A Simpler Construction of CCA2-Secure Public-Key Encryption under General Assumptions

Yehuda Lindell*

IBM T.J. Watson
19 Skyline Drive, Hawthorne,
New York 10532, USA.
lindell@us.ibm.com

Abstract. In this paper we present a simpler construction of a public-key encryption scheme that achieves adaptive chosen ciphertext security (CCA2), assuming the existence of trapdoor permutations. We build on previous works of Sahai and De Santis et al. and construct a scheme that we believe is the easiest to understand to date. In particular, it is only slightly more involved than the Naor-Yung encryption scheme that is secure against passive chosen-ciphertext attacks (CCA1). We stress that the focus of this paper is on *simplicity* only.

1 Introduction

One of the most basic tasks of cryptography is that of providing encryption schemes that enable the safe delivery of private messages on an open network. Such an encryption scheme should reveal no information about the plaintext to an eavesdropping adversary. However, it may be necessary to protect the privacy of messages from an adversary who has more power than just the ability to eavesdrop. In a chosen-ciphertext attack, the adversary has access to a decryption oracle and uses this in an attempt to “break” the encryption scheme. Such attacks come in two flavours: *passive* chosen-ciphertext attacks (CCA1), where the adversary can access the decryption oracle only up until the point that it receives a challenge ciphertext, and *adaptive* chosen-ciphertext attacks (CCA2), where the adversary can even access the decryption oracle after it receives the challenge ciphertext. (In the latter case, the adversary can query the decryption oracle for any ciphertext except for the challenge itself.) This is a very strong attack; nevertheless, there are real settings in which this level of security is required (see [2] for an example of an attack on an RSA standard that was made possible due to the fact that the encryption scheme used was not CCA2-secure). We refer the reader to [19] for a survey on the importance of CCA2 security.

Feasibility and efficiency. Two rather distinct directions of research have been considered with respect to secure encryption (and cryptography in general).

* This work was carried out while the author was at the Weizmann Institute of Science.

One direction of research focuses on proving the *feasibility* of obtaining secure schemes, while the other concentrates on doing this *efficiently*. The latter research usually relies on specific number-theoretic (or other) complexity assumptions, whereas the former prefers to only assume the existence of some general primitive, like a trapdoor permutation. This paper considers the question of the feasibility of obtaining CCA2-secure encryption, under the assumption that trapdoor permutations exist. It is well known that such a feasibility result has already been established. However, known constructions of CCA2-secure encryption schemes under general assumptions are rather complicated. Thus, despite their importance, it is hard to teach these schemes in a course on cryptography, for example. The aim of this paper is to improve this situation.

The Naor-Yung paradigm [15]. Our encryption scheme follows the Naor-Yung paradigm for constructing CCA1-secure encryption schemes. According to this paradigm, the plaintext is encrypted twice (independently), and then a non-interactive zero-knowledge proof (NIZK) is used in order to prove that both ciphertexts are encryptions to the same plaintext. The intuition behind this idea is that if the adversary manages to obtain two encryptions of the same plaintext with independent keys, then essentially it must already “know” the plaintext. Therefore, the decryption oracle that it is given is of no help.

The history of the feasibility of CCA2-encryption. The first CCA2-secure encryption scheme was presented in a breakthrough work by Dolev, Dwork and Naor (DDN) [7]. However, their construction is rather complicated, requiring many multiple encryptions and an involved key-selection technique. An important step in the simplification of CCA2-secure encryption was taken by Sahai [17] who showed that CCA2-encryption schemes can actually be constructed using the Naor-Yung paradigm. This involved introducing a stronger notion of NIZK proofs, called *one-time simulation-sound* NIZK. Loosely speaking, one-time simulation-soundness ensures that it is not feasible for an adversary to generate an accepting NIZK proof of a false statement, even if the reference string is generated by the simulator and even if a simulated proof (of a not necessarily true statement) is observed.¹ Sahai showed that the Naor-Yung encryption scheme, with the NIZK proof system replaced by one which is one-time simulation-sound, is CCA2-secure. Unfortunately, much of the complexity of the DDN construction remained in Sahai’s construction of this strong NIZK. Thus, on the one hand, the Sahai high-level construction is significantly simpler than DDN; however, when considering all the details, it is still quite involved.

Recently, De Santis et al. [6,18] presented a very elegant and far simpler construction of simulation-sound NIZK. The aim of their work was actually to strengthen the notion even further to *many-time* simulation-soundness. (In a many-time simulation-sound NIZK, the soundness is preserved even if the adversary observes many simulated proofs.) Nevertheless, their construction also yields a simpler CCA2-secure encryption scheme.

¹ This is in contrast to regular NIZK proof systems where soundness is only guaranteed relative to a uniformly distributed reference string.

