Amplification of Chosen-Ciphertext Security

Huijia Lin[1,2] and Stefano Tessaro[2]

[1] Boston University
[2] MIT
{huijia,tessaro}@csail.mit.edu

Abstract. A central question in the theory of public-key cryptography is to determine which minimal assumptions are sufficient to achieve security against chosen-ciphertext attacks (or CCA-security, for short). Following the large body of work on hardness and correctness amplification, we investigate how far we can *weaken* CCA security and still be able to efficiently transform any scheme satisfying such a weaker notion into a fully CCA-secure one.

More concretely, we consider a weak CCA-secure bit-encryption scheme with decryption error $(1-\alpha)/2$ where an adversary can distinguish encryptions of different messages with possibly large advantage $\beta < 1 - 1/\mathsf{poly}$. We show that whenever $\alpha^2 > \beta$, the weak correctness and security properties can be simultaneously amplified to obtain a fully CCA-secure encryption scheme with negligible decryption error. Our approach relies both on a new hardcore lemma for CCA security as well as on revisiting the recently proposed approach to obtain CCA security due to Hohenberger *et al* (EUROCRYPT '12).

We note that such amplification results were only known in the simpler case of security against chosen-*plaintext* attacks.

1 Introduction

1.1 Public-Key Encryption and CCA Security

The seminal work of Goldwasser and Micali [1] introduced *semantic security* as the basic security notion for public-key encryption. Semantic security demands that no polynomial-time adversary, given only the public key, can distinguish encryptions of any two messages m_0 and m_1 of its choice, except with negligible advantage. Often, this notion is also referred to as *security against a chosen plaintext attack*, or CPA security, for short. This is in contrast to the stronger notion of *(adaptive) chosen-ciphertext security* (CCA security, for short) [2], where the above indistinguishability requirement must hold true even for adversaries with the additional ability to query a decryption oracle; as CCA security is required by many applications, it is by now considered to be the golden standard for secure public-key encryption.

In contrast to the case of CPA security, where simple constructions from generic assumptions (such as trapdoor permutations (TDP)) can be given, delivering CCA-secure public-key encryption from general assumptions proved itself to be a much more challenging problem. In particular, determining whether

T. Johansson and P. Nguyen (Eds.): EUROCRYPT 2013, LNCS 7881, pp. 503–519, 2013.

CCA-secure public-key encryption can be achieved solely from CPA-secure public-key encryption remains a major longstanding open question. Constructions additionally relying on non-interactive zero-knowledge proof systems (NIZKs) are known [3,4,2,5]. But, so far, all constructions of NIZKs require the existence of (enhanced) TDPs, which are not known to be implied by CPA-secure encryption; furthermore, known constructions based on NIZKs are all non-black-box. It is in fact likely that no black-box construction of a CCA-secure scheme from a CPA-secure one exists, as confirmed at least for a certain natural class of constructions [6]. For this reason, efficient constructions have been instead given from more concrete families of assumptions, such as hash proof systems and variants thereof [7,8], lossy TDFs [9], correlated-product secure TDFs [10], adaptive TDFs [11], or using random oracles [12,13].

1.2 Our Results: From Weak to Strong CCA Security

In this paper, we ask and answer the following question:

> *"How far can we weaken CCA security and still provide a black-box construction of a CCA-secure encryption scheme from a scheme only satisfying the weaker notion?"*

Our approach follows the one of the large body of works on *security amplification*, which has considered a wide range of cryptographic primitives and was initiated by Yao [14] in the context of one-way functions. Interestingly, limited work has been devoted to amplification of *public-key encryption*. The problem was first considered by Dwork, Naor, and Reingold [15] for CPA-secure public-key encryption. Constructions achieving better parameters were later proposed by Holenstein [16] and by Holenstein and Renner [17]. However, the question of amplifying *CCA* security has remained wide open. This is the question that we tackle and solve in this work.[1]

Modeling Weak CCA Encryption. Our model of weak CCA encryption extends naturally the model of weak CPA encryption considered in [15,17]. We start from a bit-encryption[2] scheme with key generation algorithm Gen, encryption algorithm Enc, and decryption algorithm Dec, and weaken it in two different directions, allowing both for *non-negligible decryption errors* as well as for *non-negligible adversarial advantage* in a chosen-ciphertext attack. More concretely, for two given parameters $0 < \alpha, \beta \leq 1$, where $\alpha \geq 1/p(\kappa)$ and $\beta < 1 - 1/q(\kappa)$ for some polynomials p and q, we assume the following two conditions:

[1] Note that in the *secret*-key setting, amplification of CCA security is, at least in principle, known to be feasible, as any weak form of CCA security implies weak one-way functions, and these are sufficient to build CCA-secure symmetric-key encryption via standard techniques.

[2] As every meaningful encryption scheme has at least the ability to encrypt a binary value, this is the weakest possible assumption in terms of message space of the basic scheme.

(i) **α-weak decryptability:** The decryption error over a random key-pair and a random bit is at most $\frac{1-\alpha}{2}$. We stress that this is a very weak guarantee, as it is taken over *random* choices of the keys and of the bit b, as well as of the coins used to encrypt b.

(ii) **β-weak security:** We consider the usual CCA-security game where an adversary obtains first the public key, and later a challenge ciphertext encrypting a random bit b. Moreover, the adversary can ask arbitrary decryption queries, with the sole exception that after the adversary obtains the challenge ciphertext, it cannot ask for its decryption. The task of the adversary is to output a guess b', and we are going to require that $\mathsf{Pr}[b' = b] \leq \frac{1+\beta}{2}$ for all polynomial-size adversaries.

Justifying weak CCA security. There are several reasons why assuming the existence of such a weak scheme is reasonable. Let us mention some natural examples.

- Within the general agenda of achieving CCA security from general assumptions, we may envision that a construction of a weak CCA scheme is potentially much easier to find than a construction of a full-fledged CCA-secure encryption scheme.
- An existing scheme designed to be CCA-secure may end up being less secure than expected due to the discovery of a better concrete attack or due to implementation errors, as in the recently discussed case of faulty key generation for RSA-based systems [18,19].
- It may be generally easier to build a CCA-secure scheme with large decryption errors. For example, as pointed out in [20], an encryption scheme with a simple, easily learnable, decryption algorithm must have large decryption error. In contrast to CPA encryption, reducing the decryption error turns out to be a major challenge in the case of CCA-secure encryption, *even* if the scheme is already fully CCA secure.

Our main result. The question we are going to ask is whether for certain α and β, there exists a transformation which delivers a CCA-secure encryption scheme from any scheme which has α-weak decryptability and β-weak security. We provide an affirmative answer to this question.

Theorem 1 (Main theorem, informal). *If $\alpha^2 > \beta$, there exists a black-box construction transforming any scheme with α-weak decryptability and β-weak security into a CCA-secure encryption scheme with negligible decryption error.*

Unfortunately, we cannot rule out constructions achieving a wider range of parameters α and β. In fact, we remark that the problem of determining the optimal parameters is open even in the simpler case of amplifying weak CPA security. While the constraint $\alpha^2 > \beta$ is shown [17] to be necessary for a restricted class of CPA black-box amplifiers, we see little value in extending this result to CCA security, as our amplifier itself is not within this class.

1.3 Our Techniques

We now turn to a high-level overview of our techniques. In particular, our approach builds upon a number of previous works [17,21,22], which we first review. Then, we will move to a description of our two main new tools, namely hardcore lemmas for CCA-security and heavy-ciphertext pre-sampling, and of their use.

Amplification of CPA Encryption. Given a bit-encryption scheme PKE with α-weak decryptability and β-weak security with respect to chosen-plaintext attacks, the Holenstein-Renner (HR) construction [17] produces a fully CPA-secure encryption scheme with negligible decryption error. To encrypt each message m, the HR construction invokes the basic bit-encryption scheme PKE to encrypt several fresh random bits $b_1, \cdots, b_n$ under n public keys $\mathsf{pk}_1, \cdots, \mathsf{pk}_n$, producing ciphertexts $c_1, \cdots, c_n$; the bits $b_1, \cdots, b_n$ are then carefully "combined" to generate a one-time-pad k for hiding the actual message m, as well as some additional ciphertext component c'; the additional component c' is used by the legitimate receiver, given the secret keys, to reconstruct the one-time pad, but it should not leak any information about k to the adversary. The final ciphertext is $c = (c_1, \ldots, c_n, c', m \oplus k)$.

