

Detecting Dangerous Queries: A New Approach for Chosen Ciphertext Security

Susan Hohenberger^{1,*}, Allison Lewko^{2,**}, and Brent Waters^{2,***}

¹ Johns Hopkins University, susan@cs.jhu.edu

² University of Texas at Austin, {alewko,bwaters}@cs.utexas.edu

Abstract. We present a new approach for creating chosen ciphertext secure encryption. The focal point of our work is a new abstraction that we call *Detectable Chosen Ciphertext Security* (DCCA). Intuitively, this notion is meant to capture systems that are not necessarily chosen ciphertext attack (CCA) secure, but where we can detect whether a certain query CT can be useful for decrypting (or distinguishing) a challenge ciphertext CT*.

We show how to build chosen ciphertext secure systems from DCCA security. We motivate our techniques by describing multiple examples of DCCA systems including creating them from 1-bit CCA secure encryption — capturing the recent Myers-shelat result (FOCS 2009). Our work identifies DCCA as a new target for building CCA secure systems.

1 Introduction

A central goal of public key cryptography is to design encryption systems that are secure against chosen ciphertext attacks. Public key encryption systems that are chosen ciphertext attack (CCA) secure are robust against powerful adversaries that are able to leverage interaction with a decryptor. Such an attacker is modeled by allowing him to query for the decryption of any ciphertext except a challenge ciphertext for which he is trying to break. This includes ciphertexts derived from the challenge ciphertext³. Due to its robustness against powerful attackers, chosen ciphertext security has become the accepted goal for building secure encryption. For this reason, building chosen ciphertext secure systems has

* Sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under contract FA8750-11-2-0211, the Office of Naval Research under contract N00014-11-1-0470, a Microsoft Faculty Fellowship and a Google Faculty Research Award. Applying to all authors, the views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

** Sponsored by a Microsoft Research Ph.D. Fellowship.

*** Supported by NSF CNS-0915361 and CNS-0952692, AFOSR Grant No: FA9550-08-1-0352, DARPA PROCEED, DARPA N11AP20006, Google Faculty Research award, the Alfred P. Sloan Fellowship, Microsoft Faculty Fellowship, and Packard Foundation Fellowship.

³ We use “CCA” and “CCA2” interchangeably in this paper.

been a central pursuit of cryptographers for over twenty years and we have seen many distinct approaches to achieving CCA security.

Early pioneering work in chosen ciphertext security [23, 14, 26] introduced the technique of leveraging Non-Interactive Zero Knowledge Proofs (NIZKs) [5] to build CCA-secure encryption systems from chosen plaintext secure encryption systems. Roughly, a NIZK is used to prove that a ciphertext is “well-formed” or legal. Later Cramer and Shoup [12, 13] introduced the first practical CCA-secure systems that were built on specific number theoretic assumptions such as Decisional Diffie Hellman. These techniques implicitly embed a certain form of designated verifier Non-Interactive Zero Knowledge proofs in them. More recently, different methods for building chosen ciphertext security from Identity-Based Encryption [7] and Lossy Trapdoor Functions [25] have emerged. In addition, Myers and shelat [22] described general methods for amplifying CCA encryption of 1 bit to many bits.

In this work, we introduce a new approach to obtaining chosen ciphertext secure systems. The focal point of our work is a new abstraction that we call *Detectable Chosen Ciphertext Security* (DCCA). Intuitively, this notion is meant to capture systems that are not necessarily CCA secure, but where we can detect whether a certain query CT can be useful for decrypting (or distinguishing) a challenge ciphertext CT^* .

A system that is DCCA secure will be associated with a boolean function F that takes in three inputs: a public key pk , a challenge ciphertext CT^* and a query ciphertext CT. The function will output 1 if the query CT is “dangerous” for the an attacker wishing to distinguish CT^* . A DCCA secure system must have the following two properties stated here informally:

- **Unpredictability** Without seeing CT^* it should be hard to find a ciphertext CT such that $F(PK, CT^*, CT) = 1$. In other words, an attacker must first see a challenge ciphertext in order to discover a dangerous query for it.
- **Indistinguishability** The system will be secure under a detectable chosen ciphertext attack *if* the attacker is limited to decryption queries of ciphertexts CT where $F(pk, CT^*, CT) = 0$ for challenge ciphertext CT^* . I.e. the system is CCA secure if the attacker does not make dangerous queries.

The goal of our work will be to construct fully chosen ciphertext secure systems from detectable CCA-secure systems. We first motivate this goal by observing multiple DCCA systems that naturally occur:

- **Many bit encryption from 1-bit CCA** Suppose we have a 1-bit CCA-secure system and we wish to encrypt multiple bits by concatenating multiple 1-bit encryptions together. The resulting system is no longer chosen ciphertext secure, but is DCCA secure. The detecting function F is 1 iff any of the 1-bit ciphertext components between CT^* and CT are equal. This scenario is akin to the problem of showing that bit encryption is complete considered by Myers and shelat [22], where they worried about such “quoting” attacks.
- **Tag-Based Encryption Systems** MacKenzie, Reiter and Yang [21] and Kiltz [19] define a tag-based encryption scheme as an encryption scheme that

takes in an additional “tag” parameter on encryption and decryption. The security game allows an attacker to make decryption queries with any tag parameter t , except for the tag t^* that the challenge ciphertext is encrypted under. Several examples of tag-based schemes exist. Kiltz [19] gave a direct construction from the linear assumption. The CCA1-secure encryption variant of the Canetti, Halevi and Katz [7] construction where the tag is an IBE identity is an additional example. One can also view the CCA1-secure variant of Peikert and Waters [25] as a tag-based scheme, where the tag is the “branch” in an all-but-one encryption scheme.

Most of the above examples of tag-based encryption can be proven selectively secure, where an attacker must commit to the tag of the challenge ciphertext before seeing the public key. However, if we are willing to utilize complexity leveraging arguments, we can argue that these are adaptively secure. In addition, the CHK-lite transformation will be an adaptively secure tag-based scheme if used with an adaptively secure Identity-Based Encryption system. We observe that adaptively-secure tag-based encryption immediately gives rise to DCCA-secure encryption. A ciphertext of the DCCA-secure system consists of a random tag t plus a tag-based encryption of the message under the tag t . Decryption follows analogously and the function F simply tests if two ciphertexts have the same tag. Unpredictability follows from having a large tag space. Although it is already possible to transform tag-based encryption into CCA-secure encryption using a strongly unforgeable signature [19], these examples demonstrate natural DCCA systems.

- **“Sloppy” CCA Encryption** One can envision that in practice an encryption system is CCA secure, but an implementation of it is not due to certain nuances. For instance, suppose a number theoretic library had a slack bit in its representation of group elements (e.g. a bit that was supposed to be 0, but if set to 1 does not affect any computations.) A CCA attacker could exploit this weakness in an implementation, however, it is possible that the system would still be DCCA secure. One might use our techniques as a hedge against such problems. This is somewhat analogous to recent work [2] on applying deterministic encryption as a hedge against faulty random bit generation.

In addition to the examples listed above, we believe that it is useful to identify DCCA security as a new “target” for achieving chosen ciphertext security.

Overview of Our Techniques We now give an overview of our construction and proof. Our construction will build a chosen ciphertext secure system from three components: a chosen plaintext secure system, 1-bounded CCA-secure system⁴, and a detectable CCA-secure system. Since DCCA security (trivially) implies CPA, and we can build 1-bounded CCA from CPA encryption [24, 11, 10], it follows that all components are realizable from DCCA as a building block.

