

On the CCA Compatibility of Public-Key Infrastructure

Dakshita Khurana^{1*} and Brent Waters^{2**}

Abstract. In this work, we put forth the notion of *compatibility* of any key generation or setup algorithm. We focus on the specific case of encryption, and say that a key generation algorithm KeyGen is X -compatible (for $X \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$) if there exist encryption and decryption algorithms that together with KeyGen , result in an X -secure public-key encryption scheme.

We study the following question: Is every CPA-compatible key generation algorithm also CCA-compatible? We obtain the following answers:

- Every sub-exponentially CPA-compatible KeyGen algorithm is CCA1-compatible, assuming the existence of hinting PRGs and sub-exponentially secure keyless collision resistant hash functions.
- Every sub-exponentially CPA-compatible KeyGen algorithm is also CCA2-compatible, assuming the existence of non-interactive CCA2 secure commitments, in addition to sub-exponential security of the assumptions listed in the previous bullet.

Here, sub-exponentially CPA-compatible KeyGen refers to any key generation algorithm for which there exist encryption and decryption algorithms that result in a CPA-secure public-key encryption scheme *against sub-exponential adversaries*.

This gives a way to perform CCA secure encryption given any public key infrastructure that has been established with only (sub-exponential) CPA security in mind. The resulting CCA encryption makes black-box use of the CPA scheme and all other underlying primitives.

1 Introduction

Any public-key encryption scheme enables a receiver to recover the encrypted message only if they know a secret key corresponding to their public key. But what if the receiver only ever published a verification key for a digital signature scheme for which they possessed a signing key? Or published a hard puzzle for which they possessed a solution?

This question was one of the original motivations for the study of witness encryption. Garg et al. [14] showed that it is possible to encrypt a message so that it can only be opened by a recipient who knows an NP witness, assuming the existence of an appropriate witness encryption

* University of Illinois Urbana-Champaign. dakshita@illinois.edu.

** University of Texas at Austin and NTT Research. bwaters@cs.utexas.edu.

scheme. Put differently, assuming an appropriate witness encryption, almost any `KeyGen` algorithm that outputs a hard-to-invert string and a corresponding secret (such as a verification and signing key pair for a signature scheme) can be used to derive CPA-secure public key encryption.

In this work, we generalize this study. We put forth the notion of *compatibility* of any key generation or setup algorithm, while focusing on the specific case of encryption schemes. Here, recall that semantic security of (public key) encryption in [15] was only the first step towards formalizing security of encryption schemes. Semantic security, or equivalently indistinguishability-based security against chosen plaintext attacks (CPA) requires that encryptions of every pair of plaintexts appear indistinguishable to a computationally bounded attacker. Unfortunately, starting with the attacks of Bleichenbacher [4] against PKCS#1, it was quickly realized that systems that *only* satisfy CPA security often fail in practice. As a result, security against adaptive chosen ciphertext attacks (or, CCA security) has been accepted as the standard requirement from encryption schemes that need to withstand active attacks [8,26,11,29]. This guarantees security even against attackers that make oracle decryption queries to keys they do not have. If the adversary only has access to a decryption oracle *before* obtaining the challenge ciphertext, the resulting scheme is said to be CCA1 secure. On the other hand, if the adversary has access to the decryption oracle both before and after obtaining the challenge ciphertext, the resulting scheme is CCA2 secure.

We investigate whether arbitrary setup of `KeyGen` algorithms can be used to derive *CCA-secure* schemes. We will say that a key generation algorithm `KeyGen` is X -compatible (for $X \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$) if there exist encryption and decryption algorithms that together with `KeyGen`, result in an X -secure public-key encryption scheme. As already discussed, the existence of (extractable) witness encryption suffices to prove CPA-compatibility for many non-trivial `KeyGen` algorithms. The focus of our work is to take a closer look at CCA-compatibility.

Specifically, we analyze what it takes for a `KeyGen` algorithm to be CCA-compatible. Our primary result stated informally, is the following:

It is always possible to get CCA secure encryption from any `KeyGen` procedure that gives rise to (sub-exponentially secure) CPA encryption.

Combined with the CPA-compatibility of non-trivial `KeyGen`, this also implies CCA-compatibility of many non-trivial `KeyGen` algorithms.

This also means that we can always achieve CCA security using keys for cryptosystems that were originally deployed for CPA mode, without

having to modify the public key. This would allow parties with access to public key infrastructures that have been established with only CPA security in mind, to use these infrastructures to perform CCA secure encryption instead. For instance, over the years, multiple encryption schemes have been developed that satisfy IND-CPA security alone. A recent example that gained some popularity is the messaging service telegram, that supplies end-to-end encryption using a new protocol employing AES, RSA, and Diffie-Hellman key exchange. Recently, [17] showed that this protocol is not IND-CCA secure. Our result ensures that (under reasonable cryptographic assumptions) careful participants can use the same underlying infrastructure to engage in encrypted communication without worrying about CCA2 attacks. Alternatively, suppose a user or an organization sets up what is supposed to be a CCA secure system, but the underlying computational assumption turned out to be false. For example, perhaps an attack on DDH was found in a specific group [7], and the scheme is somehow adjusted to recover CPA security. Then, the scheme can also be adjusted to recover CCA security (under our assumptions), with the same infrastructure as that used for CPA security. In these settings, while one could potentially ask users to switch to using a new key from the same system, changing an entire public key infrastructure would be far more cumbersome. We note that simple key encapsulation strategies would be insufficient: for example, sampling the key for a CCA secure encryption scheme and encapsulating it using a key for the original CCA-insecure infrastructure would not lead to the resulting ciphertext being CCA secure.

Preliminary Solutions in Idealized Models. In idealized models, there are known methods that implicitly allow one to obtain CCA security from any CPA-compatible KeyGen algorithm. For instance, in the Random Oracle model, the famous Fujisaki-Okamoto transform [12] converts any CPA secure encryption scheme to a CCA secure one, with the same KeyGen algorithm. We are interested in whether a similar effect can be achieved in the plain model.

A natural approach without a random oracle would be to leverage a common reference string (CRS) and implement the Naor-Yung methodology [26,11] using simulation-sound NIZKs. We recall that the Naor-Yung (CCA secure) encryption of a message typically consists of two encryptions, under independent public keys, of the same message; and a simulation-sound NIZK proof that both ciphertexts encrypt the same message. If implemented naively, it appears that the KeyGen algorithm

for the resulting CCA mode would have to output *two* independent public keys corresponding to the underlying CPA secure scheme. Even if we found a method to get rid of the second key, this still requires participants to place their trust in a central setup assumption to enable the (simulation-sound) NIZK. Given this state of affairs, we ask if it is possible to obtain CCA secure encryption by relying on the KeyGen algorithm of any CPA secure encryption scheme:

- in the plain model without assuming setup, CRS, or a random oracle,
- with black-box use of the CPA scheme (and additional primitives),
- and while making the weakest possible cryptographic assumptions.

Our Results. We take a novel approach to obtain a plain model solution that makes black-box use of the CPA scheme, and does not resort to NIZK (or NIWI) style assumptions. Specifically, we demonstrate CCA1-compatibility of any sub-exponentially CPA-compatible KeyGen algorithm while making black-box use of hinting PRGs and sub-exponential keyless collision resistant hash functions. We also demonstrate CCA2-compatibility of any sub-exponentially CPA-compatible KeyGen algorithm while additionally making black-box use of non-interactive CCA secure commitments. Such commitments were recently obtained [13] based on black-box use of subexponential one-way functions in BQP, and sub-exponential quantum-hard one-way functions, in addition to the assumptions listed above. Alternatively, these can be based on sub-exponential time-lock puzzles [24,13] in addition to the assumptions listed above. We informally summarize our results below.

Informal Theorem 1 *Every sub-exponentially CPA-compatible KeyGen algorithm against non-uniform adversaries is also CCA1-compatible against uniform adversaries, assuming the existence of hinting PRGs and sub-exponential keyless collision-resistant hash functions against uniform adversaries.*

Informal Theorem 2 *Every sub-exponentially CPA-compatible KeyGen algorithm against non-uniform adversaries is also CCA2-compatible against uniform adversaries, assuming the existence of sub-exponential hinting PRGs, sub-exponential keyless collision-resistant hash functions against uniform adversaries and sub-exponential non-interactive CCA secure commitments.*

2 Our Technique

2.1 Background: A variant of Koppula-Waters [22]

Our starting point is a variant of the recent Koppula-Waters [22] approach to achieving CCA1 secure encryption based on CPA secure encryption and a new primitive they introduced, called a *hinting PRG*. A hinting PRG satisfies the following property: for a uniformly random short seed s , the matrix M obtained by first expanding $\text{PRG}(s) = z_0 z_1 z_2 \dots z_n$, sampling uniformly random $v_1 v_2 \dots v_n$, and setting for all $i \in [n]$, $M_{s_i, i} = z_i$ and $M_{1-s_i, i} = v_i$, should be indistinguishable from a uniformly random matrix. Hinting PRGs are known based on CDH, LWE [22]. (One can also pursue a similar path using any circular secure symmetric key encryption [20] in lieu of the Hinting PRG.) Koppula and Waters [22] also require the CPA scheme to have two properties. First, the scheme should have perfect decryption correctness for most public/secret keys and second, any ciphertext should be decryptable given the encryption randomness.

Now, the KeyGen algorithm of the CCA1 scheme constructed by [22,20] executes the CPA KeyGen setup *twice* to obtain two independent public/secret key pairs, denoted by (pk_0, sk_0) and (pk_1, sk_1) . Additionally, the CCA1 KeyGen algorithm samples and outputs the public parameters pp of an equivocal commitment scheme. To encrypt a message m , the encryption algorithm chooses a seed $s \leftarrow \{0, 1\}^n$ and computes $H(s) = z_0 z_1 \dots z_n$, where H is a hinting PRG. It uses z_0 to mask the message m ; that is, it computes $c = m \oplus z_0$. The remaining ciphertext will contain n ‘signals’ that help the decryption algorithm to recover s bit by bit, which in turn will allow it to compute z_0 and hence unmask c .

