

Strengthening Access Control Encryption*

Christian Badertscher, Christian Matt, and Ueli Maurer
Department of Computer Science, ETH Zurich, 8092 Zurich, Switzerland.
{badi, mattc, maurer}@inf.ethz.ch

Abstract

Access control encryption (ACE) was proposed by Damgård et al. to enable the control of information flow between several parties according to a given policy specifying which parties are, or are not, allowed to communicate. By involving a special party, called the *sanitizer*, policy-compliant communication is enabled while policy-violating communication is prevented, even if sender and receiver are dishonest. To allow outsourcing of the sanitizer, the secrecy of the message contents and the anonymity of the involved communication partners is guaranteed.

This paper shows that in order to be resilient against realistic attacks, the security definition of ACE must be considerably strengthened in several ways. A new, substantially stronger security definition is proposed, and an ACE scheme is constructed which provably satisfies the strong definition under standard assumptions.

Three aspects in which the security of ACE is strengthened are as follows. First, CCA security (rather than only CPA security) is guaranteed, which is important since senders can be dishonest in the considered setting. Second, the revealing of an (unsanitized) ciphertext (e.g., by a faulty sanitizer) cannot be exploited to communicate more in a policy-violating manner than the information contained in the ciphertext. We illustrate that this is not only a definitional subtlety by showing how in known ACE schemes, a single leaked unsanitized ciphertext allows for an arbitrary amount of policy-violating communication. Third, it is enforced that parties specified to receive a message according to the policy cannot be excluded from receiving it, even by a dishonest sender.

1 Introduction

1.1 Access Control Encryption—Model and Security Requirements

The concept of *access control encryption (ACE)* was proposed by Damgård, Haagh, and Orlandi [DHO16] in order to enforce information flow using cryptographic tools rather than a standard access control mechanism (e.g., a reference monitor) within an information system. If the encryption scheme provides certain operations (e.g., ciphertext sanitization) and satisfies an adequate security definition, then the reference monitor can be outsourced, as a component called the *sanitizer*, to an only partially trusted service provider. The goal of ACE is that the sanitizer learns nothing not intrinsically necessary. Security must also be guaranteed against dishonest users, whether senders or receivers of information, and against certain types of sanitizer misbehavior.

*This is the full version of the article published by Springer-Verlag in the proceedings of ASIACRYPT 2017, © IACR 2017, available at https://doi.org/10.1007/978-3-319-70694-8_18.

The information flow problem addressed by ACE is defined in a context with a set \mathcal{R} of roles corresponding, for example, to different security clearances. Each user in a system can be assigned several roles. For example the users are employees of a company collaborating on a sensitive project, and they need to collaborate and exchange information by sending messages. Since the information is sensitive, which information a party can see must be restricted (hence the term *access control*), even if some parties are dishonest. In the most general form, the specification of which role may send to which other role corresponds to a relation (a subset of $\mathcal{R} \times \mathcal{R}$) or, equivalently, to a predicate $P: \mathcal{R} \times \mathcal{R} \rightarrow \{0, 1\}$, where $s \in \mathcal{R}$ is allowed to communicate to $r \in \mathcal{R}$ if and only if $P(s, r) = 1$. The predicate P is called the (*security*) *policy*. Typical examples of such policies arise from the Bell-LaPadula [BL73] model where roles are (partially) ordered, and the so-called “no-write-down” rule specifies that it is forbidden for a user to send information to another user with a lower role. Note that for this specific example, the relation is transitive, but ACE also allows to capture non-transitive security policies.

ACE was designed to work in the following setting. Users can communicate anonymously with a sanitizer. If a user wants to send a message, it is encrypted under a key corresponding to the sender’s role. Then the ciphertext is sent (anonymously) to the sanitizer who applies a certain sanitization operation and writes the sanitized ciphertext on a publicly readable bulletin board providing anonymous read-access to the users (receivers). Users who are supposed to receive the message according to the policy (and only those users) can decrypt the sanitized ciphertext.

To ensure security in the described setting, the ACE scheme must at least provide the following guarantees:

1. The encryption must assure privacy and anonymity against dishonest receivers as well as the sanitizer, i.e., neither the sanitizer nor dishonest receivers without access allowed by the policy should be able to obtain information about messages or the sender’s role.
2. A dishonest sender must be unable to communicate with a (potentially dishonest) receiver, unless this is allowed according to the policy. In other words, the system must not provide covert channels allowing for policy-violating communication.