A public key from our system consists of three components. An “inner” public key PK_{in} which is a DCCA public key and two “outer” keys PK_A, PK_B respec-

⁴ A 1-bounded CCA-secure encryption system is secure against one chosen ciphertext query.

tively from 1-bounded CCA and CPA secure systems. To encrypt a message M , one first chooses the randomness r_A, r_B to be used for the outer encryptions and then encrypts the tuple (r_A, r_B, M) under the inner (detectable) key to compute an inner ciphertext CT_{in} . Next, the encryption algorithm encrypts CT_{in} under the outer public key PK_A using randomness r_A to get CT_A . It then analogously creates CT_B as the encryption of CT_{in} under key PK_B and randomness r_B . The output ciphertext is $CT = (CT_A, CT_B)$.

The structure of our ciphertexts is that the two outer ciphertexts both encrypt the same message — the inner ciphertext. This ciphertext itself encrypts the message and the randomness used to create the outer ciphertexts. Thus, the outer ciphertexts indirectly encrypt their own randomness.⁵ The decryption algorithm will receive $CT = (CT_A, CT_B)$ and first decrypt CT_A to get CT'_{in} and decrypt this to get (r'_A, r'_B, M') using the appropriate secret keys. Finally, it will check that the ciphertext is well formed by itself encrypting CT'_{in} under PK_A, PK_B and the respective randomness r'_A, r'_B and validating that the output matches CT_A and CT_B before accepting M' as the message. Our encryption system has elements both of the Naor-Yung [23] two key method for our two outer keys and the Myers-shelat [22] method of embedding outer randomness in inner ciphertexts.

Security of our system depends on the premise that no attacker is able to learn the message encrypted in the inner ciphertext. This will follow from the Detectable CCA security *if* we are able to guarantee that an attacker is unable to make any ciphertext queries CT_A, CT_B where the decryption of CT_A , denoted CT_{in} , is related to the inner component of our challenge ciphertext CT_{in}^* according to the DCCA function F . Intuitively, we hope to achieve this from the combination of two features of our system. First, the 1-bounded CCA security of PK_A will (hopefully) make it difficult to create an encryption under PK_A related to CT_{in}^* . Second, the embedded randomness will allow us to check that ciphertexts are well formed and thus answer multiple ciphertext queries under the Naor-Yung two key type manner.

The trickiness in proving security lies with the embedded randomness which is a two-edge sword. On one hand, forcing the attacker queries to embed randomness allows a reduction algorithm to decrypt if it knows either one of the two outer keys. On the other hand, it is not clear how such a reduction can create valid ciphertexts while playing the 1-bounded CCA game, since a reduction algorithm will not know the randomness r_A to embed. Thus, this circularity

⁵ This construction implicitly assumes that the length of the random string needed for encryption is dependent only on the security parameter and is independent (or at least smaller than) the message size of the outer ciphertexts. We can justify this assumption with the common technique of using a seed to a (variable length) Pseudo Random Generator (PRG) as the input to each encryption algorithm. The PRG can then extend the randomness to whatever length is required by the underlying encryption system. By using this justified assumption in our definitions, we are able to simplify the presentation of our construction and proofs. In contrast, Myers and shelat [22] explicitly carry the PRG technique through their exposition. This choice gives our exposition and proof an advantage in simplicity.

creates a fundamental barrier similar to difficulties encountered in attempts to create trapdoor functions from encryption [15].

We deal with this by arguing security in an indirect way that steps around this barrier. We first define a security game specific to our construction called nested indistinguishability. In this game, an attacker will receive a public key and is allowed to make decryption queries. The attacker at some point submits a single message M . The challenger will flip a coin z . If $z = 0$, the challenger creates a valid encryption of M ; otherwise, if $z = 1$ the challenger creates a encryption where the innermost message is all 0's — it neither includes the message nor the embedded randomness. The attacker continues to make decryption queries (other than the challenge ciphertext) and wins if it is successfully able to guess z . It follows that if no attacker is successful in this game, then our system is chosen ciphertext secure.

To prove security of this nested indistinguishability game, we begin by defining a “bad event”. The bad event is defined to be when the attacker submits a query (CT_A, CT_B) such that $CT_A \neq CT_A^*$ where CT_A^* is from the challenge ciphertext and the decryption of CT_A gives a ciphertext that is related to the inner challenge ciphertext according to F . If we can argue that such bad events only occur with negligible probability, then security of the nested indistinguishability game follows straightforwardly from DCCA security.

The crux of our proof is how we eliminate the possibility of a bad event. We do so in an indirect manner. We begin by arguing this event cannot happen in the case where $z = 1$, which is where all 0's are encrypted and the randomness is not embedded. In this case, we get the best of both worlds. We are able to require that the attacker's queries have the randomness embedded in them, so that we can check ciphertext well-formedness, however, *the challenge ciphertext is not required to embed the outer randomness*. We argue that the bad event does not happen by applying a set of hybrid experiments. First, we change CT_B^* to be an encryption of all 1's. Next, we change the decryption algorithm to decrypt using the secret key for PK_B . Finally, we change CT_A^* to be an encryption of all 1's. In each experiment we argue that the chance of a bad event must be very close to that of the prior experiment. For the last step we leverage the 1-bounded CCA property of the first component. Finally, we note that in the last experiment the probability of a bad event is negligible since the inner challenge ciphertext CT_{in}^* is replaced by all 1's and is not even present.

One interesting question is why is 1-bounded CCA security needed for the PK_A since at the last step in the proof we can use the secret key SK_B to execute decryption. While this is true, it is actually possible for the bad event to occur on a malformed ciphertext that will not decrypt. We need the 1-bounded CCA property to detect the occurrence of the bad event in this case during the security reduction.

We are not able to argue the lack of a bad event in a similar manner for the $z = 0$ (embedded randomness) case due to the aforementioned circularity problems. Instead, we can infer this from the lack of event in the $z = 1$ case along with DCCA security. To prove this, we can create an algorithm that plays

the DCCA indistinguishability game while simulating the nested indistinguishability game to the attacker. The simulator will choose the outer keys and outer randomness for the challenge ciphertext itself. It submits the message and outer randomness as one inner message and the 0's string as another. Then it will be able to decrypt all ciphertext queries until a bad event happens using its keys in addition to the DCCA decryption oracle. Once a bad event query is made though, it is stuck. However, it need not go any further! The fact that the attacker was able to create a bad event at all must mean that the message and randomness were embedded. It can then break the DCCA distinguishing game. Thus, we can infer that the bad event happens with negligible probability in either case. The remainder of the proof follows straightforwardly.

Comparison to Myers-shelat Myers and shelat [22] showed how to achieve many-bit chosen ciphertext security from 1-bit chosen ciphertext security and motivated us to explore the notion of detectability. They created a system using an inner/outer structure where the inner ciphertext encrypted the outer random coins. Their inner scheme, built from 1-bit CCA, is what they call “unquoteable” secure. Their concept is roughly analogous to a specific instance of a DCCA scheme. Encryptions of many-bit messages are concatenations of 1-bit encryptions; the system is chosen ciphertext secure as long as queries do not copy a 1-bit ciphertext component of the underlying scheme. For the outer scheme, they use a notion of security that is an amalgam of unquoteability and non-malleability. Their outer construction follows a specific adaptation of the Choi et. al. [10] methods applied to the 1-bit primitive. (No two key structure is used.) Their proof relies on defining quoting attacks on *both* the inner and outer layers and then establishing a certain order that outer quoting attacks must happen before inner quoting attacks.

We believe our methods offer benefits in terms of generality, simplicity, and efficiency. First, our general notion of Detectable Chosen Ciphertext Security can be realized by multiple systems. These include the 1-bit to many-bit examples, the tag-based encryption class and future systems that can leverage this as a new target path for creating CCA secure encryption.