The reason why such a combiner can exist is that the probability that the legitimate receiver, given the secret keys, can learn each individual bit b_i from c_i is $(1+\alpha)/2$, which we expect to be sufficiently larger than the probability that the adversary learns b_i from c_i *without* the secret keys. To make this intuition sound, one uses Impagliazzo's hardcore lemma [23] and its tighter version by Holenstein [16]: The lemma implies that if PKE is β-weakly CPA secure, then, for each i, with probability $1-\beta$ (over the choice of b_i, the randomness for sampling pk_i and encrypting b_i), the encryption of b_i is a "hard instance", meaning that given its encryption c_i, the bit b_i is (computationally) indistinguishable from a random independent bit. This gap between what an honest decryptor and an eavesdropper can recover can be leveraged by an information-theoretically secure one-way key-agreement protocol as in the setting of Maurer [24], which turns out to provide directly the right type of combiner.

From bit CCA Encryption to String CCA Encryption. It is well known that a CPA-secure string encryption scheme can be built from a CPA-secure bit-encryption scheme via simple parallel encryption of each bit. However, this approach does not lift to extending the message space of CCA-secure bit encryption, as an adversary can easily maul a challenge ciphertext $c_1 \cdots c_i \cdots c_n$ of a n-bit string $b_1 \cdots b_i \cdots b_n$ into another ciphertext $c_1 \cdots c_i' \cdots c_n$ of a related string $b_1 \cdots 0 \cdots b_n$, and thus win in the CCA security game—additional structure is needed to retain CCA security. Myers and shelat [21] showed that although this approach is not CCA secure, it satisfies a weaker adaptive security property—called UCCA security—which requires indistinguishability to hold for adversaries that can query a decryption oracle on any ciphertext $c_1, \cdots, c_n$ of their choice, except those that "quote" the challenge ciphertext, denoted as $c_1^*, \cdots, c_n^*$, at any of its components, that is $c_i = c_i^*$ for some i. Myers and shelat, and later Hohenberger, Lewko, and Waters (HLW) [22], showed how to construct a string CCA-secure scheme $\overline{\mathsf{PKE}}$ from such a UCCA-secure string encryption scheme

PKE_s.[3] We briefly review the HLW construction: It uses PKE_s as an *inner* encryption scheme $\mathsf{PKE}_{\mathsf{in}} = \mathsf{PKE}_s$ and two *outer* schemes $\mathsf{PKE}_{\mathsf{out},1}$, $\mathsf{PKE}_{\mathsf{out},2}$ that are CCA-1 and CPA secure respectively. To encrypt a message m, the encryption algorithm proceeds by encrypting m together with two random strings $r_{\mathsf{out},1}$ and $r_{\mathsf{out},2}$ into an *inner ciphertext* $c_{\mathsf{in}} = \mathsf{Enc}_{\mathsf{in}}(\mathsf{pk}_{\mathsf{in}}, (m, r_{\mathsf{out},1}, r_{\mathsf{out},2}))$; it then encrypts the inner ciphertext into two outer ciphertexts $(c_{\mathsf{out},1}, c_{\mathsf{out},2})$ using $r_{\mathsf{out},1}$ and $r_{\mathsf{out},2}$ respectively as the randomness for encryption, that is, $c_{\mathsf{out},i} = \mathsf{Enc}_{\mathsf{out},i}(\mathsf{pk}_{\mathsf{out},i}, c_{\mathsf{in}}; r_{\mathsf{out},i})$ for $i = 1, 2$; the final ciphertext is simply $(c_{\mathsf{out},1}, c_{\mathsf{out},2})$. At a high level, the two outer schemes prevent the adversary from issuing a decryption query for a ciphertext whose embedded inner ciphertext "quotes" that in the challenge ciphertext, thus reducing CCA to UCCA security.

Our Approach. A seemingly plausible attempt for constructing a CCA-secure encryption scheme from a weak scheme PKE with α-decryptability and β-weak CCA-security is to first try to show that the HR construction PKE', when instantiated with PKE as the basic bit-encryption scheme, is UCCA secure, and subsequently plugging PKE' as the inner encryption scheme into the HLW construction $\overline{\mathsf{PKE}}$, and show that it yields a CCA-secure encryption scheme.

Unfortunately, we encounter the following two challenges: First, it is unclear whether the weak CCA security of PKE is amplified through the construction of PKE' to UCCA security; in particular, known hardcore lemmas [23,16] only hold for games where the challenger is stateless, but the challenger in the CCA security game is stateful (it changes its behavior before and after the challenge ciphertext is generated). Second, it turns out that the security proof of the HLW construction requires the basic scheme PKE to have "unpredictability"—that is, a random cipheretxt (of a random bit) of PKE has high entropy and is almost impossible to blindly guess—which holds trivially for any fully-secure CPA encryption scheme with negligible decryption error, but is not satisfied by a weak CCA encryption scheme.

Overcoming these two difficulties turns out to be quite challenging and requires the adoption of new techniques, which we now illustrate.

Step 1: The Hardcore Lemma for CCA security and XCCA security. To overcome the first difficulty, we prove a variant of Impagliazzo's hardcore lemma which applies to CCA security (Theorem 2 below): It implies that if a scheme is weakly β-CCA-secure, then with probability $1 - \beta$ (over the randomness for choosing a random plaintext bit, for key generation, and for encryption), given an encryption of a random bit b, b is indistinguishable from a random *independent* bit even to adversaries with access to the decryption oracle. Our new hardcore lemma can be used to prove that PKE' satisfies an even stronger adaptive security property than UCCA, called XCCA (read as "cross"-CCA), which guarantees indistinguishability even for adversaries with access to decryption oracles that decrypts ciphertext of the basic scheme PKE under each individual

[3] In fact, [22] showed a more general construction of string CCA encryption schemes from any encryption scheme that is DCCA secure and unpredictable. In particular, UCCA security is a special case of DCCA security.

component key of PKE', subject to the restriction that the decryption oracle for the i-th component does not answer queries that "quote" the corresponding component in the challenge ciphertext. As we will see shortly, this stronger security guarantee is quintessential for overcoming the second difficulty.

Finally, rather than presenting a direct proof of the hardcore lemma for CCA security, we provide a general characterization of games for which hardcore lemmas exist, which extends beyond games for which such lemmas are known [23,16,25]. Our hardcore lemma for CCA-security is then simply derived as a special case. We believe this step to be of independent interest.

Step 2: From XCCA security to CCA security. We prove that the CCA security of $\overline{\mathsf{PKE}}$ can be based on the stronger XCCA security of the inner encryption PKE', even if the underlying basic scheme PKE is not sufficiently "unpredictable" – in contrast to the proof in [22]. This requires a substantially different analysis than the one of [22], and in particular a new reduction. Concretely, we overcome lack of unpredictability by introducing a new technique called *heavy-ciphertext pre-sampling*. Roughly speaking, this technique allows the security reduction (from CCA security of $\overline{\mathsf{PKE}}$ to XCCA security of PKE') to proactively predict and decrypt all highly likely ciphertexts of PKE, and the challenging task is to prove that these are the only components of the inner challenge ciphertext an adversary may indeed easily "quote" after seeing the challenge ciphertext.

2 Preliminaries

2.1 Basic Concepts and Notation

The probability distribution of a random variable X is usually denoted as P_X, and we occasionally use the shorthand $\mathsf{P}_X(x)$ for $\mathsf{Pr}[X = x]$. Adversaries are going to be modeled as non-uniform families of (randomized) circuits for ease of exposition, but all results extend with some work to the uniform setting.