Our contribution. In this paper, we use ideas from [6] and apply them to the problem of *one-time* simulation-soundness. Exploiting the fact that one-time simulation soundness is enough for CCA2-security, we obtain a significant simplification of the [6] construction for many-time simulation-sound NIZK. Our construction is both intuitive and simple, and not less importantly, has a short and easy proof of correctness. By plugging our construction into the Sahai CCA2-secure encryption scheme, we obtain a scheme that is only slightly more involved than the original CCA1-secure encryption scheme of Naor and Yung. Thus, we provide an alternative (and simpler proof) to the following theorem:

Theorem 1 *Assuming the existence of enhanced trapdoor permutations², there exists a public-key encryption scheme that is secure against adaptive chosen-ciphertext (CCA2) attacks.*

Related work. The focus of this work is the construction of public-key encryption schemes that are secure against adaptive chosen-ciphertext attacks, assuming only the existence of (enhanced) trapdoor permutations. As we have mentioned above, the first such scheme was presented by Dolev, Dwork and Naor [7]. On the other hand, our construction builds rather directly on the chain of works of Naor and Yung [15], Sahai [17] and De Santis et al. [6].

The first efficient CCA2-secure encryption scheme (proved in the standard model) was presented in a breakthrough work by Cramer and Shoup [4]. However, their construction relies on a specific complexity assumption (namely, the Decisional Diffie-Hellman assumption). Recently, they presented other CCA2-secure schemes, relying on other specific assumptions (some of which are more standard) [5]. (We stress that our work is incomparable to theirs. On the one hand, they achieve high efficiency while relying on specific complexity assumptions. On the other hand, we assume only the existence of trapdoor permutations, but obtain a scheme that is very inefficient). that is used.)

Much work on the problem of efficient CCA2-secure encryption has been carried out in the random-oracle model; the most famous of these being OAEP [1]. However, when the random oracle is replaced by a concrete hash function, the security argument becomes heuristic only. Thus, the existence of these schemes does not constitute a proof of Theorem 2.

Organization. As we have mentioned, the technical contribution of this paper is the construction of a simple one-time simulation-sound NIZK. Therefore, the main body of the paper focuses on this issue. In particular, Section 2 contains the formal definitions of simulation-sound NIZK proof systems and the cryptographic tools necessary for our construction. Then, in Section 3 we present our construction of a one-time simulation-sound NIZK and its proof of correctness. For the sake of completeness, in Section 4 we describe the Sahai CCA2-secure encryption scheme.

² Enhanced trapdoor permutations have the property that a random element generated by the domain sampler is hard to invert, even given the random coins used by the sampler. See [11, Appendix C] for a full discussion on enhanced trapdoor permutations and why they are needed.

2 Definitions and Cryptographic Tools

2.1 Definitions

In this section, we present the definitions for adaptive non-interactive zero-knowledge (NIZK) and one-time simulation-sound adaptive NIZK. Our formal definitions are essentially taken from [10,17]. We denote the security parameter by n , and an unspecified negligible function by $\mu(n)$ (i.e., $\mu(n)$ grows slower than $1/p(n)$ for every polynomial $p(\cdot)$). We often omit explicit reference to the security parameter n in our notations.

Adaptive non-interactive zero-knowledge. In the model of non-interactive zero-knowledge proofs [3], the prover and verifier both have access to the same uniformly distributed reference string. A proof in this model is a single string sent by the prover to the verifier, and the reference string is used for both generating proofs and verifying their validity. The soundness of the proof system is such that if the reference string is indeed uniformly distributed, then with overwhelming probability, no false theorem can be proved (even by an all-powerful cheating prover). On the other hand, the zero-knowledge property is formulated by stating that there exists a simulator who outputs a reference string and a proof, that are computationally indistinguishable from what is viewed by a verifier in the real setting described above. Notice that the simulator generates both the reference string and the proof and is not expected to simulate proofs relative to a uniformly distributed reference string (which would not be possible to achieve). In particular, this means that the simulator may choose the reference string to be pseudorandom, and according to some specific distribution.

The *adaptivity* of a NIZK system refers both to the soundness and zero-knowledge. In both cases, an adaptive NIZK is one where the statement to be proven is chosen only *after* the reference string has been fixed. Thus, the cheating prover first receives a uniformly distributed reference string, and then attempts to find some $x \notin L$ for which it can provide an accepting proof π for x . The adaptive soundness condition states that the probability that such a π is an accepting proof is negligible. The adaptive zero-knowledge property is formulated by having the simulator output a reference string *before* giving it a statement x for which it must generate a simulated proof. In the formal definition below, a function f is specified that “chooses” a statement x to be proven, based on the reference string R (for soundness, f chooses $x \notin L$, whereas for zero-knowledge f chooses $x \in L$).

Definition 2 (adaptive non-interactive zero-knowledge): *A pair of probabilistic machines (P, V) is called an adaptive non-interactive zero-knowledge proof system for a language L if the following holds:*

- Completeness: *For every $x \in L$,*

$$\Pr[V(x, R, P(x, R)) = 1] > 1 - \mu(|x|)$$

where R is a random variable uniformly distributed in $\{0, 1\}^{\text{poly}(|x|)}$.

- Adaptive Soundness: For every function $f : \{0, 1\}^{\text{poly}(n)} \rightarrow \{0, 1\}^n \setminus L$ and prover P^* ,

$$\Pr[V(f(R), R, P^*(R)) = 1] < \mu(n)$$

where R is a random variable uniformly distributed in $\{0, 1\}^{\text{poly}(|x|)}$.