Another key difference is that the outer layer of our scheme is built from simple 1-bounded CCA and CPA-secure parts. We argue these provide simpler concepts and are easier to work with. In addition, one can instantiate them from *any* 1-bounded encryption system. For instance, we can apply any candidate 1-bounded CCA-secure system and do not need to work through the Choi et. al. [10] construction. Instead we can apply the 1-bounded CCA system of Cramer et. al. [11], which is significantly more efficient and simpler than the non-malleable systems of either PSV [24] or Choi et. al. [10]. We also regard avoiding a combination security definition between 1-bounded CCA (or non-malleability) and detection as a benefit for simplicity. This simplification will also improve efficiency in the case where there is a candidate CPA primitive that is more efficient than the candidate DCCA primitive, since we can build the 1-bounded scheme out of the CPA primitive.

Our choice of abstractions and structure allow us to have a simple proof. We can eliminate the possibility of a bad event using a basic Naor-Yung two key argument. Then once we are able to eliminate this, the rest of the proof follows in a straightforward manner.

Why not CCA1? One intriguing possibility is to try to leverage our techniques to build full chosen ciphertext security from CCA1 security. A natural direction would be to use a CCA1 system for the inner component in place of the detectable encryption scheme. The intuitive rationale would be if the outer keys are 1-bounded CCA or non-malleable then the queries produced by the attacker should not be related to the inner challenge ciphertext and thus CCA1 might suffice. Unfortunately, we were able to create an attack oracle which breaks full CCA security in our scheme, yet does not perturb the 1-bounded CCA or CCA1 primitives, giving evidence that this approach may not work. However, the oracle we use is quite strong and “exotic”. This suggests that there might be primitives that lie somewhere in between DCCA and CCA1. One interesting example is the CCA-1 secure “Cramer-Shoup lite” [12] cryptosystem. There exists a malleability attack on a challenge CT^* that produces a query ciphertext which has the same distribution as a fresh encryption of a random message. Hence the CS-lite system is not CCA secure. However, it would be very interesting and surprising if there existed attack algorithms that matched the above oracle. We expand on this in the full version of this paper.

1.1 Related Work

Relaxations of CCA Multiple relaxations of chosen ciphertext security have been proposed in the literature.

One interesting class of relaxations is the notion of Replayable Chosen Ciphertext Security [8] and other similar works [29, 1]. These works aim to capture the concept that some malleability attacks might intuitively be benign. In particular, consider a cryptosystem where an attacker is only able to maul a ciphertext CT encrypting a message M into a different ciphertext C' that encrypts the *same* message M . If an application (or user) makes all decisions based on the decrypted plaintexts as opposed to the representation of the ciphertext such notions might be sufficient.

The primary goal of RCCA is to formally capture a form of “good enough” security under ciphertext attacks. In contrast, Detectable CCA inherently does not have good enough security on its own. In DCCA systems, it may be possible to maul ciphertexts to be encryptions of different messages or even create attack ciphertexts that each target a single bit of a target ciphertext. Thus, our primary focus is to create CCA security from a less secure DCCA building block.

We observe that DCCA does not imply RCCA. In [8], the authors gave an example of an RCCA scheme that could not be publicly detected. Conversely, not all DCCA schemes will be RCCA secure. Our bit encryption instance serves as an example. We also note that [8] discusses a notion of detectability and introduces a definition that combines replayable and detectable properties. This

combined definition is a particular instance of DCCA. However, they do not explore the notion of detectability in isolation or how to build CCA security from it. Canetti, Krawczyk, and Nielsen [8] do show how to create CCA security from RCCA security using the KEM/DEM framework.

Finally, Hofheinz and Kiltz [17] introduce a notion they call Constrained CCA security particular to developing Key Encapsulation Mechanisms. In their definition an attacker must include a predicate p along with each query ciphertext CT. The challenger will only answer the query if the predicate evaluated on the decrypted key of the ciphertext is true and the predicate is false for all but a negligible fraction of possible KEM keys. While this notion is weaker than CCA security, they show that when combined with a (symmetric) authenticated encryption scheme, the resulting system is CCA secure.

Other Related Work Goldwasser and Micali [16] gave the first formal definition of security for public key encryption systems. Naor and Yung [23] and Rackoff and Simon [26] extended this to include chosen ciphertext attacks.

Naor and Yung [23] initiated the approach of leveraging NIZKs to build chosen ciphertext security by introducing their “two key” method. A NIZK would guarantee the integrity of the ciphertext by giving a proof that the same message was encrypted to two keys. While their system gave security against lunchtime or CCA1 attacks, Dolev, Dwork and Naor [14] showed how to achieve full CCA2 security. In addition, they introduced the fundamental concept of non-malleability. Sahai [28] introduced a concept of simulation sound NIZKs that could be used to achieve CCA security through the NY two key structure. Bellare and Sahai [4] gave relations between non-malleability [14] chosen ciphertext security.

Since then, different approaches to achieving CCA security have been proposed. Cramer and Shoup [12, 13] showed techniques for proving ciphertexts were well-structured and abstracted this into projective hash functions. Several other novel cryptosystems make use of specific number-theoretic techniques (e.g. [19, 9, 18]). Boneh, Canetti, Halevi and Katz [6] showed a generic method of achieving chosen ciphertext security from IBE systems. Peikert and Waters [25] gave a new avenue for achieving CCA security with the introduction of Lossy Trapdoor Functions (TDFs). Notably, this gave the first chosen ciphertext secure systems from lattice-based assumptions. Subsequently, various refinements of weaker conditions on the trapdoor functions were introduced [27, 20].

The above techniques are proven secure in the standard model. Bellare and Rogaway [3] show that in the random oracle model chosen ciphertext security can be built from chosen plaintext security.

2 Detectable Chosen Ciphertext Security

In this section, we define *detectable chosen ciphertext security*. An encryption scheme satisfying this definition is called a *detectable* encryption system. Our discussions assume a familiarity with CPA, CCA1 and CCA2 security as well as bounded CCA security and non-malleability. A reader wishing to review these definitions can find them in the full version of this paper.

2.1 Detectable Encryption

We define a *detectable* encryption scheme as having the usual algorithms (KeyGen, Enc, Dec) together with an efficiently-computable boolean function F . Informally, F tests for a “detectable” relationship between two ciphertexts. The security game will mirror that of CCA2 security, except that decryption queries in the second phase will not be answered for ciphertexts detectably-related to the challenge ciphertext. Our formal definition follows below.

Definition 1 (Detectable Encryption System). A detectable encryption system is a tuple of probabilistic polynomial-time algorithms (KeyGen, Enc, Dec, F) such that:

1. (KeyGen, Enc, Dec) have the usual input and output, although we sometimes denote $\text{Enc}(pk, m; r)$ as a deterministic function of the public key pk , the message m and randomness r , and
2. $F(pk, c', c) \rightarrow \{0, 1\}$: the detecting function F takes as input a public key pk and two ciphertexts c' and c , and outputs a bit.

Correctness is the same as a regular encryption system.

A detectable encryption system must have the following two properties.

Unpredictability of the Detecting Function F . Informally, given the description of F and a public key pk , for an unknown ciphertext c , it should be hard to find a second ciphertext c' that is “related” to c ; i.e., such that $F(pk, c', c) = 1$. We consider both a basic and a strong formalization.

Basic Unpredictability Experiment. Consider the experiment $\text{Exp}_{\mathcal{A}, \Pi}^{\text{predict.basic}}(\lambda)$ defined for a detectable encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, F)$ and an adversary \mathcal{A} :

1. Setup: $\text{KeyGen}(1^\lambda)$ is run to obtain keys (pk, sk) .
2. Queries: Adversary \mathcal{A} is given pk and access to a decryption oracle $\text{Dec}(sk, \cdot)$. The adversary outputs a message m in the message space associated with pk and a ciphertext c in the ciphertext space associated with pk .
3. Challenge: A ciphertext $c^* \leftarrow \text{Enc}(pk, m)$ is computed.
4. Output: The output of the experiment is defined to be 1 if $F(pk, c^*, c)$, and 0 otherwise.