2.2 Weak and Strong CCA-secure Encryption

A *public-key encryption scheme* with message space $\mathcal{M} \subseteq \{0,1\}^*$ is a triple $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, where (i) Gen is the (randomized) *key generation algorithm*, outputting a pair $(\mathsf{pk}, \mathsf{sk})$ consisting of a *public-* and a *secret-key*, respectively (ii) Enc is the (randomized) *encryption algorithm* outputting a ciphertext $c = \mathsf{Enc}(\mathsf{pk}, m)$ for any message $m \in \mathcal{M}$ and a valid public key pk; and (iii) Dec is the deterministic *decryption algorithm* such that $\mathsf{Dec}(\mathsf{sk}, c) \in \mathcal{M} \cup \{\perp\}$. All algorithms additionally take (implicitly) as input the security parameter 1^κ in unary form, and the message space $\mathcal{M}$ may also depend on the security parameter κ. Whenever $\mathcal{M} = \{0,1\}$, we say that the scheme is a *bit-encryption* scheme. We sometimes need to make the randomness used by Gen and Enc explicit: In these cases, we write $\mathsf{Gen}(r)$ and $\mathsf{Enc}(\mathsf{pk}, m; r)$ to highlight the fact that random coins r are used to generate keys by Gen and to encrypt the message m, respectively.

CORRECTNESS OF PKE. Throughout this paper, we say that the encryption scheme PKE with message space $\mathcal{M}$ has *decryption error* δ if

$$\mathsf{Pr}\left[(\mathsf{pk},\mathsf{sk}) \xleftarrow{\$} \mathsf{Gen},\ m \xleftarrow{\$} \mathcal{M} :\ \mathsf{Dec}(\mathsf{sk},\mathsf{Enc}(\mathsf{pk},m)) \neq m\right] \leq \delta,$$

where the probability is additionally over the random coins of Enc. Moreover, we say that a scheme is *almost perfectly correct*, if for an overwhelming fraction of randomness r used by the key generation algorithm, for $(\mathsf{pk},\mathsf{sk}) = \mathsf{Gen}(r)$, and all messages $m \in \mathcal{M}$, we have $\mathsf{Pr}\left[\mathsf{Dec}(\mathsf{sk},\mathsf{Enc}(\mathsf{pk},m)) = m\right] = 1$.

SECURITY OF PKE. In general, security of the scheme $\mathsf{PKE} = (\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ is defined via the following security game involving a *challenger* $\mathcal{C}_{\mathsf{CCA2}}$ and an adversary $\mathcal{A}$:

Game $\mathsf{CCA2}^{\mathcal{A}}_{\mathsf{PKE}}$**:**

(i) $\mathcal{C}_{\mathsf{CCA2}}$ generates $(\mathsf{pk},\mathsf{sk}) \xleftarrow{\$} \mathsf{Gen}$ and $b \xleftarrow{\$} \{0,1\}$, and gives pk to $\mathcal{A}$.
(ii) $\mathcal{A}$ asks decryption queries c, which are answered with $\mathsf{Dec}(\mathsf{sk},c)$.
(iii) $\mathcal{A}$ outputs (m_0,m_1) with $|m_0| = |m_1|$; $\mathcal{C}_{\mathsf{CCA2}}$ sends $\mathcal{A}$ $c^* \xleftarrow{\$} \mathsf{Enc}(\mathsf{pk},m_b)$.
(iv) $\mathcal{A}$ asks decryption queries $c \neq c^*$, which are answered with $\mathsf{Dec}(\mathsf{sk},c)$.
(v) The adversary $\mathcal{A}$ outputs a bit b', and *wins* the game if $b' = b$.

We refer to decryption queries in phase **(ii)** and **(iv)** as *before-the-fact* and *after-the-fact* decryption queries, respectively. Moreover, in the case that PKE is a bit-encryption scheme we assume without loss of generality that $(m_0,m_1) = (0,1)$, and hence $\mathsf{Enc}(\mathsf{pk},b)$ is the challenge ciphertext. We also define the $\mathsf{CCA2}$-*advantage* of the adversary $\mathcal{A}$ as $\mathbf{Adv}^{\mathsf{CCA2}}_{\mathsf{PKE}}(\mathcal{A}) = 2\cdot\mathsf{Pr}\left[b' = b\right] - 1$. We say that an encryption scheme is *CCA-secure* if $\mathbf{Adv}^{\mathsf{CCA2}}_{\mathsf{PKE}}(\mathcal{A})$ is negligible for all polynomial-size adversaries $\mathcal{A}$. We say it is *q-CCA-secure* if this holds for adversaries making at most q decryption queries, whereas it is *CPA-secure* if it is 0-CCA-secure. The following notation will also be convenient.

Definition 1. *For $\alpha,\beta \in [0,1]$, a bit-encryption scheme PKE is* (α,β)-CCA-secure *if the following two properties hold:* **(i)** PKE *has decryption error* $(1-\alpha)/2$*, and* **(ii)** *For any polynomial-size adversary $\mathcal{A}$, we have* $\mathbf{Adv}^{\mathsf{CCA2}}_{\mathsf{PKE}}(\mathcal{A}) \leq \beta$.

In passing, we point out that CPA-secure encryption with negligible decryption error implies one-way functions [26], and in turn implies pseudorandom generators [27], all in a black-box way.

3 The Hardcore Lemma for CCA Security

Impagliazzo's Hardcore Lemma [23] asserts that if it is mildly hard to compute $P(x)$ for a predicate P on a random input x given side information $f(x)$ (i.e., say this can be done with probability at most $\frac{1+\varepsilon}{2}$), then there exists a sufficiently large subset $\mathcal{S}$ (the "hardcore set") of the inputs such that when sampling x' from $\mathcal{S}$, it is infeasible to predict $P(x')$ from $f(x')$ noticeably better than by

random guessing. A tight proof where the set $\mathcal{S}$ contains a $(1-\varepsilon)$-fraction of the inputs is due to Holenstein [16]. The main contribution of this section is to derive a similar statement for (weak) CCA-secure encryption to be used below.

In particular, we present a new abstraction of existing proofs of hardcore lemmas, which is of independent interest. Not only we apply it to derive the hardcore lemma for CCA security of bit-encryption, but it also yields previous more restricted statements [23,25] as special cases.

Bit-guessing games. We consider games (such as the CCA-security game) where the adversary is asked to guess a bit. Formally, a *bit-guessing game* is a tuple $G = (\mathsf{P}_X, \mathcal{C}, P)$, where P_X is a probability distribution with support $\mathcal{X}$, $\mathcal{C}$ is an interactive *stateful* machine taking an auxiliary input $x \in \mathcal{X}$, and $P : \mathcal{X} \to \{0,1\}$ is a predicate. Combined with an adversary $\mathcal{A}$, G defines the following random experiment: First, an input $x \xleftarrow{\$} \mathsf{P}_X$ is sampled. Then $\mathcal{A}$ interacts with the challenger $\mathcal{C}(x)$ and outputs a bit $b \xleftarrow{\$} \mathcal{A}^{\mathcal{C}(x)}$ (the oracle $\mathcal{C}(x)$ keeps state). The *G-advantage of $\mathcal{A}$ relative to a distribution* P is

$$\mathbf{Adv}^{G}_{\mathsf{P}}(\mathcal{A}) = 2 \cdot \mathsf{Pr}\left[x \xleftarrow{\$} \mathsf{P},\ b \xleftarrow{\$} \mathcal{A}^{\mathcal{C}(x)} :\ b = P(x)\right] - 1\ . \tag{1}$$

We say that G is (s,ε)-*hard* if $\mathbf{Adv}^{G}_{\mathsf{P}_X}(\mathcal{A}) \leq \varepsilon$ for all s-size adversaries $\mathcal{A}$.

Hardcore lemmas and measures. A *measure* $\mathcal{M}$ for a bit-guessing game G is a mapping $\mathcal{M} : \mathcal{X} \to [0,1]$, and its *density* is $\mu(\mathcal{M}) = \sum_{x\in\mathcal{X}} \mathsf{P}_X(x)\cdot\mathcal{M}(x)$. We associate with $\mathcal{M}$ the probability distribution $\mathsf{P}_{\mathcal{M}}$ such that $\mathsf{P}_{\mathcal{M}}(x) := \mathsf{P}_X(x) \cdot \mathcal{M}(x)/\mu(\mathcal{M})$ for all $x \in \mathcal{X}$. The role of a measure is that of adjoining an event $\mathcal{E}$ to the sampling of $x \xleftarrow{\$} \mathsf{P}_X$ such that $\mathsf{Pr}\left[\mathcal{E} \,\middle|\, X = x\right] = \mathcal{M}(x)$; then in particular $\mathsf{Pr}\left[\mathcal{E}\right] = \mu(\mathcal{M})$, and $\mathsf{Pr}\left[X = x \,\middle|\, \mathcal{E}\right] = \mathsf{P}_{\mathcal{M}}(x)$.