The i^{th} signal (for each $i \in [n]$) has three components $c_{0,i}, c_{1,i}, c_{2,i}$. If the i^{th} bit of s is 0, then $c_{0,i}$ is an encryption of a random string y_i using the public key pk_0 and randomness z_i , and $c_{1,i}$ is an encryption of y_i using pk_1 (encrypted using true randomness). If the i^{th} bit of s is 1, then $c_{0,i}$ is an encryption of y_i using public key pk_0 (encrypted using true randomness), $c_{1,i}$ is an encryption of y_i using public key pk_1 and randomness z_i . In both cases, $c_{2,i}$ is an equivocal commitment to s_i using randomness y_i . As a result, half the ciphertexts are encryptions with fresh randomness, while the remaining are encryptions with blocks of the hinting PRG output being used as randomness, and the positioning of the random/pseudorandom encryptions reveals the seed s .

The decryption algorithm first decrypts each $c_{0,i}$ (using secret key is sk_0) to obtain $y_1 y_2 \dots y_n$. It then checks if $c_{2,i}$ is an equivocal commitment to 0 with randomness y_i . If so, it guesses that $s_i = 0$, else it guesses that

$s_i = 1$. With this estimate for s , the decryption algorithm can compute $H(s) = z_0 z_1 \dots z_n$ and then compute $c \oplus z_0$ to recover the message m . The decryption algorithm needs to enforce additional checks to prevent malicious decryption queries (made during the CCA1 experiment). In particular, the decryption algorithm needs to check that the guess for s is indeed correct. It conducts the following checks and outputs $z_0 \oplus c$ if they all pass.

- If the i^{th} bit of s is guessed to be 0, then the decryption algorithm checks that $c_{0,i}$ is a valid ciphertext - it simply re-encrypts y_i using randomness z_i and checks if this equals $c_{0,i}$. Recall that the decryption procedure, before guessing the i^{th} bit of s to be 0, also checks that $c_{2,i}$ is a commitment to 0 with randomness y_i .
- If the i^{th} bit of s is guessed to be 1, then the decryption algorithm first recovers the message underlying ciphertext $c_{1,i}$. Note that $c_{1,i}$ should be encrypted using randomness z_i , hence using z_i , one can recover message \tilde{y}_i from $c_{1,i}$ (using the randomness recovery property of the PKE scheme). It then re-encrypts \tilde{y}_i and checks if it is equal to $c_{1,i}$, and also checks that $c_{2,i}$ is a commitment to 1 with randomness \tilde{y}_i .

Inadequacies of this transformation. At this point, we are far from having established CCA1 compatibility of arbitrary CPA infrastructure due the following reasons:

1. The transformation described so far crucially uses equivocal commitments, which require trusted setup or a common reference/random string, and this is something that we would like to avoid.
2. This transformation makes use of two public keys, and we are only guaranteed to get *a single key* from existing setup.

In fact, achieving CCA2 security is even more complex: in particular, the CCA2 setup in the transformation of [22] must also contain pairwise independent hash functions h_1, h_2, \dots, h_n . These are used to prevent the adversary from mauling the challenge ciphertext into related ciphertexts and querying the decryption oracle on these ciphertexts.

In the coming section, we discuss how to achieve CCA1 compatibility by eliminating the two problems listed above. In the section after that, we discuss the more complex case of CCA2 compatibility.

2.2 Techniques for CCA1 Compatibility

To address the first item listed above, we rely on equivocal commitments without setup, that satisfy binding against uniform adversaries. The re-

sulting CCA1 compatibility is also established against uniform PPT adversaries. We briefly recall how such equivocal commitments can be obtained based on keyless collision resistant hash functions against uniform adversaries [10,16,2,13]: the commit algorithm samples a uniformly random seed for a strong extractor, $g \leftarrow \{0, 1\}^\kappa$ and a value v in the domain of a sufficiently compressing keyless collision resistant hash function. A commitment to a bit b is given by the string $H(v), (\text{Ext}(g, v) \oplus b)$.

We use the resulting commitment scheme to generate $c_{2,i}$ values in the outline described above. Note that this commitment scheme cannot be efficiently equivocated by uniform adversaries, since that would lead to an efficient uniform algorithm that finds collisions in the hash function H . On the other hand, our proof of security will rely on the fact that most strings in the support of the commitment can be non-uniformly equivocated. Next, we discuss how to argue security when using these equivocal commitments in the transformation described above.

Arguing Security. To argue security, first observe that the equivocal commitment satisfies computational binding against PPT adversaries, which makes it infeasible for a CCA1 adversary to query the challenger on *ambiguous* ciphertexts that pass the decryption checks but potentially decrypt to different values under sk_0 and sk_1 . This is because for such ciphertexts, for some $i \in [n]$, the component $c_{2,i}$ is both a commitment to 0 with randomness y_i recovered from $c_{0,i}$ and a commitment to 1 with randomness \tilde{y}_i recovered from $c_{1,i}$: clearly violating the binding property of the commitment scheme.

At the same time, the equivocality of the commitment enables the challenger to set for every $i \in [n]$, the values $c_{2,i}$ that are both commitments to 0 with randomness y_i and 1 with randomness \tilde{y}_i . Next, via a careful hybrid argument that relies on perfect correctness of the encryption scheme, the binding property of the equivocal commitment and CPA security of the public key encryption scheme, the challenge ciphertext can be modified and made ambiguous: this means that in the challenge ciphertext for every $i \in [n]$, $c_{0,i}$ is an encryption of y_i , $c_{1,i}$ is a non-interactive commitment to \tilde{y}_i , $c_{2,i}$ an equivocal commitment: i.e., a commitment to 0 with randomness y_i and to 1 with randomness \tilde{y}_i .

At a very high level, this involves changing values encrypted under $c_{0,i}$ and $c_{1,i}$ by relying on CPA security of the encryption scheme. Values encrypted under $c_{1,i}$ can be modified relatively easily because only the secret key sk_0 is used to perform decryption queries. Arguing security when changing values encrypted under $c_{0,i}$ requires more care: in partic-

ular, such an argument is possible only if sk_0 is no longer used to answer the adversary’s decryption queries. Therefore we first switch to using an alternative decryption strategy that relies on sk_1 instead of sk_0 to decrypt the adversary’s ciphertexts. Unambiguity of the adversary’s ciphertexts helps ensure that alternative decryption yields the same outputs as the original decryption strategy. When using alternate decryption, it becomes possible to change values encrypted under $c_{0,i}$ since sk_0 is no longer being used.

At the end of this argument, information about the hinting PRG seed s has *almost* been removed from the ciphertext, except that for all i where $s_i = 0$, $c_{s_i,i}$ is encrypted using randomness r_i which came from running the hinting PRG on s ; whereas $c_{1-s_i,i}$ is encrypted using uniform randomness. These can all be replaced with uniformly random values by the property of the hinting PRG, thereby eliminating all information about s , and therefore m , from the ciphertext.

So far, the construction and security argument also relied on the use of *two* public/private key pairs. But as already pointed out, a CPA-compatible KeyGen algorithm will output a single public and private key. Next, we discuss how to eliminate the need for the second key.

Removing the second key via Non-interactive Commitments.

Here, we begin by observing that the actual decryption algorithm only makes use of the secret key sk_0 , and does not need the second secret key sk_1 . It recovers messages underlying $c_{1,i}$ for all $i \in [n]$ where $s_i = 1$, using the *randomness* z_i that was supposedly used to create $c_{1,i}$.

As a result, the actual decryption algorithm does not need to *efficiently decrypt* $c_{1,i}$ and has no use for the secret key sk_1 . Therefore, we eliminate the need for the second public key by setting the strings $c_{1,i}$ to be *non-interactive perfectly binding commitments* that do not require any public keys or public parameters, and where the committed message can be efficiently recovered given the randomness used to commit. Lombardi and Schaeffer [25] showed that such commitments can be obtained from any perfectly correct public-key encryption scheme. Specifically, we modify the encryption algorithm so that if $s_i = 0$, $c_{1,i}$ is a non-interactive commitment to 0^n using true randomness, and if $s_i = 1$, then $c_{1,i}$ is a non-interactive commitment to a random string x_i using randomness z_i . The remaining parts $\{c_{0,i}, c_{2,i}\}_{i \in [n]}$ will remain unmodified.

Now recall that the security argument outlined above points to an alternative decryption strategy that does actually use sk_1 , instead of sk_0 , to efficiently decrypt the adversary’s ciphertexts. However, this alternative

decryption algorithm is only used in a few hybrids in the proof of security, and when using non-interactive commitments, we allow these hybrids to *inefficiently* recover the values committed under $c_{1,i}$ by running an exponential time brute-force algorithm that checks all possible randomnesses values that could potentially be used to build $c_{1,i}$.

In order to make the hybrid strategy still go through, we rely on complexity leveraging: we set security parameters so that all other primitives are secure against adversaries that can run in time large enough to execute the brute-force algorithm that recovers values committed under $c_{1,i}$ for $i \in [n]$. Specifically, we assume that the CPA encryption scheme to be upgraded has security parameter k and is 2^{k^ϵ} secure for some constant $0 < \epsilon < 1$. We also assume that the keyless collision-resistant hash function responsible for the binding property of the equivocal commitments is 2^{k^ϵ} -secure for some constant $0 < \epsilon < 1$, and we set the security parameter for the non-interactive commitment to be $k^{\min(e,\epsilon)}$.

Additional Details of the Proof. We now provide additional details on the proof of CCA1 security of the resulting scheme. Recall that in the CCA1 security game, the adversary is allowed access to a decryption oracle before the challenge phase, where the adversary outputs m_0, m_1 and then obtains an encryption of m_b for b sampled uniformly at random.

