and without flooding; e.g., the sender is perhaps a person under duress calling covertly for help from any member of a group of potential receivers! Note in fact that by denoting by ℓ the size of the receiving group, if we allow the adversary to block or (temporarily) remove from the system up to $\ell - 1$ parties, the anamorphic message will be read. We call such a primitive an encryption

scheme with *anamorphic broadcast mode*. In this context, we let A denote the sender of the primary message, B denote its receiver, and $C_1, \ldots, C_\ell$ be the shadow receivers of the anamorphic message. The distress signal is an example of the usefulness of the system. Another use is to designate in the anamorphic message the status of the primary message (whether it is a forced fake message or a real and urgent, say, message); any of the receivers in the group may read the status and may wake up the receiver to read the message and act with urgency (or otherwise let the receiver continue sleeping!). As the examples show, there are scenarios where increasing the group of possible receivers over a broadcast medium makes a lot of sense (e.g., receivers are in different time zones or work in different shifts, etc.). Also, another motivating example is the case where the covert message demands some parallel action and it is important that a multitude of responding parties, rather than a single one, act fast. Yet another application may be covert coordination by the sender of actions involving a multitude of otherwise anonymous (shadow) receivers. Further, if we can deal with a number of groups, sending each a different message, then the flexibility of applications grows significantly. To summarize all the above, it is quite clear that the scenario we deal with is desirable and makes quite sense in the setting of messages hidden from the strong adversary (and receivers hidden as well!).

The above scenarios naturally motivate the need to design an encryption scheme with an anamorphic broadcast mode. It further leads to considering and blending two primitives: multi-channel broadcast encryption, introduced in [PPT13] and anonymous broadcast encryption, introduced in [BBW06] and developed in [LPQ12, FP12, LG18, DPY20, BMN+21]. Indeed, the multi-channel setting further allows the setup of multiple anamorphic channels so that the sender can establish different channels (i.e., different messages) with different sets of receivers, thus enhancing the functionality and messaging flexibility of the anamorphic system. Anonymous broadcast, in turn, aims at hiding the set of receivers, since otherwise, the dictator could attempt to control (and block) all of these receivers, say, and even knowing them without blocking may reveal information. In this work, in fact, we design an anonymous multi-channel broadcast encryption with the extra property that the ciphertexts are indistinguishable from the ciphertexts of an existing one-to-one scheme (needed for keeping even the existence of covert anamorphic messages hidden within a legitimate cryptosystem). In fact, we achieve this highly constrained setting with the Dual Regev Encryption scheme [GPV08, CHKP10], a scheme famously designed as a lattice-based scheme for the purpose of one-to-one public-key encryption. Recall that the interesting aspect of anamorphic methods is the fact that the anamorphic setting is added on top of systems which were originally designed without anamorphism in mind. In our one-to-many approach we preserve this aspect, which indicates again (this time in an extended setting) that cryptographic systems naturally possess anamorphic capacities!

1.2 Reviewing Our Contributions

The Concept. We introduce the concept of *encryption schemes with anamorphic broadcast mode*. This primitive allows a sender to broadcast an anamorphic message, inserting it in a ciphertext that is a valid ciphertext (with a corresponding valid regular plaintext) with respect to a regular encryption scheme. We then extend this concept to *multi-channel broadcast* in which the same ciphertext carries different messages to multiple disjoint sets of receivers (called *channels*). The adversary has access to the secret key of the receiver of the regular encryption scheme and can verify the validity of the regular ciphertext. We note that existing work only allows the sender to establish an anamorphic channel with *one* receiver, and if this receiver is not available or his connection is broken, then the anamorphic channel obviously does not work. The broadcast mode fortifies the anamorphic channel to a set of receivers and is highly suitable for covert signaling of a duress state of the sender, as argued above. However, this concept requires (as a necessary but seemingly hard-to-achieve condition) that the ciphertext outputted by the encryption scheme be indistinguishable from the ciphertext outputted by the added anamorphic broadcast mode. Otherwise, the adversary will trivially detect the use of the anamorphic mode. Achieving this condition, in fact, challenges us to develop a new basic technique to allow the design of the concept prescribed above.

The technique. A basic necessary condition for our approach to succeed is that the ciphertexts are independent from the public keys they refer to. Indeed, we exploited this property in our ElGamal-based "toy construction" for the one-to-one case. However, there we implicitly assumed that the two ElGamal public keys

(i.e. B's and C's) were generated with respect to the same prime-order group since, otherwise, ciphertexts from different groups can be easily distinguished. Lattice-based schemes share a common matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$. However, there is a unique property that current lattice-based encryption schemes enjoy: two instantiations (one for public-key encryption and one for broadcast encryption) can use two different matrices, $\mathbf{A}_1$ and $\mathbf{A}_2$, but since both matrices appear random in $\mathbb{Z}_q^{m \times n}$, the corresponding ciphertexts will both appear random in $\mathbb{Z}_q^{m+1}$, making them indistinguishable.

From the above crucial observation, we can extract an anamorphic system following the following script: In the regular mode, the ciphertext corresponds to the matrix $\mathbf{A}_{\mathsf{ENC}}$ associated with a one-to-one encryption ENC, while in the anamorphic mode, the broadcast ciphertext corresponds to the matrix $\mathbf{A}_{\mathsf{MCBE}}$ associated with a *multi-channel broadcast encryption* MCBE. Employing a different matrix than in the regular mode, the sender can generate $\mathbf{A}_{\mathsf{MCBE}}$ along with "a trapdoor" that enables him to generate an anamorphic key for the broadcast anamorphic mode.

In order to implement the above idea, the broadcast mode needs to be anonymous for different sets of receivers, and therefore we need to construct an *Anonymous Multi-Channel Broadcast* scheme with ciphertexts indistinguishable from those of a standard encryption scheme. Fortunately, this, in turn, can be achieved with LWE encryption.

The construction. We note that the existing MCBE schemes [PPT13, CPPT18, LTTT20, AD18, KHS+20] do not meet our two requirements. First, none of them has the same ciphertext format as any existing encryption. Second, only two schemes [AD18, KHS+20] achieve a certain level of anonymity, which is still insufficient for our setting. Indeed, the constructions by Acharya et al. [AD18] only achieve outsider anonymity which protects the privacy of receivers from outsiders (who are not in the broadcast channel). Note that in our setting of interest, insider anonymity is important as it guaranties that, even if the dictator manages to capture one member of the broadcast channel, the identity of the other members stays hidden. Moreover, these schemes do not have a ciphertext size independent of the number of channels, thus leaking the number of channels and failing to keep the ciphertext format unchanged, making them unsuitable for the anamorphic context. The construction by Kim et al. [KHS+20] only achieves static security, as the adversary is forced to choose the broadcast sets prior to setup. This is not suitable for our setting, where the sender can choose the anamorphic sets of receivers after all the parameters are set up.

Our first contribution is thus to construct an anonymous multi-channel broadcast scheme that is suited for our objective. We start from the anonymous broadcast scheme in [DPY20]. We generalize it to the multi-channel broadcast setting while preserving the ciphertext format. Finally, relying on k-LWE [LPSS14], we prove full anonymity with adversarially chosen target sets, which in turn allows us to implement an anamorphic broadcast mode. We believe that our anonymous multi-channel broadcast is of independent interest since it is the first secure scheme against adaptive adversaries and quantum adversaries. We further believe that our setting of anamorphic broadcast is a valid setting for a bounded model (where the total size of broadcast sets is limited) and can further boost research in the area of anonymous schemes within the bounded model, where the lower bound for anonymity in the general case prevents the efficient construction of anonymous broadcast and anonymous multi-channel broadcast schemes.

Our second contribution is showing that our anonymous multi-channel broadcast scheme is the anamorphic broadcast mode for the well known Dual Regev encryption scheme [GPV08, CHKP10].

It is worth noting that ours is the first example of anamorphism across different cryptographic encryption primitives.

Limitations in multi-channel and anonymous broadcast have no impact on our specific setting. Our approach inherits the limitations shared by all multi-channel broadcast encryptions [PPT13, CPPT18, LTTT20, AD18, KHS+20]: First, the sets of receivers corresponding to the channels must be disjoint. Second, there exists the following lower bound for anonymous broadcast [KS12]: for a large number of users N, the ciphertext size has to be linear to N. These limitations hinder the use of multi-channel and anonymous broadcast in numerous practical scenarios, since the typical use-case for broadcast encryption is Pay-TV or satellite transmission with potentially tens and even hundreds of millions of users. In contrast, for a bounded model where the number of users is rather limited, anonymous broadcast [DPY20] can be very efficient,

exactly as efficient as standard encryption. Even though in general it is difficult to find a setting in which disjoint channels over a bounded universe of users are not a limitation, this is not the case for our setting of anamorphic broadcast. Indeed, the objective of the sender in anamorphic mode is to set anamorphic channels with only a handful of people he totally trusts (and certainly their number is small, which is typical of a conspiracy group, namely a much fewer than the massive number of subscribers to Pay-TV!). Hence, in fact, we can easily say that our setting fits the bounded model perfectly. The sender also wants to set different multi-channels with different sets of receivers (compartmentalizing the receiving sets of, say, alerts regarding different subjects), naturally requiring that these receivers are in disjoint sets. Again, this is much unlike Pay-TV where one user may want to subscribe to many channels.

1.3 Related works

The concept of Anamorphic Encryption [PPY22] is the closest to and the central context of our work. It was invented to address adversaries that violate Kerckhoffs' principle thus demonstrating the futility of mandating escrow of keys and posing other limitations on the use of strong cryptography– a situation which was dubbed the "Crypto Wars." The prevalence of anamorphism in encryption was shown by [KPP+23b]. The concept was extended to signature schemes in [KPP+23a], where anamorphic encryption is performed over a signature scheme. The idea of realizing one hidden primitive over another was also discussed in [PY21]. The concept of anamorphic encryption has been extended by Banfi et al. [BGH+24] to include the notion of robustness. More recently, theoretical aspects of anamorphic encryption have been studied in [PPY24, CGM24b] and asymmetric versions of the notion have been proposed by [CGM24a] and, in a stronger form, by [PPY24].

In the original work on anamorphic encryption [PPY22] against an attacking dictator, the authors defined two types of anamorphic encryption: sender-anamorphic encryption and receiver-anamorphic encryption. These encryption systems ensure that they remain methods for exchanging secret messages even when the above attacker violates the sender-privacy assumption and/or the receiver-privacy assumption. Receiver-Anamorphic Encryption is a two party message sending where the message receiver and the anamorphic message receiver are the same.