We also define a stronger variant $\text{Exp}_{\mathcal{A}, \Pi}^{\text{predict.strong}}(\lambda)$ of the unpredictability experiment where the adversary is additionally given sk . We observe that strong unpredictability implies basic unpredictability since the adversary can simulate the decryption oracle using the secret key.

Indistinguishability of Encryptions. Next, we formalize the confidentiality guarantee. Consider the following experiment $\text{Exp}_{\mathcal{A}, \Pi}^{\text{indist}}(\lambda)$ defined for a detectable encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, F)$ and an adversary \mathcal{A} :

1. Setup: $\text{KeyGen}(1^\lambda)$ is run to obtain keys (pk, sk) .
2. Phase 1: Adversary \mathcal{A} is given pk and access to a decryption oracle $\text{Dec}(sk, \cdot)$. \mathcal{A} outputs a pair of messages m_0, m_1 of the same length in the message space associated with pk .
3. Challenge: A random bit $b \leftarrow \{0, 1\}$ is chosen, and then a ciphertext $c^* \leftarrow \text{Enc}(pk, m_b)$ is computed and given to \mathcal{A} . We call c^* the challenge ciphertext.
4. Phase 2: \mathcal{A} continues to have access to $\text{Dec}(sk, \cdot)$, but may not request a decryption of a ciphertext c such that $F(pk, c^*, c) = 1$. Finally, \mathcal{A} outputs a bit b' .
5. Output: The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise.

Definition 2 (Detectable Chosen Ciphertext Security). *A detectable encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, F)$ has an unpredictable detecting function and indistinguishable encryptions under a detectable chosen-ciphertext attack (or is DCCA-secure) if for all probabilistic polynomial-time adversaries \mathcal{A} there exists a negligible function negl such that:*

1. (*F is unpredictable:*) $\Pr[\text{Exp}_{\mathcal{A}, \Pi}^{\text{predict. basic}}(\lambda) = 1] \leq \text{negl}(\lambda)$ and
2. (*Encryptions are indistinguishable:*) $\Pr[\text{Exp}_{\mathcal{A}, \Pi}^{\text{indist}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda)$.

2.2 Facts about DCCA Security

For space reasons, we omit the simple proofs of the first two lemmas. We conjecture that the converse of Lemma 2 is not true. Indeed if the DDH assumption holds, then the CCA-1 secure Cramer-Shoup lite system would separate these two notions as discussed in the introduction.

Lemma 1 (CCA2 \implies DCCA). *If $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is a CCA2-secure encryption scheme, then $\Pi' = (\text{KeyGen}, \text{Enc}, \text{Dec}, F)$ is a DCCA-secure encryption scheme where F outputs 0 on all inputs except those of the form (\cdot, c, c) .*

Lemma 2 (DCCA \implies CCA1). *If $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, F)$ is a DCCA-secure encryption scheme, then $\Pi' = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is a CCA1-secure encryption scheme.*

We also claim that one-bit DCCA-secure encryption implies arbitrary-length DCCA-secure encryption. Say $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, F)$ is a detectable encryption system with plaintext space $\{0, 1\}$. We can construct a new scheme $\Pi' = (\text{KeyGen}, \text{Enc}', \text{Dec}', F')$ with plaintext space $\{0, 1\}^*$ by defining Enc' as:

$$\text{Enc}'(pk, m) = \text{Enc}(pk, m_1), \dots, \text{Enc}(pk, m_n)$$

where $m = m_1 \dots m_n$. The decryption algorithm Dec' decrypts each ciphertext piece using Dec . The function F' performs n^2 invocations of F , testing each ciphertext piece of C with each ciphertext piece of C' , and outputting 1 if any invocation of F returned 1, and 0 otherwise.

Lemma 3 (1-bit DCCA encryption implies many-bit DCCA encryption). *Let Π and Π' be as above. If Π is DCCA-secure, then so is Π' .*

We defer the proof of this lemma to the full version of the paper.

3 The Construction: CCA2 Security from DCCA Security

An overview of the techniques used for our construction is provided in Section 1.

The Construction Description We now construct a CCA2-secure public-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ using three building blocks⁶:

1. a DCCA-secure encryption scheme, denoted $\Pi_{\text{dcca}} = (\text{KeyGen}_{\text{dcca}}, \text{Enc}_{\text{dcca}}, \text{Dec}_{\text{dcca}}, F)$.
2. a 1-bounded CCA-secure encryption scheme with perfect correctness, denoted $\Pi_{1\text{b-cca}} = (\text{KeyGen}_{1\text{b-cca}}, \text{Enc}_{1\text{b-cca}}, \text{Dec}_{1\text{b-cca}})$.
3. a CPA-secure encryption scheme with perfect correctness, denoted $\Pi_{\text{cpa}} = (\text{KeyGen}_{\text{cpa}}, \text{Enc}_{\text{cpa}}, \text{Dec}_{\text{cpa}})$.

We assume that the message space of each system is $\{0, 1\}^*$ and that messages of the form (x, y, z) can be uniquely and efficiently encoded as strings in $\{0, 1\}^*$, where the encoding length is the same for all inputs of the same length. We assume that λ bits will be sufficient randomness for the encryption algorithm of each system, where 1^λ is the security parameter. We assume that $\Pi_{1\text{b-cca}}$ and Π_{cpa} have perfect correctness for decryption. Finally, we assume that for Π_{dcca} the ciphertext length is a deterministic function of the security parameter and the message length. We further justify these assumptions in the full version.

KeyGen(1^λ) Run $\text{KeyGen}_{\text{dcca}}(1^\lambda)$ to produce $(\text{PK}_{\text{in}}, \text{SK}_{\text{in}})$, $\text{KeyGen}_{1\text{b-cca}}(1^\lambda)$ to produce $(\text{PK}_A, \text{SK}_A)$, and $\text{KeyGen}_{\text{cpa}}(1^\lambda)$ to produce $(\text{PK}_B, \text{SK}_B)$. Set the public key as $\text{PK} := (\text{PK}_{\text{in}}, \text{PK}_A, \text{PK}_B)$ and the secret key as $\text{SK} := (\text{SK}_{\text{in}}, \text{SK}_A, \text{SK}_B)$.

Enc(PK, M) The encryption algorithm first chooses random strings $r_{\text{in}}, r_A, r_B \in \{0, 1\}^\lambda$. Next, it computes the ciphertext $\text{CT}_{\text{in}} := \text{Enc}_{\text{dcca}}(\text{PK}_{\text{in}}, (r_A, r_B, M); r_{\text{in}})$. It treats this ciphertext as the message and computes $\text{CT}_A := \text{Enc}_{1\text{b-cca}}(\text{PK}_A, \text{CT}_{\text{in}}; r_A)$ and $\text{CT}_B := \text{Enc}_{\text{cpa}}(\text{PK}_B, \text{CT}_{\text{in}}; r_B)$. Finally, it outputs $(\text{CT}_A, \text{CT}_B)$.

Dec(SK, CT) The decryption algorithm takes a ciphertext $\text{CT} := (\text{CT}_A, \text{CT}_B)$. It decrypts the first ciphertext as $\text{CT}_{\text{in}} := \text{Dec}_{1\text{b-cca}}(\text{SK}_A, \text{CT}_A)$. It then decrypts this output as $(r_A, r_B, M) := \text{Dec}_{\text{dcca}}(\text{SK}_{\text{in}}, \text{CT}_{\text{in}})$. It then checks that

$$\text{CT}_A = \text{Enc}_{1\text{b-cca}}(\text{PK}_A, \text{CT}_{\text{in}}; r_A) \text{ and } \text{CT}_B = \text{Enc}_{\text{cpa}}(\text{PK}_B, \text{CT}_{\text{in}}; r_B).$$

If all checks pass, it outputs M ; otherwise, it outputs \perp .