We develop a sequence of hybrid experiments where the decryption oracle as well as the challenge ciphertext are modified in small increments, and where the first hybrid corresponds to providing the adversary access to the actual decryption oracle together with an encryption of m_b and the last one corresponds to providing the adversary access to the actual decryption oracle together with an encryption of uniform randomness.

The very next hybrid experiment is an exponential time hybrid that samples equivocal commitments $\{c_{2,i}\}_{i \in [n]}$ for the challenge ciphertext, together with randomness $\{y_{0,i}\}_{i \in [n]}$ and $\{y_{1,i}\}_{i \in [n]}$ that can be used to equivocally open these commitments to 0 and 1 respectively.

The third hybrid additionally modifies the components $c_{1,i}$ to “drown” out information about s via noise. In particular, while in the real game, the values $c_{1,i}$ are always commitments to $y_{s_i,i}$, in the challenge ciphertext these values are modified to become commitments to $y_{1,i}$, irrespective of what s_i is. On the other hand, the values $c_{0,i}$ remain encryptions of $y_{s_i,i}$, exactly as in the real experiment. In spite of the fact that equivocation takes exponential time, the proof of indistinguishability between this hybrid and the previous one does not need to rely on an exponential time reduction. Instead, we observe that the equivocal commitment strings

$\{c_{2,i}\}_{i \in [n]}$ together with their openings can be fixed non-uniformly and independently of the strings $c_{1,i}$, and therefore these hybrids can be proven indistinguishable based on non-uniform hiding of the non-interactive commitment scheme. Since we must carefully manipulate the randomness used for $\{c_{1,i}\}_{i \in [n]}$ in both games, this hybrid requires a delicate argument.

The fourth hybrid modifies the decryption oracle so that instead of decrypting using the secret key of the public key encryption scheme, decryption is performed by running in time exponential in the security parameter of the commitment scheme (specifically, in time $2^{k^{\min(\epsilon, \epsilon)}}$) and performing a brute-force search for the randomness used to create the commitments $\{c_{1,i}\}_{i \in [n]}$. This hybrid is only indistinguishable from the previous one if an adversary cannot find ciphertexts that decrypt differently when using the secret key of the encryption scheme versus the brute-force algorithm discussed above. This hybrid requires a subtle argument that relies on the fact that no adversary can query the decryption oracle with “ambiguous” ciphertexts, in spite of being provided such ciphertexts in the challenge phase. Specifically, we crucially use the fact that the adversary does not observe any equivocations before obtaining the challenge ciphertext, and therefore cannot query the decryption oracle with any “ambiguous” ciphertexts (as this would lead to the adversary breaking binding of the equivocal commitment). This is the primary reason that we only obtain CCA1 security.

In the fifth hybrid, some of the challenge ciphertext values, that are independent of the message being encrypted, are chosen ahead of time. This maneuver helps us with the sixth hybrid, where in the challenge ciphertext, information about the PRG seed s is removed from the ciphertext components $\{c_{1,i}\}_{i \in [n]}$, making them all encryptions of $y_{0,i}$ instead of being encryptions of $y_{s_i,i}$. Again, since we must carefully manipulate the randomness used for $\{c_{0,i}\}_{i \in [n]}$ in both games, this hybrid requires a delicate argument.

In the seventh hybrid, we modify the decryption oracle again to go back to using the secret key of the public key encryption scheme to decrypt. Note that the only remaining information about s is in the *randomness* used to obtain $\{c_{i,0}, c_{i,1}\}_{i \in [n]}$ in the challenge ciphertext. In the seventh and eighth hybrids, we carefully re-order the randomness and rely on the security of the hinting PRG to switch to using uniform randomness everywhere. This eliminates all information about s and therefore about the message being encrypted in the challenge ciphertext.

2.3 Techniques for CCA2 Compatibility

We observe that the key barrier to proving CCA2 security in the hybrid arguments outlined above is the specific hybrid that *modifies the challenge ciphertext so it contains a commitment to both a 0 and a 1*. Given such a ciphertext, in a CCA2 game, an adversary could generate new strings that are a commitment to both a 0 and a 1, and use them to create ambiguous ciphertexts. Arguing that this cannot happen requires us to develop a much deeper technical toolkit.

Our first insight is that the requirement that an adversary, given an ambiguous ciphertext, be unable to generate *additional* ambiguous ciphertexts is reminiscent of *non-malleability*. As such, we will rely on non-interactive non-malleable (more precisely, CCA secure) commitments without setup. Up until recently, there were perceived strong barriers to obtaining non-malleable commitments with less than 3 rounds of interaction [28]. But a sequence of recent works obtained two round [19,24] and even non-interactive [24,3,18,13] based on well-studied sub-exponential hardness assumptions. In particular, a recent work [13] obtains black-box non-interactive non-malleable (and in fact CCA2 secure) commitments assuming keyless collision resistant hash functions, against uniform adversaries.

Relying on CCA2 Secure Commitments. We now discuss modifications to the CCA1 transformation discussed in the previous section. Specifically, we will replace the non-interactive commitment (used to generate ciphertext components $\{c_{1,i}\}_{i \in [n]}$) in the construction outlined above, with a CCA2 secure commitment. Intuitively, using CCA2 secure commitments ensures that no matter how we change the $\{c_{1,i}\}_{i \in [n]}$ components in the challenge ciphertext, the corresponding $\{c_{1,i}\}_{i \in [n]}$ components in the adversary's decryption queries do not change (except in a computationally indistinguishable way). Proving that the resulting protocol is actually a CCA2 secure encryption scheme, is much trickier. We encounter several technical barriers in this process, which we discuss below.

Arguing Security. Recall that in the CCA2 security game, the adversary is allowed access to a decryption oracle both before and after the challenge phase, where the adversary outputs m_0, m_1 and then obtains an encryption of m_b for b sampled uniformly at random.

We will consider a sequence of hybrid experiments similar to the CCA1 setting, where the decryption oracle as well as the challenge ciphertext are modified in small increments. The first hybrid corresponds to providing

the adversary access to the actual decryption oracle together with an encryption of m_b and the last one corresponds to providing the adversary access to the actual decryption oracle together with an encryption of uniform randomness.

The very next hybrid experiment, just like the CCA1 setting, is an exponential time hybrid that samples equivocal commitments $\{c_{2,i}\}_{i \in [n]}$ for the challenge ciphertext, together with randomness $\{y_{0,i}\}_{i \in [n]}$ and $\{y_{1,i}\}_{i \in [n]}$ that can be used to equivocally open these commitments to 0 and 1 respectively.

The third hybrid additionally modifies the components $c_{1,i}$ to “drown” out information about s via noise. In particular, while in the real game, the values $c_{1,i}$ are always commitments to $y_{s_i,i}$, in the challenge ciphertext these values are modified to become commitments to $y_{1,i}$, irrespective of what s_i is. On the other hand, the values $c_{0,i}$ remain encryptions of $y_{s_i,i}$, exactly as in the real experiment. At this point, the proof of indistinguishability of hybrids already significantly diverges from the CCA1 setting. Specifically, the proof of indistinguishability between this hybrid and the previous one, in the CCA1 setting, relied on non-uniform security of the non-interactive commitment – in order to perform the exponential time computation needed to equivocate the hash function. Here, we would ideally like to rely on CCA secure commitments which are potentially only secure against uniform adversaries (eg, the black-box construction in [13] which is only secure against uniform adversaries). One option could be to assume that the CCA2 commitment is “hard” against adversaries running in time that is sufficient to compute openings of the equivocal commitment.

In the fourth hybrid, we would like to modify the decryption oracle so that instead of decrypting using the secret key of the public key encryption scheme, decryption is performed by running in time exponential in the security parameter of the commitment scheme (specifically, in time $2^{k^{\min(\epsilon, \epsilon)}}$) and performing a brute-force search for the randomness used to create the commitments $\{c_{1,i}\}_{i \in [n]}$. This hybrid is indistinguishable from the previous one only if an adversary cannot find ciphertexts that decrypt differently when using the secret key of the encryption scheme versus the brute-force algorithm discussed above: in other words if the adversary cannot query the oracle with “ambiguous” ciphertexts.

This is where the CCA2 setting diverges most significantly from the CCA1 setting. In the CCA1 setting, we could prove that the adversary does not make ambiguous decryption queries by relying on uniform binding of the equivocal commitment, but this is no longer true in the CCA2

setting. Specifically, we need to rule out an adversary that given ambiguous ciphertexts, creates new ones.

Therefore, in the proof, we will now have to rely on CCA2 commitments to maintain an invariant across all the hybrids discussed above. The invariant is as follows: except with negligible probability, the adversary does not make any oracle query for which there exists some $i \in [n]$ such that the components $(c_{0,i}, c_{1,i})$ encrypt/commit to two different openings of the string $c_{2,i}$.

To ensure that this invariant holds in the initial hybrid that corresponds to the real CCA2 experiment, we will use any adversary that breaks the invariant to contradict the binding property of the equivocal commitment. The corresponding reduction would have to *extract* two openings for the same equivocal commitment string, from a decryption query provided by the adversary. In particular, these openings will actually be the plaintexts underlying the ciphertext $c_{0,i}$ and the commitment string $c_{1,i}$. Extracting these two openings involves decrypting $c_{0,i}$ under sk_0 , and brute-force breaking the CCA2 commitment string $c_{1,i}$. This use of brute force necessitates that the binding property of the equivocal commitment be hard to break even in time that is sufficient to break the CCA2 commitment.

But recall that arguing indistinguishability for the third hybrid actually required the exact opposite property: that the CCA2 commitment be hard to break even by adversaries running in time that is sufficient to compute openings of the equivocal commitment. It appears that we are at an impasse here, since we need the equivocal commitment and the CCA2 commitment to each take longer time to break than the other. One way to resolve this is to rely on a non-uniform reduction to argue indistinguishability between the second and third hybrids. But recall that the underlying black-box CCA commitments of [13] achieve only uniform security, at least when relying on on keyless collision resistant hash functions against uniform adversaries.