- Adaptive Zero-Knowledge: There exists a probabilistic polynomial-time simulator $S = (S_1, S_2)$ such that for every probabilistic polynomial-time function $f : \{0, 1\}^{\text{poly}(n)} \rightarrow \{0, 1\}^n \cap L$, the ensembles $\{f(R_n), R_n, P(f(R_n), R_n)\}_{n \in \mathbb{N}}$ and $\{S^f(1^n)\}_{n \in \mathbb{N}}$ are computational indistinguishable, where R_n is a random variable uniformly distributed in $\{0, 1\}^{\text{poly}(n)}$ and $S^f(1^n)$ denotes the output from the following experiment:

1. $(r, s) \leftarrow S_1(1^n)$: Simulator S_1 (upon input 1^n) outputs a reference string r and some state information s to be passed on to S_2 .
2. $x \leftarrow f(r)$: the statement x to be proven is chosen.
3. $\pi \leftarrow S_2(x, r, s)$: Simulator S_2 generates a simulated proof π that $x \in L$.
4. Output (x, r, π) .

Adaptive NIZK proof systems can be constructed from any enhanced trapdoor permutation [8]. We note that any NIZK proof system is *witness-indistinguishable*, where informally speaking, witness-indistinguishability means that proofs generated using one witness are indistinguishable from proofs generated using a different witness [9].

One-time simulation-sound adaptive NIZK. Loosely speaking, a NIZK proof system is one-time simulation-sound if the soundness condition holds even with respect to a reference string generated by the simulator (and not uniformly distributed), and even after a single simulated proof (of a not necessarily correct statement) has been observed. Of course, it is always possible for a cheating prover to simply copy the simulated proof that it observed. Therefore, the requirement is that it is infeasible to efficiently compute any *other* pair of an incorrect statement and accepting proof.

Definition 3 (one-time simulation soundness): Let (P, V) be an adaptive NIZK proof system for a language L , and let $S = (S_1, S_2)$ be a simulator for (P, V) . Then, we say that (P, V, S) is **one-time simulation-sound**, if for every probabilistic polynomial-time adversary $A = (A_1, A_2)$, it holds that the probability that A succeeds in the following experiment is negligible:

1. $(r, s) \leftarrow S_1(1^n)$.
2. $(x, a) \leftarrow A_1(r)$: Adversary A_1 receives the (simulator-generated) reference string r and outputs a statement x for which it wants to see a proof, and state information a for A_2 .
3. $\pi \leftarrow S_2(x, r, s)$.
4. $(x', \pi') \leftarrow A_2(x, r, \pi, a)$: Adversary A_2 receives the simulated proof, and outputs a statement x' and a proof π' that $x' \in L$.

5. We say that A **succeeds** if it outputs an accepting proof of a false statement (and did not copy the proof π). That is, A succeeds if $x' \notin L$, $(x', \pi') \neq (x, \pi)$ and $V(x', r, \pi') = 1$.

If there exists an S such that (P, V, S) is one-time simulation sound, then we say that the proof system (P, V) is a one-time simulation-sound adaptive NIZK.

As we have mentioned above, the notion of simulation-soundness was first introduced by [17] who also presented a construction for one-time simulation-soundness (and an extension to allow for any predetermined polynomial number of simulated proofs). Unbounded simulation-soundness (allowing any polynomial number of simulated proofs) was later demonstrated in [6].

2.2 Cryptographic Tools

In this section, we present informal definitions for the cryptographic tools that we use in constructing our one-time simulation-sound NIZK. These tools are standard; however we add minor (yet important) additional requirements. We note that it is easy to obtain these additional requirements using known techniques.

Non-interactive perfectly-binding commitment schemes. Loosely speaking, a non-interactive perfectly-binding commitment scheme is a *probabilistic* algorithm C with the following properties:

- *Hiding:* for every two strings s_1 and s_2 (such that $|s_1| = |s_2|$), it is hard to distinguish $\{C(s_1)\}$ from $\{C(s_2)\}$.
- *Binding:* for every two strings $s_1 \neq s_2$, the range of $C(s_1)$ is disjoint from the range of $C(s_2)$. (Thus, given $C(s_1)$, it is impossible to decommit to any value other than s_1 .)

We denote by $C(s; r)$ the output of the commitment scheme C upon input $s \in \{0, 1\}^n$ and using random coins $r \in_R \{0, 1\}^{\text{poly}(n)}$. (Thus, the binding property states that for $s_1 \neq s_2$, it holds that $C(s_1; r_1) \neq C(s_2; r_2)$ for every r_1 and r_2 .) In addition to the above, we require the following properties:

- *Pseudorandom range:* We require that the output of the commitment algorithm be pseudorandom. This property is fulfilled by the following commitment scheme based on one-way permutations: Let f be a one-way permutation and b a hard-core of f . Then, $C(\sigma) \stackrel{\text{def}}{=} (f(U_n), b(U_n) \oplus \sigma)$, where U_n denotes the uniform distribution over $\{0, 1\}^n$.
- *Negligible support:* We require that a random string is a valid commitment with only negligible probability. This is easily obtained from the above-described commitment scheme based on one-way permutations, by requiring that any commitment to s is preceded by a commitment to 0^n . That is, define $C'(s) = (\text{Commit}(0^n), \text{Commit}(s))$, where each bit is separately committed to using $C(\sigma) = (f(U_n), b(U_n) \oplus \sigma)$.³

³ We remark that if it suffices to obtain soundness for polynomial-time provers only, then the commitment scheme used need not have negligible support.

