Bounded-Collusion Identity-Based Encryption from Semantically-Secure Public-Key Encryption: Generic Constructions with Short Ciphertexts*

Stefano Tessaro[1]** and David A. Wilson[2]

[1] University of California, Santa Barbara
[2] MIT

Abstract. Identity-based encryption (IBE) is a special case of public-key encryption where user identities replace public keys. Every user is given a corresponding secret key for decryption, and encryptions for his or her identity must remain confidential even to attackers who learn the secret keys associated with other identities. Several IBE constructions are known to date, but their security relies on specific assumptions, such as quadratic residuosity, as well as different pairing-based and lattice-based assumptions.

To circumvent the lack of generic constructions, Dodis *et al.* (EUROCRYPT '02) introduced the notion of *bounded-collusion IBE* (BC-IBE), where attackers only learn secret keys of an a-priori bounded number t of identities. They provided a *generic* BC-IBE construction from any semantically-secure encryption scheme which, however, suffers from a $\omega(t)$ blow-up in ciphertext size. Goldwasser *et al.* (TCC 2012) recently presented a generic construction with no ciphertext-length blow-up. Their construction requires an underlying public-key scheme with a key homomorphism, as well as a hash-proof-style security definition that is strictly stronger than semantic security. This latter requirement in particular reduces the applicability of their construction to existing schemes.

In this paper, we present the *first* generic constructions of BC-IBE from *semantically-secure* encryption schemes with no ciphertext-length blow-up. Our constructions require different degrees of key-homomorphism and malleability properties that are usually easy to verify. We provide concrete instantiations based on the DDH, QR, NTRU, and LWE assumptions. For all of these assumptions, our schemes present the smallest BC-IBE ciphertext size known to date. Our NTRU-based construction is particularly interesting, due to the lack of NTRU-based IBE constructions as well as the fact that it supports fully-homomorphic evaluation. Our results also yield new constructions of bounded CCA-secure cryptosystems.

Keywords. Public-key cryptography, Identity-based encryption, Bounded-CCA security.

* An extended abstract of this paper appears in the proceedings of PKC 2014. This is the full version.
** Work done while the author was a research scientist at MIT CSAIL.

1 Introduction

Public-key encryption. One of the classic and best-studied models of secure communication is that of *public-key encryption (PKE)* [DH76], in which each individual independently generates a *public-key / secret-key* pair. Anyone possessing the public key can encrypt a message such that only the individual with the associated secret key can decrypt. To date, there are innumerable PKE constructions proven secure based on a wide variety of hardness assumptions.

However, the basic public-key model lacks a well-developed structure for public key verification. One can encrypt messages using a public key, but the model implies a trust that the public key belongs to a specific individual, unless an expensive public-key infrastructure is in place. In order to make explicit these assumptions and avoid potential difficulties with key distribution, cryptographers have explored other models of encryption.

Identity-based Encryption. The identity-based encryption (IBE) model, introduced by Shamir in 1984 [Sha85], attempts to alleviate the above concerns. In this model, a trusted center generates a master secret key and public parameters for the entire system. Anyone can encrypt a message to any user of the system using only these global public parameters and the user's *identity*. To decrypt, a user must obtain the secret key for their identity from the trusted center (who presumably authenticates the user before distributing the key).

The security model for IBE assumes that the adversary can adaptively obtain an arbitrary number of secret keys for users in the system, and requires that messages encrypted to any other user still be indistinguishable to the adversary. This models the idea that an individual's messages are still secure even if an arbitrary number of other users of the system collude against that user.

The first constructions of IBE came in 2001, by Boneh and Franklin [BF01] and Cocks [Coc01]. Both of these constructions assumed the existence of random oracles; however, subsequent work by Boneh and Boyen [BB04] and Waters [Wat05] achieved IBE in the standard model. There now exist a number of IBE constructions in both the random oracle and standard models, under hardness assumptions of problems in bilinear groups (e.g. [BF01,CHK03,BB04,Wat05]), various forms of the Quadric Residuosity (QR) problem (e.g. [Coc01,BGH07]), and the Learning With Errors problem (e.g. [GPV08,CHKP10,ABB10]). Some of these, moreover, and in particular all those based on the standard QR problem, additionally require random oracles. *However, no constructions of IBE are known from generic primitives.*

Bounded-collusion IBEs. As an attempt to come up with constructions under a wider range of assumptions, cryptographers began looking at a variant of IBE known as *Bounded-Collusion IBE* (BC-IBE). In this model, one only guarantees security against an adversary who obtains secret keys associated with at most t identities, where the size of the parameters of the system are allowed to depend on t. Falling short of achieving full security, the bounded-collusion model can be a realistic assumption in many settings, and is in fact a necessary restriction to achieve the more general notion of functional encryption [GVW12]. Additionally, it has been studied in other settings, notably broadcast encryption and revocation (e.g. [FN94,GSY99,GSW00,KRS99,NNL01,HS02,DF03]).

The first construction of BC-IBE came in the context of key-insulated systems in [DKXY02]. This paper gave a general reduction from any semantically secure public-key cryptosystem to a BC-IBE scheme. However, their construction suffers from a large ciphertext-size blowup – the resulting ciphertext length is a factor $\omega(t)$ larger than that of the underlying encryption scheme. To mitigate this, this work was recently followed by that of Goldwasser *et al.* [GLW12]: They provide a new construction that relies on a public-key encryption scheme which exhibits *key-homomorphic* properties, i.e., secret keys and public keys are elements of respective groups (with possibly different operations, which we denote by $+$ and $\cdot$), and there exists a homomorphism μ such that $\mu(\mathsf{sk} + \mathsf{sk}') = \mu(\mathsf{sk}) \cdot \mu(\mathsf{sk}')$, where $\mu(\mathsf{sk})$ and $\mu(\mathsf{sk}')$ are valid public keys for which sk and sk' yield correct decryption, respectively. More concretely, the *GLW construction* generates multiple public-key /

secret-key pairs $(\mathsf{pk}_1, \mathsf{sk}_1), \ldots, (\mathsf{pk}_n, \mathsf{sk}_n)$, letting the public-parameters and the master secret key of the scheme be $\mathsf{pp} = (\mathsf{pk}_1, \ldots \mathsf{pk}_n)$ and $\mathsf{msk} = (\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$, respectively. Then, an efficient map ϕ associates every identity ID with a vector $[\mathrm{id}_1, \ldots, \mathrm{id}_n]$, and a message m is encrypted for an identity ID as the ciphertext $c = \mathsf{Enc}(\mathsf{pk}_{\mathrm{ID}}, m)$, where $\mathsf{pk}_{\mathrm{ID}} = \prod_{i=1}^{n} \mathsf{pk}_i^{\mathrm{id}_i}$. By the existence of μ, this ciphertext can be decrypted using $\mathsf{sk}_{\mathrm{ID}} = \sum_{i=1}^{n} \mathrm{id}_i \cdot \mathsf{sk}_i$, since the homomorphism guarantees that $\mathsf{pk}_{\mathrm{ID}} = \mu(\mathsf{sk}_{\mathrm{ID}})$. The map ϕ is subjected to a combinatorial requirement that disallows computing $\mathsf{sk}_{\mathrm{ID}}$ given $\mathsf{sk}_{\mathrm{ID}'}$ for t different $\mathrm{ID}' \neq \mathrm{ID}$. The GLW construction *preserves* the ciphertext size of the underlying encryption, but its security requires the latter to satisfy a property which is *strictly* stronger than semantic security. This property is inspired by the security of hash-proof systems [CS02], and in particular does not allow the homomorphism μ to be one-to-one. This somewhat hinders the applicability of their framework to existing encryptions schemes not designed with this security goal in mind.

Our contributions. In this paper, we seek for generic constructions of BC-IBE which rely on encryption schemes that solely satisfy the *standard* security notion of semantic security in addition to some syntactical, non-security-related, properties which can be easily verified. Our constructions have the added benefit of conceptual simplicity, and the resulting instantiations from concrete assumptions either outperform or abstract existing BC-IBE constructions along different axes.

In summary, this paper makes three main contributions:

1. As our first contribution, we revisit the GLW approach in the context of *selective* security. The latter security notion only demands security for attackers attempting to break the confidentiality of messages encrypted for an *a-priori specified identity* (in particular, independently of the parameters of the scheme). We prove that the GLW approach is *selectively* secure for every *semantically secure* encryption scheme with key-homomorphic properties whenever ϕ satisfies a slightly stronger property that the one used in [GLW12], namely that of cover-freeness introduced in [EFF85] and used in several other works (e.g. [KRS99,CHH+07,DHT12], and others). While being strictly weaker than the notion of full security, selective security is sufficient for some applications, as discussed below.
2. Whenever the underlying semantically-secure scheme satisfies an additional new property – which we call *weak multi-key malleability* – we prove that the GLW construction achieves *full* BC-IBE security, i.e., confidentiality holds even with respect an identity chosen adaptively after learning the parameters of the schemes as well as secret keys for at most t other identities. Roughly, our malleability property states that given the encryption of $c = \mathsf{Enc}(\mathsf{pk}, m)$ of an *unknown* message m under a known public-key pk, and given an additional public-key / secret-key pair $(\mathsf{pk}', \mathsf{sk}')$, we can efficiently produce a ciphertext which is indistinguishable from an encryption of m under $\mathsf{pk} \cdot \mathsf{pk}'$. An example scheme with this property is ElGamal encryption – hence we directly obtain a DDH-based BC-IBE scheme from ElGamal encryption.
3. As our third contribution, we provide a new, alternative construction that relies on a different form of malleability (which we simply call multi-key malleability), and does not require any explicit key-homomorphic structure. Intuitively, our notion requires that given $c = \mathsf{Enc}(\mathsf{pk}, m)$ for an unknown message m, and another public key pk', we can obtain a new ciphertext c which decrypts to m under a combination of the secret keys sk and sk' associated with pk and pk'. We provide an efficient instantiation based on NTRU [HPS98], exploiting it multi-key homomorphic properties recently observed by Lopez-Alt *et al.* [LATV12]. This is of particular interest due to the fact that no fully-secure NTRU-based IBE scheme is known to date. Moreover, our constructions support homomorphic evaluation of ciphertexts, and this is the only construction of identity-based fully homomorphic encryption beyond the recent result by Gentry, Sahai, and Waters [GSW13].

Construction	Assumptions	Ciphertext size	PK size
[DKXY02]	Semantic-secure PKE	$\Theta(t \log \lvert\mathcal{ID}\rvert)$ PKE ciphertexts	$\Theta(t^2 \log \lvert\mathcal{ID}\rvert)$ PKE PKs
[GLW12]	PKE w/linear hash proof and key homomorphism	Same as underlying PKE	$\Theta(t \log \lvert\mathcal{ID}\rvert)$ PKE PKs
This work	Semantic-secure PKE; key homomorphism, weak multi-key malleability	Same as underlying PKE	$\Theta(t^2 \log \lvert\mathcal{ID}\rvert)$ PKE PKs
This work	Semantic-secure PKE; multi-key malleability	Same as underlying PKE	$\Theta(t^2 \log \lvert\mathcal{ID}\rvert)$ PKE PKs
[DKXY02]	DDH	3 group elements	$\Theta(t)$ group elements
[GLW12]	DDH	3 group elements	$\Theta(t \log \lvert\mathcal{ID}\rvert)$ group elements
This work	DDH	2 group elements	$\Theta(t^2 \log \lvert\mathcal{ID}\rvert)$ group elements
[GLW12]	QR	2 RSA group elements	$\Theta(t \log \lvert\mathcal{ID}\rvert)$ group elements
This work	LWE	Same as GPV [GPV08]	$\Theta(t^2 \log \lvert\mathcal{ID}\rvert)$ GPV PKs
This work	NTRU	Same as NTRUEncrypt [HPS98]	$\Theta(t^2 \log \lvert\mathcal{ID}\rvert)$ NTRU PKs

Table 1. Comparison with previous works on BC-IBE. Here t is the collusion parameter and $|\mathcal{ID}|$ is the total number of identities in the system. PK and ciphertext size implicitly include the security parameter. The above portion of the table considers generic constructions, whereas the lower section describes existing constructions from concrete assumptions. Note that linear hash proof property implies semantic security, while being strictly stronger than it.