We ask the question of which bit-guessing games admit a *hardcore measure*: Assuming the game G is ε-hard for some $\varepsilon \in [0,1]$, we seek for a measure $\mathcal{M}$ with large density (e.g. $\mu(\mathcal{M}) \geq 1-\varepsilon$) such that conditioned on the associated event $\mathcal{E}$, the game G is very hard to win. In [25], a proof that this is true for the case where $\mathcal{C}(x)$ is stateless for each x was given. Our new approach extends this to possibly stateful challengers, as in the case of CCA security.

Abstract hardcore lemmas. We give a sufficient condition on a bit-guessing game $G = (\mathsf{P}_X, \mathcal{C}, P)$ to admit a hardcore lemma – informally, this condition corresponds to the ability, for any given and possibly unknown x, to estimate the probability that a binary-output adversary for G, sampled according to a given distribution over circuits, outputs one when run on $\mathcal{C}(x)$. In particular, we call an oracle O a *size s circuit sampler* for G if, upon each invocation, it returns the description of a valid adversary $\mathcal{A}$ for G of size s. For each such O, we define $p_1^{G,\mathsf{O}}(x)$ as the probability that a randomly sampled adversary $\mathcal{A} \xleftarrow{\$} \mathsf{O}$ outputs one when run with $\mathcal{C}(x)$, i.e., $p_1^{G,\mathsf{O}}(x) := \mathsf{Pr}\left[\mathcal{B} \xleftarrow{\$} \mathsf{O}, b' \xleftarrow{\$} \mathcal{B}^{\mathcal{C}(x)} : b' = 1\right]$. The following definition captures the notion of a good estimation algorithm for $p_1^{G,\mathsf{O}}(x)$ which can only interact with $\mathcal{C}(x)$ and obtain samples from O, but does not learn x and must be equally successful on all such x.

Definition 2 (p_1-estimator). *A (s, s', q, γ, η)-p_1-estimator for a bit-guessing game $G = (\mathsf{P}_X, \mathcal{C}, P)$ is a size s circuit $\mathcal{E}$ with output in $[0,1]$ such that*

$$\mathsf{Pr}\left[\mathcal{B}_1, \ldots, \mathcal{B}_q \xleftarrow{\$} \mathsf{O},\ \overline{p_1} \xleftarrow{\$} \mathcal{E}^{\mathcal{C}(x)}(\mathcal{B}_1, \ldots, \mathcal{B}_q) :\ \left|\overline{p_1} - p_1^{G,\mathsf{O}}(x)\right| > \gamma\right] < \eta$$

for all size-s' circuit samplers O and for all x.

Note that in particular $q \cdot s' \leq s$. The following theorem relates the existence of a hardcore lemma for a certain game G with the existence of a p_1-sampler for G. Its proof abstracts the ones of [16,25] and is found in the full version.

Proposition 1 (The Abstract Hardcore Lemma). *Let $s \in \mathbb{N}$ and $\varepsilon \in [0,1]$. Let $G = (\mathsf{P}_X, \mathcal{C}, P)$ be a bit-guessing game which is (s,ε)-hard. Then, for all $\gamma > 0$, if for some $s' = s'(\gamma)$ there exists an $(s, s', q, \gamma(1-\varepsilon)/4, \gamma(1-\varepsilon)/4)$-$p_1$-estimator for G, then there exists a measure $\mathcal{M} = \mathcal{M}_\gamma$ such that:*

(i) $\mu(\mathcal{M}) \geq 1 - \varepsilon$, **(ii)** $\mathbf{Adv}^G_{\mathsf{P}_\mathcal{M}}(\mathcal{B}) \leq \gamma$ *for all s'-size $\mathcal{B}$.*

The hardcore lemma for CCA-security. We are now going to show a hardcore lemma for CCA-security as an application of Proposition 1. Let $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a public-key *bit* encryption scheme such that Gen and Enc take randomness of lengths ρ_{Gen} and ρ_{Enc}, respectively. Formally, we consider the bit-guessing game $\mathsf{CCA2}[\mathsf{PKE}] = (\mathsf{P}_X, \mathcal{C}_{\mathsf{CCA2}}, P)$ where P_X is the uniform distribution on $\{0,1\}^{\rho_{\mathsf{Gen}}} \times \{0,1\}^{\rho_{\mathsf{Enc}}} \times \{0,1\}$, whereas $\mathcal{C}_{\mathsf{CCA2}}(r_{\mathsf{Gen}}, r_{\mathsf{Enc}}, b)$ is the challenger for the CCA-security game for PKE with challenge bit b, public key and secret key $(\mathsf{pk}, \mathsf{sk}) = \mathsf{Gen}(r_{\mathsf{Gen}})$, and challenge ciphertext $c^* = \mathsf{Enc}(\mathsf{pk}, b; r_{\mathsf{Enc}})$. Moreover, we define $P(r_{\mathsf{Gen}}, r_{\mathsf{Enc}}, b) = b$. The following lemma gives an appropriate p_1-estimator for $\mathsf{CCA2}[\mathsf{PKE}]$.

Lemma 1. *For all $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with message space $\{0,1\}$, and all $s' \in \mathbb{N}$, $\gamma, \eta \in (0,1]$, there exists a (s, s', q, γ, η)-p_1-estimator for $\mathsf{CCA2}[\mathsf{PKE}]$ with $q = O(\log(1/\eta)/\gamma^2)$ and $s = s' \cdot q + O(1)$.*

Proof. The estimator $\mathcal{E}$, given pk from $\mathcal{C}_{\mathsf{CCA2}}$, runs sequentially each of $\mathcal{B}_1, \ldots, \mathcal{B}_q$ on input pk until they output their query $(0,1)$. All before-the-fact decryption queries are answered using the challenger $\mathcal{C}_{\mathsf{CCA2}}$. It then obtains a challenge ciphertext c^*, and then resumes the execution of $\mathcal{B}_i$'s from the last state before outputting $(0,1)$, again using the challenger to reply to decryption queries. Finally, let b'_i be the output of $\mathcal{B}_i$; the estimator $\mathcal{E}$ outputs the average $z = (1/q) \cdot \sum_{i=1}^{q} b'_i$. The error is at most γ with probability at most η by the Chernoff bound. □

The above proof crucially relies on the scheme encrypting one-bit messages: For a larger set of messages, each $\mathcal{B}_i$ could ask a different message pair, and the above estimation technique would fail.

The following theorem is a simple combination of Proposition 1 and Lemma 1.

Theorem 2 (Hardcore Lemma for CCA Security). *Let $\alpha, \beta \in [0,1]$, and let $s \in \mathbb{N}$. Moreover, let $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a public-key encryption*

scheme with message space $\{0,1\}$*, and assume that* $\mathbf{Adv}^{\mathsf{CCA2}}_{\mathsf{PKE}}(\mathcal{A}) \leq \beta$ *for all* s*-size adversaries* $\mathcal{A}$*. Then, for all* $\gamma > 0$*, there exists a measure* $\mathcal{M}$ *such that* $\mu(\mathcal{M}) \geq 1-\beta$*, and* $\mathbf{Adv}^{\mathsf{CCA2}[\mathsf{PKE}]}_{\mathsf{P}_{\mathcal{M}}}(\mathcal{B}) \leq \gamma$ *for all adversaries* $\mathcal{B}$ *with size* s'*, where* $s = O(s' \cdot \log(1/\gamma(1-\varepsilon))/\gamma^2(1-\varepsilon)^2)$.

In the full version, we provide a more detailed discussion about related results and extensions to the uniform setting, which we here omit due to lack of space.

4 From Weak to Strong CCA Security

We present our construction to transform an (α,β)-CCA encryption scheme into a fully CCA-secure encryption scheme. First, we review some tools underlying our construction, before turning to its description and security.