⁶ A 1-bounded CCA-secure encryption system is secure if an attacker makes at most one decryption query. One-bounded CCA security can be constructed from CPA security [24, 10]. CPA security is trivially implied by DCCA security. Thus, there is really only one necessary building block: a DCCA-secure system.

4 Proof of Security

We will now argue that the Section 3 construction is CCA2 secure, assuming the respective security properties of the underlying building blocks. To do so, it will be easier to consider a slight variant of the CCA2 security game, which we call *nested indistinguishability*, where the challenger either encrypts one of the two challenge messages or encrypts a string of zeros. The experiment involves three encryption schemes and combines them in the same manner as our main construction.

Nested Indistinguishability. Consider the experiment $\text{Exp}_{\mathcal{A}, \Pi_{\text{dcca}}, \Pi_{1\text{b-cca}}, \Pi_{\text{cpa}}}^{\text{nested}}(\lambda)$ defined for detectable encryption scheme Π_{dcca} , encryption schemes $\Pi_{1\text{b-cca}}, \Pi_{\text{cpa}}$ and an adversary \mathcal{A} :

1. Setup: Run $\text{KeyGen}_{\text{dcca}}, \text{KeyGen}_{1\text{b-cca}}$ and $\text{KeyGen}_{\text{cpa}}$ to obtain key pairs $(\text{PK}_{\text{in}}, \text{SK}_{\text{in}})$, $(\text{PK}_A, \text{SK}_A)$ and $(\text{PK}_B, \text{SK}_B)$ respectively. Set $pk := (\text{PK}_{\text{in}}, \text{PK}_A, \text{PK}_B)$ and $sk := (\text{SK}_{\text{in}}, \text{SK}_A, \text{SK}_B)$.
2. Phase 1: Adversary \mathcal{A} is given pk and access to a decryption oracle $\text{Dec}(sk, \cdot)$, which executes the decryption algorithm as defined in Section 3. \mathcal{A} outputs a pair of messages m_0, m_1 of the same length in the message space associated with pk .
3. Challenge: Randomness $\beta, z \leftarrow \{0, 1\}$ and $r_A, r_B \leftarrow \{0, 1\}^\lambda$ are chosen. Let ℓ denote the length of the encoding of (r_A, r_B, m_β) . Then compute:

$$\text{CT}_{\text{in}}^* := \begin{cases} \text{Enc}_{\text{dcca}}(\text{PK}_{\text{in}}, (r_A, r_B, m_\beta)) & \text{if } z = 0; \\ \text{Enc}_{\text{dcca}}(\text{PK}_{\text{in}}, 0^\ell) & \text{if } z = 1. \end{cases} \quad (1)$$

Next compute $\text{CT}_A^* := \text{Enc}_{1\text{b-cca}}(\text{PK}_A, \text{CT}_{\text{in}}^*; r_A)$ and $\text{CT}_B^* := \text{Enc}_{\text{cpa}}(\text{PK}_B, \text{CT}_{\text{in}}^*; r_B)$. Return to \mathcal{A} the ciphertext $\text{CT}^* := (\text{CT}_A^*, \text{CT}_B^*)$.

4. Phase 2: \mathcal{A} continues to have access to $\text{Dec}(sk, \cdot)$, but may not request a decryption of the challenge ciphertext CT^* . Finally, \mathcal{A} outputs a bit z' .
5. Output: The output of the experiment is defined to be 1 if $z' = z$, and 0 otherwise.

Definition 3 (Nested Indistinguishability). *A tuple of encryption systems $(\Pi_{\text{dcca}}, \Pi_{1\text{b-cca}}, \Pi_{\text{cpa}})$ has nested indistinguishable encryptions under a chosen-ciphertext attack if for all probabilistic polynomial-time adversaries \mathcal{A} there exists a negligible function negl such that:*

$$\Pr[\text{Exp}_{\mathcal{A}, \Pi_{\text{dcca}}, \Pi_{1\text{b-cca}}, \Pi_{\text{cpa}}}^{\text{nested}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda).$$

It is important to observe that the nested indistinguishability experiment combines $\Pi_{\text{dcca}}, \Pi_{1\text{b-cca}}, \Pi_{\text{cpa}}$ in exactly the same manner as the Section 3 construction. When $z = 1$, it encrypts “properly” and when $z = 0$, it encrypts all zeros.

With a goal of proving CCA2 security, our main task is to argue that our Section 3 construction provides nested indistinguishability. To do this, we must

first establish that a certain event does not happen, except with negligible probability. We define this event as follows.

Definition 4 (The Bad Query Event). Let Π_{dcca} , $\Pi_{1\text{b-cca}}$, and Π_{cpa} be the schemes parameterizing the experiment $\text{Exp}^{\text{nested}}$. Let PK_{in} be the public key output by running $\text{KeyGen}_{\text{det}}$ during the course of the experiment. We say that a bad query event has occurred during an execution of this experiment if in Phase 2, the adversary \mathcal{A} makes a decryption query of the form $\text{CT} := (\text{CT}_A, \text{CT}_B)$ such that

- (Query inner is “related” to challenge inner):
 $F(\text{PK}_{\text{in}}, \text{CT}_{\text{in}}^*, \text{Dec}_{1\text{b-cca}}(\text{SK}_A, \text{CT}_A)) = 1$, and
- (Query ciphertext differs from challenge ciphertext in first half):
 $\text{CT}_A^* \neq \text{CT}_A$.

where $\text{CT}^* := (\text{CT}_A^*, \text{CT}_B^*)$ is the challenge ciphertext and CT_A^* is an encryption of CT_{in}^* . We note that this event is well defined in both the cases where $z = 0$ and $z = 1$.

4.1 Proof that Bad Query Event Does Not Happen

Lemma 4 (No Bad Query Event when $z = 1$ (all zeros encrypted)). Suppose that Π_{dcca} is DCCA secure, $\Pi_{1\text{b-cca}}$ is 1-bounded CCA secure, and Π_{cpa} is CPA secure, all with perfect correctness. Then for all probabilistic polynomial-time adversaries \mathcal{A} , during a run of experiment $\text{Exp}_{\mathcal{A}, \Pi_{\text{dcca}}, \Pi_{1\text{b-cca}}, \Pi_{\text{cpa}}}^{\text{nested}}(\lambda)$ with $z = 1$, a bad query event does not take place except with negligible probability in λ where the probability is taken over the coins of the adversary and the experiment.

Proof. We proceed via a series of hybrids. Let BQE denote a bad query event.

Step 1: $\Pr[\text{BQE in Nested}] \sim \Pr[\text{BQE in Right-Erased}]$ from CPA-security of Π_{cpa} . We first define a variation of the nested indistinguishability experiment with $z = 1$, which we call the *right-erased* experiment. In this experiment, CT_B^* is formed as $\text{CT}_B^* := \text{Enc}_{\text{cpa}}(\text{PK}_B, 1^k; r_B)$ where k denotes the length of CT_{in}^* . CT_A^* is formed the same as in the nested indistinguishability experiment with $z = 1$. We suppose there exists a PPT adversary \mathcal{A} for the nested indistinguishability experiment which causes the bad query event to occur with non-negligibly different probability in the usual experiment with $z = 1$ compared to the right-erased experiment. We construct a PPT algorithm \mathcal{B} which violates the CPA-security of Π_{cpa} .