To conclude, we stress that our instantiation of the GLW approach is somewhat orthogonal to the one by Goldwasser *et al.*: Our instantiation requires indeed somewhat *larger* public-parameters at the cost of a weaker assumption on the underlying encryption scheme, hence leading to wider applicability and often smaller ciphertexts. Nonetheless, we believe that large ciphertexts are generally a more limiting factor than large parameters, especially in settings where many messages are encrypted with the same parameters.

A summary of our instantiations and their parameters is given in Table 1 comparing them to previously known best constructions. For LWE and NTRU, the best previously known construction was obtained by using the construction of [DKXY02]. We also provide a construction based on QR which does not outperform the one of [GLW12], even though we find it conceptually simpler.

From IBE to CCA-security. A somewhat related problem is that of building bounded-CCA secure public-key encryption [CHH+07]: Concretely, for t-bounded CCA security, semantic security must hold also for attackers which can decrypt up to t ciphertexts other than the challenge ciphertext for which we attempt to break confidentiality. We note that by re-interpreting a result of Boneh *et al.* [BCHK07], every construction of a BC-IBE scheme *selectively* secure against t-collusions directly yields a t-bounded CCA secure PKE. Hence, our BC-IBE constructions also directly yield better bounded-CCA-secure constructions, in terms of ciphertext size and/or conceptual simplicity. When applying our framework to ElGamal, for example, we obtain a construction which is equivalent to the one proposed in [CHH+07], for which a direct security proof was given. Moreover, our instantiation from NTRU is indeed more efficient than the best fully CCA-secure construction from NTRU given by Steinfeld *et al.* [SLP+12].

Outline of this paper. In Section 2, we revisit some standard notation, as well as traditional security definitions for public-key and identity-based encryption. In Section 3, we revisit the construction of Goldwasser *et al.*, assuming a cover-free identity mapping and only semantic security. We give multiple new constructions of BC-IBE under this framework. In Section 4, we define the property of multi-key malleability, prove that this property can also yield constructions of BC-IBE, and give a specific example based on the NTRU assumption. Finally, in Section 5, we discuss the role of BC-IBE systems in constructing bounded-CCA-secure schemes.

Game CPA for $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$:	**Game** CCA for $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$:
$(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{Gen}$	$(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{Gen}$
$b \xleftarrow{\$} \{0,1\}$	$b \xleftarrow{\$} \{0,1\}$
$(m_0, m_1, \mathsf{st}) \xleftarrow{\$} \mathcal{A}(\mathsf{pk})$	$(m_0, m_1, \mathsf{st}) \xleftarrow{\$} A^{\mathsf{Dec}(\mathsf{sk},\cdot)}(\mathsf{pk})$
$c^* \xleftarrow{\$} \mathsf{Enc}(\mathsf{pk}, m_b)$	$c^* \xleftarrow{\$} \mathsf{Enc}(\mathsf{pk}, m_b)$
$b' \xleftarrow{\$} \mathcal{A}(c^*, \mathsf{st})$	$b' \xleftarrow{\$} \mathcal{A}^{\mathsf{Dec}(\mathsf{sk},\cdot)}(c^*, \mathsf{st})$
Win iff $b' = b$	Win iff $b' = b$.

Fig. 1. Semantic Security for Public-Key Encryption. Security games defining semantic security of public-key encryption against chosen-plaintext attacks (left) and against chosen-ciphertext attacks (right). On the right, we tacitly assume that in the second phase of the game, the decryption oracle $\mathsf{Dec}(\mathsf{sk}, \cdot)$ answers to queries c^* with $\perp$. Also, in both games, the challenge ciphertext c^* is set to $\perp$ if $|m_0| \neq |m_1|$.

2 Preliminaries

This section reviews basic notion about pubic-key encryption and identity-based encryption used throughout the paper.

2.1 Public-key Encryption

PKE Syntax. As usual, a *public-key encryption (PKE) scheme* is a triple of efficient algorithms $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ satisfying the following synactic properties:

- Gen is the (randomized) *key generation algorithm*: it takes no input (other than the security parameter 1^k , which is implicit and generally omitted), and outputs a public-key / secret-key pair $(\mathsf{pk}, \mathsf{sk})$.
- Enc and Dec are the (randomized) *encryption* and the (deterministic) *decryption* algorithms, such that for all valid public-key / secret-key pairs $(\mathsf{pk}, \mathsf{sk})$ output by Gen, and all messages m, the probability $\mathsf{P}[\mathsf{Dec}(sk, \mathsf{Enc}(\mathsf{pk}, m)) \neq m]$ is negligible, where the probability is taken over the random coins of the encryption algorithm Enc.

Often, we allow public-key encryption schemes to additionally depend on explicit *public parameters* pp (randomly generated in an initial phase and shared across multiple instances of the PKE scheme) on which all of Gen, Enc, and Dec are allowed to depend. Examples include the description of a group G with its generator g. We will often omit them in the descriptions of generic constructions from PKE schemes.

Security of PKE. We define *security against chosen-plaintext attacks* (for short, *IND-CPA security*) [GM84,BDPR98] for a PKE scheme $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ via the security game depicted on the left of Figure 1. It involves an adversary $\mathcal{A}$ which is initially given the public key pk, and subsequently outputs a pair of equal-length messages m_0, m_1. The adversaries continues after receiving a *challenge ciphertext* $c^* \xleftarrow{\$} \mathsf{Enc}(\mathsf{pk}, m_b)$ for a random secret bit b, and then finally outputs a guess b' for b. We say that PKE is (τ, ε)-*ind-cpa-secure* if all attackers $\mathcal{A}$ with time complexity at most τ guess the right bit (i.e., $b' = b$) with probability at most $\frac{1+\varepsilon}{2}$. Moreover, it is simply *ind-cpa secure* if for all polynomials p, there exists a negligible function ν such that the scheme is $(p(k), \nu(k))$-ind-cpa-secure for all values of the security parameter k. We also consider *security against chosen ciphertext attacks* (for short, *IND-CCA security*), where the adversary is additionally able to decrypt ciphertexts under the constraint that a decryption query for the challenge ciphertext is never asked. (The associated game is on the right of Figure 1.) We say that PKE is (τ, t, ε)-ind-cca-secure

Game selective-CPA for IBE:	**Game** CPA for IBE:
$(\mathsf{pp}, \mathsf{msk}) \xleftarrow{\$} \mathsf{IBEGen}$	$(\mathsf{pp}, \mathsf{msk}) \xleftarrow{\$} \mathsf{IBEGen}$
$b \xleftarrow{\$} \{0,1\}$	$b \xleftarrow{\$} \{0,1\}$
$(\mathrm{ID}^*, \mathsf{st}) \xleftarrow{\$} \mathcal{A}$	$(m_0, m_1, \mathrm{ID}^*, \mathsf{st}) \xleftarrow{\$} \mathcal{A}^{\mathsf{IBEExtract}(\mathsf{msk},\cdot)}(\mathsf{pp})$
$(m_0, m_1, \mathsf{st}') \xleftarrow{\$} \mathcal{A}^{\mathsf{IBEExtract}(\mathsf{msk},\cdot)}(\mathsf{pp}, \mathsf{st})$	$c^* \xleftarrow{\$} \mathsf{IBEEnc}(\mathsf{pp}, \mathrm{ID}^*, m_b)$
$c^* \xleftarrow{\$} \mathsf{IBEEnc}(\mathsf{pp}, \mathrm{ID}^*, m_b)$	$b' \xleftarrow{\$} \mathcal{A}^{\mathsf{IBEExtract}(\mathsf{msk},\cdot)}(c^*, \mathsf{st})$
$b' \xleftarrow{\$} \mathcal{A}^{\mathsf{IBEExtract}(\mathsf{msk},\cdot)}(c^*, \mathsf{st}')$	

Fig. 2. Semantic Security for IBE. Security games defining semantic security of identity-based encryption, for selective attacks (left) and for full security (right). In both games, we assume that extraction queries for ID^* are answered by $\perp$, and that $c^* = \perp$ if $|m_0| \neq |m_1|$, or (for full security) if ID^* was previously queried to $\mathsf{Extract}(\mathsf{msk}, \cdot)$.

if any attacker with time complexity τ and making at most t decryption queries guesses b with probability at most $\frac{1+\varepsilon}{2}$. The asymptotic notion of *t-ind-cca-secure* is defined accordingly.

2.2 Identity-based Encryption

Recall that an *identity-based encryption (IBE)* scheme for identity set $\mathcal{ID}$ is a 4-tuple of algorithms $\mathsf{IBE} = (\mathsf{IBEGen}, \mathsf{IBEExtract}, \mathsf{IBEEnc}, \mathsf{IBEDec})$ satisfying the following syntactical properties:

- IBEGen is the randomized *parameter generator algorithm* which returns a pair $(\mathsf{msk}, \mathsf{pp})$, where msk is the so-called *master secret key*, and pp are the *public parameters*.
- The *extraction algorithm* $\mathsf{IBEExtract}$, on input the master secret-key msk and a valid identity $\mathrm{ID} \in \mathcal{ID}$ returns a secret key $\mathsf{sk}_{\mathrm{ID}} \xleftarrow{\$} \mathsf{IBEExtract}(\mathsf{msk}, \mathrm{ID})$ associated with this identity.
- The encryption algorithm IBEEnc takes as inputs the public parameters pp, an identity $\mathrm{ID} \in \mathcal{ID}$, and a message m, and returns a ciphertext $c \xleftarrow{\$} \mathsf{IBEEnc}(\mathsf{pp}, \mathrm{ID}, m)$ such that for the associated deterministic algorithm IBEDec, $\mathsf{IBEDec}(\mathsf{sk}_{\mathrm{ID}}, \mathsf{IBEEnc}(\mathsf{pp}, \mathrm{ID}, m)) = m$ with overwhelming probability for each $(\mathsf{pp}, \mathsf{msk})$ output by Gen and sk_{ID} output by $\mathsf{IBEExtract}(\mathsf{msk}, \mathrm{ID})$.

The notion of IND-CPA security is extended to the setting of IBE. The adversary, given the public parameters pp, can obtain keys $\mathsf{sk}_{\mathrm{ID}}$ for identities ID of its choice (via so-called *extraction queries*), and outputs at some point a pair of equal-length challenge messages m_0, m_1, together with a *challenge identity* ID^* for which no extraction query has been issued. It then obtains an encryption of m_b for the challenge identity ID^* and for a random bit b. The adversary is asked to guess b, constrained on not asking a key extraction query for ID^*. The game is given in Figure 2, on the right. We also consider a weaker security notion, called *selective IND-CPA security*, for which the corresponding security game is given on the left of Figure 2. Here, the adversary is required to choose its challenge identity *beforehand*, and only subsequently learns the public parameters and is given access to the $\mathsf{IBEExtract}$ oracle.