We note that the Naor commitment scheme [13] as is, has both of these above properties. (Although the [13] commitment scheme is interactive, the receiver message can be hardwired into the common reference string, and so suffices for our needs here.)

“Strong” one-time signature schemes. Loosely speaking, a one-time signature scheme is an existentially unforgeable signature scheme (secure against a chosen-message attack), with the restriction that the signer must only sign a single message with any key. Thus, such a signature scheme is defined as a triplet of algorithms $(G, \text{Sign}, \text{Verify})$, where G is a probabilistic generator that outputs a signing-key sk and a verification-key vk . The validity of the signature scheme fulfills that for every message m , $\text{Verify}(vk, m, \text{Sign}(sk, m)) = 1$, where $(vk, sk) \leftarrow G(1^n)$ (i.e., honestly generated signatures are always accepted). A signature scheme is said to be *secure* if the probability that an efficient forging algorithm S^* succeeds in generating a forgery given a single chosen signature is negligible. More formally, the following experiment is defined: The generator G is run, outputting a key-pair (vk, sk) . Then, S^* is given vk and chooses a message m for which it receives a signature $\sigma = \text{Sign}(sk, m)$. Following this, S^* outputs a pair (m', σ') and it is required that the probability that $\text{Verify}(vk, m', \sigma') = 1$ and $m' \neq m$ is negligible.

As with the commitment scheme, here we also require an additional property. The standard definition of security requires that the forger cannot generate a valid signature on any different message. We strengthen this and require that the forger cannot even generate a different valid signature on the same message. That is, we modify the above (informal) definition and require that the probability that $\text{Verify}(vk, m', \sigma') = 1$ and $(m', \sigma') \neq (m, \sigma)$, is negligible. (Thus, the only valid signature the forger can present is the exact one that it has seen.) Such a signature scheme can be constructed using universal one-way hash functions [14] and 1-1 one way functions (in a similar fashion to the standard construction based on any one-way function [16]). By using 1-1 one-way functions, we ensure that each message has a unique signature and therefore the above strengthening is achieved.

3 Simple One-Time Simulation-Sound NIZK

The motivation behind our construction is as follows. Following [8], (and similarly to [6]) the reference string for the non-interactive proof is divided into two parts. The first part of the string is used by the simulator to simulate a proof, while the second is really used for proving (according to a given NIZK scheme). Typically, in order to prove that $x \in L$, a compound statement of the following structure is proved: either the first part of the reference string has some special property or $x \in L$. Now, when the reference string is chosen at random, the first part will *not* have this special property (except with negligible probability). Therefore, if the proof is accepting it must be that $x \in L$, and soundness holds for the proof system. Zero-knowledge is derived from the fact that the simulator

is able to generate a pseudorandom string that *does* have the special property, enabling it to “cheat” in the proof.

In our scheme, the special property used is that the first part of the reference string is a commitment to a verification-key vk of a one-time signature scheme. That is, the prover sends a verification-key vk along with the proof and proves that: either the first part of the reference string is a commitment to this verification-key vk , or $x \in L$. Furthermore, the prover signs on the proof using the associated signing-key sk (and the verifier checks the validity of this signature using vk). A real prover chooses a random pair of signature keys and generates a proof based on the fact that $x \in L$.

In contrast, the simulator works by generating the reference string so that indeed the first part is a commitment to a verification-key vk (for which it knows the associated signing-key sk). Then, the simulator proves the proof using this fact (rather than the fact that $x \in L$). Notice that the simulator is also able to sign on the proof, as required, because it knows the associated secret-key. The fact that the scheme is simulation-sound follows from the observation that any accepting proof to a false statement must use the property that the first part of the reference string is a commitment to vk . In particular, this means that such a proof is accompanied with a signature using sk (where sk is known only to the simulator). Thus, it is only possible to generate an accepting proof to a false statement if it is possible to forge a signature.

We now formally present the proof system. In the presentation, we refer to an adaptive NIZK proof system, denoted (P, V) , and to commitment and signature schemes. These commitment and signature schemes have the additional properties described in Section 2.2.

Protocol 1 (NIZK Scheme Π)

- **Common reference string:** (r_1, r_2)
- **Prover protocol** (upon input $x \in L$ and a witness w for x):
 1. Choose a random pair of signature keys (vk, sk) for a strong one-time signature scheme.
 2. Let L' be the following language:

$$L' = \{(x, r_1, vk) \mid x \in L \text{ or } r_1 = \text{Commit}(vk)\}$$

Then, generate a non-interactive proof (using reference string r_2) that $(x, r_1, vk) \in L'$. That is, invoke the NIZK prover P for L' on input (x, r_1, vk) , auxiliary-input w and reference string r_2 , obtaining a proof p .