Fortunately for us, it turns out that [13] prove a much stronger property than uniform CCA security – they actually establish computation enabled CCA security. The computation enabled property allows the attacker to submit a randomized turing machine P at the beginning of the game. The challenger can run the program P and output the result for the attacker at the beginning of the game: crucially, the running time of P can be much larger than the uniform running time allowed to the adversary. This added property actually achieves a flavor of non-uniformity that helps our argument go through, by allowing us to perform special

heavy computation at the beginning of the reduction between hybrids 2 and 3, while at the same time, allowing the binding property of the equivocal commitment to be hard to break even in time that is sufficient to break the CCA2 commitment.

Once we have these ingredients in place, we still need to ensure that the invariant continues to hold in all the other hybrids described above. This is tricky because checking the invariant involves decrypting $\{c_{0,i}\}_{i \in [n]}$ under sk_0 , and also finding the messages committed via the CCA2 commitment strings $\{c_{1,i}\}_{i \in [n]}$, which may not necessarily be an efficient process. Recall that in the very next hybrid, we simply sample the commitment strings in an equivocal way – this hybrid is statistically indistinguishable from the previous one, and therefore the invariant also holds in this hybrid. In the hybrid after that, the commitment strings $c_{1,i}$ are modified in the challenge ciphertext to drown out information about s . Here, in order to prove that the invariant holds, we rely on CCA2 security of the commitment to find the messages committed via the CCA2 commitment strings $\{c_{1,i}\}_{i \in [n]}$ in all of the adversary’s queries.

In the fourth hybrid, we change the way the adversary’s queries are decrypted: here, we can prove (this time, by relying on the invariant) that the adversary does not make decryption queries that decrypt differently, except with negligible probability. In the next hybrid, we modify the decryption oracle again to go back to using the secret key sk_0 of the public key encryption scheme to decrypt. At this point, we are no longer able to argue that the invariant holds, but note that we only needed the invariant to argue that the way the adversary’s queries are decrypted can be changed without affecting the adversary’s advantage. Therefore, this point on, we will not make any changes to how the adversary’s queries are decrypted, and so all we will need to do is argue indistinguishability of the subsequent hybrids. At this point, the only remaining information about s is in the *randomness* used to obtain $\{c_{i,0}, c_{i,1}\}_{i \in [n]}$ in the challenge ciphertext. In the next couple of hybrids, we carefully re-order the randomness and rely on the security of the hinting PRG to switch to using uniform randomness everywhere. All this while, we decrypt the adversary’s oracle queries by breaking the CCA commitments (via brute-force). As a result, our reductions run in superpolynomial time, and we rely on sub-exponential hardness of the hinting PRG. This is different from the CCA1 setting where we could first go back to decrypting the adversary’s oracle queries in polynomial time and then rely on polynomial hardness of the hinting PRG.

We provide some additional technical details about how we implement the invariant discussed in this overview. Specifically, we insert a hybrid after the first hybrid, where the experiment aborts (and the adversary wins) if he makes an oracle query that breaks the invariant: that is, if the adversary makes an oracle query for which there exists $i \in [n]$ such that the components $(c_{0,i}, c_{1,i})$ encrypt/commit to two different openings of the string $c_{2,i}$. This (inefficient) check is performed in all subsequent hybrids up until the fourth one, where we change the way the adversary’s queries are decrypted. We perform careful reductions to argue indistinguishability of these hybrids while performing this inefficient check (as described above). After the fourth hybrid, we no longer need the invariant and we therefore remove this check before proceeding with subsequent hybrids. This concludes an overview of our construction and proof of security.

2.4 On Security against Non-Uniform Adversaries

Very recently, the security of keyless hash functions against adversaries with *non-uniform advice* has also been explored; in particular, [1,2,21] defined and constructed keyless collision-resistant hash functions that satisfy the following property: there exists a polynomial $p(\cdot)$ such that for any polynomial $s(\cdot)$, any PPT adversary with $s(\kappa)$ bits of non-uniform advice cannot find more than $p(s(\kappa))$ pairs of collisions. Subsequently, [3] used these hash functions and (sub-exponential) NIWIs to obtain one-message zero-knowledge without trusted setup and a weak form of soundness against provers with non-uniform advice.

We observe that relying on non-uniform secure primitives; more specifically substituting keyless collision-resistant hash functions against uniform adversaries with keyless collision-resistant hash functions against adversaries with non-uniform advice as described above, helps make our CCA constructions secure against non-uniform adversaries. In other words, we can make a stronger assumption on the underlying keyless hash function, to obtain stronger (non-uniform) security. The only difference would be the observation that an adversary with polynomial advice can only find polynomially many collisions, and therefore query the decryption oracle with only polynomially many ambiguous ciphertexts – the answers to which can be non-uniformly fixed and hardwired into the oracle.

2.5 On Setting Parameters for CCA Compatibility

For both the CCA1 and CCA2 transformations, our non-interactive commitment scheme used to create $\{c_{1,i}\}_{i \in [n]}$ needs to be easier to break

“along some axis of hardness” than the PKE scheme so that there is a way to open it, while the PKE scheme is still hard. Our axis of choice in this paper, is basic computation time. As a result, our theorem statement requires the KeyGen algorithm to be sub-exponentially CPA compatible, i.e. to give rise to a sub-exponentially secure CPA encryption scheme. This could also potentially lead to issues if the original PKE scheme had parameters “on edge” of being broken: since we would need commitment scheme to be even easier to break in terms of computation time.

We point out that in these cases, there could be other different axes of hardness (eg,time-lock puzzles [24,3]) that could be exploited to achieve the same effect. As another example, following [18], one could show that any KeyGen that gives rise to *polynomially hard* PKE scheme secure against quantum adversaries can be combined with a commitment scheme that is quantum *in*-secure, to achieve CCA compatibility. As a result, there is still a way to open the commitments in BQP, while the CPA-secure PKE scheme is still hard. These approaches could improve the concrete parameters that one would need to use to instantiate these transformations, and the exact axis of hardness can be chosen depending upon the specifics of the application.

In the coming sections, we first discuss some key building blocks used by our transformations in Section 3, and define the notion of compatibility in Section 4. Next, we describe our CCA2 compatibility construction in Section 5, with analysis and proof of security deferred to the full version. We can use simpler assumptions and a simpler construction to achieve the weaker goal of CCA1 compatibility, as discussed above. This construction and analysis are deferred to the full version due to lack of space.

3 Preliminaries

In this section we will provide notions and security definitions for public key encryption, keyless collision resistant hash functions and non-interactive perfectly binding commitments. For public key encryption we will formulate a definition that can capture IND-CPA, IND-CCA1 and IND-CCA2 security. For all of our definitions we will be explicit to whether we are describing security against uniform or non-uniform adversaries as our results will be sensitive to this nuance.

We will use κ to denote the security parameter. We will denote by $\text{negl}(\kappa)$ a function that is asymptotically smaller than the inverse of every polynomial in κ .

Public Key Encryption

A public key encryption scheme is specified by a triple of algorithms $(\text{KeyGen}, \text{Enc}, \text{Dec})$, where $\text{KeyGen}(1^\kappa; r_{\text{KeyGen}}) \rightarrow (\text{sk}, \text{pk})$, $\text{Enc}(\text{pk}, \text{msg}; r_{\text{Enc}}) \rightarrow \text{ct}$ and $\text{Dec}(\text{sk}, \text{ct}) \rightarrow \text{msg}$. These algorithms satisfy (perfect) correctness, and IND-CPA/CCA1/CCA2 security, which we will describe below. In addition, we require the following additional properties.

Security Parameter Retrievability. A PKE scheme is *security parameter retrievable* if there exists a polynomial time algorithm `RetrieveParam` that can extract the security parameter used to generate a public key. More formally $\forall \kappa, r_{\text{KeyGen}}$ it must be that $\text{RetrieveParam}(\text{pk}) = \kappa$ where $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\kappa; r_{\text{KeyGen}})$.

Message Recovery from Randomness. We will additionally assume a message recovery from randomness property as given in [22]. Suppose that ct is an encryption of message msg under a (valid) public key cpa.pk and randomness r . Then there exists an algorithm `CPA.Recover` where $\text{CPA.Recover}(\text{cpa.pk}, \text{ct}, r) = \text{msg}$.

The encryption algorithm of any IND-CPA secure PKE scheme can be modified to include this property, as follows. Assume that messages are n bits long. Then one can use n additional random coins r' during encryption and append $\text{msg} \oplus r'$ to the end of the ciphertext. The message can then be recovered from the random coins by a simple XOR operation with r' . Moreover, since the r' portion of the coins are not used elsewhere in encryption, IND-CPA security is preserved. This simple transformation only modifies the encryption algorithm and not the public key. Thus, from a compatibility perspective the setup algorithm remains the same. Therefore in our presentation we will assume that the public key encryption scheme has this property.

Security. We now describe security for public key encryption schemes. We will present a single game of (full) chosen ciphertext security and then derive IND-CCA-1 and IND-CPA security. We define the following security game between a challenger \mathcal{C} and a *stateful* attacker \mathcal{A} .

1. \mathcal{C} runs $\text{KeyGen}(1^\kappa; r_{\text{KeyGen}}) \rightarrow (\text{sk}, \text{pk})$ and gives pk to \mathcal{A} .
2. \mathcal{A} then is allowed to make oracle queries to the function $\text{Dec}(\text{sk}, \cdot)$
3. \mathcal{A} submits two messages $\text{msg}_0, \text{msg}_1 \in \mathcal{M} \times \mathcal{M}$ to \mathcal{C} .
4. \mathcal{C} chooses a coin $b \in \{0, 1\}$ and outputs $\text{ct}^* \leftarrow \text{Enc}(\text{pk}, \text{msg}_b; r_{\text{Enc}})$ for random r_{Enc} .