4.1 Information-Theoretically Secure Key-Agreement

We consider the problem of two parties, Alice and Bob, agreeing on a secret key with *unconditional security* in a setting where they each hold values $X_1,\ldots X_n$ and $Y_1\ldots,Y_n$, respectively, in presence of an adversary obtaining correlated values $Z_1,\ldots,Z_n$; in particular, (X_i,Y_i,Z_i) are sampled independently from a given tripartite probability distributions P_{XYZ} for all $1 \leq i \leq n$. That is, (X_i,Y_i,Z_i) are correlated for each i, but independent across distinct indices $i \neq j$. Moreover, Alice and Bob are connected via an authenticated channel, allowing them to exchange messages, which is however wiretapped by the adversary. Secret-key agreement in this setting was first considered by Maurer [24]. Here, we consider the special case where the channel only allows *one-way* communication from Alice to Bob. The following definition captures protocols for this setting.

Definition 3 (One-way key-agreement). *Let* $\varepsilon,\delta : \mathbb{N} \to [0,1]$*, and let* $n,\ell : \mathbb{N} \to \mathbb{N}$ *be monotonically increasing. Also, let* $\mathcal{P} = \{\mathcal{P}_\kappa\}_{\kappa\in\mathbb{N}}$ *be a family of sets of probability distribution* P_{XYZ}*. A* $(\mathcal{P},\varepsilon,\delta,n,\ell)$-one-way key-agreement (OKA) protocol *is a pair of probabilistic polynomial-time algorithms* $\mathsf{OKA} = (\mathsf{KAEnc},\mathsf{KADec})$ *such that for all* $\kappa \in \mathbb{N}$ *and* $\mathsf{P}_{XYZ} \in \mathcal{P}_\kappa$*, the following two properties hold when sampling* $(X_1,Y_1,Z_1),\ldots,(X_n,Y_n,Z_n) \xleftarrow{\$} \mathsf{P}_{XYZ}$ *(where* $n = n(\kappa)$*),* $(C,K) \xleftarrow{\$} \mathsf{KAEnc}(1^\kappa,X_1,\ldots,X_n)$*,* $K' \xleftarrow{\$} \mathsf{KADec}(1^\kappa,Y_1,\ldots,Y_n;C)$*, and* $K'' \xleftarrow{\$} \{0,1\}^{\ell(\kappa)}$*:* **(1)** $K = K'$ *with probability at least* $1-\delta(\kappa)$*, and* **(2)** $(C,K,Z_1,\ldots,Z_n)$ *and* $(C,K'',Z_1,\ldots,Z_n)$ *have statistical distance at most* $\varepsilon(\kappa)$.

The following set of distributions was introduced in [17].

Definition 4 ([17]). *Let* $\alpha,\beta : \mathbb{N} \to [0,1]$*. Let* $\mathcal{D}(\alpha,\beta) = \{\mathcal{D}_\kappa(\alpha,\beta)\}_{\kappa\in\mathbb{N}}$ *be such that for all* $\kappa \in \mathbb{N}$*,* $\mathsf{P}_{XYZ} \in \mathcal{D}_\kappa(\alpha,\beta)$ *if* $(X,Y,Z) \xleftarrow{\$} \mathsf{P}_{XYZ}$ *satisfies* **(i)** $\mathsf{Pr}[X=0] = \mathsf{Pr}[X=1] = \frac{1}{2}$*, i.e.,* X *is uniform,* **(ii)** $\mathsf{Pr}[X=Y] \geq \frac{1+\alpha(\kappa)}{2}$*,* **(iii)** *there exists an event* $\mathcal{E}$*, defined on* (X,Z)*, such that* $\mathsf{Pr}\left[X=0 \,\middle|\, Z=z,\mathcal{E}\right] = \mathsf{Pr}\left[X=1 \,\middle|\, Z=z,\mathcal{E}\right] = \frac{1}{2}$ *for all* z*, and* $\mathsf{Pr}[\mathcal{E}] \geq 1-\beta(\kappa)$.

The following two propositions show feasibility of OKA protocols for $\mathcal{D}(\alpha, \beta)$ for certain values of α and β. The first proposition was proved by Holenstein and Renner [17], the second is proved in the full version. We note that there is no a-priori reason why α^2 and β could not be closer, yet no better gap can be proven given existing constructions of capacity-achieving error-correcting codes.

Proposition 2. *Let $\alpha, \beta : \mathbb{N} \to [0,1]$ be such that $\alpha^2 > \beta + \Omega(1)$, and let $\ell : \mathbb{N} \to \mathbb{N}$ be a polynomial function. Then, there exists a polynomial-time $(\mathcal{D}(\alpha,\beta), \varepsilon, \delta, n, \ell)$-OKA protocol such that $n(\kappa) = \frac{1}{7} \cdot \ell(\kappa) \cdot (\alpha^2 - \beta - O(1))$ and moreover, $\varepsilon(\kappa)$ is negligible in $n(\kappa)$, and $\delta(\kappa) = 2^{-\Theta(n(\kappa))}$.*

Proposition 3. *Let $p, \ell : \mathbb{N} \to \mathbb{N}$ be polynomially bounded and let $\varepsilon' : \mathbb{N} \to [0,1]$. Then, there exists a $\mathcal{D}(1, 1 - \frac{1}{p(\kappa)}, \varepsilon, \delta, n, \ell)$-OKA protocol where $n(\kappa) = 2/(1-\beta(\kappa)) \cdot (\ell(\kappa) + 2\log(1/\varepsilon'(\kappa)) + O(1))$, $\varepsilon(\kappa) \leq O(\sqrt{\varepsilon'(\kappa)})$, and $\delta(\kappa) = 0$.*

4.2 The Construction

Let $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a bit-encryption scheme which is (α, β)-secure. Assuming the existence of an information-theoretically secure one-way key agreement protocol for $\mathcal{D}(\alpha, \beta)$, we present a construction of a CCA-secure public-key encryption scheme $\overline{\mathsf{PKE}} = (\overline{\mathsf{Gen}}, \overline{\mathsf{Enc}}, \overline{\mathsf{Dec}})$, with message length $\ell = \ell(\kappa)$ and negligible decryption error, which makes black-box use of the basic scheme PKE.

At the highest level, our construction $\overline{\mathsf{PKE}}$ follows the paradigm recently proposed by Hohenberger, Lewko, and Waters [22]. In particular, it consists of an inner scheme $\mathsf{PKE}_{\mathsf{in}} = (\mathsf{Gen}_{\mathsf{in}}, \mathsf{Enc}_{\mathsf{in}}, \mathsf{Dec}_{\mathsf{in}})$ and two outer schemes $\mathsf{PKE}_{\mathsf{out},1} = (\mathsf{Gen}_{\mathsf{out},1}, \mathsf{Enc}_{\mathsf{out},1}, \mathsf{Dec}_{\mathsf{out},1})$ and $\mathsf{PKE}_{\mathsf{out},2} = (\mathsf{Gen}_{\mathsf{out},2}, \mathsf{Enc}_{\mathsf{out},2}, \mathsf{Dec}_{\mathsf{out},2})$, all three of which will be built from PKE, and specified below. For $\star \in \{\mathsf{in}, (\mathsf{out},1), (\mathsf{out},2)\}$, let us further denote by $\ell_\star, \rho_\star$ and $t_\star$ the message, randomness, and ciphertext lengths of $\mathsf{PKE}_\star$, respectively. We are going to require $\ell_{\mathsf{in}} = \ell + \rho_{\mathsf{out},1} + \rho_{\mathsf{out},2}$ as well as $\ell_{\mathsf{out},1} = \ell_{\mathsf{out},2} = t_{\mathsf{in}}$. A formal description of $\overline{\mathsf{PKE}}$ is given in Figure 1, on top: We encrypt the message m, together with two random values $r_{\mathsf{out},1}$ and $r_{\mathsf{out},2}$, obtaining an *inner* ciphertext c_{in}, which is then encrypted twice with the two outer schemes, using $r_{\mathsf{out},1}$ and $r_{\mathsf{out},2}$ as the respective random coins. Decryption recovers the message by decrypting the ciphertext via $\mathsf{Dec}_{\mathsf{out},1}$ and $\mathsf{Dec}_{\mathsf{in}}$ using the corresponding secret keys, and then checks validity of the ciphertext by re-encrypting the inner ciphertext using the public keys and the recovered random coins.

We now turn to describing the construction of the component schemes $\mathsf{PKE}_{\mathsf{in}}$, $\mathsf{PKE}_{\mathsf{out},1}$ and $\mathsf{PKE}_{\mathsf{out},2}$ from the basic scheme PKE.