In our work, we extend the anamorphic approach to a one-to-many setting in which the anamorphic message is received over a public communication channel by a shadow receiver different from the real receiver (akin to sender anamorphic), where shadow receiver is a set of groups of receivers, with each group (channel) receiving a different message. Unlike sender anamorphism, the adversary gets the key of the public communication (and therefore the regular message) but does not get the anamorphic messages, nor does he know who received it, and whether it exists at all. Furthermore, in our setting, even if the adversary blocks a number of parties, say t, of a group of size $k > t$, the group will still get the anamorphic message (in fact, all groups will get their messages off the public communication channel).

In [KPP+23b] and [BGH+24], the concept of a multi-receiver setting or multiple cover model for receiver-anamorphic encryption was explored. However, we note that these are broadcasts with a different meaning from our anamorphic broadcast mode. In [KPP+23b], Multiple-Receiver Anamorphic denotes decryption relying solely on the double secret key dsk, without necessitating the anamorphic secret key ($\mathtt{ask}$). Thus, any user who has the double secret key can decrypt the anamorphic message without compromising security. In [BGH+24], a multiple covert model is considered, allowing a user to generate many double keys for its public key. However, in both scenarios, encryption requires the double key, which is then needed for decryption, thus still being a one-to-one communication. Extending the receiver-anamorphic model to enable encryption with a set of double keys, then anyone can decrypt with one of these double keys, seems an open interesting question. Our work relies on and extends previous work on broadcast encryption, introduced by Fiat and Naor [FN94]. Broadcast encryption is a useful cryptographic technique that allows a broadcaster to efficiently transmit encrypted content through a public channel to a specific group of users while preventing access to those outside the group. However, the repeated use of broadcast encryption to send different ciphertexts to multiple user groups over multiple channels significantly increases both computation and bandwidth. To address this issue, Phan et al. [PPT13] put forth the multi-channel broadcast encryption (MCBE) concept and provided an efficient construction. In an MCBE scheme, there is a single header for multiple user groups, enabling all members of a specific group to extract the same session key. Each group has a distinct session

key, and the single header contains all the necessary encrypted information to allow each group to recover their respective session keys.

An essential performance criterion for MCBE is that the size of the header/ciphertext remains constant, regardless of the number of channels in the system. Furthermore, it would be a desired criterion if an MCBE could protect user privacy, and such a system is called an anonymous MCBE.

1.4 Paper organization

In Section 2 the basic Lattice-based cryptographic background is reviewed. In Section 3 Anonymous Multi-Channel Broadcast Encryption is defined, designed, and its properties are proven, while in Section 4 we use the construction to build our scheme of "Public Key Encryption with an Anamorphic Broadcast Mode," and prove its properties.

2 Preliminaries

2.1 Notations

We denote the security parameter by λ, the set of real numbers by $\mathbb{R}$, and the integers by $\mathbb{Z}$. Let $[n]$ denote the set $\{1, 2, \ldots, n\}$ for an integer n. For any $q \geq 2$, we let $\mathbb{Z}_q$ denote the ring of integers with addition and multiplication modulo q. Let $\lfloor x \rfloor$ denote the largest integer less than or equal to x, and let the rounding function for a real number x, denoted by $\lfloor x \rceil$, be the closest integer to x. If x is exactly halfway between two integers, the rounding function rounds to the nearest even integer. Denote $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ as the group of reals $[0, 1)$ with mod 1 addition.

A probabilistic polynomial time (PPT) algorithm $\mathcal{A}$ runs in time polynomial in the (implicit) security parameter λ. A function $f(\lambda)$ is negligible in λ if it is $O(\lambda^{-c})$ for every $c \in \mathbb{N}$ and we write $f = \mathsf{negl}(\lambda)$ for short. We say that a probability is overwhelming if it is $1 - \mathsf{negl}(\lambda)$. Similarly, we write $f = \mathsf{poly}(\lambda)$ if $f(\lambda))$ is a polynomial with variable λ. We let $\langle \mathbf{a}, \mathbf{b} \rangle$ denote the inner product of the vectors $\mathbf{a}$ and $\mathbf{b}$ of the same length.

If D is a probability distribution, $x \leftarrow D$ means that x is sampled from D and if S is a countable set, $x \xleftarrow{\$} S$ means that x is sampled uniformly and independently at random from S. We also let $U(S)$ denote the uniform distribution over S. If D_1 and D_2 are distributions over a countable set X, their statistical distance is defined as $\frac{1}{2}\sum_{x \in X} |D_1(x) - D_2(x)|$ and will be denoted by $\Delta(D_1, D_2)$. We say that two distributions (two ensembles of distributions indexed by n) are statistically close if their statistical distance is negligible in n. Furthermore, we say that two distributions D_0 and D_1 are computationally indistinguishable, denoted $D_0 \approx D_1$, if for all PPT adversaries $\mathcal{A}$, we have

$$|\mathsf{Pr}[x_0 \leftarrow D_0 : \mathcal{A}(1^n, x_0) = 1] - \mathsf{Pr}[x_1 \leftarrow D_1 : \mathcal{A}(1^n, x_1) = 1]| = \mathsf{negl}(n).$$

We review the basic notion of $\mathsf{IND-CPA}$ security and the Dual Regev Encryption scheme in the following Section

2.2 Public Key Encryption

We start by defining the notion of a public key encryption scheme.

Definition 1. *A* public key encryption scheme $\mathcal{E}$ *is a tuple* $\mathcal{E} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ *of* PPT *algorithms with the following syntax*

1. *the* key-generator algorithm KG *takes as input the security parameter* 1^λ*, and returns a* public key pk *and a* secret key sk*;*
2. *the* encryption algorithm Enc *takes as input a public key* pk *and a message* m*, and returns a ciphertext* ct*;*

3. *the* decryption algorithm Dec *takes as input a ciphertext* $\mathtt{ct}$ *and a secret key* sk*, and returns a message* m *or the symbol* $\perp$*.*

and that satisfy the following correctness requirement:

- *there exists a negligible function* $\mathsf{negl}(\lambda)$ *such that for all* m

$$\mathsf{Pr}\left[\begin{array}{ll}\mathsf{Dec}(\mathsf{sk},\mathtt{ct}) \neq m: & (\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KG}(1^\lambda) \\ & \mathtt{ct} = \mathsf{Enc}(\mathsf{pk}, m)\end{array}\right] \leq \mathsf{negl}(\lambda).$$

We will consider the $\mathsf{IND-CPA}$ notion of security for $\mathcal{E}$.

Definition 2. *A public key encryption scheme* $\mathcal{E}$ *is called* $\mathsf{IND-CPA}$ secure *if, for every* PPT *adversary* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, there exists a negligible function* $\mathsf{negl}(\lambda)$ *such that*

$$\mathsf{Pr}\left[\begin{array}{ll} & (\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KG}(1^\lambda) \\ & (m_0, m_1, \mathsf{st}) \leftarrow \mathcal{A}_1(\mathsf{pk}) \\ b = b': & b \xleftarrow{\$} \{0,1\} \\ & \mathtt{ct}^* \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b) \\ & b' \leftarrow \mathcal{A}_2(\mathsf{st}, \mathtt{ct}^\star)\end{array}\right] - \frac{1}{2} \leq \mathsf{negl}(\lambda).$$

2.3 Dual Regev Public Key Encryption

We next recap the Dual Regev Encryption scheme, the 1-bit public key encryption scheme from [GPV08] whose security relies on the hardness of the LWE problem.

The scheme is parameterized by two integers, m and q, a positive real number r and a 1-dimensional Gaussian distribution ν_α with standard deviation α. Following [GPV08], we set the parameters α, q, r, and m in the following way to guarantee both security and correctness: $q \geq 5r(m+1)$ is a prime, $r \geq \omega(\sqrt{\log m})$, $\alpha \leq 1/(r\sqrt{m+1} \cdot \omega(\sqrt{\log n}))$, and $m \geq 2n \log q$.

$\mathsf{KG}(1^\lambda)$**:** Set $n = \lambda$ and randomly select matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$. $\mathbf{A}$ is the index of the function $f_{\mathbf{A}}(\mathbf{e}) = \mathbf{A}^T\mathbf{e}$ mod q. Choose an error vector $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, r}$. The secret key is $\mathbf{e}$ and the public key is the pair $(\mathbf{A}, \mathbf{u})$ where $\mathbf{u} = f_{\mathbf{A}}(\mathbf{e})$ is the syndrome.

$\mathsf{Enc}((\mathbf{A}, \mathbf{u}), b)$**:** Sample $\mathbf{s} \xleftarrow{\$} U(\mathbb{Z}_q^n)$ and set $\mathbf{p} = \mathbf{A}\mathbf{s} + \mathbf{x} \in \mathbb{Z}_q^m$, where $\mathbf{x} \leftarrow \nu_\alpha^m$. Output ciphertext $\mathtt{ct} = (\mathbf{p}; c = \mathbf{u}^T\mathbf{s} + x + b\lfloor \frac{q}{2} \rfloor) \in \mathbb{Z}_q^m \times \mathbb{Z}_q$, where $x \leftarrow \nu_\alpha$.

$\mathsf{Dec}((\mathbf{p}, c), \mathbf{e})$**:** Compute $z = c - \mathbf{e}^T\mathbf{p} \in \mathbb{Z}_q$. Output 0 if $z \mod q$ is closer to 0 than to $\lfloor \frac{q}{2} \rfloor$; otherwise output 1.

2.4 Lattice and k-LWE problem

For two matrices $\mathbf{A}, \mathbf{B}$ of compatible dimensions, let $(\mathbf{A} \| \mathbf{B})$ (or sometimes $\left(\frac{\mathbf{A}}{\mathbf{B}}\right)$) denote vertical concatenations of $\mathbf{A}$ and $\mathbf{B}$. For $\mathbf{x} \in \mathbb{Z}^m$, define $\mathbf{x}^+ = (1 \| \mathbf{x}) \in \mathbb{Z}^{m+1}$. For $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, define $\mathrm{Im}(\mathbf{A}) = \{\mathbf{A}\mathbf{s} \mid \mathbf{s} \in \mathbb{Z}_q^n\} \subseteq \mathbb{Z}_q^m$. For $\mathbf{X} \subseteq \mathbb{Z}_q^m$, let $\mathrm{Span}(\mathbf{X})$ denote the set of all linear combinations of elements of $\mathbf{X}$ and define $\mathbf{X}^\perp$ to be $\{\mathbf{b} \in \mathbb{Z}_q^m \mid \forall \mathbf{c} \in \mathbf{X}, \langle \mathbf{b}, \mathbf{c} \rangle = 0\}$.