3. Compute $\sigma = \text{Sign}_{sk}(x, p)$.
 4. Output $\pi = (vk, x, p, \sigma)$.
- **Verifier protocol** (upon input x and a proof $\pi = (vk, x, p, \sigma)$):
 1. Check the signature using vk . That is, check that $\text{Verify}_{vk}((x, p), \sigma) = 1$.

2. Invoke the NIZK verifier V and check that p constitutes a correct proof that $(x, r_1, vk) \in L'$ when the reference string equals r_2 (i.e., check that $V((x, r_1, vk), r_2, p) = 1$).
3. Output 1 if and only if the above two checks succeed.

We now proceed to prove the correctness of Protocol 1:

Theorem 4 *Assume that (P, V) is a secure adaptive NIZK proof system for $\mathcal{NP}$, and that the signature and commitment schemes meet the requirements as described in Section 2.2. Then, Protocol 1 constitutes a one-time simulation-sound adaptive NIZK proof system for $\mathcal{NP}$.*

Proof: We begin by proving that Protocol 1 is an adaptive non-interactive proof system. Completeness is immediate. Soundness follows from the fact that the commitment scheme used has negligible support (see Section 2.2), and thus a random string is a valid commitment with only negligible probability. Therefore, when r_1 is uniformly chosen, $x \notin L$ implies that $(x, vk, r_1) \notin L'$, except with negligible probability. Adaptive soundness then follows from the adaptive soundness of the NIZK proof system (P, V) for which proofs are generated using the uniformly distributed reference string r_2 .

We now proceed to demonstrate the zero-knowledge property. As we mentioned in the motivating discussion, intuitively, zero-knowledge holds because a simulator can set r_1 to be a commitment to a verification-key vk , for which it knows the associated signing-key sk . Then, the simulator proves that $(x, vk, r_1) \in L'$ based on the fact that $r_1 = \text{Commit}(vk)$, and without any witness to the fact that $x \in L$. Since, the commitment scheme used has a pseudorandom range, such a r_1 is indistinguishable from a random string. Furthermore, the NIZK proof system (P, V) is witness indistinguishable and therefore the simulated proof cannot be distinguished from a real one. We now provide the exact description of the simulator. The simulator is divided into two parts: S_1 who chooses the reference string and S_2 who generates simulated proofs.

1. Simulator S_1 :

- a) Choose a random pair of signature keys (vk, sk) for a strong one-time signature scheme.
- b) Compute $r_1 = \text{Commit}(vk) = C(vk; r_c)$ for a random r_c .
- c) Choose a *uniformly* distributed string r_2 .
- d) Output (r_1, r_2) and $s = (vk, sk, r_c)$ where s is S_1 's output state information to be given to S_2 .

2. Simulator S_2 upon input $x, (r_1, r_2)$ and $s = (vk, sk, r_c)$:

- a) Invoke the NIZK prover for L' (as defined in Protocol 1) on input (x, r_1, vk) , auxiliary-input r_c and reference string r_2 , and obtain a proof p . (Note that the witness provided to the NIZK prover is for $r_1 = \text{Commit}(vk)$, and not the witness for $x \in L$.)
- b) Compute $\sigma = \text{Sign}_{sk}(x, p)$.
- c) Output $\pi = (vk, x, p, \sigma)$.

There are two differences between a real proof and that provided by the simulator (where we consider the joint distribution $\{x, (r_1, r_2), \pi\}$ of the reference string and the proof). Firstly, the string r_1 generated by S_1 is only pseudorandom (and not random as in a real setting). Secondly, the proof p provided by S_2 is based on the witness for the fact that $r_1 = \text{Commit}(vk)$, rather than being based on the witness for $x \in L$. Intuitively, since r_1 is pseudorandom and the NIZK is witness indistinguishable, these distributions cannot be distinguished. Formally, one defines a hybrid distribution in which $r_1 = \text{Commit}(vk)$ and yet the proof p is based on the witness for $x \in L$. Then, the hybrid is indistinguishable from a real proof by the indistinguishability of r_1 from a random string (everything else is exactly the same). Furthermore, the hybrid is indistinguishable from a simulated proof due to the witness indistinguishability of the NIZK. (Notice that the reference string r_2 for this NIZK is uniformly distributed, and thus the witness indistinguishability property holds.) The indistinguishability of a simulated proof from a real one follows. (We note that from the above proof it follows that the underlying non-interactive proof need not be zero-knowledge; rather, adaptive witness indistinguishability suffices.)

One-time simulation-soundness. Until now, we have shown that Protocol 1 constitutes a non-interactive zero-knowledge proof system. It remains to show that it is also one-time simulation-sound. Intuitively, this holds for the following reason. Let (vk, sk) be the signature keys chosen by the simulator S_1 . Then, if an adversary generates a proof based on the fact that $r_1 = \text{Commit}(vk)$, it must sign on the proof using the key sk (otherwise, the verification of the signature will fail). This constitutes a successful forgery of the signature scheme and therefore can only occur with negligible probability. Details follow.