The Inner Scheme. Let $\mathsf{OKA} = (\mathsf{KAEnc}, \mathsf{KADec})$ be a $(\mathcal{D}(\alpha,\beta), \varepsilon, \delta, n, \ell_{\mathsf{in}})$-one-way key agreement protocol such that ε and δ are negligible, and known (recall that PKE is (α,β)-CCA secure). We define $\mathsf{PKE}_{\mathsf{in}} = (\mathsf{Gen}_{\mathsf{in}}, \mathsf{Enc}_{\mathsf{in}}, \mathsf{Dec}_{\mathsf{in}})$ as in Figure 1, at the bottom: It encrypts random bits $b_1, \ldots, b_n$ with the basic scheme, and then generates a session key k via $\mathsf{KAEnc}(b_1, \ldots, b_n)$, and a ciphertext c', and uses the key k as an one-time pad. Decryption via KADec is then obvious. It is easy to see that the decryption error of this scheme is inherited from OKA, i.e., it is upper bounded by exactly δ.

Scheme $\overline{\mathsf{PKE}} = (\overline{\mathsf{Gen}}, \overline{\mathsf{Enc}}, \overline{\mathsf{Dec}})$:

Key generation $\overline{\mathsf{Gen}}(1^\kappa)$**:** Sample $(\mathsf{pk}_{\mathsf{in}}, \mathsf{sk}_{\mathsf{in}}) \xleftarrow{\$} \mathsf{Gen}_{\mathsf{in}}(1^\kappa)$ and for $i = 1, 2$, $(\mathsf{pk}_{\mathsf{out},i}, \mathsf{sk}_{\mathsf{out},i}) \xleftarrow{\$} \mathsf{Gen}_{\mathsf{out},i}(1^\kappa)$. Return $(\overline{\mathsf{pk}} = (\mathsf{pk}_{\mathsf{in}}, \mathsf{pk}_{\mathsf{out},1}, \mathsf{pk}_{\mathsf{out},2}), \overline{\mathsf{sk}} = (\mathsf{sk}_{\mathsf{in}}, \mathsf{sk}_{\mathsf{out},1}, \mathsf{pk}_{\mathsf{out},1}, \mathsf{pk}_{\mathsf{out},2}))$.

Encryption $\overline{\mathsf{Enc}}(\overline{\mathsf{pk}}, m)$**,** $m \in \{0,1\}^\ell$**:** Sample $r_{\mathsf{out},i} \xleftarrow{\$} \{0,1\}^{\rho_{\mathsf{out},i}}$ for $i = 1, 2$. Generate $c_{\mathsf{in}} \xleftarrow{\$} \mathsf{Enc}_{\mathsf{in}}(\mathsf{pk}_{\mathsf{in}}, m \,\|\, r_{\mathsf{out},1} \,\|\, r_{\mathsf{out},2})$ and $c_{\mathsf{out},i} \leftarrow \mathsf{Enc}_{\mathsf{out},i}(\mathsf{pk}_{\mathsf{out},i}, c_{\mathsf{in}}; r_{\mathsf{out},i})$ for $i = 1, 2$. Output ciphertext $c_{\mathsf{out},1} \,\|\, c_{\mathsf{out}_2}$.

Decryption $\overline{\mathsf{Dec}}(\overline{\mathsf{sk}}, c = c_{\mathsf{out},1} \,\|\, c_{\mathsf{out}_2})$**:** Decrypt $c'_{\mathsf{in}} \leftarrow \mathsf{Dec}_{\mathsf{out},1}(\mathsf{sk}_{\mathsf{out},1}, c_{\mathsf{out},1})$ and $m' \,\|\, r'_{\mathsf{out},1} \,\|\, r'_{\mathsf{out},2} \leftarrow \mathsf{Dec}_{\mathsf{in}}(\mathsf{sk}_{\mathsf{in}}, c'_{\mathsf{in}})$. If $\mathsf{Enc}_{\mathsf{out},i}(\mathsf{pk}_{\mathsf{out},i}, c'_{\mathsf{in}}; r'_{\mathsf{out},i}) = c_{\mathsf{out},i}$ for $i = 1, 2$ then return m, else return $\perp$.

Scheme $\mathsf{PKE}_{\mathsf{in}} = (\mathsf{Gen}_{\mathsf{in}}, \mathsf{Enc}_{\mathsf{in}}, \mathsf{Dec}_{\mathsf{in}})$:

Key generation $\mathsf{Gen}_{\mathsf{in}}(1^\kappa)$**:** Sample $(\mathsf{pk}_1, \mathsf{sk}_1), \ldots, (\mathsf{pk}_n, \mathsf{sk}_n) \xleftarrow{\$} \mathsf{Gen}(1^\kappa)$. Return $(\mathbf{pk} = (\mathsf{pk}_1, \ldots, \mathsf{pk}_n), \mathbf{sk} = (\mathsf{sk}_1, \ldots, \mathsf{sk}_n))$.

Encryption $\mathsf{Enc}_{\mathsf{in}}(\mathbf{pk}, m)$**,** $m \in \{0,1\}^{\ell_{\mathsf{in}}}$**:** For all $i \in [n]$, sample $b_i \xleftarrow{\$} \{0,1\}$ and generate $c_i \xleftarrow{\$} \mathsf{Enc}(\mathbf{pk}[i], b_i)$. Compute $(k, c') \xleftarrow{\$} \mathsf{KAEnc}(b_1, \ldots, b_n)$. Return ciphertext $(c_1, \ldots, c_n, c', m \oplus k)$.

Decryption $\mathsf{Dec}_{\mathsf{in}}(\mathbf{sk}, c = (c_1, \ldots, c_n, c', c''))$**:** Decrypt $b'_i \leftarrow \mathsf{Dec}(\mathbf{sk}[i], c_i)$ for $i = [n]$ and $k' \leftarrow \mathsf{KADec}(b'_1, \ldots, b'_n; c')$. Return plaintext $m' = c'' \oplus k'$.

Fig. 1. Descriptions of public-key encryption schemes $\overline{\mathsf{PKE}}$ and $\mathsf{PKE}_{\mathsf{in}}$

THE OUTER SCHEMES. We now instantiate the two outer schemes. The following description is fairly high-level, but sufficient to fully specify the construction. We refer the reader unfamiliar with the basic components to the full version for a more detailed description.

We first derive a CPA-secure public-key encryption scheme $\mathsf{PKE}^{\ell,\rho}_{\mathsf{out}}$ with message length $\ell = \mathsf{poly}(\kappa)$ and randomness length $\rho = \omega(\log(\kappa))$ from the basic scheme PKE which also enjoys almost-perfect correctness:[4]

1. We use the same construction as in $\mathsf{PKE}_{\mathsf{in}}$ to achieve a CPA-secure scheme $\mathsf{PKE}'_{\mathsf{out}}$, with message length truncated to 1-bit. CPA-security of the resulting scheme follows from the proof in [17] or from the stronger Lemma 2 below. Let ρ be the randomness length of $\mathsf{PKE}'_{\mathsf{out}}$.
2. We apply the transformation by Dwork, Naor, and Reingold [15] to enhance correctness of $\mathsf{PKE}'_{\mathsf{out}}$ with negligible decryption error to almost-perfect correctness, via sparsification of the randomness space. Let δ be the decryption error of $\mathsf{PKE}'_{\mathsf{out}}$. The transformation of [15] reduces randomness length to $\rho' = \frac{1}{4} \cdot \log(1/\delta(\kappa)) = \omega(\log(\kappa))$ via a PRG $G : \{0,1\}^{\rho'} \to \{0,1\}^{\rho}$, whose existence is implied by the existence of $\mathsf{PKE}'_{\mathsf{out}}$ in a black-box fashion [26,27].

[4] In the following, we are not going to optimize the complexity of the scheme; it is clear that some modifications can be done to save on complexity.

3. We then use parallel repetition of ℓ copies of $\mathsf{PKE}''_{\mathsf{out}}$ to obtain $\mathsf{PKE}^{\ell,\rho}_{\mathsf{out}}$, possibly using a PRG again to shorten the overall randomness length to ρ.