Let $\mathbf{B} = \{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n\} \subset \mathbb{R}^n$ consists of n linearly independent vectors. The n-dimensional lattice Λ generated by the basis $\mathbf{B}$ is $\Lambda = L(\mathbf{B}) = \{\mathbf{B}\mathbf{c} = \sum_{i \in [n]} c_i \cdot \mathbf{b}_i \mid \mathbf{c} \in \mathbb{Z}^n\}$. The length of a matrix $\mathbf{B}$ is defined as the norm of its longest column: $\|\mathbf{B}\| = \max_{1 \leq i \leq n} \|\mathbf{b}_i\|$. Here we view a matrix as simply the set of its column vectors.

For a lattice $L \subseteq \mathbb{R}^m$ and an invertible matrix $\mathbf{S} \in \mathbb{R}^{m \times m}$, we define the Gaussian distribution of parameters L and $\mathbf{S}$ by $D_{L,\mathbf{S}}(\mathbf{b}) = \exp(-\pi \|\mathbf{S}^{-1}\mathbf{b}\|^2)$ for all $\mathbf{b} \in L$.

The q-ary lattice associated with a matrix $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$ is defined as $\Lambda^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{x}^t \cdot \mathbf{A} = \mathbf{0} \bmod q\}$. It has dimension m, and a basis can be computed in polynomial-time from $\mathbf{A}$. For $\mathbf{u} \in \mathbb{Z}_q^n$, we define $\Lambda_{\mathbf{u}}^\perp(\mathbf{A})$ as the coset $\{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{x}^t \cdot \mathbf{A} = \mathbf{u}^t \bmod q\}$ of $\Lambda^\perp(\mathbf{A})$.

Let ν_α denote the 1-dimensional Gaussian distribution with standard deviation α, $\alpha \in (0, 1)$.

Lemma 1 (Theorem 3.1, [AP11]). *There is a probabilistic polynomial-time algorithm that, on input positive integers $n, m, q \geq 2$, outputs two matrices $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$ and $\mathbf{T} \in \mathbb{Z}^{m\times m}$ such that the distribution of $\mathbf{A}$ is within statistical distance $2^{-\Omega(n)}$ from $U(\mathbb{Z}_q^{m\times n})$; the rows of $\mathbf{T}$ form a basis of $\Lambda^\perp(\mathbf{A})$; each row of $\mathbf{T}$ has norm $\leq 3mq^{n/m}$.*

Lemma 2 (GPV, [GPV08, LPSS14]). *Given $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$, an invertible $\mathbf{S} \in \mathbb{R}^{m\times m}$, and $\mathbf{u} \in \mathbb{Z}_q^n$, there exists a probabilistic polynomial-time algorithm that, on input a trapdoor matrix $\mathbf{T} \in \mathbb{Z}^{m\times m}$ for the lattice $\Lambda^\perp(\mathbf{A})$, it outputs a sample from a distribution that is statistically close to $D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}),\mathbf{S}}$.*

Definition 3 (The LWE problem). *Let $\alpha > 0$, n and $m = m(n)$ be integers, $q = q(n)$ a prime. The $\mathsf{LWE}_{\alpha,q,m}(n)$ problem is defined as follows. Given $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{m\times n})$, distinguish between the following two distributions over $\mathbb{T}^m$*

$$\frac{1}{q} \cdot U\Big(\mathrm{Im}(\mathbf{A})\Big) + \nu_\alpha^m \quad \textit{and} \quad \frac{1}{q} \cdot U\Big(\mathbb{Z}_q^m\Big) + \nu_\alpha^m,$$

where ν_α denotes the one-dimensional Gaussian distribution with standard deviation α.

The $\mathsf{LWE}_{\alpha,q,m}$ problem is conjectured to be hard for any $\alpha > 0$ and prime q such that $\alpha \cdot q \geq n$, and for any $m = \mathsf{poly}(\lambda)(n)$. Next we introduce a new problem whose hardness implies the hardness of the LWE problem.

Definition 4 (k-LWE problem, [LPSS14]). *Let $k \leq m, \alpha > 0$, and $\mathbf{S} \in \mathbb{R}^{m\times m}$ be an invertible matrix. We denote $\mathbb{T}^{m+1} = (\mathbb{R}/\mathbb{Z})^{m+1}$. The $(k, \mathbf{S}) - \mathsf{LWE}$ problem is: given $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{m\times n})$, $\mathbf{u} \leftarrow U(\mathbb{Z}_q^n)$ and $\mathbf{x}_i \leftarrow D_{\Lambda_{-\mathbf{u}}^\perp(\mathbf{A}),\mathbf{S}}$ for $i \leq k \leq m$, the goal is to distinguish between the distributions (over $\mathbb{T}^{m+1}$)*

$$\frac{1}{q} \cdot U\Big(\mathrm{Im}\Big(\frac{\mathbf{u}^t}{\mathbf{A}}\Big)\Big) + \nu_\alpha^{m+1} \quad \textit{and} \quad \frac{1}{q} \cdot U\Big(\mathrm{Span}_{i\leq k}(1\|\mathbf{x}_i)^\perp\Big) + \nu_\alpha^{m+1},$$

where ν_α denotes the one-dimensional Gaussian distribution with standard deviation α.

In [LPSS14], it was shown that this problem can be polynomially reduced to LWE problem for a specific class of diagonal matrices $\mathbf{S}$. In our work, we only need an arbitrary $\mathbf{S}$ such that $(k, \mathbf{S})$-LWE is hard, and thus the use of $\mathbf{S}$ is implicit. For simplicity in the exposition, we use k-LWE instead of $(k, \mathbf{S})$-LWE.

We restate an important result that is frequently used in our proofs.

Lemma 3 (Theorem 25, [LPSS14]). *Under the k-LWE assumption, for $k > t$, given t small vectors $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_t$, which are defined as in Definition 4, for any $j \notin [t]$, the distributions*

$$U\left(\mathrm{Span}_{i\in[t]}(\mathbf{x}_i^+)^\perp\right) + \lfloor \nu_{\alpha q} \rceil^{m+1}, U\left(\mathrm{Span}_{i\in[t]\cup\{j\}}(\mathbf{x}_i^+)^\perp\right) + \lfloor \nu_{\alpha q} \rceil^{m+1},$$

are indistinguishable.

Remark 1. Lemma 3 states that we can add a constraint to the distribution, with $\mathbf{x}_j^+$, without the adversary noticing, as long as the adversary does not possess the key $\mathbf{x}_j^+$. Another interpretation of this lemma can be expressed as follows (by removing a constraint, if the adversary does not have the corresponding key):

- If $\mathbf{y} \in \mathbb{Z}_q^{m+1}$ is randomly sampled in $\mathbf{y} \in \mathbb{Z}_q^{m+1}$ and satisfies $t + 1$ (for $t < k$) constraints $\langle \mathbf{y}, \mathbf{x}_i^+ \rangle = 0, i \in [t + 1]$, then under the $k - \mathsf{LWE}$ assumption, we can remove one constraint $\langle \mathbf{y}, \mathbf{x}_j^+ \rangle = 0$ if the adversary

does not have the corresponding key $\mathbf{x}_j^+$ and $\mathbf{y}+\mathbf{e}$, for $\mathbf{e}$ is sampled in $\lfloor \nu_{\alpha q} \rceil^{m+1}$, remains in the same distributions under the view of the adversary. The proof of this lemma remains valid if we replace the constant 0 in the constraints with any constant c_i, resulting in constraints of the form $\langle \mathbf{y}, \mathbf{x}_i^+ \rangle = c_i$. We will use this form in our proof.

- In the particular case where the adversary does not have any key $\mathbf{x}_i^+$, we can remove all the constraints, and the samples can be taken uniformly and independently from $U(\mathbb{Z}_q^{m+1})$.

3 Anonymous Multi-Channel Broadcast Encryption

We start with the formal definition of a broadcast system, and then we proceed to formalize its security properties. We will consider two notions of security: IND−CPA *security*, that protects the privacy of the encrypted message, and *anonymity*, that protects the privacy of the channels. More specifically, a broadcast system is called *anonymous* (AnoBE for short) if it allows addressing a message to a subset of the users, without revealing this privileged set even to users who successfully decrypt the message. This type of anonymity is sometimes called *insider anonymity*.

3.1 Broadcast systems

A *broadcast system* allows addressing a message to a subset of users. In this paper, we will look at *multi-channel broadcast encryption* in which the same ciphertext can be used to encrypt different messages to different disjoint sets of users, called *channels*. We consider the setting when the number N of users in the system is known during setup and the users themselves are identified with the integers in $[N]$. Thus a *channel* $\mathcal{S}$ is simply a subset of $[N]$. We follow the definition in [LPQ12].

Definition 5 (MCBE). *A* multi-channel broadcast encryption *(*MCBE*, in short) is a quadruple of algorithms* $\mathsf{MC} = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ *such that*

$\mathsf{Setup}(1^\lambda, 1^N)$**:** *Takes as input the security parameter λ and the maximum number of users N, and it outputs a* master secret key msk.

$\mathsf{KeyGen}(\mathsf{msk}, i)$**:** *Takes as input the master secret key* msk*, and a user index $i \in [N]$, and it outputs the* decryption key dk_i *for user with index i.*
We denote by dkset *the set containing all the user's decryption keys, dk_i, for $i = 1, \ldots, N$.*

$\mathsf{Encrypt}(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \mathcal{S}_1), \ldots, (\beta_\ell, \mathcal{S}_\ell)))$**:** *Takes as input the master secret key* msk*, the set of all decryption keys* dkset*, and the* broadcast message structure *$((\beta_1, \mathcal{S}_1), \ldots, (\beta_\ell, \mathcal{S}_\ell))$, consisting of ℓ pairs of message and channel, and it outputs the ciphertext* ct.