5. \mathcal{A} then is allowed to make oracle queries to the function $\text{Dec}(\text{sk}, \cdot)$ with the restriction that ct^* is not given as input to the oracle.
6. \mathcal{A} outputs a bit b' .

We refer to the above security game as the IND-CCA2 security game. We define IND-CCA1 security as above, with the exception that the attacker is not allowed any decryption oracle queries in Step 5. We define the IND-CPA security game as above with the exception that the attacker is not allowed any decryption oracle queries in Step 2 and none in Step 5.

Definition 1 (Secure Public Key Encryption). *We will say that a public key encryption scheme is (IND-CCA2, IND-CCA1, IND-CPA) secure if for all non-uniform poly-time attackers \mathcal{A} there exists a negligible function negl such that $\Pr[b' = b] \leq \frac{1}{2} + \text{negl}(\kappa)$ in the (IND-CCA2, IND-CCA1, IND-CPA) security game.*

Definition 2 (Uniform Secure Public Key Encryption). *We will say that a public key encryption scheme is (IND-CCA2, IND-CCA1, IND-CPA) secure if for all poly-time uniform attackers \mathcal{A} we have that there exists a negligible function negl such that $\Pr[b' = b] \leq \frac{1}{2} + \text{negl}(\kappa)$ in the (IND-CCA2, IND-CCA1, IND-CPA) security game.*

In our construction we will also need to consider more fined-grained notions of security where we will specify a time function T that the attacker is allowed to run in. Typically, this will be used to specify security against an attacker that runs in time subexponential in the security parameter.

Definition 3 (Non-uniform T -secure Public Key Encryption). *We will say that a public key encryption scheme is T -(IND-CCA2, IND-CCA1, IND-CPA) secure if for every polynomial $p(\cdot)$, all non-uniform attackers \mathcal{A} running in time at most $p(T(\kappa))$ and with at most $p(T(\kappa))$ bits of advice there exists a negligible function negl such that $\Pr[b' = b] \leq \frac{1}{2} + \text{negl}(\kappa)$ in the (IND-CCA2, IND-CCA1, IND-CPA) security game.*

Definition 4 (Uniform T -secure Public Key Encryption). *We will say that a public key encryption scheme is T -(IND-CCA2, IND-CCA1, IND-CPA) secure if for every polynomial $p(\cdot)$ and all uniform attackers \mathcal{A} running in time at most $p(T(\kappa))$ we have that there exists a negligible function negl such that $\Pr[b' = b] \leq \frac{1}{2} + \text{negl}(\kappa)$ in the (IND-CCA2, IND-CCA1, IND-CPA) security game.*

Non-Interactive Perfect Binding Commitments

Definition 5 (Non-Interactive Perfectly Binding Commitments with Non-Uniform Security). A non-interactive perfectly binding commitment is specified by a poly-time computable randomized algorithm Com that on input $(1^\kappa, \text{msg}; r)$ outputs a commitment string c of length $\ell(\kappa)$, where $\ell(\cdot)$ is a polynomially bounded function, satisfying:

- **Perfect Binding:** For all $c \in \{0, 1\}^*$, κ , $\nexists(\text{msg}_0, \text{msg}_1, r_0, r_1)$ such that $\text{msg}_0 \neq \text{msg}_1$, and $c = \text{Com}(1^\kappa, \text{msg}_0; r_0)$ and $c = \text{Com}(1^\kappa, \text{msg}_1; r_1)$.
- **Computational Hiding:** There exists a negligible function $\text{negl}(\cdot)$ such that $\forall \text{msg}_0, \text{msg}_1 \in \{0, 1\}^*$ s.t. $|\text{msg}_0| = |\text{msg}_1|$, \forall non-uniform PPT \mathcal{A} ,

$$\left| \Pr[\mathcal{A}(\text{Com}(1^\kappa, \text{msg}_0, r)) = 1] - \Pr[\mathcal{A}(\text{Com}(1^\kappa, \text{msg}_1, r)) = 1] \right| \leq \text{negl}(\kappa)$$
 where \mathcal{M} denotes message space and the probability is over r .

We will also assume a property of message recovery from randomness for our commitment scheme. Suppose that c is a commitment of message msg under randomness r . Then there exists an algorithm Com.Recover where $\text{Com.Recover}(c, r) = \text{msg}$. A similar argument to the one given above for public key encryption shows how one can derive a commitment scheme with the message recover from randomness property from any ordinary one. Finally, we will implicitly assume that any message m committed to using security parameter 1^κ can be retrieved with probability 1 by an algorithm running in time $2^\kappa q(\kappa)$ for some polynomial function q . We denote Com.Dec as the algorithm for doing this.

Equivocal Commitments without Setup

Equivocal commitments were proposed by DiCrescenzo, Ishai and Ostrovsky [9] as a bit commitment scheme with a trusted setup algorithm. During normal setup, the bit commitment scheme is statistically binding. However, there exists an alternative setup which produces public parameters along with a trapdoor, that produces commitments which can be opened to either 0 or 1. Moreover, the public parameters of the normal and alternative setup are computationally indistinguishable.

Here we will define a similar primitive, but without utilizing a trusted setup algorithm. In order for such a notion to be meaningful, we will require the commitment scheme to be computationally binding for any uniform T -time attacker, but there will exist an algorithm running in time $\text{poly}(2^\kappa)$ that can be opened to 0 or 1. Moreover, such a commitment with one of the openings should be statistically indistinguishable from a commitment created in the standard manner.

An equivocal commitment scheme without setup consists of 3 algorithms:

$\text{Equiv.Com}(1^\kappa, b) \rightarrow (c, d)$ is a randomized PPT algorithm that on input a bit b and the 1^κ outputs a commitment c and decommitment d .

$\text{Equiv.Decom}(c, d) \rightarrow \{0, 1, \perp\}$ is a deterministic polytime algorithm that takes in part of the commitment and it's opening and reveals the bit that it was committed to or \perp to indicate failure.

$\text{Equiv.Equivocate}(1^\kappa) \rightarrow (c, d_0, d_1)$ is an (inefficient) randomized algorithm that in input 1^κ outputs decommitments to both 0 and 1.

An equivocal commitment is perfectly correct if $\forall b \in \{0, 1\}$

$$\Pr \left[\begin{array}{l} (c, d) \leftarrow \text{Equiv.Com}(1^\kappa, b) \\ b' \leftarrow \text{Equiv.Decom}(c, d) \\ b' = b \end{array} \right] = 1$$

An equivocal commitment is efficient if Equiv.Com and Equiv.Decom run in $\text{poly}(\kappa)$ time, and Equiv.Equivocate runs in time 2^κ .

An equivocal commitment without setup scheme (Equiv.Com , Equiv.Decom , Equiv.Equivocate) is said to be $T(\cdot)$ binding secure if for any *uniform* adversary \mathcal{A} running in time $p(T(\kappa))$ for some polynomial p , there exists a negligible function $\text{negl}(\cdot)$ such that

$$\Pr [(c, d_0, d_1) \leftarrow \mathcal{A}(1^\kappa) : \text{Equiv.Decom}(c, d_0) = 0 \wedge \text{Equiv.Decom}(c, d_1) = 1] \leq \text{negl}(\kappa).$$

We sat that a scheme is equivocal if for all $b \in \{0, 1\}$ the statistical difference between the following two distributions is negligible in κ .

- $\mathcal{D}_0 = (c, d)$ where $\text{Equiv.Com}(1^\kappa, b) \rightarrow (c, d)$.
- $\mathcal{D}_1 = (c, d_b)$ where $\text{Equiv.Equivocate}(1^\kappa) \rightarrow (c, d_0, d_1)$.

We observe that our security definitions do not include an explicit hiding property of a committed bit. This property is actually implied by our equivocal property, and hiding will not be explicitly needed by our proof.

Hinting PRGs

We now provide the definition of hinting PRGs taken from [22]. Let $n(\cdot, \cdot)$ be a polynomial. A n -hinting PRG scheme consists of two PPT algorithms Setup , Eval with the following syntax.

$\text{Setup}(1^\kappa, 1^\ell)$: The setup algorithm takes as input the security parameter κ , and length parameter ℓ , and outputs public parameters pp and input length $n = n(\kappa, \ell)$.

Eval ($\mathbf{pp}, s \in \{0, 1\}^n, i \in [n] \cup \{0\}$): The evaluation algorithm takes as input the public parameters \mathbf{pp} , an n bit string s , an index $i \in [n] \cup \{0\}$ and outputs an ℓ bit string y .

Definition 6. A *hinting PRG scheme* ($\mathbf{Setup}, \mathbf{Eval}$) is said to be secure if for any PPT adversary \mathcal{A} , polynomial $\ell(\cdot)$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\left| \Pr \left[\beta \leftarrow A \left(\mathbf{pp}, \left(y_0^\beta, \{y_{i,b}^\beta\}_{i \in [n], b \in \{0,1\}} \right) \right) \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

where the probability is over $(\mathbf{pp}, n) \leftarrow \mathbf{Setup}(1^\kappa, 1^{\ell(\lambda)})$, $s \leftarrow \{0, 1\}^n$, $\beta \leftarrow \{0, 1\}$, $y_0^0 \leftarrow \{0, 1\}^\ell$, $y_0^1 = \mathbf{Eval}(\mathbf{pp}, s, 0)$, $y_{i,b}^0 \leftarrow \{0, 1\}^\ell \forall i \in [n], b \in \{0, 1\}$, and $y_{i,s_i}^1 = \mathbf{Eval}(\mathbf{pp}, s, i)$, $y_{i,\bar{s}_i}^1 \leftarrow \{0, 1\}^\ell \forall i \in [n]$.