\mathcal{B} is given PK_B . \mathcal{B} then runs $\text{KeyGen}_{\text{dcca}}$ and $\text{KeyGen}_{1\text{b-cca}}$ for itself to produce $\text{PK}_{\text{in}}, \text{SK}_{\text{in}}$ and PK_A, SK_A respectively. It gives \mathcal{A} $pk = (\text{PK}_{\text{in}}, \text{PK}_A, \text{PK}_B)$. \mathcal{B} can simulate the decryption oracle $\text{Dec}(sk, \cdot)$ for \mathcal{A} by running the usual decryption algorithm (note that this does not require SK_B).

The adversary \mathcal{A} outputs a pair of messages m_0, m_1 of the same length in the message space associated with pk . \mathcal{B} chooses $r_A \in \{0, 1\}^\lambda$ and computes $\text{CT}_{\text{in}}^* = \text{Enc}_{\text{dcca}}(\text{PK}_{\text{in}}, 0^\ell)$, where ℓ is the length of the encoding of (r_A, r_A, m_0) .

It then computes $\text{CT}_A^* = \text{Enc}_{1b\text{-cca}}(\text{PK}_A, \text{CT}_{\text{in}}^*; r_A)$. It submits CT_{in}^* and 1^k to its challenger as its two messages. It receives CT_B^* as the ciphertext. It gives $\text{CT}^* := (\text{CT}_A^*, \text{CT}_B^*)$ to \mathcal{A} .

To respond to remaining decryption queries \mathcal{A} makes, \mathcal{B} runs the usual decryption algorithm (after checking that the query is not equal to the challenge ciphertext). In addition, \mathcal{B} checks for the bad query event by first checking if $\text{CT}_A \neq \text{CT}_A^*$ and then computing $F(\text{PK}_{\text{in}}, \text{CT}_{\text{in}}^*, \text{Dec}_{1b\text{-cca}}(\text{SK}_A, \text{CT}_A))$. We recall that \mathcal{B} generated SK_A, PK_A for itself, so it can compute $\text{Dec}_{1b\text{-cca}}(\text{SK}_A, \text{CT}_A)$.

If CT_B^* is an encryption of CT_{in}^* , then \mathcal{B} has properly simulated the usual experiment with $z = 1$. If it is instead an encryption of 1^k , then \mathcal{B} has properly simulated the right-erased experiment. We note that the bad query event occurs in the simulation if and only if it is detected by \mathcal{B} .

We let ϵ denote the probability that the bad query event occurs in the usual experiment with $z = 1$ and δ denote this probability in the right-erased experiment. We suppose $\epsilon - \delta$ is positive and non-negligible (the opposite case is analogous). Now, if \mathcal{B} detects the bad query event, it guesses that CT_A^* is an encryption of CT_{in}^* . Otherwise, it guesses the opposite. \mathcal{B} 's probability of guessing correctly in the CPA security game for Π_{cpa} is then equal to $\frac{\epsilon}{2} + \frac{1}{2}(1 - \delta) = \frac{1}{2} + \frac{1}{2}(\epsilon - \delta)$. The quantity $\epsilon - \delta$ is non-negligible, so \mathcal{B} violates the CPA-security of Π_{cpa} . Hence we may conclude that the probability of the bad query event happening in the usual experiment with $z = 1$ is the same (up to a negligible difference) as the probability of the bad query event happening in the right-erased experiment for any PPT adversary.

Step 2: Pr[BQE in Full-Erased] is negligible from the unpredictability of the detecting function of Π_{dcca} . We now define an additional variation of the experiment, which we call the *full-erased* experiment. This is like the right-erased experiment, except that CT_A^* is also an encryption of 1^k , instead of an encryption of CT_{in}^* . We claim that in the full-erased experiment, the bad query event can only occur with negligible probability. To see this, we suppose we have a PPT adversary \mathcal{A} which causes the bad query event to occur with non-negligible probability in the full-erased experiment. We will build a PPT adversary \mathcal{B} for the basic unpredictability experiment which violates unpredictability of the detecting function for Π_{dcca} .

\mathcal{B} is given PK_{in} and access to a decryption oracle $\text{Dec}(\text{SK}_{\text{in}}, \cdot)$. It runs $\text{KeyGen}_{1b\text{-cca}}$ and $\text{KeyGen}_{\text{cpa}}$ for itself to produce PK_A, SK_A and PK_B, SK_B . It gives $(\text{PK}_{\text{in}}, \text{PK}_A, \text{PK}_B)$ to \mathcal{A} . \mathcal{B} can simulate the decryption oracle for \mathcal{A} using SK_A and its own decryption oracle. \mathcal{A} outputs m_0, m_1 . \mathcal{B} then computes $\text{CT}_A^* = \text{Enc}_{1b\text{-cca}}(\text{PK}_A, 1^k)$ and $\text{CT}_B^* = \text{Enc}_{\text{cpa}}(\text{PK}_B, 1^k)$ and gives $\text{CT}^* = (\text{CT}_A^*, \text{CT}_B^*)$ to \mathcal{A} . We let q denote the number of Phase 2 queries made by \mathcal{A} . \mathcal{B} can respond to these queries as before. \mathcal{B} chooses a random $i \in \{1, 2, \dots, q\}$ and a random bit $b \in \{0, 1\}$. It takes the i^{th} Phase 2 query of \mathcal{A} , denoted by $(\text{CT}_A^i, \text{CT}_B^i)$, and computes $\text{CT}_{\text{in}}^i = \text{Dec}_{1b\text{-cca}}(\text{SK}_A, \text{CT}_A^i)$. It submits m_b and CT_{in}^i to its challenger. Then, the distribution of $c^* = \text{Enc}_{\text{dcca}}(\text{PK}_{\text{in}}, m_b)$ in the basic unpredictability experiment is precisely the distribution of CT_{in}^* . Hence, the bad query event for query i corresponds to an output of 1 for basic unpredictability

experiment. Thus, if the bad query event occurs with some non-negligible probability ϵ , \mathcal{B} will cause an output of 1 in the basic unpredictability experiment with probability at least $\frac{\epsilon}{q}$, which is non-negligible.

Step 3: $Pr[\text{BQE in Right-Erased}] \sim Pr[\text{BQE in Full-Erased}]$ from the 1-bounded CCA security of Π_{1b-cca} . We now return to considering a PPT adversary \mathcal{A} in the right-erased experiment. We let q denote the number of Phase 2 queries made by \mathcal{A} . We suppose that \mathcal{A} causes the bad query event with non-negligible probability. Then there exists some index $i \in \{1, \dots, q\}$ such that \mathcal{A} causes the bad query event to occur with non-negligible probability on its i^{th} Phase 2 query. In other words, if there exists a PPT adversary \mathcal{A} for which the bad query event occurs with non-negligible probability in the right-erased experiment, then for each value of the security parameter, there exists an index i such that \mathcal{A} causes the BQE to occur on its i^{th} Phase 2 query with non-negligible probability. We note that for *any* i , the probability that \mathcal{A} causes the BQE to occur on its i^{th} Phase 2 query in the full-erased experiment is negligible, as we proved above.

We fix such an i , and we define a PPT algorithm \mathcal{B} which violates the 1-bounded CCA security of Π_{1b-cca} . \mathcal{B} receives PK_A from its challenger. It runs $\text{KeyGen}_{\text{dcca}}$ and $\text{KeyGen}_{\text{cpa}}$ for itself to produce $\text{PK}_{\text{in}}, \text{SK}_{\text{in}}$ and PK_B, SK_B . It gives $(\text{PK}_{\text{in}}, \text{PK}_A, \text{PK}_B)$ to \mathcal{A} as the public key.