As usual in a context with dishonest senders, the first goal requires security against chosen-ciphertext attacks (CCA) because dishonest users can send a ciphertext for which they do not know the contained message and by observing the effects the received message has on the environment, potentially obtain information about the message. This corresponds to the availability of a decryption oracle, as in the CCA-security definition.

Note that the second goal is only achievable if users cannot directly write to the repository or communicate by other means bypassing the sanitizer, and if the sanitizer is not actively dishonest because a dishonest sanitizer can directly write any information received from a dishonest sender to the repository. The assumption that a user cannot bypass the sanitizer and communicate to another party outside of the system can for example be justified by assuming that users, even if dishonest, want to avoid being caught communicating illegitimately, or if only a user’s system (not the user) is corrupted, and the system can technically only send message to the sanitizer.

Since the sanitizer is not fully trusted in our setting, one should consider the possibility that an unsanitized ciphertext is leaked (intentionally or unintentionally) to a dishonest party. This scenario can be called (*unsanitized*) *ciphertext-revealing attack*. Obviously, all information contained in this ciphertext gets leaked to that party. While this cannot be avoided, such an attack should not enable dishonest parties to violate the security requirements beyond that.

We point out that previously proposed encryption techniques (before ACE), such as attribute-based encryption [SW05; GPSW06] and functional encryption [BSW11], enable the design of schemes where a sender can encrypt messages such that only designated receivers (who possess the required key) can read the message. This captures the access control aspects of *read* permissions, but it does not allow to capture the control of *write/send* permissions. In other words, such schemes only achieve the first goal listed above, not the second one.

1.2 Contributions of this Paper

While the proposal of the ACE-concept and of efficient ACE-schemes were important first steps toward outsourcing access control, the existing security definition turns out to be insufficient for several realistic attack scenarios. The main contributions of this paper consist of uncovering issues with existing definitions and schemes, fixing these issues by proposing stronger security notions, and constructing a scheme satisfying our stronger notions.

Issues with existing definitions and schemes. As argued above, chosen-ciphertext attacks should be considered since the use case for ACE includes dishonest senders. Existing definitions, however, do not take this into account, i.e., the adversary does not have access to a decryption oracle in the security games.

Furthermore, existing notions do not consider ciphertext-revealing attacks. Technically speaking, the security game that is supposed to prevent dishonest senders from transmitting information to dishonest receivers (called no-write game), gives the adversary only access to an encryption oracle that sanitizes ciphertexts before returning them. This means that the adversary has no access to unsanitized ciphertexts. This is not only a definitional subtlety, but can completely break down any security guarantees. We demonstrate that existing ACE schemes allow the following attack: Assume there are three users A , M , and E in the system, where A is honest and by the policy allowed to send information to E , and M and E are dishonest and not allowed to communicate. If A sends an (innocent) message to E and the corresponding unsanitized ciphertext is leaked to M , malleability of the ciphertext can be exploited by M to subsequently communicate an arbitrary number of arbitrary messages chosen by M to E . Note that while this attack crucially exploits malleability of ciphertexts, it is not excluded by CCA security for two reasons: first, CCA security does not prevent an adversary from producing valid ciphertexts for *unrelated* messages, and second, the integrity should still hold if the adversary has the decryption key (but not the encryption key).

Finally, existing security definitions focus on preventing dishonest parties from communicating if disallowed by the policy, but they do not enforce information flow. For example, if user A only has a role such that according to the policy, users B and C can read what A sends, existing schemes do not prevent A from sending a message that can be read by B but not by C , or sending a message such that B and C receive different messages. This is not as problematic as the two issues above, and one can argue that A could anyway achieve something similar by additionally encrypting the message with another encryption scheme. Nevertheless, for some use cases, actually precisely enforcing the policy can be required (consider, e.g., a logging system), and one might intuitively expect that ACE schemes achieve this.

New security definitions. We propose new, stronger security definitions for ACE that exclude all issues mentioned above. First, we give the adversary access to a decryption oracle. More

precisely, the oracle first sanitizes the given ciphertext and then decrypts it, since this is what happens in the application if a dishonest party sends a ciphertext to the sanitizer. Second, we incorporate ciphertext-revealing attacks by giving the adversary access to an encryption oracle that returns unsanitized ciphertexts for arbitrary roles. Finally, we introduce a new security game in which an adversary can obtain encryption keys and decryption keys from an oracle and has to output a ciphertext such that one of the following events occur: either the set of roles that can successfully decrypt the ciphertext (to an arbitrary message) is inconsistent with the policy for all sender roles for which the adversary has an encryption key (in this case, we say the adversary is not *role-respecting*); or the ciphertext can be successfully decrypted with two keys such that two different messages are obtained (in this case, we say the *uniform-decryption* property is violated).