We let $\mathsf{PKE}_{\mathsf{out},2} = \mathsf{PKE}^{\ell_{\mathsf{out},2},\rho_{\mathsf{out},2}}_{\mathsf{out}}$. To obtain the first outer scheme $\mathsf{PKE}_{\mathsf{out},1}$, we rely on the result by Cramer *et al* [28] that transforms a CPA-secure encryption scheme into a 1-CCA secure one in a black-box way, which preserves the almost perfect correctness property of the underlying CPA-secure scheme. By applying their transformation to $\mathsf{PKE}^{\ell_{\mathsf{out},1},\rho}_{\mathsf{out}}$ (for some $\rho = \mathsf{poly}(\kappa)$), and then finally using a PRG to reduce the randomness length to $\rho_{\mathsf{out},1}$, we obtain a 1-CCA secure encryption scheme that is almost-perfectly correct.

4.3 CCA Security of $\overline{\mathsf{PKE}}$

We turn to our main result and show that our construction $\overline{\mathsf{PKE}}$ is CCA secure.

Theorem 3. *Let ε and δ be two negligible functions. Assume that PKE is (α,β)-CCA-secure, and OKA is a $(\mathcal{D}(\alpha,\beta),\varepsilon,\delta,n,\ell_{\mathsf{in}})$-one-way key-agreement protocol. Then, $\overline{\mathsf{PKE}}$ is a CCA-secure encryption scheme with negligible decryption error.*

In particular, by Propositions 2 and 3, we achieve amplification whenever $\alpha^2 > \beta + \Omega(1)$, and whenever $\alpha = 1$ and $\beta < 1 - \frac{1}{p(\kappa)}$ for some polynomial p.

Overview of the Security Proof. Towards showing the CCA security of $\overline{\mathsf{PKE}}$, we first show that it follows from Theorem 2 that the inner encryption scheme $\mathsf{PKE}_{\mathsf{in}}$ satisfies a strong adaptive security property, which we refer to as XCCA (to be read as "cross"-CCA) security. We are then going to reduce the CCA security of $\overline{\mathsf{PKE}}$ to the XCCA security of $\mathsf{PKE}_{\mathsf{in}}$ using the 1-CCA security of $\mathsf{PKE}_{\mathsf{out},1}$ and the CPA security of $\mathsf{PKE}_{\mathsf{out},2}$, combined with their almost perfect correctness. This second step resembles the proof of [22] only at a first glance, as it will require a completely different technique to handle the fact that ciphertexts of the basic scheme PKE are not sufficiently unpredictable.

Before proceeding to describing the two steps in more details, we first describe the XCCA security game. For simplicity, here we only define the XCCA game w.r.t. the concrete scheme $\mathsf{PKE}_{\mathsf{in}}$; one can easily generalize the definition to a larger class of encryption schemes whose ciphertext contains multiple component ciphertexts of a base encryption scheme, similarly to [21]; we omit the details here. The game proceeds almost identically to the CCA game except that instead of having access to the decryption oracle for the whole encryption scheme, the adversary has access to the decryption oracles of the basic encryption scheme PKE using each of the component secret keys; the i'th decryption oracle using the i'th component secret key is denoted as $\mathsf{Dec}(\mathbf{sk}[i],\cdot)$. As a result, the adversary cannot make any after-the-fact decryption queries that is the same as any of the component ciphertexts encrypted using one of the component public keys $\mathbf{pk}[i]$ in the challenge ciphertext. Similar to the CCA game, we define the XCCA-*advantage* of the adversary $\mathcal{A}$ as $\mathbf{Adv}^{\mathsf{XCCA}}_{\mathsf{PKE}_{\mathsf{in}}}(\mathcal{A}) = 2 \cdot \Pr[b' = b] - 1$.

We say that $\mathsf{PKE}_{\mathsf{in}}$ is XCCA-secure if no polynomial sized adversary can achieve a non-negligible advantage in the XCCA game.

We remark that the XCCA game is closely related to the notion of UCCA security defined in [21], and the similar notion of DCCA security in [22]: In comparison, in the UCCA security game w.r.t. $\mathsf{PKE}_{\mathsf{in}}$, the adversary only has access to the decryption oracle of the *whole encryption scheme*, but is not allowed to make any after-the-fact query that quotes any of the component ciphertexts in the challenge ciphertext (in DCCA a more fine grained control on disallowed queries is considered). As we will see shortly, the stronger security guarantee given by XCCA is crucial for our proof to succeed.

With the definition of the XCCA game in mind, the remainder of the proof proceeds via the following two lemmas, for which we give a proof sketch. (A formal proof is given in the full version.)

Lemma 2. *Let ε and δ be two negligible functions. Assume that* PKE *is* (α, β)*-secure, and* OKA *is a* $(\mathcal{D}(\alpha, \beta), \varepsilon, \delta, n, \ell_{\mathsf{in}})$*-one-way KA protocol. Then,* $\mathsf{PKE}_{\mathsf{in}}$ *is XCCA-secure.*

Lemma 3. *Assume that* $\mathsf{PKE}_{\mathsf{in}}$*,* $\mathsf{PKE}_{\mathsf{out},1}$ *and* $\mathsf{PKE}_{\mathsf{out},2}$ *are respectively XCCA, 1-CCA and CPA secure, and* $\mathsf{PKE}_{\mathsf{out},1}$ *and* $\mathsf{PKE}_{\mathsf{out},2}$ *have almost-perfect correctness, then* $\overline{\mathsf{PKE}}$ *is CCA secure.*

Proof Sketch of Lemma 2: We are going to use the hardcore lemma for CCA-security (Theorem 2) to show that $\mathsf{PKE}_{\mathsf{in}}$ is XCCA secure. Informally speaking, in the XCCA game, with respect to each random bit b_i used to generate the component c_i of the challenge ciphertext, the adversary is participating in an independently and randomly executed CCA game for PKE: Indeed, each random bit b_i is encrypted using an independently and randomly chosen public key $\mathbf{pk}[i]$ and random coins, and the adversary has access to the decryption oracle $\mathsf{Dec}(\mathbf{sk}[i], \cdot)$. Thus, by the hardcore lemma, each of these CCA games has probability $1 - \beta$ of delivering an "hard instance", and thus the corresponding bit b_i remains hidden to the adversary, i.e., it looks (pseudo-)random with probability $1 - \beta$. More precisely, each triple $(b_i, \mathsf{Dec}(\mathbf{sk}[i], c_i), c_i)$, with $c_i \xleftarrow{\$} \mathsf{Enc}(\mathbf{pk}[i], b_i)$ is computationally indistinguishability from a sample from a valid distribution from $\mathcal{D}(\alpha, \beta)$. In this case, then it simply follows from the fact that OKA is a $(\mathcal{D}(\alpha, \beta), \varepsilon, \delta, n, \ell_{\mathsf{in}})$-one-way key agreement scheme that the key k output by $\mathsf{KAEnc}(b_1, \cdots, b_n)$ remains random and thus the message m_b is hidden.

Proof Sketch of Lemma 3: We base the CCA security of $\overline{\mathsf{PKE}}$ on the XCCA security of $\mathsf{PKE}_{\mathsf{in}}$ via a black-box reduction. The reduction $\mathcal{B}$ participates in the XCCA game for $\mathsf{PKE}_{\mathsf{in}}$ and internally emulates the CCA game for $\overline{\mathsf{PKE}}$ to a CCA-adversary $\mathcal{A}$ succeeding with non-negligible advantage γ as follows:

- It receives the public key $\mathbf{pk}$ in the XCCA game and internally generates the public key $\overline{\mathsf{pk}}$ by sampling key pairs $(\mathsf{pk}_{\mathsf{out},1}, \mathsf{sk}_{\mathsf{out},1})$ and $(\mathsf{pk}_{\mathsf{out},2}, \mathsf{sk}_{\mathsf{out},2})$ for the two outer schemes and gives $\overline{p}k = (\mathbf{pk}, \mathsf{pk}_{\mathsf{out},1}, \mathsf{pk}_{\mathsf{out},2})$ to $\mathcal{A}$.