Assume by contradiction that there exists an adversary $A = (A_1, A_2)$ (as in Definition 3) who sees a simulated proof π for a statement x , and with non-negligible probability outputs a pair $(x', \pi') \neq (x, \pi)$ where $x' \notin L$ and π' is an accepting proof. That is, let (r_1, r_2) be the reference string output by S_1 , let x be the statement that A_1 outputs upon receiving (r_1, r_2) , and let $\pi = (vk, x, p, \sigma)$ be the simulated proof of $x \in L$ that is supplied by S_2 . Then, by the contradicting assumption, with non-negligible probability A_2 outputs an accepting proof $\pi' = (vk', x', p', \sigma')$, such that $(x', \pi') \neq (x, \pi)$, and $x' \notin L$. We consider two different cases:

1. *Case 1 – $vk' \neq vk$:* First recall that by the definition of the simulator S_1 , the string r_1 is such that $r_1 = \text{Commit}(vk)$. However, $vk' \neq vk$ and therefore we have that $r_1 \neq \text{Commit}(vk')$ (by the perfect binding property of the commitment scheme). Therefore, $x' \notin L$ implies that $(x, r_1, vk) \notin L'$. By the (unconditional) soundness of the underlying NIZK scheme, we have that the probability that p' (and therefore π') is an accepting proof is at most negligible.
2. *Case 2 – $vk' = vk$:* In this case, we use A to contradict the (strong) security of the signature scheme. Recall that A 's proof π' is only accepting if $\text{Verify}_{vk'}((x', p'), \sigma') = 1$. Since $(x', \pi') \neq (x, \pi)$ and $vk' = vk$, it holds that

$(x', p', \sigma') \neq (x, p, \sigma)$. Therefore, we have that A received a message and signature $((x, p), \sigma)$ and generated a valid message and signature $((x', p'), \sigma')$, where $((x', p'), \sigma') \neq ((x, p), \sigma)$. By the strong security of the signature scheme, A can succeed in doing this with only negligible probability.

Formally, we construct a forger A' who receives vk and a single oracle query to $\text{Sign}_{sk}(\cdot)$ and successfully forges a signature. A' works exactly like the simulator S except that in Step (b) of S_2 's specification, it “computes” the signature by consulting its oracle. Notice that S_1 and S_2 need no knowledge of sk in order to complete all their other steps. Thus, A' can perfectly emulate the simulation setting for A . Therefore, if A outputs $\pi' = (vk, x', p', \sigma')$, where $(x', p', \sigma') \neq (x, p, \sigma)$ and σ' is a valid signature on (x', p') , then this constitutes a successful forgery of the signature scheme. This implies that A succeeds with at most negligible probability.

This completes the proof. ■

We remark that enhanced trapdoor permutations suffice for securely obtaining all the building blocks required in the construction of Protocol 1. We therefore have the following result:

Proposition 5 *Assuming the existence of enhanced trapdoor permutations, there exists a one-time simulation-sound adaptive NIZK proof system.*

4 The Encryption Scheme

In this section, we describe the CCA2-secure public-key encryption scheme of Sahai [17]. This scheme is exactly the scheme of Naor-Yung [15], with the modification that the NIZK used is one-time simulation-sound. We stress that our contribution is in Section 3, where we present a simple one-time simulation-sound NIZK. Thus, we directly plug our NIZK into the construction (and proof) of [17], obtaining a new (and simpler) CCA2-secure public-key encryption scheme. (For the exposition below, we omit the formal definitions of encryption and CCA2 security and assume that the reader is familiar with them.)

The Naor-Yung paradigm. As we have mentioned, the [17] encryption scheme is based on the Naor-Yung paradigm [15]. According to this paradigm, the plaintext is encrypted twice with independent keys (from an encryption scheme that is secure against chosen-plaintext attacks) and then a NIZK proof is provided to ensure that both encryptions are indeed of the same plaintext. Passive chosen-ciphertext security (CCA1) or adaptive chosen-ciphertext security (CCA2) are achieved by applying NIZKs with certain “special” properties. For CCA1, the NIZK is such that soundness holds even with respect to the pseudorandom string output by the simulator (as long as a simulated proof has not been observed).⁴ For CCA2, the NIZK must be one-time simulation-sound. From here on, we focus

⁴ This is in contrast with standard NIZK proof systems, where soundness is guaranteed only if the reference string is uniformly distributed.

on the CCA2 case. However, we stress that the CCA2 scheme and its proof of security are almost identical to that of the CCA1 scheme. This highlights one of the conceptual advantages of the [17] approach; both CCA1 and CCA2-secure encryption schemes can be presented and proved together (and for almost the price of one).