Construction of an ACE scheme for our stronger notions. Our construction proceeds in three steps and follows the general structure of the generic construction by Fuchsbauer et al. [FGKO17]. Since we require much stronger security notions in all three steps, our constructions and proofs are consequently more involved than existing ones. First, we construct a scheme for a primitive we call *enhanced sanitizable public-key encryption (sPKE)*. Second, we use an sPKE scheme to construct an ACE scheme satisfying our strong security notion for the equality policy, i.e., for the policy that allows s to send to r if and only if $r = s$. Third, we show how to lift an ACE scheme for the equality policy to an ACE scheme for the disjunction of equalities policy. This policy encodes roles as vectors $\mathbf{x} = (x_1, \dots, x_\ell)$ and allows role \mathbf{x} to send to role \mathbf{y} if and only if $x_1 = y_1 \vee \dots \vee x_\ell = y_\ell$. As shown by Fuchsbauer et al. [FGKO17], useful policies including the inequality predicate corresponding to the Bell-LaPadula model can efficiently be implemented using this policy by encoding the roles appropriately.

Enhanced sanitizable PKE. An sPKE scheme resembles public-key encryption with an additional setup algorithm that outputs sanitizer parameters and a master secret key. The master secret key is needed to generate a public/private key pair and the sanitizer parameters can be used to sanitize a ciphertext. A sanitized ciphertext cannot be linked to the original ciphertext without the decryption key. We require the scheme to be CCA secure (with respect to a sanitize-then-decrypt oracle) and anonymous. Sanitization resembles rerandomization [Gro04; PR07], also called universal re-encryption [GJJS04], but we allow sanitized ciphertexts to be syntactically different from unsanitized ciphertexts. This allows us to achieve full CCA security, which is needed for our ACE construction and unachievable for rerandomizable encryption.

Our scheme is based on ElGamal encryption [Elg85], which can easily be rerandomized and is anonymous. We obtain CCA security using the technique of Naor and Yung [NY90], i.e., encrypting the message under two independent keys and proving in zero-knowledge that the ciphertexts are encryptions of the same message, which was shown by Sahai to achieve full CCA security if the zero-knowledge proof is simulation-sound [Sah99]. A technical issue is that if the verification of the NIZK proof was done by the decrypt algorithm, the sanitization would also need to sanitize the proof. Instead, we let the sanitizer perform the verification. Since we want to preserve anonymity, this needs to be done without knowing under which public keys the message was encrypted. Therefore, the public keys are part of the witness in the NIZK proof. Now the adversary could encrypt the same message under two different public keys that were not produced together by the key-generation, which would break the reduction. To prevent this, the pair of public keys output by the key-generation is signed using a signature key that is

contained in the master secret key and the corresponding verification key is contained in the sanitizer parameters.

ACE for equality. The basic idea of our ACE scheme for the equality policy is to use for each role, encryption and decryption keys of an sPKE scheme as the encryption and decryption keys of the ACE scheme, respectively. Since we need to prevent dishonest senders without an encryption key for some role from producing valid ciphertexts for that role even after seeing encryptions of other messages for this role and obtaining encryption keys for other roles, we add a signature key to the encryption key, sign this pair using a separate signing key, where the corresponding verification key is part of the sanitizer parameters, and let senders sign their ciphertexts. To preserve anonymity, this signature cannot be part of the ciphertext. Instead, senders prove in zero-knowledge that they know such a signature and that the encryption was performed properly.

ACE for disjunction of equalities. The first step of our lifting is identical to the lifting described by Fuchsbaauer et al. [FGKO17]: for each component of the role-vector, the encryption and decryption keys contain corresponding keys of an ACE scheme for the equality policy. To encrypt a message, this message is encrypted under each of the key-components. In a second step, we enforce role-respecting security with the same trick we used in our ACE scheme for equality; that is, we sign encryption key-vectors together with a signing key for that role, and senders prove in zero-knowledge that they have used a valid key combination to encrypt and that they know a signature of the ciphertext vector.

1.3 Related Work

The concept of access control encryption has been introduced by Damgård et al. [DHO16]. They provided the original security definitions and first schemes. Subsequent work by Fuchsbaauer et al. [FGKO17], by Tan et al. [TZMT17], and by Kim and Wu [KW17] focused on new schemes that are more efficient, based on different assumptions, or support more fine grained access control policies. In contrast to our work, they did not attempt to strengthen the security guarantees provided by ACE.