In analogy to the case of conventional PKE, we say that IBE is (τ, t, ε)-cpa-secure if all τ-time adversaries $\mathcal{A}$ making t extraction queries output b with probability at most $\frac{1+\varepsilon}{2}$ in the CPA-security game above. Similarly, we define (τ, t, ε)-selective-cpa-secure likewise for the selective-CPA game above, as well as the asymptotic notions of t-cpa and t-selective-cpa security.

3 Revisiting the GLW Construction

In the first part of this paper, we revisit the IBE construction for bounded-collusion security proposed by Goldwasser, Lewko, and Wilson [GLW12] – henceforth, we refer to this construction as the *GLW construction*. We show two generic results, the first one for selective security and the

second one for full ibe security. Then, we discuss two new instantiations of this paradigm based on DDH and LWE. A third instantiation based on the QR assumption is deferred to Appendix A.

3.1 The GLW Construction

Secret-key to Public-key Homomorphisms. Throughout this section, we (tacitly) consider only public-key cryptosystems $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with the property that secret and public keys are elements of groups G and H, respectively. For convenience and ease of distinction, we will denote the group operations on G and H as $+$ and $\cdot$, respectively.

Definition 1 (Secret-key to public-key homomorphism). *We say that* PKE *admits a* secret-key to public-key homomorphism *if there exists a map* $\mu : G \to H$ *such that:*

- **(i)** *μ is a homomorphism, i.e., for all* $\mathsf{sk}, \mathsf{sk}' \in G$*, we have* $\mu(\mathsf{sk} + \mathsf{sk}') = \mu(\mathsf{sk}) \cdot \mu(\mathsf{sk}')$*;*
- **(ii)** *Every output* $(\mathsf{sk}, \mathsf{pk})$ *of* Gen *satisfies* $\mathsf{pk} = \mu(\mathsf{sk})$*.*

We stress that we are *not* requiring that every element $\mathsf{sk} \in G$ is a valid secret key output by Gen. This will be important in our LWE instantiation below. In this case, we still want to make sure that decryption is correct: In particular, we say below that μ *satisfies* n*-correctness* if for any $n' \leq n$ valid secret keys $\mathsf{sk}_1, \ldots, \mathsf{sk}_{n'}$ output by Gen, the probability $\mathsf{P}[\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mu(\mathsf{sk}), m)) \neq m]$ is negligible for all messages m, where the probability is over the coins of Enc and where $\mathsf{sk} = \mathsf{sk}_1 + \cdots + \mathsf{sk}_{n'}$. (This property is implicitly satisfied for all n if all elements of G are valid secret keys.)

Also note that the map μ does *not* need to be efficiently computable for our applications, even though the map is often very efficient. Additionally, we observe that in case the scheme depends on some explicit public parameter (like a generator or a matrix, as will be the case in our examples below), μ is indeed allowed to be parameter-dependent.

The GLW Construction. Goldwasser, Lewko, and Wilson [GLW12] presented a generic approach to build a bounded-collusion secure IBE from a public-key encryption scheme admitting a secret-key to public-key homomorphism. Specifically, let $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be such a public-key encryption scheme with homomorphism $\mu : G \to H$ satisfying n-correctness, and let $\phi : \mathcal{ID} \to \{0,1\}^n$ be a polynomial-time computable function, called the *identity map*. (With a slight abuse of notation, it will be convenient to consider the output ϕ as a subset of $\{1, \ldots, n\}$, encoded in the canonical way as an n-bit string.) Then, the GLW construction for PKE and ϕ gives rise to the following IBE scheme $\mathsf{IBE} = (\mathsf{IBEGen}, \mathsf{IBEExtract}, \mathsf{IBEEnc}, \mathsf{IBEDec})$ with identities from the set $\mathcal{ID}$ defined as follows:

$\underline{\mathsf{IBEGen}:}$	$\underline{\mathsf{IBEExtract}(\mathsf{msk} = \mathbf{sk}, \mathrm{ID}):}$	$\underline{\mathsf{IBEEnc}(\mathsf{pp} = \mathbf{pk}, \mathrm{ID}, m):}$	$\underline{\mathsf{IBEDec}(\mathsf{sk}_{\mathrm{ID}}, c):}$
$(\mathbf{pk}, \mathbf{sk}) \xleftarrow{\$} \mathsf{Gen}^n$	$\mathsf{sk}_{\mathrm{ID}} = \sum_{i \in \phi(\mathrm{ID})} \mathbf{sk}[i]$	$\mathsf{pk}_{\mathrm{ID}} = \prod_{i \in \phi(\mathrm{ID})} \mathbf{pk}[i]$	$m' \leftarrow \mathsf{Dec}(\mathsf{sk}_{\mathrm{ID}}, c)$
$\mathsf{msk} \leftarrow \mathbf{sk}$	Return $\mathsf{sk}_{\mathrm{ID}}$	$c \xleftarrow{\$} \mathsf{Enc}(\mathsf{pk}_{\mathrm{ID}}, m)$	Return m'
$\mathsf{pp} \leftarrow \mathbf{pk}$		Return c	
Return $(\mathsf{msk}, \mathsf{pp})$			

The notation $(\mathbf{pk}, \mathbf{sk}) \xleftarrow{\$} \mathsf{Gen}^n$ denotes running Gen n times, with independent random coins, and $\mathbf{pk}, \mathbf{sk}$ are vectors such that $(\mathbf{pk}[i], \mathbf{sk}[i])$ is the output of the i-th execution of Gen. First note that correctness of IBE follows trivially from the correctness of PKE and the existence of a secret-key to public-key homomorphism μ with n-correctness, since $\mathsf{pk}_{\mathrm{ID}} = \mu(\mathsf{sk}_{\mathrm{ID}})$ holds for all IDs and $\mathsf{sk}_{\mathrm{ID}}$ is the sum of at most n valid secret keys. We stress that a central advantage of the above construction is that IBE ciphertexts are ciphertexts of the underlying encryption scheme PKE. Also, note that if PKE relies on some public parameters, these are generated once and used across all uses of Gen, Enc, and Dec.

INSTANTIATING THE IDENTITY MAP. We still need to discuss how the map ϕ is instantiated. In all constructions of this paper, we rely on constructions based on *cover-free sets*, following previous work on bounded-collusion IBE [DKXY02], bounded-CCA security [CHH+07], and bounded security for FDH signatures [DHT12]. Concretely, let $2^{[n]}$ be the set of subsets of $[n] := \{1, \ldots, n\}$.

Definition 2 (Cover-free sets). *We say that $\phi : \mathcal{ID} \to 2^{[n]}$ is (t, s)-*cover free *if $|\phi(x)| = s$ for all $x \in \mathcal{ID}$, and moreover $\phi(x_t) \setminus \bigcup_{i=1}^{t-1} \phi(x_i) \neq \emptyset$ for all $x_1, \ldots, x_t \in \mathcal{ID}$, i.e., the set $\phi(x_t)$ is not covered by the union of $\phi(x_1), \ldots, \phi(x_{t-1})$.*

In general, we will equivalently think of ϕ as a map $\mathcal{ID} \to \{0,1\}^n$, where we output the characteristic vector of the associated set, instead of the set itself. The following gives the currently best-known construction of cover-free sets.

Theorem 1 ([CHH+07]). *For all integers $t \geq 1$, there exists a polynomial-time computable (t, s)-cover-free map $\phi : \mathcal{ID} \to \{0,1\}^n$, where $n = 16t^2 \log |\mathcal{ID}|$ and $s = 4t \log |\mathcal{ID}|$.*

We note that Goldwasser, Lewko, and Wilson used a weaker requirement of ϕ that only requires linear independence of the vectors $\phi(x_1), \ldots, \phi(x_t)$. In this case, the output length n can be reduced to $O(t \log |\mathcal{ID}|)$, or even $O(t)$ if we allow both identities as well as components of $\phi(x)$ to be elements of $\mathbb{Z}_p$ for some large prime p. However, the price they pay compared to our results below is that the underlying encryption scheme is required to satisfy a harder to show notion than in our results given below assuming cover-freeness, and this is often reflected in instantiations with larger ciphertexts.

3.2 Selective Security of the GLW Construction

We start with selective security, which will be important to obtain bounded CCA-secure cryptosystems with short ciphertexts, as we explain below in Section 5. In the following, let $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be an arbitrary public-key encryption scheme which admits secret-key to public-key homomorphism, and let IBE be the IBE scheme resulting from the above construction, using an underlying identity map ϕ.

Theorem 2 (Selective ID Security of GLW). *Assume that PKE is ind-cpa-secure, and that ϕ is $(t+1, s)$-cover free. Then, the GLW construction is t-selective-cpa-secure.*

Proof. Let $\mathcal{A}$ be a selective-cpa adversary for IBE which outputs $b' = b$ with probability at least $(1 + n\varepsilon)/2$, and which makes at most t extraction queries. We construct an ind-cpa adversary $\mathcal{B}$ for PKE from $\mathcal{A}$, guessing the bit b with probability $\frac{1+\varepsilon}{2}$. Concretely, the adversary $\mathcal{B}$ first runs $\mathcal{A}$, obtaining the challenge identity ID^*, and chooses an index i^* uniformly at random from the set $S^* = \{i : \mathsf{id}_i^* = 1\}$, where $\phi(\mathrm{ID}^*) = [\mathsf{id}_1^*, ..., \mathsf{id}_n^*]$. It then gets a public key pk^* from the underlying CPA game, and computes $(\mathbf{pk}[j], \mathbf{sk}[j]) \xleftarrow{\$} \mathsf{Gen}$ for all $j \in [n] \setminus \{i^*\}$. Finally, it sets $\mathbf{pk}[i^*] = \mathsf{pk}^* \cdot \left(\prod_{j \neq i^*} \mathbf{pk}[j]^{-\mathsf{id}_j^*}\right)$.

The adversary $\mathcal{B}$ then gives $\mathsf{pp} = \mathbf{pk}$ to $\mathcal{A}$ and runs it until it outputs a pair (m_0, m_1). In particular, $\mathcal{A}$'s extraction queries for $\mathrm{ID} \neq \mathrm{ID}^* \in \mathcal{ID}$ are replied by computing $[\mathsf{id}_1, \ldots, \mathsf{id}_n] = \phi(\mathrm{ID})$ and, if $\mathsf{id}_{i^*} = 0$, returning $\mathsf{sk}_{ID} := \sum_i \mathsf{id}_i \cdot \mathbf{sk}[i]$. Note that if $\mathsf{id}_{i^*} = 1$, then $\mathcal{B}$ cannot answer the extraction query, as it does not know any corresponding $\mathbf{sk}[i^*]$. In this case, it returns $\perp$, and sets a flag bad to $\mathtt{true}$. When the adversary $\mathcal{A}$ outputs a pair (m_0, m_1) of messages of equal length, $\mathcal{B}$ forwards them to the CPA, obtaining a challenge ciphertext c^* , which it then gives back to $\mathcal{A}$, and its simulated execution is continued until it outputs a bit b'. To conclude, $\mathcal{B}$ outputs the bit b' if bad is not set to $\mathtt{true}$, and returns a random bit otherwise. Note that we have $\mathsf{pk}_{\mathrm{ID}^*} = \mathsf{pk}^*$ by our definition.