- To emulate the challenge ciphertext c^* of $\overline{\mathsf{PKE}}$ that encrypts either m_0 or m_1 chosen by $\mathcal{A}$ in the emulated CCA game, $\mathcal{B}$ first chooses random $r_{\mathsf{out},1}$ and $r_{\mathsf{out},2}$, and obtains the challenge ciphertext c^*_{in} of $\mathsf{PKE}_{\mathsf{in}}$ that encrypts $m_b \| r_{\mathsf{out},1} \| r_{\mathsf{out},2}$ for a random $b \in \{0,1\}$ chosen in the XCCA game. It then produces c^* honestly by encrypting $c^*_{\mathsf{out},1} = \mathsf{Enc}_{\mathsf{out},1}(c^*_{\mathsf{in}}; r_{\mathsf{out},1})$ and $c^*_{\mathsf{out},2} = \mathsf{Enc}_{\mathsf{out},2}(c^*_{\mathsf{in}}; r_{\mathsf{out},2})$.
- Finally, it emulates the decryption oracle $\overline{\mathsf{Dec}}(\overline{\mathsf{sk}}, \cdot)$ for $\mathcal{A}$ by using the secret key $\mathsf{sk}_{\mathsf{out},1}$ and the decryption oracles $\{\mathsf{Dec}(\mathsf{sk}[i], \cdot)\}_{i \in [n]}$ in the XCCA game.

It is easy to see that as long as $\mathcal{A}$ does not ask any after-the-fact queries whose inner ciphertext (embedded in the first outer ciphertext) "quotes" the inner challenge ciphertexts c^*_{in}, i.e., it does not share a common component ciphertext, $\mathcal{B}$ always decrypts queries from $\mathcal{A}$ perfectly and consequently also emulates the view of $\mathcal{A}$ perfectly.

It is therefore tempting to try to show that the probability that $\mathcal{A}$ "quotes" is negligible. Indeed, this is the approach taken by [21,22]. The rationale in their proof is that if the basic scheme PKE has unpredictability — a random ciphertext of a random bit has high entropy and is hard to blindly guess — then the fact that $\mathcal{A}$ manages to quote would violate the 1-CCA security of the first outer scheme or the CPA-security of the second outer scheme. In [22], a series of hybrids is used to remove the circular dependence between the inner challenge ciphertext and the randomness used in its two outer encryptions, and move to a setting where $\mathcal{A}$'s view is *statistically* independent from the inner challenge ciphertext, but the quoting probability is negligibly close to the original one. One can then easily show that unpredictability of PKE yields that quoting occurs with negligible probability only.

Unfortunately, this approach fails completely in our setting, as our basic encryption scheme PKE does not ensure unpredictability; in fact, it is possible to build an (α, β)-CCA-secure scheme where ciphertexts have very low min-entropy. We address this via a new technique, called *heavy ciphertext pre-sampling*: We observe that if $\mathcal{A}$ can blindly guess some component ciphertext c_i in c^*_{in}, then c_i is a ciphertext value which appears with sufficiently large probability when encrypting a random bit under $\mathbf{pk}[i]$. Hence, we can hope that the same value is hit by the reduction $\mathcal{B}$ by simply generating a large number of random encryptions (of random bits) of PKE under $\mathbf{pk}[i]$; call these pre-sampled ciphertexts. Since the component ciphertexts in c^*_{in} are generated identically to the pre-sampled ciphertexts, the probability that $\mathcal{A}$'s guess collides with the former is the same as the probability it collides with any of the pre-sampled ciphertexts. Setting the size of the pre-sampling large enough, say $\mathsf{poly}(1/\varepsilon)$, the reduction can exhaust all the component ciphertexts that $\mathcal{A}$ may "quote" with probability $1 - \varepsilon$, for any ε. Furthermore, due to the strong security provided by the XCCA game, the reduction $\mathcal{B}$, with access to the decryption oracles of the component ciphertexts, can obtain the decrypted values of these pre-sampled ciphertexts before-the-fact. This is crucial, since even if *we know* that a ciphertext is obtained by encrypting some bit d, its actual decryption could well be equal $1 - d$ due to the weak α-correctness.

Intuitively this solves the problem, as whenever $\mathcal{A}$ makes an after-the-fact query that "quotes" c^*_{in}, $\mathcal{B}$ can still decrypt by using either the external decryption oracles (for components that do not quote) *or* the decrypted values of the pre-sampled ciphertexts (for these that quote). This will allow us to show that $\mathcal{B}$ succeeds in emulating the view of $\mathcal{A}$ with high probability, and thus the CCA security of $\overline{\mathsf{PKE}}$ reduces to the XCCA security of $\mathsf{PKE}_{\mathsf{in}}$.

Acknowledgments. We wish to thank Russell Impagliazzo and Thomas Ristenpart for insightful discussions at early stages of this work.

This work was partially supported by NSF grant CCF-1018064. This material is based on research sponsored by DARPA under agreement number 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

1. Goldwasser, S., Micali, S.: Probabilistic encryption. Journal of Computer and System Sciences 28(2), 270–299 (1984)
2. Rackoff, C., Simon, D.R.: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 433–444. Springer, Heidelberg (1992)
3. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC. ACM Press (May 1990)
4. Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography. In: 23rd ACM STOC, pp. 542–552. ACM Press (May 1991)
5. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: 40th FOCS, pp. 543–553. IEEE Computer Society Press (October 1999)
6. Gertner, Y., Malkin, T., Myers, S.: Towards a separation of semantic and CCA security for public key encryption. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 434–455. Springer, Heidelberg (2007)
7. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002)
8. Wee, H.: Efficient chosen-ciphertext security via extractable hash proofs. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 314–332. Springer, Heidelberg (2010)
9. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC, pp. 187–196. ACM Press (May 2008)
10. Rosen, A., Segev, G.: Chosen-ciphertext security via correlated products. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 419–436. Springer, Heidelberg (2009)
11. Kiltz, E., Mohassel, P., O'Neill, A.: Adaptive trapdoor functions and chosen-ciphertext security. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 673–692. Springer, Heidelberg (2010)

12. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 1993, pp. 62–73. ACM Press (November 1993)
13. Bellare, M., Rogaway, P.: Optimal asymmetric encryption. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 92–111. Springer, Heidelberg (1995)
14. Yao, A.C.: Theory and applications of trapdoor functions. In: 23rd FOCS, pp. 80–91. IEEE Computer Society Press (November 1982)
15. Dwork, C., Naor, M., Reingold, O.: Immunizing encryption schemes from decryption errors. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 342–360. Springer, Heidelberg (2004)
16. Holenstein, T.: Key agreement from weak bit agreement. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC, pp. 664–673. ACM Press (May 2005)
17. Holenstein, T., Renner, R.S.: One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 478–493. Springer, Heidelberg (2005)
18. Lenstra, A.K., Hughes, J.P., Augier, M., Bos, J.W., Kleinjung, T., Wachter, C.: Public keys. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 626–642. Springer, Heidelberg (2012)
19. Heninger, N., Durumeric, Z., Wustrow, E., Halderman, J.A.: Mining your ps and qs: Detection of widespread weak keys in network devices. In: Proceedings of the 21st USENIX Security Symposium (2012)
20. Kearns, M.J., Valiant, L.G.: Cryptographic limitations on learning boolean formulae and finite automata. J. ACM 41(1), 67–95 (1994)
21. Myers, S., Shelat, A.: Bit encryption is complete. In: 50th FOCS, pp. 607–616. IEEE Computer Society Press (October 2009)
22. Hohenberger, S., Lewko, A., Waters, B.: Detecting dangerous queries: A new approach for chosen ciphertext security. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 663–681. Springer, Heidelberg (2012)
23. Impagliazzo, R.: Hard-core distributions for somewhat hard problems. In: FOCS 1995, pp. 538–545 (1995)
24. Maurer, U.M.: Protocols for secret key agreement by public discussion based on common information. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 461–470. Springer, Heidelberg (1993)
25. Tessaro, S.: Security amplification for the cascade of arbitrarily weak pRPs: Tight bounds via the interactive hardcore lemma. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 37–54. Springer, Heidelberg (2011)
26. Impagliazzo, R., Luby, M.: One-way functions are essential for complexity-based cryptography. In: 30th FOCS, pp. 230–235. IEEE Computer Society Press (October / November 1989)
27. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM Journal on Computing 28(4), 1364–1396 (1999)
28. Cramer, R., Hanaoka, G., Hofheinz, D., Imai, H., Kiltz, E., Pass, R., Shelat, A., Vaikuntanathan, V.: Bounded CCA2-secure encryption. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 502–518. Springer, Heidelberg (2007)