2 Preliminaries

2.1 Notation

We write $x \leftarrow y$ for assigning the value y to the variable x . For a finite set X , $x \leftarrow X$ denotes assigning to x a uniformly random value in X . For $n \in \mathbb{N}$, we use the convention

$$[n] := \{1, \dots, n\}.$$

By \mathbb{Z}_n we denote the ring of integers modulo n , and by \mathbb{Z}_n^* its multiplicative group of units. The probability of an event A in an experiment E is denoted by $\Pr^E[A]$, e.g., $\Pr^{x \leftarrow \{0,1\}}[x = 0] = \frac{1}{2}$. If the experiment is clear from the context, we omit the superscript. The conditional probability of A given B is denoted by $\Pr[A \mid B]$ and the complement of A is denoted by $\neg A$. For a probabilistic algorithm \mathcal{A} and $r \in \{0, 1\}^*$, we denote by $\mathcal{A}(x; r)$ the execution of \mathcal{A} on input x with randomness r . For algorithms \mathcal{A} and \mathcal{O} , $\mathcal{A}^{\mathcal{O}(\cdot)}(x)$ denotes the execution of \mathcal{A} on input x , where \mathcal{A} has oracle access to \mathcal{O} .

2.2 Security Definitions, Advantages, Efficiency, and Negligibility

We define the security of a scheme via a random experiment (or game) involving an adversary algorithm \mathcal{A} . For a given scheme \mathcal{E} and adversary \mathcal{A} , we define the advantage of \mathcal{A} , which is a function of the security parameter κ . To simplify the notation, we omit the security parameter when writing the advantage, e.g., we write $\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{Sig-EUF-CMA}}$ instead of $\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{Sig-EUF-CMA}}(\kappa)$ for the advantage of \mathcal{A} in the existential unforgeability game for the signature scheme \mathcal{E} . Such a scheme is considered *secure* if $\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{Sig-EUF-CMA}}$ is *negligible* for all *efficient* \mathcal{A} . An algorithm \mathcal{A} is *efficient* if it runs in *probabilistic polynomial time (PPT)*, i.e., \mathcal{A} has access to random bits and there is a polynomial p such that $\mathcal{A}(x)$ terminates after at most $p(|x|)$ steps (on some computational model, e.g., Turing machines) for all inputs x , where $|x|$ denotes the bit-length of x . A function f is *negligible* if for every polynomial p , there exists $n_0 \in \mathbb{N}$ such that $f(n) < 1/p(n)$ for all $n \geq n_0$. While these asymptotic definitions yield concise statements, we will in all proofs derive precise bounds on the advantages, following a concrete security approach.

2.3 Access Control Encryption

We recall the definition of access control encryption by Damgård et al. [DHO16]. For definitions of other cryptographic primitives used in this paper, see [Appendix A](#). Following Fuchsbauer et al. [FGKO17], we do not have sanitizer keys and require Gen to be deterministic. The set of roles is assumed to be $\mathcal{R} = [n]$.

Definition 2.1. An *access control encryption (ACE) scheme* \mathcal{E} consists of the following five PPT algorithms:

Setup: The algorithm Setup on input a security parameter 1^κ and a *policy* $P: [n] \times [n] \rightarrow \{0, 1\}$, outputs a *master secret key* msk and *sanitizer parameters* sp . We implicitly assume that all keys include the finite *message space* \mathcal{M} and the *ciphertext spaces* $\mathcal{C}, \mathcal{C}'$.

Key generation: The algorithm Gen is deterministic and on input a master secret key msk , a role $i \in [n]$, and the type \mathbf{sen} , outputs an *encryption key* ek_i ; on input msk , $j \in [n]$, and the type \mathbf{rec} , outputs a *decryption key* dk_j .

Encryption: The algorithm Enc on input an encryption key ek_i and a message $m \in \mathcal{M}$, outputs a ciphertext $c \in \mathcal{C}$.

Sanitization: The algorithm San on input sanitizer parameters sp and a ciphertext $c \in \mathcal{C}$, outputs a *sanitized ciphertext* $c' \in \mathcal{C}' \cup \{\perp\}$.

Decryption: The algorithm Dec on input a decryption key dk_j and a sanitized ciphertext $c' \in \mathcal{C}'$, outputs a message $m \in \mathcal{M} \cup \{\perp\}$; on input dk_j and \perp , it outputs \perp .