Since ϕ is $(t+1, s)$-cover-free, we know that there exists at least one i^* such that $\mathsf{id}_{i^*}^* = 1$, but $\mathsf{id}_{i^*} = 0$ for all vectors $\phi(\mathrm{ID})$ corresponding to the (at most t) extraction queries $\mathrm{ID} \neq \mathrm{ID}^*$.

Intuitively, such an index i^* is chosen hence with probability at least $1/|S^*| = 1/s \geq 1/n$, and conditioned on this, the simulation is easily seen to be perfect. Formally, we let $\mathsf{Win}_{\mathsf{PKE}}$ and $\mathsf{Win}_{\mathsf{IBE}}$ be the events that $\mathcal{B}$ and $\mathcal{A}$ guess the bit in the respective security games. Then,

$$\begin{aligned}\mathsf{P}\left[\mathsf{Win}_{\mathsf{PKE}}\right] &= \mathsf{P}\left[\mathsf{Win}_{\mathsf{PKE}} \wedge \mathsf{bad} = \mathtt{false}\right] + \mathsf{P}\left[\mathsf{Win}_{\mathsf{PKE}} \wedge \mathsf{bad} = \mathtt{true}\right] \\ &\geq \mathsf{P}\left[\mathsf{bad} = \mathtt{false}\right] \cdot \mathsf{P}\left[\mathsf{Win}_{\mathsf{PKE}} \,\middle|\, \mathsf{bad} = \mathtt{false}\right] + \mathsf{P}\left[\mathsf{bad} = \mathtt{true}\right] \cdot \mathsf{P}\left[\mathsf{Win}_{\mathsf{PKE}} \,\middle|\, \mathsf{bad} = \mathtt{true}\right] .\end{aligned}$$

Now, clearly, $\mathsf{P}\left[\mathsf{bad} = \mathtt{true}\right] = 1 - \mathsf{P}\left[\mathsf{bad} = \mathtt{false}\right]$, and $\mathsf{P}\left[\mathsf{Win}_{\mathsf{PKE}} \,\middle|\, \mathsf{bad} = \mathtt{true}\right] \geq \frac{1}{2}$, since $\mathcal{B}$ outputs a random bit if bad is $\mathtt{true}$. Moreover, one can verify that $\mathsf{P}\left[\mathsf{bad} = \mathtt{false}\right] \geq \frac{1}{n}$, and, as the simulation is perfect, $\mathsf{P}\left[\mathsf{Win}_{\mathsf{PKE}} \,\middle|\, \mathsf{bad} = \mathtt{false}\right] = \mathsf{P}\left[\mathsf{Win}_{\mathsf{IBE}}\right]$. Formalizing these last two argument actually requires some (standard) extra work, using the fact that all random coins are independent of the choice of i^*, but we dispense with the details here. Plugging in terms into the above concludes the proof. □

3.3 Full Security of GLW

We note that the above proof strategy used in Theorem 2 fails when we do not know the challenge identity ID^* at the point in time when the reduction $\mathcal{B}$ sets the public parameters pp. However, an additional syntactic requirement on the underlying cryptosystem PKE yields full security, as we show below. This requirement is captured by the following definition.

Definition 3 (Weak Multi-Key Malleability). *We say that* PKE *is* weakly n-key malleable *if there exists an efficient algorithm* $\mathsf{Simulate}$ *such that for all messages m, all $I \subseteq [n]$, and all $i \in I$, the probability distributions D_0 and D_1 are computationally indistinguishable, where with $(\mathbf{pk}, \mathbf{sk}) \xleftarrow{\$} \mathsf{Gen}^n$, D_b consists of $(\mathbf{pk}, \mathbf{sk}[[n] \setminus \{i\}], c_b)$ such that*

(1) $c_0 \xleftarrow{\$} \mathsf{Enc}(\prod_{i \in I} \mathbf{pk}[i], m)$;
(2) $c \xleftarrow{\$} \mathsf{Enc}(\mathbf{pk}[i], m)$, $c_1 \xleftarrow{\$} \mathsf{Simulate}(i, I, c, \mathbf{pk}, \mathbf{sk}[[n] \setminus \{i\}])$.

In other words, given a ciphertext c encrypting with public key $\mathbf{pk}[i]$ (where i is part of some set I) an arbitrary *unknown* message m, we can efficiently generate a ciphertext c' encrypting the same message m under the *product* of the keys $\mathbf{pk}[j]$ for $j \in I$ without knowing the secret key $\mathbf{sk}[i]$, but still possibly using $\mathbf{sk}[j]$ for $j \neq i$. The resulting ciphertext has the right distribution in the eyes of a computationally bounded distinguisher.

The proof of the following theorem follows a similar approach to the one of Theorem 2, and is deferred to Appendix B.1.

Theorem 3 (Full Security of GLW). *Assume that* PKE *is ind-cpa-secure and weakly n-key malleable, and that ϕ is $(t+1, s)$-cover free. Then, the GLW construction is t-cpa-secure.*

3.4 Instantiation from DDH

We present a simple instantiation of the above paradigm based on the DDH assumptions and the ElGamal cryptosystem. The resulting scheme has smaller ciphertexts than earlier BC-IBE schemes [GLW12,DKXY02], both requiring *three* group elements.

Concretely, let G be a group with prime order $|G| = q$. Recall that the *Decisional Diffie-Hellman (DDH)* assumption demands that the distributions (g, g^a, g^b, g^{ab}) (for $g \xleftarrow{\$} G$, $a, b \xleftarrow{\$} \mathbb{Z}_q$) and (g, g^a, g^b, g^c) (where $c \xleftarrow{\$} \mathbb{Z}_q$) are computationally indistinguishable. For the same group G, the *ElGamal cryptosystem* has as a public parameter an element $g \xleftarrow{\$} G$, secret key $\mathsf{sk} \xleftarrow{\$} \mathbb{Z}_q$, and public key $\mathsf{pk} = g^{sk}$. For a message $m \in G$, the encryption algorithm is $\mathsf{Enc}(\mathsf{pk}, m) = (g^r, m \cdot \mathsf{pk}^r)$, where $r \xleftarrow{\$} Z_q$, whereas $\mathsf{Dec}(\mathsf{sk}, (c_1, c_2)) = c_2 \cdot c_1^{-\mathsf{sk}}$. ElGamal is easily shown to be ind-cpa-secure under the DDH assumption. Moreover, we observe the following two properties of the ElGamal cryptosystem:

1. ElGamal admits a secret-key to public-key homomorphism $\mu : \mathbb{Z}_q \to G$ where $\mu(x) = g^x$, and n-correctness is satisfied for any n.
2. Moreover, it satisfies (perfect) weak n-key malleability: Namely, just consider the algorithm that for all $I \subseteq [n]$, $i \in I$, and secret- and public-key vectors $\mathbf{sk}$ and $\mathbf{pk}$, outputs

$$c^* = \mathsf{Simulate}(i, I, \mathbf{pk}, \mathbf{sk}[[n] \setminus \{i\}], (c_1, c_2)) = (c_1, c_2 \cdot c_1^{\sum_{j \neq i} \mathbf{sk}[j]}) \ . \tag{1}$$

In particular, the resulting IBE scheme with identities $\mathcal{ID}$ obtained by plugging ElGamal into the GLW construction, for any $(t+1, s)$-cover-free map $\phi : \mathcal{ID} \to \{0,1\}^n$, is as follows, and Theorem 3 implies its t-ibe-cpa security under the DDH assumption. (The decryption algorithm remains the same as in the original ElGamal scheme.)

$\underline{\mathsf{IBEGen}}$:	$\underline{\mathsf{IBEExtract}(\mathsf{msk} = \mathbf{sk}, \mathrm{ID})}$:	$\underline{\mathsf{IBEEnc}(\mathsf{pp} = (g, \mathbf{pk}), \mathrm{ID}, m)}$:
$g \xleftarrow{\$} G$	$[\mathsf{id}_1, \ldots, \mathsf{id}_n] \leftarrow \phi(\mathrm{ID})$	$[\mathsf{id}_1, \ldots, \mathsf{id}_n] \leftarrow \phi(\mathrm{ID})$
$\mathbf{sk} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{pk}[i] \leftarrow g^{\mathbf{sk}[i]}$	$\mathsf{sk}_{\mathrm{ID}} \leftarrow \sum_{i=1}^n \mathsf{id}_i \cdot \mathbf{sk}[i]$	$r \xleftarrow{\$} \mathbb{Z}_q$
$\mathsf{pp} \leftarrow (g, \mathbf{pk})$, $\mathsf{msk} \leftarrow \mathbf{sk}$	Return $\mathsf{sk}_{\mathrm{ID}}$	$c \leftarrow (g^r, m \cdot \prod_{i=1}^n \mathbf{pk}[i]^{r \cdot \mathsf{id}_i})$
Return $(\mathsf{pp}, \mathsf{msk})$		Return c

3.5 Instantiation from LWE

We now turn to a somewhat more involved example based on the GPV cryptosystem [GPV08]. For ease of exposition, we omit a too-detailed discussion of parameters in the following.

The LWE assumption. Let us first recall the *learning with errors (LWE)* problem, introduced by Regev [Reg05]. Let n, q be parameters. For any noise distribution χ on $\mathbb{Z}_q$, and vector $\mathbf{s} \in \mathbb{Z}_q^n$, the oracle $\mathsf{LWE}_{q,n,\chi}(\mathbf{s})$ samples a fresh random n-dimensional vector $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n$, as well as noise $e \xleftarrow{\$} \chi$, and returns $(\mathbf{a}, \mathbf{a}^T\mathbf{s} + e)$. The *LWE assumption* with noise χ states that for every PPT distinguisher D,

$$\mathsf{P}\left[\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n : D^{\mathsf{LWE}_{q,n,\chi}(\mathbf{s})} = 1\right] - \mathsf{P}\left[\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n : D^{\mathsf{LWE}_{q,n,U}(\mathbf{s})} = 1\right] = \mathsf{negl}(n) \ , \tag{2}$$

where U is the uniform distribution on $\mathbb{Z}_q$. (In other words, in addition to $\mathbf{a}$, the oracle $\mathsf{LWE}_{q,n,U}(\mathbf{s})$ simply returns uniform random samples, independent of $\mathbf{s}$.) In general, the error distribution χ is chosen to be a discrete Gaussian on $\mathbb{Z}_q$.

The GPV Cryptosystem. The following is a variant of the cryptosystem suggested by Gentry, Peikert, and Vaikuntanathan [GPV08] (and is in fact the dual of Regev's PKE [Reg05]): For a public random parameter $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, where $m \geq n \cdot \log q + \omega(\log n)$, the secret key $\mathsf{sk} = \mathbf{s}$ consists of a vector $\mathbf{s} \xleftarrow{\$} \{0,1\}^n$ (i.e., the secret key is selected among *binary* vectors), whereas the public key is $\mathsf{pk} = \mathbf{As}$. Then, encryption of a message $b \in \{0,1\}$ is by first computing $\mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^m$, $\mathbf{e} \xleftarrow{\$} \chi^n$, $e' \xleftarrow{\$} \chi'$ (for some distribution χ' related to χ and to be specified below), and then outputting the ciphertext

$$\mathbf{c} = (\mathbf{r}^T\mathbf{A} + \mathbf{e}^T, \mathbf{r}^T\mathsf{pk} + e' + b \cdot (q-1)/2) \ .$$

For decryption of a ciphertext $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2)$ one takes the secret key $\mathbf{s}$, compute $\mathbf{c}_2 - \mathbf{c}_1\mathbf{s}$, and checks whether the outcome is closer to 0 or $(q-1)/2$ (modulo q). The noise distribution must guarantee that this is indeed true with overwhelming probability. We omit the details of this discussion, but this essentially accounts to showing that $e' - \mathbf{e}^T\mathbf{s}$ is not too large.