$\mathsf{Decrypt}(\mathsf{dk}, \mathtt{ct})$**:** *Takes as input a decryption key* dk *and a ciphertext* ct*, and it outputs either a message β or the symbol $\perp$.*

and that satisfies the following correctness *requirement*

Correctness: *For all $N = \mathsf{poly}(\lambda)$, for all messages $\beta_1, \ldots, \beta_\ell$, for all collections $\mathcal{S}_1, \ldots, \mathcal{S}_\ell$ of disjoint channels, if* $\mathsf{msk} \leftarrow \mathsf{Setup}(1^\lambda, 1^N)$*,* $\mathsf{dkset} = \{\mathsf{dk}_i \leftarrow \mathsf{KeyGen}(\mathsf{msk}, i)\}_{i \in [N]}$*, and* $\mathtt{ct} \leftarrow \mathsf{Encrypt}(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \mathcal{S}_1), \ldots, (\beta_\ell, \mathcal{S}_\ell)))$*, then we have that for all channels $\mathcal{S}_j$, all users $i \in \mathcal{S}_j$ get the same message $\beta_j = \mathsf{Decrypt}(\mathsf{dk}_i, \mathtt{ct})$, except with probability negligible in λ.*

We remark that the Encrypt algorithm takes as input the set of decryption keys dkset and thus the scheme is private-key. This fits our setting, since we do not want the adversary **D** to be aware of the public key for the MCBE for otherwise, by the key rule, it can request the corresponding secret key. More precisely, let us go back to our example of A wishing to communicate with $C_1, \ldots, C_t$. In this setting, A establishes a hidden channel with the C_i's by means of an MCBE and in doing so A generates and shares the key dk_i with each member of the channel. Obviously, being an MCBE *multi-channel* scheme, A can establish ℓ channels with distinct groups of users and send different messages to different groups with only one ciphertext.

3.2 IND−CPA security

The notion of IND−CPA security of a MCBE scheme MC is defined based on the following game $\mathsf{CPAGame}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N)$ between an adversary $\mathbf{D}$ and a challenger $\mathcal{C}$

- The challenger $\mathcal{C}$ runs $\mathsf{msk} \leftarrow \mathsf{Setup}(1^\lambda, 1^N)$, initializes the list of indices of corrupted users $\mathsf{L} = \emptyset$, and constructs the set $\mathsf{dkset} = \{\mathsf{dk}_i, i \in [N]\}$ of all decryption keys by setting $\mathsf{dk}_i \leftarrow \mathsf{KeyGen}(\mathsf{msk}, i)$, for $i \in [N]$.
- The adversary $\mathbf{D}$ adaptively issues queries for user decryption keys. For each query $i \in [N]$ issued by $\mathbf{D}$, the challenger $\mathcal{C}$ returns dk_i computed during the first step and adds i to L.
- Challenge query:
  - The adversary $\mathbf{D}$ outputs $\ell + 1$ messages $\beta_1, \ldots, \beta_{j,0}, \beta_{j,1}, \ldots, \beta_\ell$, ℓ disjoint channels $\mathcal{S}^\star_1, \ldots, \mathcal{S}^\star_\ell$, and an index j, which specifies the *target* channel $\mathcal{S}^\star_j$.
  - The challenger $\mathcal{C}$ picks a random $b \xleftarrow{\$} \{0,1\}$ and runs
$$\mathsf{Encrypt}\left(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \mathcal{S}^\star_1), \ldots, (\beta_{j,b}, \mathcal{S}^\star_j), \ldots, (\beta_\ell, \mathcal{S}^\star_\ell))\right)$$
    and gets $\mathtt{ct}^\star$.
  - Finally, $\mathcal{C}$ sends $\mathtt{ct}^\star$ to $\mathbf{D}$.
- The adversary $\mathbf{D}$ continues by asking for adaptively chosen user decryption keys. Queries are handled by $\mathcal{C}$ as in the previous step.
- Finally, $\mathbf{D}$ returns its guess $b' \in \{0,1\}$ for the b chosen by the challenger.

This concludes the description of $\mathsf{CPAGame}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N)$. We say that $\mathbf{D}$ wins game $\mathsf{CPAGame}^{\mathbf{D}}_{\mathsf{MC}}$ if $b' = b$ and $\mathcal{S}^\star_j \cap \mathsf{L} = \emptyset$. We let $\mathsf{Succ}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N)$ denote its probability of winning.

Definition 6. *We say that an* MCBE MC *is* IND−CPA secure *if, for all polynomial-time adversaries* $\mathbf{D}$ *and for all* $N \leq \mathsf{poly}(\lambda)$*, we have that*
$$\left|\mathsf{Succ}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N) - \frac{1}{2}\right| \leq \mathsf{negl}(\lambda).$$

3.3 Anonymity

Next, we define the anonymity property of an MCBE MC through the following game $\mathsf{AnonGame}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N)$ between an adversary $\mathbf{D}$ and a challenger $\mathcal{C}$.

- The challenger $\mathcal{C}$ runs $\mathsf{msk} \leftarrow \mathsf{Setup}(1^\lambda, 1^N)$, initializes the list of indices of corrupted users $\mathsf{L} = \emptyset$, and constructs the set $\mathsf{dkset} = \{\mathsf{dk}_i, i \in [N]\}$ of all decryption keys by setting $\mathsf{dk}_i \leftarrow \mathsf{KeyGen}(\mathsf{ek}, \mathsf{msk}, i)$, for $i \in [N]$.
- The adversary $\mathbf{D}$ adaptively issues queries for user decryption keys. For each query $i \in [N]$ issued by $\mathbf{D}$, the challenger $\mathcal{C}$ returns dk_i computed during the first step and adds i to L.
- Challenge query:
  - The adversary $\mathbf{D}$ outputs an integer j, ℓ messages $\beta_1, \ldots, \beta_\ell$ and $\ell+1$ channels $\mathcal{S}^\star_1, \ldots, \mathcal{S}^\star_{j,0}, \mathcal{S}^\star_{j,1}, \ldots, \mathcal{S}^\star_\ell$.
  - The challenger $\mathcal{C}$ picks a random $b \xleftarrow{\$} \{0,1\}$ and sets
$$\mathtt{ct}^\star = \mathsf{Encrypt}\left(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \mathcal{S}^\star_1), \ldots, (\beta_j, \mathcal{S}^\star_{j,b}), \ldots, (\beta_\ell, \mathcal{S}^\star_\ell))\right),$$
    and then returns $\mathtt{ct}^\star$ to $\mathbf{D}$.
- The adversary $\mathbf{D}$ continues by asking for adaptively chosen user decryption keys. Queries are handled by $\mathcal{C}$ as in the previous step.
- $\mathbf{D}$ returns its guess $b' \in \{0,1\}$ for the b chosen by the challenger.

This concludes the description of $\mathsf{CPAGame}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N)$. To define the winning condition for $\mathsf{AnonGame}$, we set $\mathcal{S}^\star_j := \mathcal{S}^\star_{j,0} \bigtriangleup \mathcal{S}^\star_{j,1} = (\mathcal{S}^\star_{j,0} \setminus \mathcal{S}^\star_{j,1}) \cup (\mathcal{S}^\star_{j,1} \setminus \mathcal{S}^\star_{j,0})$. We say the adversary wins the game if $b' = b$ and $\mathcal{S}^\star_j \cap \mathsf{L} = \emptyset$, and we let $\mathsf{anonSucc}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N)$ denote its success probability.

Definition 7. *We say that* MCBE MC *is* anonymous *if for all* $N = \mathsf{poly}(\lambda)$ *and for all polynomial-time adaptive adversaries* $\mathbf{D}$*, we have that*

$$\left| \mathsf{anonSucc}^{\mathbf{D}}_{\mathsf{MC}}(\lambda, N) - \frac{1}{2} \right| \leq \mathsf{negl}(\lambda).$$

3.4 Construction

In this section we describe AnonMCBE, an Anonymous MCBE whose security relies on the k-LWE problem. Let N be the maximum number of users (receivers are implicitly represented by the integers in $[N]$). Given a security parameter λ, we assert that parameters $q, m, \alpha, \mathbf{S}$ are chosen so that the k-LWE problem is hard to solve as presented in [LPSS14]. Since the adversary can corrupt any user, we require that $N \leq k$ (the system's bounded universe constraint).

Next, we formally describe the 4 algorithms Setup, KeyGen, Encrypt, and Decrypt that make up AnonMCBE.

$\mathsf{Setup}(1^\lambda, 1^N)$**:** Takes as input the security parameter λ and the maximum number of users N. It sets $n = \lambda$ and generates 2 matrices $(\mathbf{A}, \mathbf{T}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}^{m \times m}$. Moreover, it picks $\mathbf{u}$ uniformly in $\mathbb{Z}_q^n$ and sets $\mathbf{A}^+ = (\mathbf{u}^t \| \mathbf{A})$. Finally, it outputs the master secret key $\mathsf{msk} = (\mathbf{T}, \mathbf{A}^+)$.

$\mathsf{KeyGen}(\mathsf{msk}, j)$**:** Takes as input the master secret key msk and a user index $j \in [N]$. It calls the GPV algorithm (see Lemma 2) using the basis $\Lambda^\perp(\mathbf{A})$ consisting of the rows of $\mathbf{T}$ and the standard deviation matrix $\mathbf{S}$. It obtains a sample $\mathbf{x}_j$ from $D_{\Lambda^\perp_{-\mathbf{u}}(\mathbf{A}), \mathbf{S}}$. Finally, the algorithm outputs decryption key $\mathsf{dk}_j = \mathbf{x}_j^+ := (1 \| \mathbf{x}_j) \in \mathbb{Z}^{m+1}$ for user j.

We denote by dkset the set containing all the user's decryption key dk_i.

$\mathsf{Encrypt}(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \mathcal{S}_1), \ldots, (\beta_\ell, \mathcal{S}_\ell)))$**:** Takes as input the master secret key msk, the set $\mathsf{dkset}\{\mathsf{dk}_i | i \in [N]\}$, and broadcast message structure $((\beta_j, \mathcal{S}_j))_{j \in [\ell]}$, where $\beta_j \in \{0, 1\}$ and $\mathcal{S}_j \subset [N]$, for $j \in [\ell]$. It first sets $\mathbf{r}_j = (\beta_j, 0, \ldots, 0)$, then it randomly samples $\mathbf{y} = (y_1, \ldots, y_{m+1})$ in

$$U(\mathrm{Span}_{j=1,\ldots,\ell, i \in \mathcal{S}_j} (\mathbf{x}_i^+ - \mathbf{r}_j)^\perp).$$

Finally, Encrypt outputs the ciphertext $\mathsf{ct} = \mathbf{y} + \mathbf{e}$, where $\mathbf{e} \leftarrow \lfloor \nu_{\alpha q} \rceil^{m+1}$.