Computation Enabled CCA Commitments

We now define “computation enabled” CCA secure commitments [13]. Intuitively, these are tagged commitments where a commitment to message m under tag \mathbf{tag} and randomness r is created as $\text{CCA.Com}(\mathbf{tag}, m; r) \rightarrow \text{Com}$. The scheme is statistically binding in that for all $\mathbf{tag}_0, \mathbf{tag}_1, r_0, r_1$ and $m_0 \neq m_1$ we have that $\text{CCA.Com}(\mathbf{tag}_0, m_0; r_0) \neq \text{CCA.Com}(\mathbf{tag}_1, m_1; r_1)$.

Our hiding property follows along the lines of chosen commitment security definitions [6] where an attacker gives a challenge tag \mathbf{tag}^* along with messages m_0, m_1 and receives a challenge commitment Com^* to either m_0 or m_1 from the experiment. The attacker’s job is to guess the message that was committed to with the aid of oracle access to an (inefficient) value function CCACom.Val where $\text{CCACom.Val}(\text{Com})$ will return m if $\text{CCA.Com}(\mathbf{tag}, m; r) \rightarrow \text{Com}$ for some r . The attacker is allowed oracle access to $\text{CCACom.Val}(\cdot)$ for any $\mathbf{tag} \neq \mathbf{tag}^*$.

The computation enabled property allows the attacker to submit a randomized turing machine P at the beginning of the game. The challenger will run the program P and output the result for the attacker at the beginning of the game. This added property will be useful in our proof of security. In addition, we require a message recovery from randomness property, which allows one to open the commitment given all the randomness used to generate said commitment.

A computation enabled CCA secure commitment is parameterized by a tag space of size $N = N(\kappa)$ and tags in $[1, N]$. It consists of 3 algorithms:

$\text{CCA.Com}(1^\kappa, \mathbf{tag}, m; r) \rightarrow \text{Com}$ is a randomized PPT algorithm that takes as input the security parameter κ , a tag $\mathbf{tag} \in [N]$, a mes-

sage $m \in \{0, 1\}^*$ and outputs a commitment Com , including the tag Com.tag . We denote the random coins explicitly as r .

$\text{CCACom.Val}(\text{Com}) \rightarrow m \cup \perp$ is a deterministic inefficient algorithm that takes in a commitment Com and outputs either a message $m \in \{0, 1\}^*$ or a reject symbol \perp .

$\text{CCACom.Recover}(\text{Com}, r) \rightarrow m$ is a deterministic algorithm which takes a commitment Com and the randomness r used to generate Com and outputs the underlying message m .

We now define the correctness, efficiency properties, as well as the security properties of perfectly binding and message hiding.

A computation enabled CCA secure commitment scheme is perfectly correct if the following holds. $\forall m \in \{0, 1\}^*$, $\text{tag} \in [N]$ and r we have that

$$\text{CCACom.Val}(\text{CCA.Com}(1^\kappa, \text{tag}, m; r)) = m.$$

A computation enabled CCA secure commitment scheme is efficient if CCA.Com , CCACom.Recover run in time $\text{poly}(|m|, \kappa)$, while CCACom.Val runs in time $\text{poly}(|m|, 2^\kappa)$.

A computation enabled CCA secure commitment is perfectly binding if $\forall m_0, m_1 \in \{0, 1\}^*$ s.t. $m_0 \neq m_1$ there does not exist $\text{tag}_0, \text{tag}_1, r_0, r_1$ such that $\text{CCA.Com}(1^\kappa, \text{tag}_0, m_0; r_0) = \text{CCA.Com}(1^\kappa, \text{tag}_1, m_1; r_1)$.

Remark 1. We remark that this is implied by correctness, as we know that if $\text{CCA.Com}(1^\kappa, \text{tag}_0, m_0; r_0) = \text{CCA.Com}(1^\kappa, \text{tag}_1, m_1; r_1)$, then

$$\begin{aligned} m_0 &= \text{CCACom.Val}(\text{CCA.Com}(1^\kappa, \text{tag}_0, m_0; r_0)) \\ &= \text{CCACom.Val}(\text{CCA.Com}(1^\kappa, \text{tag}_1, m_1; r_1)) = m_1, \end{aligned}$$

but $m_0 \neq m_1$, a contradiction.

We define our message hiding game between a challenger and an attacker. The game is parameterized by a security parameter κ .

1. The attacker sends a randomized and inputless Turing Machine algorithm P . The challenger runs the program on random coins and sends the output to the attacker. If the program takes more than 2^{2^κ} time to halt, the outputs halts the evaluation and outputs the empty string.¹
2. The attacker sends a “challenge tag” $\text{tag}^* \in [N]$.

¹ The choice of 2^{2^κ} is somewhat arbitrary as the condition is in place so that the game is well defined on all P .

3. The attacker makes repeated commitment queries Com . If $\text{Com.tag} = \text{tag}^*$ the challenger responds with \perp . Otherwise it responds as

$$\text{CCACom.Val}(\text{Com}).$$

4. For some w , the attacker sends two messages $m_0, m_1 \in \{0, 1\}^w$.
5. The challenger flips a coin $b \in \{0, 1\}$ and sends $\text{Com}^* = \text{CCA.Com}(\text{tag}^*, m_b; r)$ for randomly chosen r .
6. The attacker again makes repeated queries of commitment Com . If $\text{Com.tag} = \text{tag}^*$ the challenger responds with \perp . Otherwise it sends

$$\text{CCACom.Val}(\text{Com}).$$

7. The attacker finally outputs a guess b' .

We define the attacker's advantage in the game to be $\Pr[b' = b] - \frac{1}{2}$ where the probability is over all the attacker and challenger's coins.

Definition 7. *An attack algorithm \mathcal{A} is said to be e -conforming for some real value $e > 0$ if:*

1. \mathcal{A} is a (randomized) uniform algorithm.
2. \mathcal{A} runs in polynomial time.
3. The program P output by \mathcal{A} in Step 1 of the game terminates in time $p(2^{\kappa^e})$ and outputs at most $q(\kappa)$ bits for some polynomial functions p, q (For all possible random tapes given to the program P).

Definition 8. *A computation enabled CCA secure commitment scheme given by algorithms $(\text{CCA.Com}, \text{CCACom.Val}, \text{CCACom.Recover})$ is said to be e -computation enabled CCA secure if for any e -conforming adversary \mathcal{A} there exists a negligible function $\text{negl}(\cdot)$ such that the attacker's advantage in the game is $\text{negl}(\kappa)$.*

Definition 9. *We say that our CCA secure commitment scheme can be recovered from randomness if the following holds. For all $m \in \{0, 1\}^*$, $\text{tag} \in [N]$, and r , $\text{CCACom.Recover}(\text{CCA.Com}(1^\kappa, \text{tag}, m; r), r) = m$.*

Claim. Let $(\text{CCA.Com}, \text{CCACom.Val})$ be a set of algorithms which satisfy the correctness, efficiency, binding and Definition 8. Then there exists a set of algorithms $(\text{CCA}'.\text{Com}, \text{CCA}'.\text{Val}, \text{CCA}'.\text{Recover})$ which satisfy the same properties as well as Definition 9.

Proof. Consider the following transformation:

$$\begin{aligned} \text{RecoverRandom}(\text{NM} = (\text{CCA.Com}, \text{CCACom.Val})) &\rightarrow \text{NM}' = \\ &(\text{CCA}'.\text{Com}, \text{CCA}'.\text{Val}, \text{CCA}'.\text{Recover}) : \end{aligned}$$

$\text{CCA}'.\text{Com}(\text{tag}, m; r = (r_0, r_1))$: Let $\text{Com} = \text{CCA}.\text{Com}(\text{tag}, r_0)$, and $c = r_1 \oplus m$. Output (Com, c) .
 $\text{CCA}'.\text{Val}(\text{Com}' = (\text{Com}, c))$: Output $\text{CCACom}.\text{Val}(\text{Com})$.
 $\text{CCA}'.\text{Recover}(\text{Com}' = (\text{Com}, c), r = (r_0, r_1))$: Output $c \oplus r_1$.

We can see that correctness, efficiency and binding all hold if they do in the underlying scheme as they call the underlying $\text{CCA}.\text{Com}$, $\text{CCACom}.\text{Val}$ once. To see that Definition 8 still holds, we can consider an attacker \mathcal{A} against NM' . We can construct an attacker for NM by taking the challenge commitment Com , appending w uniformly random bits c' to it, and running \mathcal{A} on (Com, c') . Let m be the underlying message in Com . Since c' is independent and uniformly random, so is $c' \oplus m$, meaning that (Com, c') produces a distribution of Com' identical to $\text{CCA}'.\text{Com}$. Finally, we can see that our transformation satisfies Definition 9 as $c \oplus r_1 = m \oplus r_1 \oplus r_1 = m$.

Connecting to Standard Security We now connect our computation enabled definition of security to the standard notion of chosen commitment security. In particular, the standard notion of chosen commitment security is simply the computation enabled above, but removing the first step of submitting a program P . We prove two straightforward lemmas. The first showing that any computation enabled CCA secure commitment scheme is a standard secure one against uniform attackers. The second is that any non-uniformly secure standard scheme satisfies e -computation enabled security for any constant $e \geq 0$.

Definition 10. *A commitment $(\text{CCA}.\text{Com}, \text{CCACom}.\text{Val}, \text{CCACom}.\text{Recover})$ is said to be CCA secure against uniform/non-uniform attackers if for any poly-time uniform/non-uniform adversary \mathcal{A} there exists a negligible function $\text{negl}(\cdot)$ such that \mathcal{A} 's advantage in the above game with Step 1 removed is $\text{negl}(\kappa)$.*

Claim. If $(\text{CCA}.\text{Com}, \text{CCACom}.\text{Val}, \text{CCACom}.\text{Recover})$ is an e -computation enabled CCA secure commitment scheme for some e as per Definition 8, then it is also a scheme that achieves standard CCA security against uniform poly-time attackers as per Definition 10.