One can show the above cryptosystem to be secure if the LWE assumption with distributions χ and χ' is true. (Note that the original GPV cryposystem has $\chi = \chi'$, but this distinction will be necessary for our analysis below.)

Secret-key to public-key homomorphism. In order to build an IBE scheme via the GLW-construction, we first observe that the cryptosystem admits a secret-key to public-key homomorphism $\mu : \mathbb{Z}_q^n \to \mathbb{Z}_q^m$ such that $\mu(\mathbf{s}) = \mathbf{As}$. Note that for any two valid secret keys $\mathsf{sk}, \mathsf{sk}' \in \{0,1\}^n$ it is *not* necessarily true that $\mathsf{sk} + \mathsf{sk}'$ is still a valid secret key. However, for any ℓ, it is still true that μ satisfies ℓ-correctness as long as χ and χ' are appropriately bounded.

Weak multi-key-malleability. For weak ℓ-key malleability, we specify the algorithm $\mathsf{Simulate}$ such that, for all $I \subseteq [\ell]$, $i \in I$, $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2)$, $\mathbf{sk} = (\mathbf{s}_1, , ..., \mathbf{s}_\ell)$ and $\mathbf{pk} = [\mathbf{As}_1, ..., \mathbf{As}_\ell]$, we have

$$\mathbf{c} = \mathsf{Simulate}(i, I, \mathbf{pk}, \mathbf{sk}[[\ell] \setminus \{i\}], (\mathbf{c}_1, \mathbf{c}_2)) = \left(\mathbf{c}_1, \mathbf{c}_2 + \mathbf{c}_1 \cdot \sum_{j \in I \setminus \{i\}} \mathbf{s}_j \right) .$$

Note that in contrast to the above DDH-based example, simulation is not perfect. Indeed, the output of $\mathsf{Simulate}$ can indeed be rewritten as

$$\overline{\mathbf{c}}_1 = \left(\mathbf{r}^T \mathbf{A} + \mathbf{e}^T, \mathbf{r}^T \mathbf{A} \sum_{j \in I} \mathbf{s}_j + \mathbf{e}^T \sum_{j \in I \setminus \{i\}} \mathbf{s}_j + e' + b(q-1)/2) \right) .$$

whereas the term $\mathbf{e}^T \sum_{j \in I \setminus \{i\}} \mathbf{s}_j$ is missing in the real ciphertext $\overline{\mathbf{c}}_0$.

However, statistical indistinguishability of $(\mathbf{pk}, \mathbf{sk}[[\ell] \setminus \{i\}], \overline{\mathbf{c}}_b)$ for $b = 0, 1$ is achieved by choosing χ' to be a distribution with a much larger variance than χ. If elements sampled by χ are bounded by B with overwhelming probability, then $\mathbf{e}^T \sum_{j \in I \setminus \{i\}} \mathbf{s}_j$ is at most $|I| \cdot n \cdot B$. We know that if χ' is a discrete Gaussian distribution with standard deviation βq, then the statistical distance of χ' and $\chi' + \mathbf{e}^T \sum_{j \in I \setminus \{i\}} \mathbf{s}_j$ is at most $|I| \cdot n \cdot B/(\beta q)$ [DGK+10, Lemma 3]. Thus, we wish to choose q large enough such that this factor is negligible, yet the LWE problem with distributions χ and χ' is still hard. If we choose $q = 2^{n^\varepsilon}$, $\beta = 2^{n^{-\varepsilon/2}}$, and $B = 2^{n^{\varepsilon/4}}$ for some constant $\varepsilon > 0$, and $|I| = \mathsf{poly}(n)$, we can make the statistical distance smaller than any inverse polynomial in n while retaining subexponential hardness in the LWE assumption. We can thus reduce the security of this PKE scheme to the hardness of subexponential approximations of certain lattice problems [Pei09].

The final LWE-based IBE Scheme. Consequently, every $(s, t+1)$-cover-free map $\phi : \mathcal{ID} \to \{0,1\}^\ell$, every $n \geq m \log q + \omega(n)$, and noise distributions χ, χ' as above yield the following scheme with identity set $\mathcal{ID}$, which, by Theorem 3, is t-ibe-cpa secure under the LWE assumption for distribution χ:

$\underline{\mathsf{IBEGen}}$:	$\underline{\mathsf{IBEExtract}(\mathsf{msk} = \mathbf{sk}, \mathrm{ID})}$:	$\underline{\mathsf{IBEEnc}(\mathsf{pp} = (\mathbf{A}, \mathbf{pk}), \mathrm{ID}, b)}$:
$\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$	$[\mathsf{id}_1, \ldots, \mathsf{id}_\ell] \leftarrow \phi(\mathrm{ID})$	$[\mathsf{id}_1, \ldots, \mathsf{id}_n] \leftarrow \phi(\mathrm{ID})$
$\mathbf{sk}[1], \ldots, \mathbf{sk}[\ell] \xleftarrow{\$} \mathbb{Z}_q^n$	$\mathsf{sk}_{\mathrm{ID}} \leftarrow \sum_{i=1}^{\ell} \mathsf{id}_i \cdot \mathbf{sk}[i]$	$\mathsf{pk} = \sum_{i=1}^{\ell} \mathsf{id}_i \cdot \mathbf{pk}[i]$
$\mathbf{pk}[i] \leftarrow \mathbf{A}\,\mathbf{sk}[i]$ $(i = 1, \ldots, \ell)$	Return $\mathsf{sk}_{\mathrm{ID}}$	$\mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^m$, $\mathbf{e} \xleftarrow{\$} \chi^n$, $e' \xleftarrow{\$} \chi'$
$\mathsf{pp} \leftarrow (\mathbf{A}, \mathbf{pk})$, $\mathsf{msk} \leftarrow \mathbf{sk}$		$\mathbf{c}_1 \leftarrow \mathbf{r}^T \cdot \mathbf{A} + \mathbf{e}$
Return $(\mathsf{pp}, \mathsf{msk})$		$\mathbf{c}_2 \leftarrow \mathbf{r}^T \mathsf{pk} + e' + b \cdot (q-1)/2$
		Return c

Also, we note that if we are only interested in proving selective security using Theorem 2, then weak ℓ-key malleability is unnecessary, and we can fix $\chi = \chi'$. This allows to choose a polynomial modulus, avoiding the subexponential LWE assumption. We omit more detailed discussion of parameters, which follows from standard techniques.

4 Construction from Multi-Key Malleability

4.1 Bounded-IBE Construction

We present a further construction of BC-IBE from PKE schemes which satisfy a different notion of key malleability than the one given above, which we first introduce. Our notion requires that

given an encryption of a message under one public key, we are asking for the ability to produce a new ciphertext of the same message which decrypts under a combination of secret keys (e.g., the product) for which we only know the corresponding public keys. Note that we are only asking for decryptability under the combination of the secret keys. In particular, in contrast to the above notion of weak key-malleability, the distribution of the resulting ciphertext may not be a valid encryption under some well-defined combination of the corresponding public keys, and moreover, we require the ability to compute this ciphertext without knowledge of any secret keys.

Definition 4 (Multi-Key Malleability). *Let* PKE *be a public-key encryption scheme. We say that* PKE *is* n-key malleable *if there exist algorithms* Modify *and* $\mathsf{Combine}$ *such that the following properties hold:*

(i) *For all valid messages* m*, all* $I \subseteq [n]$*, and all* $i \in I$*, the following probability is negligible (taken over the coins of* Enc*):*

$$\mathsf{P}\left[(\mathbf{pk}, \mathbf{sk}) \xleftarrow{\$} \mathsf{Gen}^n, c \xleftarrow{\$} \mathsf{Enc}(\mathbf{pk}[i], m), c' \xleftarrow{\$} \mathsf{Modify}(i, I, \mathbf{pk}, c) : \mathsf{Dec}(\mathsf{Combine}(I, \mathbf{sk}), c') \neq m\right] .$$

(ii) *For all* $I \subseteq [n]$, $\mathsf{Combine}(I, \mathbf{sk})$ *does* not *depend on* $\mathbf{sk}[i]$ *for* $i \notin I$.

(iii) *For all* $I \subseteq [n]$ *and all valid public-key / secret-key vectors* $(\mathbf{pk}, \mathbf{sk})$*, for all* $i, j \in I$*, the values* $\mathsf{Modify}(i, I, \mathbf{pk}, \mathsf{Enc}(\mathbf{pk}[i], m))$ *and* $\mathsf{Modify}(j, I, \mathbf{pk}, \mathsf{Enc}(\mathbf{pk}[j], m))$ *are equally distributed.*

We note that Property **(iii)** above is not really necessary (a computational relaxation would suffice), but will make the presentation somewhat simpler and is true in the only instantiation we give below.

The IBE construction and its security. For an identity map $\phi : \mathcal{ID} \to \{0,1\}^n$, we now propose a construction of an identity-based encryption scheme $\mathsf{IBE} = (\mathsf{IBEGen}, \mathsf{IBEExtract}, \mathsf{IBEEnc}, \mathsf{IBEDec})$ from an n-key malleable encryption scheme $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. The decryption algorithm is unaltered, i.e., $\mathsf{IBEDec} = \mathsf{Dec}$, and moreover the construction consists of the following algorithms. (Note that the choice of i as $\min\{\phi(\mathrm{ID})\}$ below within IBEEnc is purely arbitrary.)

$\underline{\mathsf{IBEGen}:}$	$\underline{\mathsf{IBEExtract}(\mathsf{msk} = \mathbf{sk}, \mathrm{ID}):}$	$\underline{\mathsf{IBEEnc}(\mathsf{pp} = \mathbf{pk}, \mathrm{ID}, m):}$
$(\mathbf{pk}, \mathbf{sk}) \xleftarrow{\$} \mathsf{Gen}^n$	$\mathsf{sk}_{\mathrm{ID}} \leftarrow \mathsf{Combine}(\phi(\mathrm{ID}), \mathbf{sk})$	$i \leftarrow \min\{\phi(\mathrm{ID})\}$
$\mathsf{msk} \leftarrow \mathbf{sk}$	Return $\mathsf{sk}_{\mathrm{ID}}$	$c' \xleftarrow{\$} \mathsf{Enc}(\mathbf{pk}[i], m)$
$\mathsf{pp} \leftarrow \mathbf{pk}$		$c \leftarrow \mathsf{Modify}(i, \phi(\mathrm{ID}), \mathbf{pk}, c')$
Return $(\mathsf{msk}, \mathsf{pp})$		Return c

Correctness of the scheme follows by Property **(i)** above. The following theorem establishes security of our new construction. The proof is deferred to Appendix B.2.