Formal definition of the scheme. We now present the construction of the encryption scheme. Let (G, E, D) be a public-key encryption scheme that is secure against chosen-plaintext attacks. Furthermore, let (P, V) be a one-time simulation-sound adaptive NIZK proof system for the following NP-language:

$$L = \{(c_1, c_2, pk_1, pk_2) \mid \exists m \text{ s.t. } c_1 = E_{pk_1}(m) \ \& \ c_2 = E_{pk_2}(m)\}$$

That is, L is the language of pairs of ciphertexts (and public-keys), such that both ciphertexts are encryptions of the same message (we denote $c = E_{pk}(m)$, if c is an encryption of m). Then, the CCA2-secure scheme, denoted $(\mathcal{G}, \mathcal{E}, \mathcal{D})$, is defined as follows:

Construction 2 (Adaptive chosen-ciphertext encryption scheme $(\mathcal{G}, \mathcal{E}, \mathcal{D})$):

- **Key Generation:** *Obtain two independent key sets from G . That is, obtain $(pk_1, sk_1) \leftarrow G(1^n)$ and $(pk_2, sk_2) \leftarrow G(1^n)$. Furthermore, choose a uniformly distributed reference string r of the correct length for the NIZK proof system (P, V) . The public key is defined by $\mathcal{PK} = (pk_1, pk_2, r)$ and the secret key is $\mathcal{SK} = (sk_1, sk_2)$.*
- **Encryption:** *In order to encrypt a plaintext m , compute $c_1 = E_{pk_1}(m; r_1)$ and $c_2 = E_{pk_2}(m; r_2)$, for random strings r_1 and r_2 . Then, invoke the NIZK prover P upon (c_1, c_2, pk_1, pk_2) with reference string r , obtaining a proof π . Notice that P can prove this statement efficiently when it is given the witness (m, r_1, r_2) . Finally, output $\mathcal{E}(m) = (c_1, c_2, \pi)$.*
- **Decryption:** *In order to decrypt (c_1, c_2, π) , first verify that π is an accepting proof for the statement (c_1, c_2, pk_1, pk_2) with reference string r . If yes, then decrypt c_1 and output the decryption value m .⁵*

The fact that the above encryption scheme is secure against adaptive chosen-ciphertext attacks has been proven in [17]. That is,

Theorem 6 (Sahai [17]): *Assume that (G, E, D) is a public-key encryption scheme secure against chosen-plaintext attacks, and that (P, V) is a one-time simulation-sound adaptive NIZK proof system. Then, the encryption scheme $(\mathcal{G}, \mathcal{E}, \mathcal{D})$ of Construction 2 is secure against adaptive chosen-ciphertext attacks.*

Combining Theorem 6 with Proposition 5, we obtain the existence of CCA2-secure encryption assuming enhanced trapdoor permutations only (i.e., we obtain Theorem 2). For the sake of completeness, we describe the main ideas behind this proof.

⁵ Our choice of decrypting the first ciphertext c_1 is arbitrary; equivalently, one could define the decryption algorithm by having it decrypt c_2 .

Motivation for the proof of security. The basic idea underlying Construction 2 of $(\mathcal{G}, \mathcal{E}, \mathcal{D})$ is that it is enough to use only one secret-key in order to decrypt ciphertexts. This is because anyone can verify the validity of a NIZK proof. Therefore, given knowledge of any of the two secret-keys, decryption can be carried out by verifying the NIZK and then decrypting. Since the NIZK proof ensures that both encryptions are to the *same* plaintext, it does not matter which secret-key is used. Now, consider an adversary A_{cpa} who carries out a chosen-plaintext attack (CPA) on the scheme (G, E, D) . (Recall that in a CPA attack, the adversary gets no decryption oracle.) Adversary A_{cpa} receives a public-key pk_1 and proceeds to generate a simulated public-key for the scheme $(\mathcal{G}, \mathcal{E}, \mathcal{D})$ which incorporates pk_1 . Specifically, A_{cpa} chooses a second key-pair (pk_2, sk_2) and a NIZK reference string r , thereby obtaining a public-key (pk_1, pk_2, r) for $(\mathcal{G}, \mathcal{E}, \mathcal{D})$. Furthermore, A_{cpa} knows *one* of the two secret-keys (namely sk_2). Therefore, as we have discussed above, A_{cpa} is able to correctly decrypt ciphertexts. The important point is that A_{cpa} is able to correctly simulate a *decryption oracle* for a CCA2-adversary $\mathcal{A}$ who attacks $(\mathcal{G}, \mathcal{E}, \mathcal{D})$. In other words, given only chosen-plaintext ability, A_{cpa} can simulate an adaptive chosen-ciphertext attack for a CCA2-adversary $\mathcal{A}$.