$\mathsf{Decrypt}(\mathsf{dk}_i, \mathsf{ct})$**:** Takes as input the decryption key $\mathsf{dk}_i = \mathbf{x}_i^+$ of user i and a ciphertext ct. It computes $z = \langle \mathsf{dk}_i, \mathsf{ct} \rangle$ and outputs 0 if z is closer to 0 than to $y_1 \mod q$; otherwise, it outputs 1.

Correctness of AnonMCBE**.** We next show that, for any subset $\mathcal{S}_j \subseteq [N]$ and for all users $i \in \mathcal{S}_j$, if $\mathsf{ct} \leftarrow \mathsf{Encrypt}(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \mathcal{S}_1), \ldots, (\beta_\ell, \mathcal{S}_\ell)))$ and dk_i is the decryption key for user i appearing in dkset, then $\mathsf{Decrypt}(\mathsf{dk}_i, \mathsf{ct})$ returns the bit message β_j. Indeed, since $\mathbf{y} \in \mathrm{Span}_{j=1,\ldots,\ell, i \in \mathcal{S}_j} (\mathbf{x}_i^+ - \mathbf{r}_j)^\perp$, for each user $i \in \mathcal{S}_j$, we have $\langle \mathbf{x}_i^+, \mathbf{y} \rangle = \langle \mathbf{r}_j, \mathbf{y} \rangle$. Therefore,

$$\begin{aligned} z = \langle \mathsf{dk}_i, \mathsf{ct} \rangle &= \langle \mathbf{x}_i^+, \mathbf{y} + \mathbf{e} \rangle = \langle \mathbf{x}_i^+, \mathbf{y} \rangle + \langle \mathbf{x}_i^+, \mathbf{e} \rangle \\ &= \langle \mathbf{x}_i^+, \mathbf{e} \rangle + \langle \mathbf{r}_j, \mathbf{y} \rangle \\ &= \langle \mathbf{x}_i^+, \mathbf{e} \rangle + y_1 \beta_j. \end{aligned}$$

We note that $\mathbf{e} \leftarrow \lfloor \nu_{\alpha q} \rceil^{m+1}$ and therefore, according to [LPSS14], the quantity $\langle \mathbf{x}_i^+, \mathbf{e} \rangle$ is small modulo q compared to y_1 with high probability. Therefore, every user $i \in \mathcal{S}_j$ can decrypt successfully with high probability and gets the same plaintext bit β_j.

Remark 2. We note that, similar to the anonymous broadcast scheme in [DPY20], we provide a description for a scheme with a bounded universe of N users and allow full collusion: the adversary can corrupt all

N users. Another equivalent view of these schemes is that the number of users N is not bounded (as the key generation KeyGen can generate as many keys as we want) but the adversary can only corrupt up to k users (for which the $k-\mathsf{LWE}$ problem is hard). We use the term *bounded model* to refer to both views of the scheme: bounded universe with full collusion or unbounded universe with bounded-size collusion. In the context of anamorphic broadcast mode, we use the bounded universe model as described above.

3.5 Security Proof

In this section we prove the security properties of our construction $\mathsf{AnonMCBE}$. Specifically, Theorem 1 will establish the $\mathsf{IND-CPA}$ security and Theorem 2 will establish anonymity.

Theorem 1. *Under the assumption that the* $k-\mathsf{LWE}$ *problem is hard, the* $\mathsf{AnonMCBE}$ *scheme constructed above is* $\mathsf{IND-CPA}$ *secure for any number of users* $N \leq k$.

Proof. We proceed using the hybrid games between a challenger $\mathcal{C}$ and an attacker $\mathbf{D}$.
Game G_0**:** This is the $\mathsf{IND-CPA}$ game and the interaction between the challenger $\mathcal{C}$ and the adversary $\mathbf{D}$ takes place as follows.
Setup. $\mathcal{C}$ generates matrix $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{m\times n})$ along with a trapdoor $\mathbf{T} \in \mathbb{Z}^{m\times m}$. It also samples $\mathbf{u} \leftarrow U(\mathbb{Z}_q^n)$ and sets $\mathbf{A}^+ = (\mathbf{u}^t \| \mathbf{A})$ and sets $\mathsf{msk} = (\mathbf{T}, \mathbf{A}^+)$. $\mathcal{C}$ constructs the set dkset consisting of dk_i computed as follows. $\mathcal{C}$ samples, $\mathbf{x}_i \leftarrow D_{\Lambda^{\perp}_{-\mathbf{u}}(\mathbf{A}),\mathbf{S}}$ and sets $\mathsf{dk}_i := (1\|\mathbf{x}_i) \in \mathbb{Z}^{m+1}$. Finally, $\mathcal{C}$ initializes the list of indices of corrupted users $\mathsf{L} = \emptyset$,
Query Phase 1. When $\mathbf{D}$ queries decryption keys for users $i \in [N]$, $\mathcal{C}$ returns dk_i^+ computed by Setup and adds i to L.
Challenge phase. The adversary $\mathbf{D}$ outputs ℓ disjoint sets $\mathcal{S}_1^\star, \ldots, \mathcal{S}_\ell^\star \subseteq [N]$, an index j, which specifies the attacked target set $\mathcal{S}_j^\star$, and two sequences of 1-bit messages $(\beta_1, \ldots, \beta_j = 0, \ldots, \beta_\ell)$ and $(\beta_1, \ldots, \beta_j = 1, \ldots, \beta_\ell)$ that differ in the j-th position. The challenger computes ciphertext $\mathtt{ct}^\star$ as follows:

- Randomly sample $\beta^\star \xleftarrow{\$} \{0,1\}$ and set

$$\mathbf{r}_1 = \begin{pmatrix} \beta_1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \ldots, \mathbf{r}_j = \begin{pmatrix} \beta^\star \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \ldots, \mathbf{r}_\ell = \begin{pmatrix} \beta_\ell \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$

- Randomly sample $\mathbf{y} = (y_1, \ldots, y_{m+1})$ in

$$U(\mathrm{Span}_{t=1,\ldots,\ell, i\in\mathcal{S}_t^\star}(\mathbf{x}_i^+ - \mathbf{r}_t)^\perp).$$

- Sample $\mathbf{e} \leftarrow \lfloor \nu_{\alpha q} \rceil^{m+1}$ and set $\mathtt{ct}^\star = \mathbf{y} + \mathbf{e}$.

Finally, the challenger sends $\mathtt{ct}^\star$ to the adversary $\mathbf{D}$.
Query Phase 2. $\mathbf{D}$ continues to query for decryption keys for users $i \in [N]$, $\mathcal{C}$ returns dk_i^+ computed by Setup and adds i to L. Recall that the list L in the $\mathsf{IND-CPA}$ game verifies $\mathcal{S}_j^\star \cap \mathsf{L} = \emptyset$, otherwise $\mathbf{D}$ loses the game.
Guess. $\mathbf{D}$ gives a guess β for $\beta^\star$.

This concludes the description of Game G_0. In each of the $c := |\mathcal{S}_j^\star|$ subsequent games, $\mathbf{D}$ will receive ciphertexts in which the challenger $\mathcal{C}$ has progressively revoked the decryption capability of the users in $\mathcal{S}_j^\star$ until, in the last game, the j-th channel contains no users. Specifically, let us fix an arbitrary ordering $i_1, i_2, \ldots, i_c$ of the c users in $\mathcal{S}_j^\star$ and define $\mathcal{S}_{j,h}^\star = \mathcal{S}_j^\star \setminus \{i_1, \ldots, i_h\}$. In Game G_h the ciphertext is computed with respect to the channels $\mathcal{S}_1^\star, \ldots, \mathcal{S}_{j,h}^\star, \ldots, \mathcal{S}_\ell^\star$. This is reflected in the computation of $\mathtt{ct}^\star$ that will consider set $\mathcal{S}_{j,h}^\star$ instead of $\mathcal{S}_j^\star$ in the sampling of $\mathbf{y}$ from

$$U(\mathrm{Span}_{t=1,\ldots,\ell, i\in\mathcal{S}_t^\star}(\mathbf{x}_i^+ - \mathbf{r}_t)^\perp).$$

As already observed, since $\mathcal{S}^\star_{j,c} = \emptyset$, in Game G_c, $\mathbf{D}$'s view is independent from $\beta^\star$. Let us now argue that consecutive games are indistinguishable. Note that $\mathbf{y} = (y_1, \ldots, y_{m+1})$ is chosen randomly from $U(\mathrm{Span}_{t\in[\ell], i\in\mathcal{S}^\star_t}(\mathbf{x}^+_i - \mathbf{r}_t)^\perp)$, which signifies that we randomly choose $\mathbf{y} = (y_1, \ldots, y_{m+1})$ under fewer than k constraints (since the total number of $t \in [\ell], i \in \mathcal{S}^\star_t$ is at most $N < k$): $\langle \mathbf{y}, \mathbf{x}^+_i \rangle = y_1\beta_j, \forall i \in \mathcal{S}^\star_j$. By Lemma 3 and Remark 1, we can remove one constraint as the adversary has no key in $\mathcal{S}^\star_j$, which implies that the distribution $U(\mathrm{Span}_{t\in[\ell], i\in\mathcal{S}^\star_t}(\mathbf{x}^+_i - \mathbf{r}_t)^\perp) + \lfloor \nu_{\alpha q} \rceil^{m+1}$ and the equivalent in which in one of the sets one user is removed are computational indistinguishable.

To summarize, we have a sequence of games where, in the final game Game G_c, the adversary has zero advantage, and the difference in the adversary's advantage between each two successive games Game G_{h-1} and Game G_h, is negligible. Since c is polynomial, the proposed scheme is $\mathsf{IND-CPA}$ secure.

□

Anonymity. We now state the theorem that shows that the proposed $\mathsf{AnonMCBE}$ scheme also enjoys anonymity. This is crucial for achieving the anamorphism developed in the next section.