Proof. This follows from the fact that any uniform attacker \mathcal{A} in the standard security game with advantage $\epsilon(\kappa) = \epsilon$ immediately implies an e -conforming attacker \mathcal{A}' with the same advantage where \mathcal{A}' outputs a program P that immediately halts and then runs \mathcal{A} .

Claim. If $(\text{CCA.Com}, \text{CCACom.Val}, \text{CCACom.Recover})$ achieves standard CCA security against *non-uniform* poly-time attackers as per Definition 10, then it is an e -computation enabled CCA secure commitment scheme for any e as per Definition 8.

Proof. Suppose \mathcal{A} is an e -conforming attacker for some e with some advantage $\epsilon = \epsilon(\kappa)$. Then our non-uniform attacker \mathcal{A}' can fix the random coins of \mathcal{A} and to maximize its probability of success. Since now \mathcal{A} is deterministic save for randomness produced by the challenger in step 5, this deterministically fixes the P \mathcal{A} sends, so \mathcal{A}' can fix the coins of P to maximize success. Thus, \mathcal{A}' can simulate \mathcal{A} given the above aforementioned random coins of \mathcal{A} and the output of P , both of which are poly-bounded by the fact that \mathcal{A} is e -conforming. Since all non-challenger randomness was non-uniformly fixed to maximize success, \mathcal{A}' has at least advantage ϵ as well. By our definition of standard security hiding, the advantage of \mathcal{A}' must be negligible, so \mathcal{A} 's advantage must be as well.

Decryption in exponential time. We will implicitly assume that any message m committed to using security parameter 1^κ can be retrieved with probability 1 by an algorithm running in time $2^\kappa q(\kappa)$ for some polynomial function q . We denote CCACom.Dec as the algorithm for doing this.

4 Defining CCA Compatibility

In this section we provide formal definitions of what it means for a scheme to be CPA/CCA compatible. This will be a property of any KeyGen algorithm, and our main technical result will establish that CPA compatibility implies CCA compatibility (under appropriate hardness assumptions).

Definition 11 (CPA Compatibility). *An algorithm KeyGen is said to be non-uniform (resp., uniform) T -CPA-compatible for $T = T(\kappa)$ and message space $\mathcal{M}(\kappa)$, if there exist poly-time algorithms Enc, Dec such that $(\text{KeyGen}, \text{Enc}, \text{Dec})$ comprise a public key encryption scheme for message space $\mathcal{M}(\kappa)$, that satisfies $p(T)$ -IND-CPA according to Definition 3 (resp., Definition 4), for every polynomial function $p(\cdot)$.*

Definition 12 (CCA1 Compatibility). *An algorithm KeyGen is said to be non-uniform (resp., uniform) T -CCA1-compatible for message space $\mathcal{M}(\kappa)$, if there exist poly-time algorithms Enc, Dec such that $(\text{KeyGen}, \text{Enc}, \text{Dec})$ comprise a public key encryption scheme for message space $\mathcal{M}(\kappa)$, that satisfies $p(T)$ -IND-CCA1 according to Definition 3 (resp., Definition 4), for every polynomial function $p(\cdot)$.*

Definition 13 (CCA2 Compatibility). *An algorithm KeyGen is said to be non-uniform (resp., uniform) T -CCA2-compatible for message space $\mathcal{M}(\kappa)$, if there exist poly-time algorithms Enc, Dec such that (KeyGen, Enc, Dec) comprise a public key encryption scheme for message space $\mathcal{M}(\kappa)$, that satisfies $p(T)$ -IND-CCA2 according to Definition 3 (resp., Definition 4), for every polynomial function $p(\cdot)$.*

Our main result is that any KeyGen that is non-uniform $T(\lambda)$ -CPA-compatible where $T = 2^{\lambda^c}$ for any constant $c > 0$, is uniform λ -CCA1-compatible and uniform λ -CCA2-compatible, under appropriate computational hardness assumptions.

5 On CCA2 Compatibility

Our Construction

Let κ denote the security parameter, $0 < \delta < 1$ be a constant and $\kappa' = \kappa^\delta$. We now provide our construction of an IND-CCA2 secure encryption system that uses any $2^{\kappa'}$ -CPA compatible KeyGen algorithm, according to Definition 11. Our construction relies on a hinting PRG, non-interactive computation enabled CCA commitments and subexponentially secure equivocal commitments.

Let (CPA.Enc, CPA.Dec) be the encryption and decryption algorithms of the non-uniform $2^{\kappa'}$ -IND-CPA secure public key encryption scheme with randomness-recoverable ciphertexts and perfect decryption correctness, that is guaranteed to exist by Definition 11. We will also assume that the following exist:

- An equivocal commitment (Equiv.Com, Equiv.Decom, Equiv.Equivocate) that is $T = 2^{\kappa'}$ binding secure.
- A $2^{\kappa'}$ -secure hinting PRG scheme HPRG = (HPRG.Setup, HPRG.Eval) against non-uniform adversaries.
- A non-interactive e -computation enabled CCA commitment scheme represented by algorithms (CCA.Com, CCACom.Val, CCACom.Recover), with security parameter κ' and with $e = 1/\delta$ (for the same δ), such that the commitment scheme can be broken in brute force in time $2^{\kappa'}$.
- An existentially unforgeable under chosen message attack (EUF-CMA) signature (Signature.Setup, Sign, Verify) with security parameter κ' .

We will now describe our CCA secure public key encryption scheme $\text{PKE}_{\text{CCA}} = (\text{KeyGen}, \text{CCA.Enc}, \text{CCA.Dec})$ with message space $\{0, 1\}^{\ell(\kappa)}$. For simplicity of notation, we will skip the dependence of ℓ on κ . We will

also assume that the CPA scheme has message space $\{0, 1\}^{\kappa+1}$ and uses $\ell(\kappa)$ bits of randomness for encryption.

KeyGen(1^κ): The **KeyGen** algorithm outputs a public key **cca.pk**.

CCA.Enc(**cca.pk**, $m \in \{0, 1\}^\ell$): The encryption algorithm is as follows:

1. It runs **RetrieveParam**(**cca.pk**) $\rightarrow \kappa$ and then calculates $\kappa' = \kappa^\delta$.
2. It samples $(\text{HPRG.pp}, 1^n) \leftarrow \text{HPRG.Setup}(1^\lambda, 1^\ell)$.
3. It then chooses $s \leftarrow \{0, 1\}^n$.
4. For each $i \in [n]$, it chooses random $r_i \leftarrow \{0, 1\}^\ell$ and sets $\tilde{r}_i = \text{HPRG.Eval}(\text{HPRG.pp}, s, i)$.
5. For each $i \in [n]$, it chooses $v_i \leftarrow \{0, 1\}^\kappa$. It sets $\sigma_i = \text{Equiv.Com}(1^\kappa, s_i; v_i)$, and $y_i = s_i | v_i$.
6. It sets $c = \text{HPRG.Eval}(\text{HPRG.pp}, s, 0) \oplus m$ and for each $i \in [n]$
 - If $s_i = 0$, $c_{0,i} = \text{CPA.Enc}(\text{cpa.pk}, y_i; \tilde{r}_i)$, $c_{1,i} = \text{CCA.Com}(1^{\kappa'}, vk, y_i; r_i)$.
 - If $s_i = 1$, $c_{0,i} = \text{CPA.Enc}(\text{cpa.pk}, y_i; r_i)$, $c_{1,i} = \text{CCA.Com}(1^{\kappa'}, vk, y_i; \tilde{r}_i)$.²
7. It sets $\alpha = \left(\text{HPRG.pp}, 1^n, c, (c_{0,i}, c_{1,i}, \sigma_i)_{i \in [n]} \right)$.
8. It samples $(vk, sk) \leftarrow \text{Signature.Setup}(1^{\kappa'})$.
9. Finally, it computes $\tau = \text{Sign}(sk, \alpha)$, and outputs (vk, α, τ) as the ciphertext.

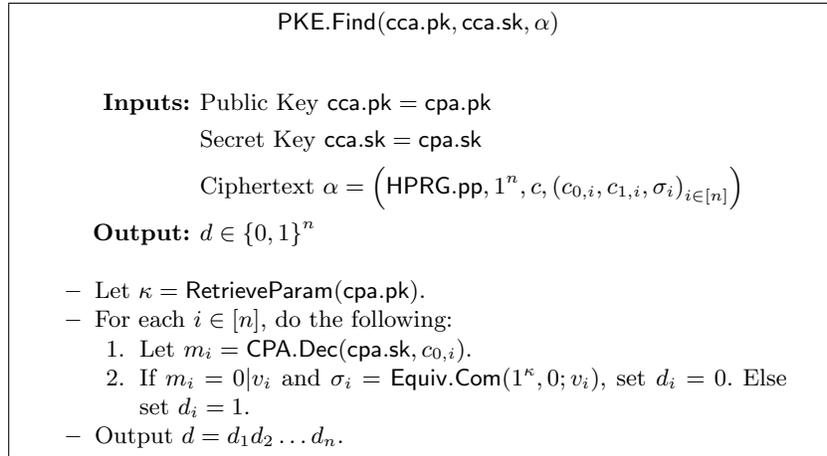


Fig. 1. Routine **PKE.Find**

² For ease of exposition we assume that ℓ coins are both used for encryption with security parameter κ as well as a commitment with security parameter κ' . In practice if one is less than then other the extraneous bits can be truncated.