Theorem 4. *Assume that* PKE *is ind-cpa-secure and* n*-key malleable, and that* ϕ *is* $(t+1, s)$*-cover free. Then,* IBE *is* t*-ibe-cpa-secure.*

4.2 NTRU Based Instantiation and Fully-Homomorphic IBE

We provide an instantiation of the above constructing using the multi-key homomorphic properties of NTRU-based public-key encryption [LATV12], which we first review. For some parameters r, n and q (where q is a prime), consider the ring of polynomials $R = \mathbb{Z}[x]/(x^r + 1)$, and let χ be a B-bounded distribution on R, i.e., with overwhelming probability, χ samples a polynomial from R whose coefficients are all at most B in absolute value. All operations on polynomials are to be understood as over the ring $R_q = R/qR$. The NTRU cryptosystem is such that key generation Gen samples $f, g \xleftarrow{\$} \chi$ subject to the constraint that $f \equiv 1 \pmod 2$, and sets $\mathsf{pk} = 2g/f$ and $\mathsf{sk} = f$.

(Possibly, f needs to be resampled until it admits an inverse in R_q, and χ is such that this happens with good probability.) The message $b \in \{0,1\}$ is encrypted as

$$\mathsf{Enc}(\mathsf{pk}, m) = h \cdot \mathsf{pk} + 2e + b \;,$$

where $h, e \xleftarrow{\$} \chi$. Finally, decryption, given c, outputs $\mathsf{Dec}(\mathsf{sk}, c) = \mathsf{sk} \cdot c \pmod 2$. To see why decryption is correct, note that

$$\mathsf{sk} \cdot c \equiv f \cdot (2h \cdot g/f + 2e + b) \equiv 2h \cdot g + 2e \cdot f + f \cdot b \pmod q \;.$$

If $B \leq \sqrt{q/2}/r$, then all coefficients from $h \cdot g$ and $e \cdot f$ are of size at most $r^2 B^2 < q/2$. Consequently, $2hg$ and $2ef$ only have even coefficients, and are 0 modulo 2. And finally, $f \cdot b$ clearly always equals b modulo 2.

The scheme was proven ind-cpa-secure under a fairly ad-hoc assumption in [LATV12], where it was also shown to have strong homomorphic properties we address below, and which we exploit for our construction.

The IBE Scheme. We turn now to building an IBE scheme from the above NTRU-based PKE scheme PKE using the above generic approach. In the following, we assume that r is our security parameter, $q = 2^{n^\varepsilon}$ for some constant $\varepsilon < 1$, $B = \mathsf{poly}(r)$, and $n = \Theta(r^\delta)$ for some constant $\delta < 1$.

We first show ℓ-key malleability exploiting the multi-key homomorphic properties of NTRU shown in [LATV12]. To this end, we define the algorithm $\mathsf{Combine}$ which given $I \subseteq [\ell]$ and $\mathbf{sk} \in R_q^\ell$ outputs

$$\mathsf{Combine}(I, \mathbf{sk}) = \prod_{i \in I} \mathbf{sk}[i] \;.$$

Moreover, we also define the (randomized) function Modify, which given $I \subseteq [\ell]$, $i \in I$, $c \in R_q$, and $\mathbf{pk} \in R_q^\ell$, outputs

$$\mathsf{Modify}(i, I, c, \mathbf{pk}) = c + \sum_{j \in I \setminus \{i\}} h_j \cdot \mathbf{pk}[j] \;,$$

where h_j for $j \in I \setminus \{i\}$ are sampled independently from the B-bounded distribution χ as above. Now, Properties (ii) and (iii) in Definition 4 are immediate to verify. Moreover, for Property (i), fix $I \subseteq [\ell]$ and $i \in I$, and $\mathbf{pk}, \mathbf{sk} \in R_q^\ell$, each consisting of ℓ B-bounded polynomials as components, then define c as

$$c = \mathsf{Modify}(i, I, \mathsf{Enc}(\mathbf{pk}[i], b), \mathbf{pk}) = \sum_{j \in I} h_j \cdot \mathbf{pk}[j] + 2e + b \;,$$

and observe that

$$\mathsf{Dec}(\mathsf{Combine}(I, \mathbf{sk}), c) = \left(\prod_{i \in I} \mathbf{sk}[i]\right) \cdot \left(\sum_{j \in I} h_j \cdot \mathbf{pk}[j] + 2e + b\right) \pmod 2 \;.$$

In particular,

$$\left(\prod_{i \in I} \mathbf{sk}[i]\right) \cdot \left(\sum_{j \in I} h_j \cdot \mathbf{pk}[j] + 2e + b\right) \equiv \sum_{j \in I} 2h_j \cdot g_j \cdot \prod_{i \in I \setminus \{j\}} f_\ell + \left(2e \cdot \prod_{i \in I} f_\ell\right) + b \cdot \left(\prod_{i \in I} f_\ell\right) \;.$$

Note that in the above sum, only products of at most $|I| + 1$ B-bounded polynomials occurs. The coefficients of the resulting products have size at most $r^{|I|} \cdot B^{|I|+1}$, which (given previous parameter choices) is smaller than $q/2$ as long as $|I| = o(n^\varepsilon / \log n)$. This yields correct decryption as no wraparound (modulo q) occurs.

THE FINAL SCHEME. Overall, this yields to the following scheme, for any identity mapping $\phi : \mathcal{ID} \to \{0,1\}^\ell$ which is $(s, t+1)$-cover-free for some $s = o(n^\varepsilon)$, which is t-ind-cpa secure by Theorem 4.

IBEGen:	IBEExtract($\mathsf{msk} = \mathbf{sk}, \mathrm{ID}$):	IBEEnc($\mathsf{pp} = \mathbf{pk}, \mathrm{ID}, m$):
$f_1, \ldots, f_n \xleftarrow{\$} \chi$, $f_i \equiv 1 \pmod 2$, $f_i \in R_q^*$ $g_1, \ldots, g_n \xleftarrow{\$} \chi$ $\mathsf{msk} \leftarrow (f_1, \ldots, f_n)$ $\mathsf{pp} \xleftarrow{\$} (2g_1/f_1, \ldots, 2g_n/f_n)$ Return $(\mathsf{msk}, \mathsf{pp})$	$\mathsf{sk}_{\mathrm{ID}} \leftarrow \prod_{i \in \phi(\mathrm{ID})} \mathbf{sk}[i]$ Return $\mathsf{sk}_{\mathrm{ID}}$	$h_1, \ldots, h_n, e \xleftarrow{\$} \chi$ $c \leftarrow \sum_{i \in \phi(\mathrm{ID})} \mathbf{pk}[i] \cdot h_i + 2e + m$ Return c

FULLY-HOMOMORPHIC IBE. The above instantiation has additionally the property of being fully-homomorphic in the following sense: Given encryptions $\mathsf{IBEEnc}(\mathrm{ID}, m_1), \ldots, \mathsf{IBEEnc}(\mathrm{ID}, m_t)$, and a function $f : \{0,1\}^t \to \{0,1\}$, we can compute a ciphertext which decrypts to $f(m_1, \ldots, m_t)$ under $\mathsf{sk}_{\mathrm{ID}}$ using the homomorphic-evaluation procedures given in [LATV12].

We note that in general one can provide a construction, along the lines given above, from multi-key fully-homomorphic encryption to fully-homomorphic identity-based encryption for bounded collusions. We defer a full discussion, noting in passing that the above is the only instantiation of this paradigm we are aware of.

5 Applications: Bounded CCA Security with Short Ciphertexts

In this section, we revisit the generic transform by Boneh, Canetti, Halevi, and Katz [BCHK07] in the context of BC-IBE, and use it to obtain constructions of bounded-CCA2 secure encryption schemes with short ciphertexts from any semantically secure scheme with a secret-key to public-key homomorphism.

ONE-TIME SIGNATURES. Recall that a one-time signature scheme $\mathsf{SS} = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ consists of a parameter generator algorithm Gen, outputting a pair consisting of the signing key sk and the verification key vk, and a signing algorithm Sign and verification algorithm Verify such that for any message m, Sign returns a signature $\sigma = \mathsf{Sign}(\mathsf{sk}, m)$ such that $\mathsf{Verify}(\mathsf{vk}, \sigma) = \mathtt{true}$ with overwhelming probability for the associated verification key vk. We say that SS is *strongly one-time secure* if for all PPT adversaries $\mathcal{A}$, there exists a negligible function ν such that

$$\mathsf{P}\left[\begin{array}{l}(\mathsf{sk}, \mathsf{vk}) \xleftarrow{\$} \mathsf{Gen}, (m, \mathsf{st}) \xleftarrow{\$} \mathcal{A}(\mathsf{vk}), \\ \sigma = \mathsf{Sign}(\mathsf{sk}, m), (m', \sigma') \xleftarrow{\$} \mathcal{A}(\mathsf{st}, \sigma)\end{array} : (m, \sigma) \neq (m', \sigma') \wedge \mathsf{Verify}(\mathsf{vk}, m', \sigma') = \mathtt{true}\right] \leq \nu(k)\ ,$$

where k is the corresponding (implicit) security parameter.

THE BCHK TRANSFORM . In the following, let $\mathsf{IBE} = (\mathsf{IBEGen}, \mathsf{IBEExtract}, \mathsf{IBEEnc}, \mathsf{IBEDec})$ be an IBE scheme and let $\mathsf{SS} = (\mathsf{Gen}_{\mathsf{SS}}, \mathsf{Sign}, \mathsf{Verify})$ be a strong signature scheme. Boneh et al [BCHK07] presented the following construction of an encryption scheme $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ from IBE:

Gen:	Enc(pk, m):	Dec($\mathsf{sk}, (\mathsf{vk}, c, \sigma)$):

$(\mathsf{pp}, \mathsf{msk}) \xleftarrow{\$} \mathsf{IBEGen}$
$\mathsf{pk} \leftarrow \mathsf{pp}$
$\mathsf{sk} \leftarrow \mathsf{msk}$
Return $(\mathsf{pk}, \mathsf{sk})$.
	$(\mathsf{sk}', \mathsf{vk}') \xleftarrow{\$} \mathsf{Gen}_{\mathsf{SS}}$
	$c \xleftarrow{\$} \mathsf{IBEEnc}(\mathsf{pp}, \mathsf{vk}', m)$
	$\sigma \xleftarrow{\$} \mathsf{Sign}(\mathsf{sk}', c)$
	Return $(\mathsf{vk}', c, \sigma)$.
		If $\mathsf{Verify}(\mathsf{vk}, c, \sigma) = \mathtt{false}$ then
		$m \leftarrow \perp$
		Else
		$\mathsf{sk}_{\mathsf{vk}} \leftarrow \mathsf{IBEExtract}(\mathsf{sk}, \mathsf{vk})$
		$m \leftarrow \mathsf{IBEDec}(\mathsf{sk}_{\mathsf{vk}}, c)$
		Return m

Chosen-ciphertext security of this construction was proven in [BCHK07, Theorem 1]. Their proof considers two cases: in order to defeat the CCA security of the above construction, an adversary must either forge a signature for SS or defeat the selective security of IBE. They thus provide reductions in two different security games: one with a signature oracle (for the former case), and one with an IBE challenger (for the latter).

Of note is that in their reduction for the IBE case, the reduction makes at most one IBEExtract query for each decryption query it receives from the adversary, and no other parameters change. Thus, their proof carries through exactly in the bounded-collusion case, yielding:

Theorem 5. *If* IBE *is t-selective-ibe-cpa-secure, and if* SS *is strongly one-time secure, then* PKE *is t-CCA secure.*

Applications. Using previous results, we directly obtain bounded CCA PKE constructions from DDH, QR, NTRU, and (standard) LWE using the constructions of the previous sections. In particular, note that only standard LWE is required as we only need selective security to instantiate the above paradigm. Moreover, the resulting DDH construction is essentially equivalent to the one presented in [CHH+07], and our construction thus provides an abstraction to obtain the same construction.