Theorem 2. *Under the assumption that the* $k - \mathsf{LWE}$ *problem is hard,* $\mathsf{AnonMCBE}$ *is an anonymous* MCBE*, for any* $N \leq k$*.*

Proof. We only give an overview of the proof strategy as the proof is very similar to the one of Theorem 1.

We start from two games G^0_0 and G^1_0 between challenger $\mathcal{C}$ and adversary $\mathbf{D}$ consisting of the anonymity game and that differ only in the way the challenge ciphertext is computed. Specifically, if $\mathbf{D}$ outputs index j, ℓ bits $\beta_1, \ldots, \beta_\ell$ and $\ell+1$ channels $\mathcal{S}^\star_1, \ldots, \mathcal{S}^\star_{j,0}, \mathcal{S}^\star_{j,1}, \ldots, \mathcal{S}^\star_\ell$, the challenge ciphertext $\mathtt{ct}^\star$ of G^0_0 is computed with respect to the broadcast message structure $((\beta_1, \mathcal{S}^\star_1), \ldots, (\beta_j, \mathcal{S}^\star_{j,0}), \ldots, (\beta_\ell, \mathcal{S}^\star_\ell))$, whereas in G^1_0 it is computed with respect to $((\beta_1, \mathcal{S}^\star_1), \ldots, (\beta_j, \mathcal{S}^\star_{j,1}), \ldots, (\beta_\ell, \mathcal{S}^\star_\ell))$. We wish to prove that G^0_0 and G^1_0 are indistinguishable.

We fix an ordering $i_1, \ldots, i_c$ of the users in the symmetric difference of $\mathcal{S}^\star_{j,0}$ and $\mathcal{S}^\star_{j,1}$. For $h = 1, \ldots, c$, we define games G^0_h to use as j-th channel the set $\mathcal{S}^\star_{j,0} \cup \{i_1, \ldots, i_h\}$ and games G^1_h to use as j-th channel the set $\mathcal{S}^\star_{j,1} \cup \{i_1, \ldots, i_h\}$. Note that last games G^0_c and G^1_c both use the same channel $\mathcal{S}^\star_{j,0} \cup \mathcal{S}^\star_{j,1}$ and are thus indistinguishable. The proof is then completed by showing that pairs of consecutive games G^0_{h-1} and G^0_h, and G^1_{h-1} and G^1_h are indistinguishable for $h = 1, \ldots, c$.

Consider games G^0_{h-1} and G^0_h (a similar argument holds for games G^1_{h-1} and G^1_h) and distinguish two cases. If $i_h \in \mathcal{S}^\star_{j,0}$ then the two games are trivially identical. Suppose now that $i_h \notin \mathcal{S}^\star_{j,0}$. Then it is sufficient to prove that the following two distributions are indistinguishable:

$$U(\mathrm{Span}_{\substack{t\in([\ell]\setminus\{j\}), i\in\mathcal{S}^\star_t \\ t=j,\ i\in S_{h-1}}}(\mathbf{x}^+_i - \mathbf{r}_t)^\perp) + \lfloor \nu_{\alpha q} \rceil^{m+1},$$

and

$$U(\mathrm{Span}_{\substack{t\in([\ell]\setminus\{j\}), i\in\mathcal{S}^\star_t \\ t=j,\ i\in S_{h-1}\cup\{i_h\}}}(\mathbf{x}^+_i - \mathbf{r}_t)^\perp) + \lfloor \nu_{\alpha q} \rceil^{m+1},$$

where S_{h-1} denotes the set $\mathcal{S}^\star_{j,0} \cup \{i_1, \ldots, i_{h-1}\}$. This is proved by using an argument similar to the one used to prove $\mathsf{IND-CPA}$ security (see Theorem 1) by applying Lemma 3 and Remark 1.

We can thus conclude that $\mathsf{AnonMCBE}$ is an anonymous MCBE scheme. □

Indistinguishability of $\mathsf{LWE-PKE}$ and $\mathsf{AnonMCBE}$ ciphertexts. The following corollary is important for enabling broadcast anamorphic mode. It establishes that an adversary will not be able to detect that an anamorphic broadcast mode is employed, as the ciphertexts look exactly the same as in the regular $\mathsf{LWE-PKE}$. We note that the dictator is not aware of the $\mathsf{AnonMCBE}$, and thus has no key in $\mathsf{AnonMCBE}$, but in the corollary we allow the adversary to see the matrix $\mathbf{A}$ and to choose the channels/broadcast sets, as this will be used in our later proofs with hybrid arguments.

Corollary 1. *Under the assumption that the* $k - \mathsf{LWE}$ *problem is hard, for any* $N \leq k$*, given the public parameters, public key and secret key* $(\mathbf{A}_{\mathsf{PKE}} \in \mathbb{Z}_q^{m\times n}, \mathbf{u}, \mathsf{sk})$ *of the* $\mathsf{LWE-PKE}$ *and the public parameter*

$\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ *of the* AnonMCBE*, no PPT adversary, with no any user's key in* AnonMCBE *can distinguish a ciphertext sampled from* AnonMCBE *for* N *users, for adversarial choice of broadcast message structure, and a ciphertext of a random bit message in* $\mathsf{LWE} - \mathsf{PKE}$.

Proof. The proof follows easily from anonymity. First, because the adversary has no any keys in the AnonMCBE system, the encryption oracle of AnonMCBE can be simulated by randomly sample in $U(\mathbb{Z}_q^{m+1})$, following Lemma 3 and Remark 1. Now, given adversarially chosen broadcast message structure $((\beta_1, \mathcal{S}_1), \ldots, (\beta_\ell, \mathcal{S}_\ell))$, from the anonymity established in Theorem 2, we can replace $\mathcal{S}_1$ with $\emptyset$, and the adversary is not able to distinguish between

$$\mathsf{Encrypt}\left(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \mathcal{S}_1), \ldots, (\beta_\ell, \mathcal{S}_\ell))\right),$$

$$\mathsf{Encrypt}\left(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \emptyset), \ldots, (\beta_\ell, \mathcal{S}_\ell))\right),$$

because the adversary has no secret key $\mathcal{S}_1$.

We continue in the same way to replace $\mathcal{S}_2$ with $\emptyset$, and so on, until $\mathcal{S}_\ell$ is replaced by $\emptyset$. In the end, the adversary will receive

$$\mathsf{Encrypt}\left(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \emptyset), \ldots, (\beta_\ell, \emptyset))\right),$$

without noticing the change.

Finally, by definition,

$$\mathsf{Encrypt}\left(\mathsf{msk}, \mathsf{dkset}, ((\beta_1, \emptyset), \ldots, (\beta_\ell, \emptyset))\right),$$

is $\mathbf{y} + \mathbf{e}$, for a random $\mathbf{y} \in \mathbb{Z}_q^{m+1}$, which is thus indistinguishable from a ciphertext of a random bit message in $\mathsf{LWE} - \mathsf{PKE}$, by the security of $\mathsf{LWE} - \mathsf{PKE}$.

□

4 Public Key Encryption with an Anamorphic Broadcast Mode

In this section we, finally, present and prove the properties of our PKE system with anamorphic broadcast properties. We start by defining the anamorphic property we want to prove.

4.1 Formal Definition

Let $\mathcal{E}$ be public key encryption scheme and MCBE a multi-channel broadcast encryption. We wish to capture the fact that it is possible to produce a ciphertext for multi-channel broadcast encryption that encrypts a given broadcast message structure consisting of ℓ channels and messages $((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell))$ into a ciphertext that is also a valid ciphertext for encryption scheme $\mathcal{E}$; here validity is to be taken in the sense that the decryption with $\mathcal{E}$'s secret key will be successful. We require that such a ciphertext be indistinguishable from a ciphertext computed by encrypting a random message with respect to $\mathcal{E}$.

We present our definition for one-bit encryption schemes.

Definition 8. *Let* $\mathcal{E} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ *be a one-bit public key encryption scheme. An* anamorphic broadcast encryption *mode for* $\mathcal{E}$ *is a tuple*

$$\mathsf{aMCBE} = (\mathsf{aSetup}, \mathsf{aKG}, \mathsf{aEnc}, \mathsf{aDec})$$

of algorithms with the following syntax:

1. aSetup*, on input security parameter* 1^λ*, maximum number of users* 1^N*, and* forced *public key* fpk *for* $\mathcal{E}$*, outputs public key* ek *and master secret key* msk.
2. aKG*, on input* ek *and* msk*, outputs a set* dkset *which contains* N *anamorphic secret keys* $\mathtt{ask}_1, \ldots, \mathtt{ask}_N$.

3. aEnc, *on input the forced public key* fpk, *master secret key* msk, *a set* dkset *of anamorphic secret keys, and* broadcast message structure *consisting of* ℓ *pairs of anamorphic messages and disjoint channels*

$$((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)),$$

produces an anamorphic ciphertext $\mathtt{act}$.
4. aDec *is an algorithm that on input an anamorphic secret key* $\mathtt{ask}$ *and an anamorphic ciphertext* $\mathtt{act}$, *outputs a message* $\mathtt{am}$ *or the symbol* $\perp$.

and that satisfies the following correctness *requirement*

- *For* $(\mathsf{fpk}, \mathsf{fsk}) \leftarrow \mathsf{KG}(1^\lambda)$, *for all* $N = \mathsf{poly}(\lambda)(\lambda)$, $(\mathsf{ek}, \mathsf{msk}) \leftarrow \mathsf{aSetup}(1^\lambda, 1^N, \mathsf{fpk})$, $\mathsf{dkset} \leftarrow \mathsf{aKG}(\mathsf{ek}, \mathsf{msk})$, *and for all* ℓ *pairs of anamorphic messages and disjoint channels* $((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell))$, *and for* i *and* j *such that* $i \in \mathtt{aset}_j$, *if*

$$\mathtt{act} \leftarrow \mathsf{aEnc}(\mathsf{fpk}, \mathsf{ek}, \mathsf{msk}, \mathsf{dkset}, (\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)),$$

then we have $\mathtt{am}_j = \mathsf{aDec}(\mathtt{ask}_i, \mathtt{act})$, *except with probability negligible in* λ.