\mathcal{B} simulates the decryption oracle for \mathcal{A} as follows. Upon receiving a ciphertext $(\text{CT}_A, \text{CT}_B)$, \mathcal{B} decrypts CT_B using Dec_{cpa} with SK_B , and we let CT_{in} denote the output. It then decrypts CT_{in} using Dec_{dcca} with SK_{in} , and parses the output as r_A, r_B, M . It checks if $\text{CT}_A = \text{Enc}_{1b-cca}(\text{PK}_A, \text{CT}_{\text{in}}; r_A)$ and if $\text{CT}_B = \text{Enc}_{\text{cpa}}(\text{PK}_B, \text{CT}_{\text{in}}; r_B)$. If both checks pass, it outputs M . Else, it outputs \perp .

We claim that this matches the output of the usual decryption algorithm, even though \mathcal{B} is first decrypting CT_B instead of CT_A . To see this, note that the outputs are the same whenever $\text{Dec}_{1b-cca}(\text{CT}_A, \text{SK}_A) = \text{Dec}_{\text{cpa}}(\text{CT}_B, \text{SK}_B)$. Whenever these are unequal, both decryption methods will output \perp . This is because $\text{CT}_A = \text{Enc}_{1b-cca}(\text{PK}_A, \text{CT}_{\text{in}}; r_A)$ and $\text{CT}_B = \text{Enc}_{\text{cpa}}(\text{PK}_B, \text{CT}_{\text{in}}; r_B)$ imply that $\text{Dec}_{1b-cca}(\text{CT}_A, \text{SK}_A) = \text{CT}_{\text{in}} = \text{Dec}_{\text{cpa}}(\text{CT}_B, \text{SK}_B)$. (Recall here that we have assumed Π_{1b-cca} and Π_{cpa} have perfect correctness.)

At some point, \mathcal{A} outputs m_0, m_1 . \mathcal{B} forms $\text{CT}_{\text{in}}^* = \text{Enc}_{\text{dcca}}(\text{PK}_{\text{in}}, 0^\ell)$ and $\text{CT}_B^* = \text{Enc}_{\text{cpa}}(\text{PK}_B, 1^k)$. It outputs the messages CT_{in}^* and 1^k to its challenger, and receives a ciphertext which it sets as CT_A^* . It gives the ciphertext $(\text{CT}_A^*, \text{CT}_B^*)$ to \mathcal{A} . It can then respond to \mathcal{A} 's Phase 2 decryption queries in the same way as before. When it receives the i^{th} Phase 2 query of \mathcal{A} , denoted by $(\text{CT}_A^i, \text{CT}_B^i)$, \mathcal{B} checks for the bad query event by first checking if $\text{CT}_A^i \neq \text{CT}_A^*$ and if so, submitting CT_A^i as its one decryption query to its decryption oracle for PK_A . It can compute $F(\text{PK}_{\text{in}}, \text{CT}_{\text{in}}^*, \text{Dec}(\text{SK}_A, \text{CT}_A^i))$. This equals 1 if and only if the bad query event has occurred for query i , and in this case \mathcal{B} guesses that CT_A^* is an encryption of CT_{in}^* . Otherwise, \mathcal{B} guesses the opposite.

We observe that when CT_A^* is an encryption of CT_{in}^* , then \mathcal{B} has properly simulated the right-erased experiment, and when CT_A^* is an encryption of 0^k , then \mathcal{B} has properly simulated the full-erased experiment. We let ϵ denote the

non-negligible probability that \mathcal{A} causes the bad query event to occur on (Phase 2) query i in the right-erased experiment, and we let δ denote the corresponding probability for the full-erased experiment. We know that δ must be negligible, therefore $\epsilon - \delta$ is positive and non-negligible. The probability that \mathcal{B} guesses correctly is: $\frac{1}{2}(1 - \delta) + \frac{1}{2}\epsilon = \frac{1}{2} + \frac{1}{2}(\epsilon - \delta)$, so \mathcal{B} achieves a non-negligible advantage in the 1-bounded CCA security game for Π_{1b-cca} .

Thus, it must be the case that for all PPT algorithms \mathcal{A} , the BQE occurs with only negligible probability in the right-erased experiment, and hence also in the nested experiment with $z = 1$.

Lemma 5 (No Bad Query Event when $z = 0$ (real message encrypted)). *As a consequence of Lemma 4 and the DCCA security of Π_{dcca} , it holds that for all probabilistic polynomial-time adversaries \mathcal{A} , during a run of experiment $\text{Exp}_{\mathcal{A}, \Pi_{dcca}, \Pi_{1b-cca}, \Pi_{cpa}}^{\text{nested}}(\lambda)$ with $z = 0$, a bad query event does not take place except with negligible probability in λ where the probability is taken over the coins of the adversary and the experiment.*

Proof. In Lemma 4, we established that bad query events happen with at most negligible probability when $z = 1$. We will use this fact to argue that they cannot happen much more frequently when $z = 0$. Suppose to the contrary that there exists a PPT adversary \mathcal{A} that forces bad query events to happen with non-negligible probability ϵ when $z = 0$. We create an PPT adversary \mathcal{B} who interacts with \mathcal{A} in a run of the nested indistinguishability experiment to break the DCCA security of Π_{dcca} with detecting function F with probability negligibly-close to $\frac{1}{2} + \frac{\epsilon}{2}$ as follows:

1. Setup: \mathcal{B} obtains PK_{in} from the $\text{Exp}^{\text{indist}}$ challenger. It runs KeyGen_{1b-cca} to obtain $(\text{PK}_A, \text{SK}_A)$ and KeyGen_{cpa} to obtain $(\text{PK}_B, \text{SK}_B)$.
2. Phase 1: \mathcal{B} gives to \mathcal{A} the public key $\text{PK} = (\text{PK}_{\text{in}}, \text{PK}_A, \text{PK}_B)$. When \mathcal{A} queries the decryption oracle on CT , \mathcal{B} can simulate the normal decryption algorithm using SK_A and the phase 1 oracle $\text{Dec}(\text{SK}_{\text{in}}, \cdot)$. Eventually, \mathcal{A} outputs a pair of messages m_0, m_1 .
3. Challenge: Choose random $\beta \in \{0, 1\}$ and $r_A, r_B \in \{0, 1\}^\lambda$. Send to the $\text{Exp}^{\text{indist}}$ challenger the messages $M_0 = (r_A, r_B, m_\beta)$ and $M_1 = 0^{|M_0|}$, and obtain from the challenger CT_{in}^* . Compute $\text{CT}_A^* := \text{Enc}_{1b-cca}(\text{PK}_A, \text{CT}_{\text{in}}^*; r_A)$ and $\text{CT}_B^* := \text{Enc}_{cpa}(\text{PK}_B, \text{CT}_{\text{in}}^*; r_B)$. Return $\text{CT}^* := (\text{CT}_A^*, \text{CT}_B^*)$ to \mathcal{A} .
4. Phase 2: When \mathcal{A} queries the decryption oracle on $\text{CT} := (\text{CT}_A, \text{CT}_B)$, compute $\text{CT}_{\text{in}} := \text{Dec}_{1b-cca}(\text{SK}_A, \text{CT}_A)$. If
 - (a) Case 1 (a bad query event): $\text{CT}_A \neq \text{CT}_A^*$ and yet $F(\text{PK}_{\text{in}}, \text{CT}_{\text{in}}^*, \text{CT}_{\text{in}}) = 1$, then abort and output the bit 0.
 - (b) Case 2 (partial match with challenge): $\text{CT}_A = \text{CT}_A^*$, then return \perp to \mathcal{A} .
 Otherwise, query the phase 2 oracle, $\text{Dec}(\text{SK}_{\text{in}}, \cdot)$, to decrypt CT_{in} , and return its response to \mathcal{A} .
5. Output: When \mathcal{A} outputs a bit, \mathcal{B} echos the bit as its output.