As an example, we give the t-CCA PKE based on the NTRU assumption that comes from applying Theorem 5 to the BC-IBE of Section 4.2. (Here the parameters q, χ, R_q^* are defined as in that section.)

Gen:	Enc(pk, m):	Dec($\mathsf{sk}, (\mathsf{vk}, c, \sigma)$):

$f_1, \ldots, f_n \xleftarrow{\$} \chi$, $f_i \equiv 1 \pmod 2$, $f_i \in R_q^*$
$g_1, \ldots, g_n \xleftarrow{\$} \chi$
$\mathsf{sk} \leftarrow (f_1, \ldots, f_n)$
$\mathsf{pk} \leftarrow (2g_1/f_1, \ldots, 2g_n/f_n)$
Return $(\mathsf{pk}, \mathsf{sk})$.
	$(\mathsf{sk}_{\mathsf{SS}}, \mathsf{vk}_{\mathsf{SS}}) \xleftarrow{\$} \mathsf{Gen}_{\mathsf{SS}}$
	$h_1, \ldots, h_n, e \xleftarrow{\$} \chi$
	$c \leftarrow \sum_{i \in \phi(\mathsf{vk}_{\mathsf{SS}})} \mathbf{pk}[i] \cdot h_i + 2e + m$
	$\sigma \xleftarrow{\$} \mathsf{Sign}(\mathsf{sk}_{\mathsf{SS}}, c)$
	Return $(\mathsf{vk}_{\mathsf{SS}}, c, \sigma)$.
		If $\mathsf{Verify}(\mathsf{vk}, c, \sigma) = \texttt{false}$ then
		$m \leftarrow \perp$
		Else
		$\mathsf{sk}_{\mathsf{vk}} \leftarrow \prod_{i \in \phi(\mathsf{vk})} \mathbf{sk}[i]$
		$m \leftarrow \mathsf{sk}_{\mathsf{vk}} \cdot c \pmod 2$
		Return m

The ciphertext size of the CCA scheme generated by the BCHK transform is the same as the ciphertext size of the IBE scheme (and hence of the NTRU encryption scheme), plus a verification key and signature. Steinfeld *et al.* [SLP+12] show a (fully) CCA-secure construction based on NTRU; their ciphertext contains k ciphertexts of the underlying NTRUEncrypt algorithm (where $k = \Theta(1)$ is a parameter that depends on the hardness assumption used, but is at least 4), and additionally a verification key, a signature, and a blinded message. (Since the NTRUEncrypt ciphertexts are polynomials in R_q, they will typically be much larger than the other values.) Thus, we obtain a constant-factor improvement in ciphertext size by moving to the bounded-query model, in addition to the conceptual simplicity of the proof.

Acknowledgments

The authors wish to thank Shafi Goldwasser for insightful feedback and motivating us to write the present paper.

The research of this paper was partially supported by NSF Contract CCF-1018064. Moreover, this material is based on research sponsored by DARPA under agreement numbers FA8750-11-C-0096 and FA8750-11-2-0225. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

ABB10. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 98–115. Springer, August 2010.

BB04. Dan Boneh and Xavier Boyen. Secure identity based encryption without random oracles. In Matthew Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 443–459. Springer, August 2004.

BCHK07. Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM Journal on Computing*, 36(5):1301–1328, 2007.

BDPR98. Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. Relations among notions of security for public-key encryption schemes. In Hugo Krawczyk, editor, *Advances in Cryptology – CRYPTO'98*, volume 1462 of *Lecture Notes in Computer Science*, pages 26–45. Springer, August 1998.

BF01. Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 213–229. Springer, August 2001.

BGH07. Dan Boneh, Craig Gentry, and Michael Hamburg. Space-efficient identity based encryption without pairings. In *48th Annual Symposium on Foundations of Computer Science*, pages 647–657. IEEE Computer Society Press, October 2007.

CHH+07. Ronald Cramer, Goichiro Hanaoka, Dennis Hofheinz, Hideki Imai, Eike Kiltz, Rafael Pass, Abhi Shelat, and Vinod Vaikuntanathan. Bounded CCA2-secure encryption. In Kaoru Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007*, volume 4833 of *Lecture Notes in Computer Science*, pages 502–518. Springer, December 2007.

CHK03. Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 255–271. Springer, May 2003.

CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 523–552. Springer, May 2010.

Coc01. Clifford Cocks. An identity based encryption scheme based on quadratic residues. In Bahram Honary, editor, *8th IMA International Conference on Cryptography and Coding*, volume 2260 of *Lecture Notes in Computer Science*, pages 360–363. Springer, December 2001.

CS02. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 45–64. Springer, April / May 2002.

DF03. Yevgeniy Dodis and Nelly Fazio. Public key broadcast encryption for stateless receivers. In Joan Feigenbaum, editor, *Digital Rights Management*, volume 2696 of *Lecture Notes in Computer Science*, pages 61–80. Springer Berlin Heidelberg, 2003.

DGK+10. Yevgeniy Dodis, Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Public-key encryption schemes with auxiliary inputs. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 361–381. Springer, February 2010.

DH76. Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.

DHT12. Yevgeniy Dodis, Iftach Haitner, and Aris Tentes. On the instantiability of hash-and-sign RSA signatures. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 112–132. Springer, March 2012.

DKXY02. Yevgeniy Dodis, Jonathan Katz, Shouhuai Xu, and Moti Yung. Key-insulated public key cryptosystems. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 65–82. Springer, April / May 2002.

EFF85. P. Erdös, P. Frankel, and Z. Furedi. Families of finite sets in which no set is covered by the union of r others. *Israeli Journal of Mathematics*, 51:79–89, 1985.

FN94. Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO'93*, volume 773 of *Lecture Notes in Computer Science*, pages 480–491. Springer, August 1994.

GLW12. Shafi Goldwasser, Allison B. Lewko, and David A. Wilson. Bounded-collusion IBE from key homomorphism. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 564–581. Springer, March 2012.

GM84. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.

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 Annual ACM Symposium on Theory of Computing*, pages 197–206. ACM Press, May 2008.

GSW00. Juan A. Garay, Jessica Staddon, and Avishai Wool. Long-lived broadcast encryption. In Mihir Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 333–352. Springer, August 2000.

GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. Cryptology ePrint Archive, Report 2013/340, 2013. http://eprint.iacr.org/.

GSY99. Eli Gafni, Jessica Staddon, and Yiqun Lisa Yin. Efficient methods for integrating traceability and broadcast encryption. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 372–387. Springer, August 1999.

GVW12. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 162–179. Springer, August 2012.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. Ntru: A ring-based public key cryptosystem. In *Lecture Notes in Computer Science*, pages 267–288. Springer-Verlag, 1998.

HS02. Dani Halevy and Adi Shamir. The LSD broadcast encryption scheme. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 47–60. Springer, August 2002.

KRS99. Ravi Kumar, Sridhar Rajagopalan, and Amit Sahai. Coding constructions for blacklisting problems without computational assumptions. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 609–623. Springer, August 1999.

LATV12. Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In Howard J.

Karloff and Toniann Pitassi, editors, *44th Annual ACM Symposium on Theory of Computing*, pages 1219–1234. ACM Press, May 2012.

NNL01. Dalit Naor, Moni Naor, and Jeffery Lotspiech. Revocation and tracing schemes for stateless receivers. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 41–62. Springer, August 2001.

Pei09. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Computing*, pages 333–342. ACM Press, May / June 2009.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93. ACM Press, May 2005.

Sha85. Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *Advances in Cryptology – CRYPTO'84*, volume 196 of *Lecture Notes in Computer Science*, pages 47–53. Springer, August 1985.

SLP+12. Ron Steinfeld, San Ling, Josef Pieprzyk, Christophe Tartary, and Huaxiong Wang. NTRUCCA: How to strengthen NTRUEncrypt to chosen-ciphertext security in the standard model. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012: 15th International Workshop on Theory and Practice in Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 353–371. Springer, May 2012.

Wat05. Brent R. Waters. Efficient identity-based encryption without random oracles. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 114–127. Springer, May 2005.

A Instantiation from Quadratic Residuosity

We now discuss an instantiation from the QR assumption. This application does not achieve any improvement with respect to parameter size of the construction of [GLW12]. This is mostly due to the fact that the best semantically secure QR-based scheme with key-homomorphism is a variation of the scheme given in [GLW12]. However, as we only need semantic security, the corresponding security proof is simpler and allows the underlying PKE scheme to have shorter secret keys.

The Quadratic Residuosity assumption. The *Quadratic Residuosity assumption (QR) modulo* N states that it is compuationally infeasible, given composite number N and a uniform element y of Jacobi symbol 1 modulo N, to determine whether there exists an x such that $x^2 \equiv y \pmod N$ with probability nonnegligibly better than $1/2$. We will let $\mathbb{QR}_N$ denote the group of quadratic residues modulo N. Note that if N is a Blum integer (that is, if $N = pq$ where primes $p, q \equiv 3 \pmod 4$), then -1 has Jacobi symbol 1 modulo N, but $-1 \notin \mathbb{QR}_N$.

A QR-based scheme. We consider the QR-based PKE of Goldwasser *et al.* [GLW12], which is defined as follows:

$\underline{\mathsf{Gen}}$:	$\underline{\mathsf{Enc}(\mathsf{pp} = (N, g), \mathsf{pk}, m \in \{1, -1\})}$:	$\underline{\mathsf{Dec}(c = (c_1, c_2), \mathsf{sk})}$:
$p, q \xleftarrow{\$} \mathbb{N}$ s.t. p, q are prime, $p \equiv q \equiv 3 \pmod 4$, and $\mathbb{QR}_{pq}$ is cyclic $N \leftarrow pq, g \xleftarrow{\$} \mathbb{QR}_N$ $\mathsf{sk} \xleftarrow{\$} [\rho(N)]$, $\mathsf{pk} \leftarrow g^{\mathsf{sk}} \pmod N$, $\mathsf{pp} \leftarrow (N, g)$ Return $(\mathsf{pp}, \mathsf{pk}, \mathsf{sk})$	$r \xleftarrow{\$} [N^2]$ s.t. r odd Return $(g^r \pmod N, m \cdot \mathsf{pk}^r \pmod N)$	Return $c_2 \cdot (c_1^{\mathsf{sk}})^{-1} \pmod N$

Goldwasser *et al.* prove security under the QR assumption modulo N where $\rho(N) = N^{\Theta(1)}$ [GLW12]. However, they require a stronger linear-hash-proof property, whereas we only require semantic security. Thus, we are able to use $\rho(N) = \varphi(N)/2$ (where φ here is Euler's totient function), and furthermore can provide a shorter, simpler security proof.

Theorem 6. *Assuming the quadratic residuosity problem modulo N is hard, the above PKE scheme, with $\rho(N) = \varphi(N)/2$, is semantically secure.*

Proof. Assume an adversary $\mathcal{A}$ that is able to distinguish encryptions of 1 and -1 under the above scheme with probability $1/2+\varepsilon$, where ε is nonnegligible. We show a reduction to the QR_N problem.