We next formalize the anamorphic properties of $\mathcal{E}$ and aMCBE by requiring experiments $\mathsf{RegularG}$ and AnamG, formally described below, to be indistinguishable. Roughly speaking, in the regular game $\mathsf{RegularG}$ the adversary is given access to an oracle that, on input a broadcast message structure, returns the encryption of a random bit with respect to $\mathcal{E}$. In the anamorphic game AnamG, instead, the adversary obtains a ciphertext computed by encrypting the broadcast message structure. In both games the adversary is given in input the pair $(\mathsf{fpk}, \mathsf{fsk})$ of forced public and secret keys and therefore, quite trivially, it has access to the production of ciphertexts carrying bits of $\mathbf{D}$'s choice (opposed to just a random bit).

$\mathsf{RegularG}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N)$

1. Generate $(\mathsf{fpk}, \mathsf{fsk}) \leftarrow \mathsf{KG}(1^\lambda)$.
2. Return $\mathbf{D}^{\mathsf{rEO}(\mathsf{fpk},\cdot)}(\mathsf{fpk}, \mathsf{fsk})$, where
   - rEO, on input message structure $((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell))$,
     - randomly selects $\tilde{b} \leftarrow \{0, 1\}$.
     - returns $\mathsf{Enc}(\mathsf{fpk}, \tilde{b})$.

$\mathsf{AnamG}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N)$

1. Generate $(\mathsf{fpk}, \mathsf{fsk}) \leftarrow \mathsf{KG}(1^\lambda)$.
2. Generate $\mathsf{msk} \leftarrow \mathsf{aSetup}(1^\lambda, 1^N, \mathsf{fpk})$.
3. Generate $\mathsf{dkset} \leftarrow \mathsf{aKG}(\mathsf{ek}, \mathsf{msk})$.
4. Return $\mathbf{D}^{\mathsf{amEO}(\mathsf{fpk},\mathsf{msk},\mathsf{dkset},\cdot)}(\mathsf{fpk}, \mathsf{fsk})$, where
   - amEO, on input message structure $((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell))$, returns $\mathsf{aEnc}(\mathsf{fpk}, \mathsf{msk}, \mathsf{dkset}, ((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)))$.

We denote by $p^{\mathsf{RegularG}}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N)$ (respectively, $p^{\mathsf{AnamG}}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N)$) the probability that $\mathbf{D}$ outputs 1 in $\mathsf{RegularG}$ (respectively, in AnamG) and present the following definition.

Definition 9. *A public key encryption scheme* $\mathcal{E}$ *is said to have a secure anamorphic broadcast mode* aMCBE *if*

1. $\mathcal{E}$ *is an* $\mathsf{IND{-}CPA}$ *public-key encryption scheme;*
2. *for all PPT adversaries* $\mathbf{D}$, *and all* $N = \mathsf{poly}(\lambda)(\lambda)$

$$\left| p^{\mathsf{RegularG}}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N) - p^{\mathsf{AnamG}}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N) \right| \leq \mathsf{negl}(\lambda)(\lambda).$$

Discussion. In the security games of Definition 9, the adversary $\mathbf{D}$ receives as a reply to their queries a valid encryptions $\mathtt{ct}$ of a random bits w.r.t to fpk. The ciphertext $\mathtt{ct}$ is computed in one of two ways: in $\mathsf{RegularG}$, the ciphertext is computed by encrypting a random bit for fpk, whereas in AnamG, the ciphertext is computed by encrypting the provided broadcast message structure. The security definition requires the two ways of producing the ciphertext to be indistinguishable. Let us map this to our example with A trying to communicate with C (actually, ℓ anonymous broadcast channels) and disguising the ciphertext as direct to B. The adversary $\mathbf{D}$ has B's secret key and is guaranteed to receive a ciphertext that is indistinguishable from one carrying a random message. This is models the setting in which A, a *client* that has no public key, initiates the interaction with B, a *server* that has a public key, by sending an ephemeral key to be used to encrypt subsequent messages. For example, this is essentially how a client in SSL and TLS initiates a session with a server.

4.2 Construction

In this section, we show how to embed an anamorphic broadcast mode into an LWE PKE, namely the Dual Regev PKE [GPV08], presented in Section 2.3. Our approach is to make use of anonymous multi-channel broadcast encryption, thus allowing a sender to open multiple anamorphic channels. The details of our construction are given below:

Let $\mathcal{E} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ be the Dual Regev PKE scheme and let $\mathsf{MC} = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ be the anonymous multi-channel broadcast encryption of Section 3.3. We will construct anamorphic broadcast mode, $\mathsf{aMCBE} = (\mathsf{aSetup}, \mathsf{aKG}, \mathsf{aEnc}, \mathsf{aDec})$ in which the algorithms are defined as follows:

aSetup**:** Takes as input the security parameter 1^λ, maximum number of users 1^N, and forced public key $\mathsf{fpk} = (\mathbf{A}, \mathbf{u}) \leftarrow \mathsf{KG}(1^\lambda)$, where $\mathbf{A} \in \mathbb{Z}_q^{m \times n}, \mathbf{u} \in \mathbb{Z}_q^n$. It calls $\mathsf{Setup}(1^\lambda, 1^N)$ to generate anamorphic master secret key which consists of an anamorphic public parameter and an anamorphic private parameter, $\mathsf{msk} = (\mathbf{A}_{\mathsf{Broad}}, \mathbf{T}_{\mathsf{Broad}}) \leftarrow \mathsf{Setup}(1^\lambda, 1^N)$, where $\mathbf{A}_{\mathsf{Broad}} \in \mathbb{Z}_q^{m \times n}$ and $\mathbf{T}_{\mathsf{Broad}} \in \mathbb{Z}^{m \times m}$ is a trapdoor matrix for $\mathbf{A}_{\mathsf{Broad}}$.

aKG**:** on input $\mathsf{msk} = (\mathbf{A}_{\mathsf{Broad}}, \mathbf{T}_{\mathsf{Broad}})$, outputs a set dkset which contains N anamorphic secret keys $\mathtt{ask}_1, \ldots, \mathtt{ask}_N$, where

$$\mathsf{ask}_i \leftarrow \mathsf{KeyGen}(\mathbf{A}_{\mathsf{Broad}}, \mathbf{T}_{\mathsf{Broad}}, i).$$

aEnc**:** on input the forced public key fpk, master secret key $\mathsf{msk} = (\mathbf{A}_{\mathsf{Broad}}, \mathbf{T}_{\mathsf{Broad}})$, a set dkset of anamorphic secret keys, and broadcast message structure consisting of ℓ pairs of anamorphic messages and disjoint channels

$$((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)),$$

produces an anamorphic ciphertext $\mathtt{act}$ by calling

$$\mathsf{act} \leftarrow \mathsf{Encrypt}(\mathsf{msk}, \mathsf{dkset}, ((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell))).$$

aDec**:** The algorithm takes as input an anamorphic secret key $\mathtt{ask}$ and an anamorphic ciphertext $\mathtt{act}$, and it calls $\mathsf{Decrypt}(\mathtt{ask}_i, \mathtt{act})$ to output a message $\mathtt{am}_j$ or the invalid symbol $\perp$.

Correctness. The correctness property of aMCBE is directly inferred from the correctness of the MC scheme: for $i \in \mathsf{aset}_j$, $\mathsf{amsg}_j = \mathsf{aDec}(\mathsf{ask}_i, \mathsf{act})$, except with negligible probability in λ.

Theorem 3. *Dual Regev PKE is a public key encryption scheme having an anamorphic broadcast encryption mode.*

Proof. We will show that for any PPT adversary $\mathbf{D}$, the advantage of $\mathbf{D}$ in distinguishing games $\mathsf{RegularG}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N)$ and $\mathsf{AnamG}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N)$ is negligible. We let $Q = \mathsf{poly}(\lambda)(\lambda)$ denote the maximum number of queries issued by $\mathbf{D}$.

We start by setting game G_0 to be the anamorphic game AnamG be the anamorphic game. For $i = 1, \ldots, Q$, we let G_i be the same as game G_{i-1} except for the reply to i-th query. Specifically, challenger $\mathcal{C}$ replaces the

i-th anamorphic ciphertext $\mathtt{act}_i$ with the regular ciphertext $\mathtt{ct}_i$, computed as $\mathtt{ct}_i \leftarrow \mathsf{Enc}(\mathsf{fpk}, b)$, for a random bit b. Thus, we have that game G_Q coincides with $\mathsf{RegularG}$. Relying on Corollary 1, we show that game G_{i-1} is indistinguishable from game G_i, for each $i \in [Q]$. Therefore, $\mathsf{G}_0 \approx \mathsf{G}_Q$, from which the theorem follows.

In more detail, our proof is through the following sequence of games:

Game G_0: this is the anamorphic game. The adversary $\mathbf{D}$ is given:

1. a forced public key fpk and a forced secret key fsk;
2. Q anamorphic ciphertexts from Q chosen queries to amEO. For each $i \in [Q]$, $\mathtt{act}_i$ is computed as

$$\mathtt{act}_i \leftarrow \mathsf{aEnc}(\mathsf{fpk}, \mathsf{msk}, \mathsf{dkset}, ((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)).$$

Game G_1: let Game G_1 be the same Game G_0 except that the first anamorphic ciphertext $\mathtt{act}_1$, which is computed as

$$\mathtt{act}_1 \leftarrow \mathsf{aEnc}(\mathsf{fpk}, \mathsf{msk}, \mathsf{dkset}, ((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)),$$

is replaced by a regular ciphertext $\mathtt{ct}_1$,

$$\mathtt{ct}_1 \leftarrow \mathsf{Enc}(\mathsf{fpk}, b),$$

for a random bit b. The hybrid transition of ciphertexts from Game G_0 to Game G_1 is depicted as below:

$$\begin{pmatrix} \mathsf{act}_1 & \mathsf{act}_2 & \ldots & \mathsf{act}_Q \end{pmatrix} \leftarrow \text{Game } \mathsf{G}_0$$
$$\begin{pmatrix} \mathsf{ct}_1 & \mathsf{act}_2 & \ldots & \mathsf{act}_Q \end{pmatrix} \leftarrow \text{Game } \mathsf{G}_1$$

Based on Corollary 1, we observe that this transition cannot be detected by the adversary. Indeed, we can construct a simulator, acting as an adversary in Corollary 1, to simulate either Game G_0 or Game G_1:

- Generate $(\mathsf{fpk}, \mathsf{fsk}) \leftarrow \mathsf{KG}(1^\lambda)$.
- For each of query $i = 2, \ldots, Q$: on input broadcast message structure

$$((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)),$$

  the anamorphic ciphertexts $\mathsf{act}_2, \ldots, \mathsf{act}_Q$ can be obtained via $Q-1$ queries to encryption oracle $\mathsf{aEnc}(\mathsf{fpk}, \mathsf{msk}, \mathsf{dkset}, ((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)))$.
- Use the challenge of the simulator (act from aEnc or ct from Enc) to answer the first query.