Analysis. We begin our analysis by arguing that \mathcal{B} correctly answers all decryption queries except when it aborts. First, we show that a partial match with the challenge, causing the \perp response in Case 2, is correct because that query must be invalid. Since a decryption query on the challenge is forbidden by the experiment, if $\text{CT}_A = \text{CT}_A^*$, then $\text{CT}_B \neq \text{CT}_B^*$. However, we argue that this must be an invalid ciphertext, i.e., one on which the main construction’s decryption algorithm would return \perp . We see this as follows. Since decryption is deterministic, we have $T := \text{Dec}_{1\text{b-cca}}(\text{SK}_A, \text{CT}_A) = \text{Dec}_{1\text{b-cca}}(\text{SK}_A, \text{CT}_A^*)$ and $(r_A, r_B, m) := \text{Dec}_{\text{dcca}}(\text{SK}_{\text{in}}, T)$. By the checks enforced by the main construction’s decryption algorithm, there is only one “second half” that matches $\text{CT}_A = \text{CT}_A^*$, that is $\text{Enc}_{\text{cpa}}(\text{PK}_B, T; r_B)$. Since the challenge is a valid ciphertext, CT_B^* must be this value and CT_B must cause an error.

When neither Case 1 or Case 2 applies in phase 2, the inner decryption query will succeed since the ciphertext is not detectably related to the challenge. This allows \mathcal{B} to respond correctly.

When a bad query event occurs in Phase 2, \mathcal{B} cannot query $\text{Exp}^{\text{indist}}$ ’s decryption oracle to decrypt the ciphertext. At first glance, one seems stuck. However, we assumed bad query events happen only when $z = 0$ with all but negligible probability. Thus, \mathcal{B} can guess that \mathcal{A} thinks $z = 0$, which corresponds to M_0 being encrypted in our reduction. Thus, \mathcal{B} can abort and guess 0 at this point.

When \mathcal{B} aborts, it causes the $\text{Exp}^{\text{indist}}$ experiment to output 1 with high probability. When \mathcal{B} does not abort, it causes $\text{Exp}^{\text{indist}}$ experiment to output 1 with probability $\frac{1}{2}$. Since \mathcal{B} aborts with non-negligible probability ϵ when $z = 0$, then \mathcal{B} causes the experiment’s output to be 1 with probability non-negligibly greater than $\frac{1}{2}$.

4.2 Putting the Proof of the Main Theorem Together

Theorem 1 (Main Construction is Nested Indistinguishable). *Our main construction in Section 3, comprised of the three building blocks Π_{dcca} , $\Pi_{1\text{b-cca}}$, Π_{cpa} , has nested indistinguishable encryptions under a chosen-ciphertext attack under the assumptions that Π_{dcca} is DCCA secure, $\Pi_{1\text{b-cca}}$ is 1-bounded CCA secure, and Π_{cpa} is CPA secure, all with perfect correctness.*

Proof of Theorem 1 appears in the full version. The crux of the argument is that bad query events do not happen (except with negligible probability). This was already established in Lemmas 4 and 5. Armed with this fact, we can prove the nested indistinguishability of the main construction based on the indistinguishability property of the DCCA-security of Π_{dcca} . The reduction and its analysis are similar to those in the proof of Lemma 5.⁷

The following corollary follows from Theorem 1. Informally, if the adversary cannot distinguish an encryption of a message from an encryption of zeros, then she also cannot distinguish between the encryptions of two different messages.

⁷ We alternatively could have merged the proofs of Lemma 5 and Theorem 1. However, we chose to keep the bad event analysis separate for pedagogical purposes at the expense of some redundancy in the description of the related reductions.

Corollary 1 (Main Construction is CCA2 Secure). *Our main construction in Section 3, comprised of the three building blocks Π_{dcca} , $\Pi_{1\text{b-cca}}$, Π_{cpa} , is CCA2 secure under the assumptions that Π_{dcca} is DCCA secure, $\Pi_{1\text{b-cca}}$ is 1-bounded CCA secure, and Π_{cpa} is CPA secure, all with perfect correctness.*

Acknowledgments

The authors thank Steven Myers and the anonymous reviewers for helpful comments.

References

1. Jee Hea An, Yevgeniy Dodis, and Tal Rabin. On the security of joint signature and encryption. In *EUROCRYPT*, volume 2332, pages 83–107, 2002.
2. Mihir Bellare, Zvika Brakerski, Moni Naor, Thomas Ristenpart, Gil Segev, Hovav Shacham, and Scott Yilek. Hedged public-key encryption: How to protect against bad randomness. In *ASIACRYPT*, pages 232–249, 2009.
3. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
4. Mihir Bellare and Amit Sahai. Non-malleable encryption: Equivalence between two notions, and an indistinguishability-based characterization. In *CRYPTO*, pages 519–536, 1999.
5. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *STOC*, pages 103–112, 1988.
6. Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM J. Comput.*, 36(5):1301–1328, 2007.
7. Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In *EUROCRYPT*, pages 207–222, 2004.
8. Ran Canetti, Hugo Krawczyk, and Jesper Buus Nielsen. Relaxing chosen-ciphertext security. In *CRYPTO*, volume 2729, pages 565–582, 2003.
9. David Cash, Eike Kiltz, and Victor Shoup. The twin diffie-hellman problem and applications. In *EUROCRYPT*, pages 127–145, 2008.
10. Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. Black-box construction of a non-malleable encryption scheme from any semantically secure one. In *TCC*, pages 427–444, 2008.
11. Ronald Cramer, Goichiro Hanaoka, Dennis Hofheinz, Hideki Imai, Eike Kiltz, Rafael Pass, Abhi Shelat, and Vinod Vaikuntanathan. Bounded cca2-secure encryption. In *ASIACRYPT*, pages 502–518, 2007.
12. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *CRYPTO*, pages 13–25, 1998.
13. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *EUROCRYPT*, pages 45–64, 2002.
14. Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *STOC*, pages 542–552, 1991.
15. Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *FOCS*, pages 126–135, 2001.

16. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.
17. Dennis Hofheinz and Eike Kiltz. Secure hybrid encryption from weakened key encapsulation. In *CRYPTO*, volume 4622, pages 553–571, 2007.
18. Dennis Hofheinz and Eike Kiltz. Practical chosen ciphertext secure encryption from factoring. In *EUROCRYPT*, pages 313–332, 2009.
19. Eike Kiltz. Chosen-ciphertext security from tag-based encryption. In *TCC*, 2006.
20. Eike Kiltz, Payman Mohassel, and Adam O’Neill. Adaptive trapdoor functions and chosen-ciphertext security. In *EUROCRYPT*, pages 673–692, 2010.
21. Philip D. MacKenzie, Michael K. Reiter, and Ke Yang. Alternatives to non-malleability: Definitions, constructions, and applications (extended abstract). In *TCC*, pages 171–190, 2004.
22. Steven Myers and Abhi Shelat. Bit encryption is complete. In *FOCS*, pages 607–616, 2009.
23. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, pages 427–437, 1990.
24. Rafael Pass, Abhi Shelat, and Vinod Vaikuntanathan. Construction of a non-malleable encryption scheme from any semantically secure one. In *CRYPTO*, pages 271–289, 2006.
25. Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196, 2008.
26. Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *CRYPTO*, pages 433–444, 1991.
27. Alon Rosen and Gil Segev. Chosen-ciphertext security via correlated products. In *TCC*, pages 419–436, 2009.
28. Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, pages 543–553, 1999.
29. Victor Shoup. A proposal for an iso standard for public key encryption. Cryptology ePrint Archive, Report 2001/112, 2001. <http://eprint.iacr.org/>.