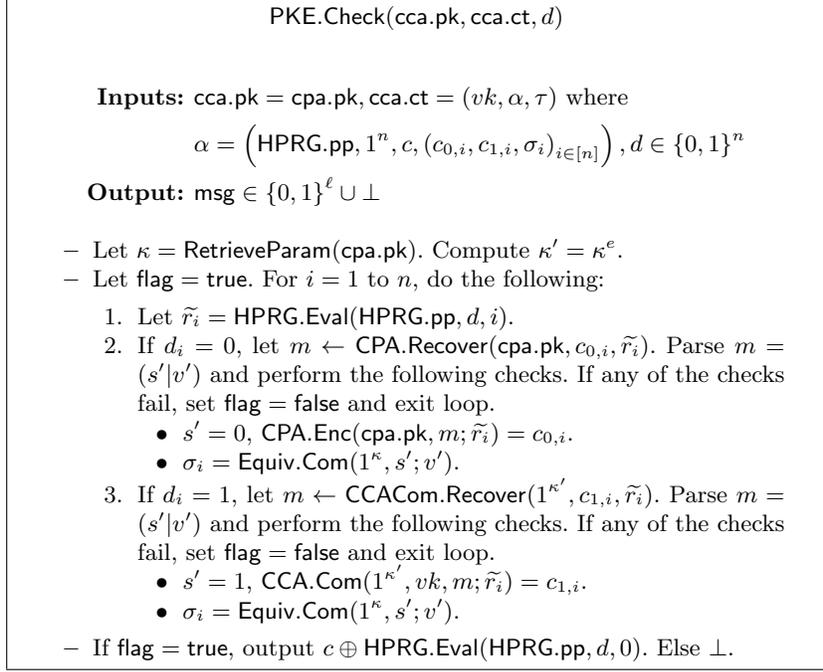


Fig. 2. Routine PKE.Check

CCA.Dec($\text{cca.sk}, \text{cca.pk}, \text{cca.ct}$): Parse ciphertext cca.ct as (vk, α, τ) where $\text{cca.sk} = \text{cpa.sk}$ and $\alpha = \left(\text{HPRG.pp}, 1^n, c, (c_{0,i}, c_{1,i}, \sigma_i)_{i \in [n]} \right)$. Output \perp if $\text{Verify}(vk, \alpha, \tau) = 0$. Otherwise, set $d = \text{PKE.Find}(\text{cca.pk}, \text{cca.sk}, \alpha)$ (where PKE.Find is defined in Figure 1), and output $\text{PKE.Check}(\text{cca.pk}, \text{cca.ct}, d)$ (where PKE.Check is defined in Figure 2).

References

1. Berman, I., Degwekar, A., Rothblum, R.D., Vasudevan, P.N.: Multi-collision resistant hash functions and their applications. In: Nielsen and Rijmen [27], pp. 133–161, https://doi.org/10.1007/978-3-319-78375-8_5
2. Bitansky, N., Kalai, Y.T., Paneth, O.: Multi-collision resistance: a paradigm for keyless hash functions. In: Diakonikolas, I., Kempe, D., Henzinger, M. (eds.) Proceedings of the 50th Annual ACM SIGACT STOC 2018, Los Angeles, CA, USA, June 25-29, 2018. pp. 671–684. ACM (2018). <https://doi.org/10.1145/3188745.3188870>
3. Bitansky, N., Lin, H.: One-message zero knowledge and non-malleable commitments. In: Beimel, A., Dziembowski, S. (eds.) Theory of Cryptography - 16th International Conference, TCC 2018, Panaji, India, November 11-14, 2018, Proceedings, Part I. Lecture Notes in Computer Science, vol. 11239, pp. 209–234. Springer (2018). https://doi.org/10.1007/978-3-030-03807-6_8

4. Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In: Krawczyk [23], pp. 1–12. <https://doi.org/10.1007/BFb0055716>, <https://doi.org/10.1007/BFb0055716>
5. Boldyreva, A., Micciancio, D. (eds.): Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part III, Lecture Notes in Computer Science, vol. 11694. Springer (2019). <https://doi.org/10.1007/978-3-030-26954-8>, <https://doi.org/10.1007/978-3-030-26954-8>
6. Canetti, R., Lin, H., Pass, R.: Adaptive Hardness and Composable Security in the Plain Model from Standard Assumptions. In: Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science. pp. 541–550. FOCS '10 (2010)
7. Castryck, W., Sotáková, J., Vercauteren, F.: Breaking the decisional diffie-hellman problem for class group actions using genus theory. Cryptology ePrint Archive, Report 2020/151 (2020), <https://eprint.iacr.org/2020/151>
8. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk [23], pp. 13–25. <https://doi.org/10.1007/BFb0055717>, <https://doi.org/10.1007/BFb0055717>
9. Crescenzo, G.D., Ishai, Y., Ostrovsky, R.: Non-interactive and non-malleable commitment. In: Proceedings of the Thirtieth Annual ACM STOC, Dallas, Texas, USA, May 23-26, 1998. pp. 141–150 (1998). <https://doi.org/10.1145/276698.276722>, <http://doi.acm.org/10.1145/276698.276722>
10. Damgård, I.B., Pedersen, T.P., Pfitzmann, B.: On the existence of statistically hiding bit commitment schemes and fail-stop signatures. In: Annual International Cryptology Conference. pp. 250–265. Springer (1993)
11. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. *SIAM J. Comput.* **30**(2), 391–437 (2000). <https://doi.org/10.1137/S0097539795291562>, <https://doi.org/10.1137/S0097539795291562>
12. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. *J. Cryptology* **26**(1), 80–101 (2013). <https://doi.org/10.1007/s00145-011-9114-1>, <https://doi.org/10.1007/s00145-011-9114-1>
13. Garg, R., Khurana, D., Lu, G., Waters, B.: Black-box non-interactive non-malleable commitments. Manuscript (2020)
14. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013. pp. 467–476. ACM (2013). <https://doi.org/10.1145/2488608.2488667>, <https://doi.org/10.1145/2488608.2488667>
15. Goldwasser, S., Micali, S.: Probabilistic encryption. *J. Comput. Syst. Sci.* **28**(2), 270–299 (1984). [https://doi.org/10.1016/0022-0000\(84\)90070-9](https://doi.org/10.1016/0022-0000(84)90070-9), [https://doi.org/10.1016/0022-0000\(84\)90070-9](https://doi.org/10.1016/0022-0000(84)90070-9)
16. Halevi, S., Micali, S.: Practical and Provably-Secure Commitment Schemes from Collision-Free Hashing. In: Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology. pp. 201–215. CRYPTO '96 (1996)
17. Jakobsen, J., Orlandi, C.: On the CCA (in)security of mtproto. In: Proc. of the 6th Workshop on Security and Privacy in Smartphones Mobile Devices, SPSM@CCS 2016. pp. 113–116 (2016), <http://dl.acm.org/citation.cfm?id=2994468>
18. Kalai, Y.T., Khurana, D.: Non-interactive non-malleability from quantum supremacy. In: Boldyreva and Micciancio [5], pp. 552–582.

- https://doi.org/10.1007/978-3-030-26954-8_18, https://doi.org/10.1007/978-3-030-26954-8_18
19. Khurana, D., Sahai, A.: How to achieve non-malleability in one or two rounds. In: Umans [30], pp. 564–575. <https://doi.org/10.1109/FOCS.2017.58>, <https://doi.org/10.1109/FOCS.2017.58>
 20. Kitagawa, F., Matsuda, T., Tanaka, K.: CCA security and trapdoor functions via key-dependent-message security. In: Boldyreva and Micciancio [5], pp. 33–64. https://doi.org/10.1007/978-3-030-26954-8_2, https://doi.org/10.1007/978-3-030-26954-8_2
 21. Komargodski, I., Naor, M., Yogev, E.: Collision resistant hashing for paranooids: Dealing with multiple collisions. In: Nielsen and Rijmen [27], pp. 162–194. https://doi.org/10.1007/978-3-319-78375-8_6, https://doi.org/10.1007/978-3-319-78375-8_6
 22. Koppula, V., Waters, B.: Realizing chosen ciphertext security generically in attribute-based encryption and predicate encryption. In: Boldyreva, A., Micciancio, D. (eds.) *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part II. *Lecture Notes in Computer Science*, vol. 11693, pp. 671–700. Springer (2019). https://doi.org/10.1007/978-3-030-26951-7_23, https://doi.org/10.1007/978-3-030-26951-7_23
 23. Krawczyk, H. (ed.): *Advances in Cryptology - CRYPTO '98*, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23–27, 1998, Proceedings, *Lecture Notes in Computer Science*, vol. 1462. Springer (1998). <https://doi.org/10.1007/BFb0055715>, <https://doi.org/10.1007/BFb0055715>
 24. Lin, H., Pass, R., Soni, P.: Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles. In: Umans [30], pp. 576–587. <https://doi.org/10.1109/FOCS.2017.59>, <https://ieeexplore.ieee.org/xpl/conhome/8100284/proceeding>
 25. Lombardi, A., Schaeffer, L.: A note on key agreement and non-interactive commitments. *Cryptology ePrint Archive*, Report 2019/279 (2019), <https://eprint.iacr.org/2019/279>
 26. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: Ortiz, H. (ed.) *Proceedings of the 22nd Annual ACM STOC*, 1990. pp. 427–437. ACM (1990). <https://doi.org/10.1145/100216.100273>, <https://doi.org/10.1145/100216.100273>
 27. Nielsen, J.B., Rijmen, V. (eds.): *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II, *Lecture Notes in Computer Science*, vol. 10821. Springer (2018), <https://doi.org/10.1007/978-3-319-78375-8>
 28. Pass, R.: Unprovable security of perfect NIZK and non-interactive non-malleable commitments. *Computational Complexity* **25**(3), 607–666 (2016). <https://doi.org/10.1007/s00037-016-0122-2>, <http://dx.doi.org/10.1007/s00037-016-0122-2>
 29. Rackoff, C., Simon, D.R.: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In: Feigenbaum, J. (ed.) *Advances in Cryptology — CRYPTO '91*. pp. 433–444. Springer Berlin Heidelberg, Berlin, Heidelberg (1992)
 30. Umans, C. (ed.): *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, Berkeley, CA, USA, October 15–17, 2017. *IEEE Computer Society* (2017), <https://ieeexplore.ieee.org/xpl/conhome/8100284/proceeding>