We receive a pair (N, x) where N is a Blum integer and x has Jacobi symbol 1 modulo N. We choose a random quadratic residue g modulo n. Note that $|\mathbb{QR}_N| = \varphi(N)/4$; thus, if we choose $\mathsf{sk} \xleftarrow{\$} [N^2]$, the distribution of $\mathsf{pk} = g^{\mathsf{sk}}$ is statistically close to uniform over $\mathbb{QR}_N$ (note that in the real experiment, since sk is chosen from $[\varphi(N)/2]$, the resulting distribution is exactly uniform). We then choose $b \xleftarrow{\$} \{0, 1\}$ and send pk and the challenge ciphertext $(x, (-1)^b \cdot x^{\mathsf{sk}} \pmod N)$ to $\mathcal{A}$.

If x is a quadratic residue, the ciphertext is a valid ciphertext for message $(-1)^b$. If x is a nonresidue, it is in the group with Jacobi symbol 1, which has size $\varphi(N)/2$. However, information-theoretically, the secret key is only determined from the public key modulo $\varphi(N)/4$. Therefore, no adversary is able to distinguish between this experiment run with sk and this experiment run with $\mathsf{sk} \pm \varphi(N)/4$ (whichever lies in the correct interval).

However, in the case where x is a nonresidue, it can be expressed as $-1 \cdot g^y$ for some y. The resulting decrypted plaintext using secret key $\mathsf{sk} \pm \varphi(N)/4$ is thus multiplied by $(-1)^{\varphi(N)/4} = -1$. Therefore, information-theoretically the adversary cannot distinguish between encryptions of 1 and -1. We can thus use $\mathcal{A}$ in the standard way, guessing that x is a quadratic residue iff $\mathcal{A}$ guesses correctly, and we will be correct with nonnegligible advantage (specifically, half the advantage of $\mathcal{A}$). □

Note that this PKE scheme, though based on the QR assumption, strongly resembles the ElGamal construction.[3] It thus exhibits properties very similar to the previous case: it admits a natural key homomorphism with mapping $x \mapsto g^x \pmod N$, and it is weakly n-key-malleable: for all $I \subseteq [n]$, $i \in I$, and secret- and public-key vectors $\mathbf{sk}$ and $\mathbf{pk}$, we have the $\mathsf{Simulate}$ algorithm

$$c^* = \mathsf{Simulate}(i, I, \mathbf{pk}, \mathbf{sk}[[n] \setminus \{i\}], (c_1, c_2)) = (c_1, c_2 \cdot c_1^{\sum_{j \neq i} \mathbf{sk}[j]} \pmod N) . \tag{3}$$

This yields the following BC-IBE construction, which is t-ind-cpa secure by Theorem 3.

$\underline{\mathsf{IBEGen}}$:	$\underline{\mathsf{IBEExtract}(\mathsf{msk} = \mathbf{sk}, \mathrm{ID})}$:	$\underline{\mathsf{IBEEnc}(\mathsf{pp} = \mathbf{pk}, \mathrm{ID}, m)}$:
$p, q \xleftarrow{\$} \mathbb{N}$ s.t. p, q are prime, $p \equiv q \equiv 3 \pmod 4$, and $\mathbb{QR}_{pq}$ is cyclic	$[\mathsf{id}_1, \ldots, \mathsf{id}_n] \leftarrow \phi(\mathrm{ID})$	$[\mathsf{id}_1, \ldots, \mathsf{id}_n] \leftarrow \phi(\mathrm{ID})$
$N \leftarrow pq, g \xleftarrow{\$} \mathbb{QR}_N$	$\mathsf{sk}_{\mathrm{ID}} \leftarrow \sum_{i=1}^n \mathsf{id}_i \cdot \mathbf{sk}[i]$	$r \xleftarrow{\$} [N^2]$ s.t. r odd
$\mathbf{sk} \xleftarrow{\$} [\rho(N)]^n$, $\mathbf{pk}[i] \leftarrow g^{\mathbf{sk}[i]} \pmod N$	Return $\mathsf{sk}_{\mathrm{ID}}$	$c \leftarrow (g^r \pmod N, m \cdot \prod_{i=1}^n \mathsf{id}_i \cdot \mathbf{pk}[i]^r) \pmod N$
$\mathsf{pp} \leftarrow (N, g, \mathbf{pk}), \mathsf{msk} \leftarrow \mathbf{sk}$		Return c
Return $(\mathsf{pp}, \mathsf{msk})$		

[3] ElGamal relies on the DDH assumption to hide values in the exponent, ensuring that the secret key and encryption randomness are hidden. This scheme does not require that the DDH or discrete logarithm problem be hard, but relies on the fact that both the secret key and the randomness are chosen from a range of size larger than N; thus, expressing these values in the exponent modulo N information-theoretically loses information, thus hiding the message bit.

B Missing Proofs

B.1 Full Security of GLW (Theorem 3)

Proof (Sketch). We only sketch the reduction – the analysis is similar to the one in the proof of Theorem 2. Let $\mathcal{A}$ be an ibe-cpa adversary for IBE which guesses the underlying bit b with probability at least $(1+n\cdot\varepsilon)/2$, and which makes at most t extraction queries. We construct an ind-cpa adversary $\mathcal{B}$ for PKE from $\mathcal{A}$, winning with probability at least $\frac{1+\varepsilon}{2}$. Concretely, the adversary $\mathcal{B}$ starts by choosing an index i^* uniformly at random from the set $[n]$. It then gets a public key pk^* from the underlying IND-CPA game, and computes $(\mathbf{pk}[j], \mathbf{sk}[j]) \xleftarrow{\$} \mathsf{Gen}$ for all $j \in [n]\backslash\{i^*\}$. Finally, it sets $\mathbf{pk}[i^*] = \mathsf{pk}^*$. It then gives $\mathsf{pp} = \mathbf{pk}$ to $\mathcal{A}$ and starts its execution, until it outputs a target identity ID^* as well as a message pair (m_0, m_1). In particular, $\mathcal{A}$'s extraction queries for $\mathrm{ID} \in \{0,1\}$ are replied by computing $[\mathsf{id}_1, \ldots, \mathsf{id}_n] = \phi(\mathrm{ID})$ and, if $\mathsf{id}_{i^*} = 0$, returning $\mathsf{sk}_{\mathrm{ID}} = \sum_i \mathsf{id}_i \mathbf{sk}[i]$. Note that if $\mathsf{id}_{i^*} = 1$, then $\mathcal{B}$ cannot answer the extraction query. In this case, it returns $\perp$, and sets the flag bad to $\mathtt{true}$. Given ID^* and (m_0, m_1), let $\phi(\mathrm{ID}^*) = [\mathsf{id}_1^*, \ldots, \mathsf{id}_n^*]$. If $\mathsf{id}_{i^*}^* = 0$, then bad is set to $\mathtt{true}$, and $\perp$ is returned to $\mathcal{A}$. Otherwise, $\mathcal{B}$ forwards (m_0, m_1) to the ind-cpa game, obtaining a challenge ciphertext c. Given this, it outputs $c^* \xleftarrow{\$} \mathsf{Simulate}(i^*, \phi(\mathrm{ID}^*), c, \mathbf{pk}, \mathbf{sk}[[n] \setminus \{i^*\}])$, which it then gives back to $\mathcal{A}$, continuing it execution, until $\mathcal{A}$ outputs a bit b'. To conclude, $\mathcal{B}$ outputs the bit b' if bad is $\mathtt{false}$, and returns a random bit otherwise. Unlike in Theorem 2, this simulation is not (necessarily) perfect, since we only require computational indistinguishability from $\mathsf{Simulate}$. However, this will only affect $\mathcal{A}$'s output with negligible probability (or else $(\mathcal{B}, \mathcal{A})$ would serve as an efficient distinguisher for the output of $\mathsf{Simulate}$), and thus $\mathcal{B}$'s overall success probability is still nonnegligibly greater than $\frac{1}{2}$. □

B.2 Proof of Theorem 4

Proof. Let $\mathcal{A}$ be an ibe-cpa-adversary for IBE making t extraction queries and which succeeds with probability $\frac{1+n\varepsilon}{2}$ in guessing the right bit b. We build an ind-cpa adversary $\mathcal{B}$ for PKE as follows: The adversary $\mathcal{B}$ initially chooses an index $i^* \xleftarrow{\$} [n]$, and is then given pk^*. It sets $\mathbf{pk}[i^*] = \mathsf{pk}$, and then samples $(\mathbf{pk}[i], \mathbf{sk}[i]) \xleftarrow{\$} \mathsf{Gen}$ for all $i \neq i^*$, and then runs $\mathcal{A}$ with public parameters $\mathsf{pp} = \mathbf{pk}$. Then, $\mathcal{A}$'s extraction queries for ID are simulated as follows: If $i^* \notin \phi(\mathrm{ID})$, then it returns $\mathsf{Combine}(\phi(\mathrm{ID}), \mathbf{sk})$. (Note that this can be done by Property (ii), since the output of $\mathsf{Combine}$ does not depend on $\mathsf{sk}[i^*]$.) Else it returns $\perp$ if $i^* \in \phi(\mathrm{ID})$, and sets a flag bad to $\mathtt{true}$. Moreover, on input a triple $(m_0, m_1, \mathrm{ID}^*)$, $\mathcal{B}$ forwards m_0, m_1 to the ind-cpa game, obtaining $c \xleftarrow{\$} \mathsf{Enc}(\mathsf{pk}^*, m_b)$, and then, if $i^* \in \phi(ID^*)$ and $c \neq \perp$, it gives $c^* = \mathsf{Modify}(i^*, \phi(\mathrm{ID}^*), \mathbf{pk}, c)$ back to $\mathcal{A}$. If $c = \perp$, then it sets $c^* = \perp$. Otherwise, it gives simply $\perp$ back and sets bad to $\mathtt{true}$. Finally, $\mathcal{B}$ outputs $\mathcal{A}$'s final output b' if bad was never set, and a random bit is returned otherwise.

We now turn to the analysis of the success probability of $\mathcal{B}$ in winning the ind-cpa game for PKE. Let bad be the event that the bad flag is set, and let good be its complement. First note that $\mathsf{P}\left[b' = b \,\middle|\, \mathsf{good}\right] \geq \frac{1+n\cdot\varepsilon}{2}$, because as long as the bad flag is never set, all extraction queries have been replied as in the original ibe-cpa game. Also note that $\mathsf{P}\left[b' = b \,\middle|\, \mathsf{bad}\right] \geq 1/2$, since in this case $\mathcal{B}$ outputs a uniform bit b'.

It remains to prove a lower bound on the probability of good happening. Note that since ϕ is $(t+1, s)$-cover-free, then there must exist an index $i \in \phi(\mathrm{ID}^*)$ such that $i \notin \phi(\mathrm{ID})$ for all (at most t) extraction queries $\mathrm{ID} \neq \mathrm{ID}^*$. If i^* takes such a value, then the bad flag is never set, i.e., good holds. It is not hard to show that the probability that such an index is hit is at least $\frac{1}{n}$, even though this requires some (standard) work (which we omit) due to the fact that i^* is chosen *before* $\mathcal{A}$'s execution starts. To conclude, we obtain

$$\mathsf{P}\left[b' = b\right] > \frac{1}{n} \cdot \frac{1+n\varepsilon}{2} + \left(1 - \frac{1}{n}\right) \cdot \frac{1}{2} = \frac{1+\varepsilon}{2} ,$$

which contradicts ind-cpa security of PKE. $\square$