The above shows how the decryption oracle in a CCA2-attack can be simulated by A_{cpa} . However, A_{cpa} must also be able to generate a challenge ciphertext for $\mathcal{A}$ from $(\mathcal{G}, \mathcal{E}, \mathcal{D})$, given its own challenge ciphertext from (G, E, D) . That is, during its attack, A_{cpa} receives some challenge ciphertext c_1 . Based on c_1 , A_{cpa} must provide $\mathcal{A}$ with a challenge. Furthermore, it must be shown that if $\mathcal{A}$ can distinguish ciphertexts in $(\mathcal{G}, \mathcal{E}, \mathcal{D})$ with non-negligible probability, then A_{cpa} can use this to also distinguish ciphertexts in (G, E, D) . Loosely speaking, A_{cpa} generates the needed ciphertext by simply computing $c_2 = E_{pk_2}(0^n)$ (i.e., c_2 is an encryption to garbage) and then providing a proof π that c_1 and c_2 are encryptions of the same message. Of course, this statement may not be true (since A_{cpa} does not know if c_1 is an encryption of 0^n or of some other message). Nevertheless, such a proof can be generated using the NIZK simulator, and this will be indistinguishable from a real ciphertext. Thus, $\mathcal{A}$ receives the challenge ciphertext (c_1, c_2, π) in its CCA2-attack. The point is that $\mathcal{A}$'s challenge ciphertext contains c_1 and therefore any “information” learned by $\mathcal{A}$ about its challenge ciphertext (c_1, c_2, π) can be used by A_{cpa} to derive information about its own challenge ciphertext c_1 . Observe, however, that by the way A_{cpa} constructs the challenge ciphertext, it follows that $\mathcal{A}$ receives a *simulated* NIZK proof π during its attack. Furthermore, $\mathcal{A}$ is able to ask for more decryptions of ciphertexts after seeing this simulated proof, and these ciphertexts contain NIZK proofs. In order for the decryption simulation of A_{cpa} described above to be correct, it must hold that $\mathcal{A}$ cannot generate accepting proofs of false statements, even in such a setting. This is where the one-time simulation-soundness of the NIZK is utilized. Thus we have that A_{cpa} can simulate a complete CCA2-attack for $\mathcal{A}$.

A full proof of Theorem 6 can be found in [17] and in [12, Appendix A].

Acknowledgements. We thank Amit Sahai for pointing out to us that the soundness of the simulation-sound NIZK can be made unconditional by using a commitment scheme with negligible support. We also thank Oded Goldreich for helpful discussions and for encouraging us to write this work.

References

1. M. Bellare and P. Rogaway. Optimal asymmetric encryption – How to encrypt with RSA. In *EUROCRYPT'94*, Springer-Verlag (LNCS 950), pages 92–111, 1994.
2. D. Bleichenbacher. Chosen Ciphertext Attacks Against Protocols Based on the RSA Encryption Standard PKCS#1. In *CRYPTO'98*, Springer-Verlag (LNCS 1462), pages 1–12, 1998.
3. M. Blum, P. Feldman and S. Micali. Non-interactive zero-knowledge and its applications. In *20th STOC*, pages 103–112, 1988.
4. R. Cramer and V. Shoup. A practical public-key cryptosystem provably secure against adaptive chosen ciphertext attack. In *CRYPTO'98*, Springer-Verlag (LNCS 1462), pages 13–25, 1998.
5. R. Cramer and V. Shoup. Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption. In *EUROCRYPT 2002*, Springer-Verlag (LNCS 2332), pages 45–64, 2002.
6. A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano and A. Sahai. Robust Non-interactive Zero-Knowledge. In *CRYPTO 2001*, Springer-Verlag (LNCS 2139), pages 566–598, 2001.
7. D. Dolev, C. Dwork and M. Naor. Non-malleable Cryptography. In *SICOMP*, 30(2):391–437, 2000.
8. U. Feige, D. Lapidot and A. Shamir. Multiple Non-Interactive Zero-Knowledge Proofs Under General Assumptions. In *SICOMP*, 29(1):1–28, 1999.
9. U. Feige and A. Shamir. Witness Indistinguishability and Witness Hiding Protocols. In *22nd STOC*, pages 416–426, 1990.
10. O. Goldreich. *Foundation of Cryptography – Basic Tools*. Cambridge University Press, 2001.
11. O. Goldreich. *Foundations of Cryptography: Volume 2 – Basic Applications*. To be published. Available from <http://www.wisdom.weizmann.ac.il/~oded>.
12. Y. Lindell. A Simpler Construction of CCA2-Secure Public-Key Encryption Under General Assumptions. *Cryptology ePrint Archive*, Report 2002/057, <http://eprint.iacr.org/>, 2002.
13. M. Naor. Bit Commitment using Pseudorandom Generators. *Journal of Cryptology*, 4(2):151–158, 1991.
14. M. Naor and M. Yung. Universal One-Way Hash Functions and their Cryptographic Applications. In *21st STOC*, pages 33–43, 1989.
15. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd STOC*, pages 427–437, 1990.
16. J. Rompel. One-way functions are necessary and efficient for secure signatures. In *22nd STOC*, pages 387–394, 1990.
17. A. Sahai. Non-Malleable Non-Interactive Zero Knowledge and Adaptive Chosen-Ciphertext Security. In *40th FOCS*, pages 543–553, 1999.
18. A. Sahai. Simulation-Sound Non-Interactive Zero Knowledge. Manuscript, 2000.
19. V. Shoup. Why chosen ciphertext security matters. *IBM Research Report RZ 3076*, November, 1998.