Depending on the challenge (act or ct, respectively) the simulator simulates Game G_0 or Game G_1 respectively for $\mathbf{D}$. Thus, under the view of $\mathbf{D}$, if Game G_0 and Game G_1 are distinguishable (the difference of winning probability between the two games for $\mathbf{D}$ is non-negligible), then the simulator can distinguish act and ct, which contradicts Corollary 1.

Game G_i, for $i = 2, \ldots, Q$: let Game G_i be the same Game G_{i-1} except that the anamorphic ciphertext at the i-th position act_i is computed as

$$\mathtt{act}_i \leftarrow \mathsf{aEnc}(\mathsf{fpk}, \mathsf{msk}, \mathsf{dkset}, ((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell))),$$

is replaced by a regular ciphertext $\mathtt{ct}_i$,

$$\mathtt{ct}_i \leftarrow \mathsf{Enc}(\mathsf{fPK}, b_i),$$

for random message b_i The hybrid transition of ciphertexts from Game G_{i-1} to Game G_i is depicted as below:

$$\begin{pmatrix} \mathsf{ct}_1 & \ldots & \mathsf{ct}_{i-1} & \mathsf{act}_i & \ldots & \mathsf{act}_Q \end{pmatrix} \leftarrow \text{Game } \mathsf{G}_{i-1}$$
$$\begin{pmatrix} \mathsf{ct}_1 & \ldots & \mathsf{ct}_{i-1} & \mathsf{ct}_i & \ldots & \mathsf{act}_Q \end{pmatrix} \leftarrow \text{Game } \mathsf{G}_i$$

Based on Corollary 1, we observe that this transition cannot be detected by the adversary. Indeed, we can construct a simulator, acting as an adversary in Corollary 1, to simulate either Game G_{i-1} or Game G_i:

1. Generate $(\mathsf{fpk}, \mathsf{fsk}) \leftarrow \mathsf{KG}(1^\lambda)$.
2. With the public key fPK, all regular ciphertexts $\mathsf{ct}_1, \ldots, \mathsf{ct}_{i-1}$ can be simulated by running $i-1$ times $\mathsf{Enc}(\mathsf{fPK}, b_i)$ for random message b_i.
3. For each of query $i+1, \ldots, Q$: on input message structure $((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell))$, the anamorphic ciphertexts $\mathsf{act}_{i+1}, \ldots, \mathsf{act}_Q$ can be obtained via $Q-1$ queries to encryption oracle $\mathsf{aEnc}(\mathsf{fpk}, \mathsf{msk}, \mathsf{dkset}, ((\mathtt{am}_1, \mathtt{aset}_1), \ldots, (\mathtt{am}_\ell, \mathtt{aset}_\ell)))$.
4. Use the challenge of the simulator ($\mathtt{act}$ from aEnc or $\mathtt{ct}$ from Enc) to answer the i-th query.

Depending on the challenge ($\mathtt{act}$ or $\mathtt{ct}$, respectively) the simulator simulates Game G_{i-1} or Game G_i respectively for $\mathbf{D}$. Thus, under the view of $\mathbf{D}$, if Game G_{i-1} and Game G_i are distinguishable (the difference of winning probability between the two games for $\mathbf{D}$ is non-negligible), then the simulator can distinguish $\mathtt{act}$ and $\mathtt{ct}$, which contradicts Corollary 1. For $i = Q$ we reach the Game G_Q in which all anamorphic ciphertexts from the previous game are replaced with regular ciphertexts. We can thus conclude that $\mathsf{AnamG}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N) \approx \mathsf{RegularG}_{\mathcal{E},\mathsf{aMCBE},\mathbf{D}}(\lambda, N)$. □

Acknowledgments

This work was supported in part by the France 2030 ANR Project ANR-22-PECY-003 SecureCompute. Xuan Thanh Do was supported by the 03/2025/KCM project (KCM/21-30 research program).

References

AD18. Kamalesh Acharya and Ratna Dutta. Constructions of secure multi-channel broadcast encryption schemes in public key framework. In Jan Camenisch and Panos Papadimitratos, editors, *CANS 18*, volume 11124 of *LNCS*, pages 495–515. Springer, Cham, September / October 2018. 6

AP11. J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. *Theor. Comput. Science*, 48(3):535–553, 2011. 10

BBW06. Adam Barth, Dan Boneh, and Brent Waters. Privacy in encrypted content distribution using private broadcast encryption. In Giovanni Di Crescenzo and Avi Rubin, editors, *FC 2006*, volume 4107 of *LNCS*, pages 52–64. Springer, Berlin, Heidelberg, February / March 2006. 5

BGH+24. Fabio Banfi, Konstantin Gegier, Martin Hirt, Ueli Maurer, and Guilherme Rito. Anamorphic encryption, revisited. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 3–32. Springer, Cham, May 2024. 4, 7

BMN+21. Olivier Blazy, Sayantan Mukherjee, Huyen Nguyen, Duong Hieu Phan, and Damien Stehlé. An anonymous trace-and-revoke broadcast encryption scheme. In Joonsang Baek and Sushmita Ruj, editors, *ACISP 21*, volume 13083 of *LNCS*, pages 214–233. Springer, Cham, December 2021. 5

CGM24a. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Anamorphic encryption: New constructions and homomorphic realizations. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 33–62. Springer, Cham, May 2024. 4, 7

CGM24b. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Limits of black-box anamorphic encryption. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part II*, volume 14921 of *LNCS*, pages 352–383. Springer, Cham, August 2024. 7

CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 523–552. Springer, Berlin, Heidelberg, May / June 2010. 5, 6

CPPT18. Sébastien Canard, Duong Hieu Phan, David Pointcheval, and Viet Cuong Trinh. A new technique for compacting ciphertext in multi-channel broadcast encryption and attribute-based encryption. *Theor. Comput. Sci.*, 723:51–72, 2018. 6

DPY20. Xuan Thanh Do, Duong Hieu Phan, and Moti Yung. A concise bounded anonymous broadcast yielding combinatorial trace-and-revoke schemes. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *ACNS 20International Conference on Applied Cryptography and Network Security, Part II*, volume 12147 of *LNCS*, pages 145–164. Springer, Cham, October 2020. 5, 6, 13

FN94. Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 480–491. Springer, Berlin, Heidelberg, August 1994. 7

FP12. Nelly Fazio and Irippuge Milinda Perera. Outsider-anonymous broadcast encryption with sublinear ciphertexts. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 225–242. Springer, Berlin, Heidelberg, May 2012. 5

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 5, 6, 9, 10, 18

KHS+20. Intae Kim, Seong Oun Hwang, Willy Susilo, Joonsang Baek, and Jongkil Kim. Efficient anonymous multi-group broadcast encryption. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *ACNS 20International Conference on Applied Cryptography and Network Security, Part I*, volume 12146 of *LNCS*, pages 251–270. Springer, Cham, October 2020. 6

KPP+23a. Miroslaw Kutylowski, Giuseppe Persiano, Duong Hieu Phan, Moti Yung, and Marcin Zawada. Anamorphic signatures: Secrecy from a dictator who only permits authentication! In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 759–790. Springer, Cham, August 2023. 4, 7

KPP+23b. Miroslaw Kutylowski, Giuseppe Persiano, Duong Hieu Phan, Moti Yung, and Marcin Zawada. The self-anti-censorship nature of encryption: On the prevalence of anamorphic cryptography. *PoPETs*, 2023(4):170–183, October 2023. 4, 7

KS12. Aggelos Kiayias and Katerina Samari. Lower bounds for private broadcast encryption. In *Information Hiding*, volume 7692 of *Lecture Notes in Computer Science*, pages 176–190. Springer, 2012. 6

LG18. Jiangtao Li and Junqing Gong. Improved anonymous broadcast encryptions - tight security and shorter ciphertext. In Bart Preneel and Frederik Vercauteren, editors, *ACNS 18International Conference on Applied Cryptography and Network Security*, volume 10892 of *LNCS*, pages 497–515. Springer, Cham, July 2018. 5

LPQ12. Benoît Libert, Kenneth G. Paterson, and Elizabeth A. Quaglia. Anonymous broadcast encryption: Adaptive security and efficient constructions in the standard model. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 206–224. Springer, Berlin, Heidelberg, May 2012. 5, 11

LPSS14. San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld. Hardness of k-LWE and applications in traitor tracing. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 315–334. Springer, Berlin, Heidelberg, August 2014. 6, 10, 13

LTTT20. Minh Ha Le, Vinh Duc Tran, Van Anh Trinh, and Viet Cuong Trinh. Compacting ciphertext in multi-channel broadcast encryption and attribute-based encryption. *Theor. Comput. Sci.*, 804:219–235, 2020. 6

PPT13. Duong Hieu Phan, David Pointcheval, and Viet Cuong Trinh. Multi-channel broadcast encryption. In Kefei Chen, Qi Xie, Weidong Qiu, Ninghui Li, and Wen-Guey Tzeng, editors, *ASIACCS 13*, pages 277–286. ACM Press, May 2013. 5, 6, 7

PPY22. Giuseppe Persiano, Duong Hieu Phan, and Moti Yung. Anamorphic encryption: Private communication against a dictator. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 34–63. Springer, Cham, May / June 2022. 3, 4, 7

PPY24. Giuseppe Persiano, Duong Hieu Phan, and Moti Yung. Public-key anamorphism in (CCA-secure) public-key encryption and beyond. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part II*, volume 14921 of *LNCS*, pages 422–455. Springer, Cham, August 2024. 7

PY21. Duong Hieu Phan and Moti Yung. Privacy in advanced cryptographic protocols: Prototypical examples. *Journal of Computer Science and Cybernetics*, 37(4):429–452, Oct. 2021. 7

WCHY23. Yi Wang, Rongmao Chen, Xinyi Huang, and Moti Yung. Sender-anamorphic encryption reformulated: Achieving robust and generic constructions. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VI*, volume 14443 of *LNCS*, pages 135–167. Springer, Singapore, December 2023. 4