For a probabilistic algorithm \mathcal{A} , consider the experiment $\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-CORR}}$ that given a security parameter 1^κ and a policy P , executes $(sp, msk) \leftarrow \text{Setup}(1^\kappa, P)$, $(m, i, j) \leftarrow \mathcal{A}^{\text{Gen}(msk, \cdot, \cdot)}(sp)$, $ek_i \leftarrow \text{Gen}(msk, i, \mathbf{sen})$, and $dk_j \leftarrow \text{Gen}(msk, j, \mathbf{rec})$. We define the *correctness advantage* of \mathcal{A} (for security parameter κ and policy P) as

$$\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{ACE-CORR}} := \Pr[P(i, j) = 1 \wedge \text{Dec}(dk_j, \text{San}(sp, \text{Enc}(ek_i, m))) \neq m],$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-CORR}}$ and the random coins of Enc , San , and Dec . The scheme \mathcal{E} is called *correct* if $\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{ACE-CORR}}$ is negligible for all efficient \mathcal{A} , and *perfectly correct* if $\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{ACE-CORR}} = 0$ for all \mathcal{A} .

Remark. Correctness of an encryption scheme is typically not defined via a game with an adversary, but by requiring that decryption of an encryption of m yields m with probability 1. This perfect correctness requirement is difficult to achieve for ACE schemes and not necessary for applications because it is sufficient if a decryption error only occurs with negligible probability in any execution of the scheme. Damgård et al. [DHO16] define correctness by requiring that for all m , i , and j with $P(i, j) = 1$, the probability that a decryption fails is negligible, where the probability is over setup, key generation, encrypt, sanitize, and decrypt. While this definition is simpler than ours, it does not guarantee that decryption errors only occur with negligible probability in any execution of the scheme. For example, a scheme could on setup choose a random message m and embed it into all keys such that decryption always fails for encryptions of this particular message. This does not violate the definition by Damgård et al. since for any fixed message, the probability that this message is sampled during setup is negligible (if the message space is large). Nevertheless, an adversary can always provoke a decryption error by sending that particular message m , which is not desirable. The above example might at first sight seem somewhat artificial, and typically, schemes do not have such a structure. However, capturing correctness via an experiment is important when thinking of composition, since we expect that the correctness guarantee still holds when the ACE scheme is run as part of a larger system. In order to meet this expectation, and to exclude the above issue, we formalize correctness via an experiment.

Additionally, Fuchsbauer et al. have defined detectability, which guarantees that decrypting with a wrong key yields \perp with high probability [FGKO17]. This allows receivers to detect whether a message was sent to them. As for correctness, we define it via an experiment. The notion is related to robustness for public-key encryption [ABN10]. We additionally define strong detectability, in which the randomness for the encryption is adversarially chosen.

Definition 2.2. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an ACE scheme and let \mathcal{A} be a probabilistic algorithm. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-DTCT}}$ that given a security parameter 1^κ and a policy P , executes $(sp) \leftarrow \text{Setup}(1^\kappa, P)$, $(m, i, j) \leftarrow \mathcal{A}^{\text{Gen}(msk, \cdot, \cdot)}(sp, msk)$, $ek_i \leftarrow \text{Gen}(msk, i, \text{sen})$, and $dk_j \leftarrow \text{Gen}(msk, j, \text{rec})$. We define the *detectability advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-DTCT}} := \Pr[P(i, j) = 0 \wedge \text{Dec}(dk_j, \text{San}(sp, \text{Enc}(ek_i, m))) \neq \perp],$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-DTCT}}$ and the random coins of Enc , San , and Dec . The scheme \mathcal{E} is called *detectable* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-DTCT}}$ is negligible for all efficient \mathcal{A} . The experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sDTCT}}$ is identical to $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-DTCT}}$ except that \mathcal{A} returns (m, r, i, j) . The *strong detectability advantage* of \mathcal{A} is defined as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sDTCT}} := \Pr[P(i, j) = 0 \wedge \text{Dec}(dk_j, \text{San}(sp, \text{Enc}(ek_i, m; r))) \neq \perp],$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sDTCT}}$ and the random coins of San and Dec . The scheme \mathcal{E} is called *strongly detectable* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sDTCT}}$ is negligible for all efficient \mathcal{A} .

2.4 Existing Security Definitions

Existing notions for ACE specify two core properties: the so-called *no-read rule* and the *no-write rule*. The no-read rule formalizes privacy and anonymity: roughly, an honestly generated ciphertext should not leak anything about the message, except possibly its length, or about the

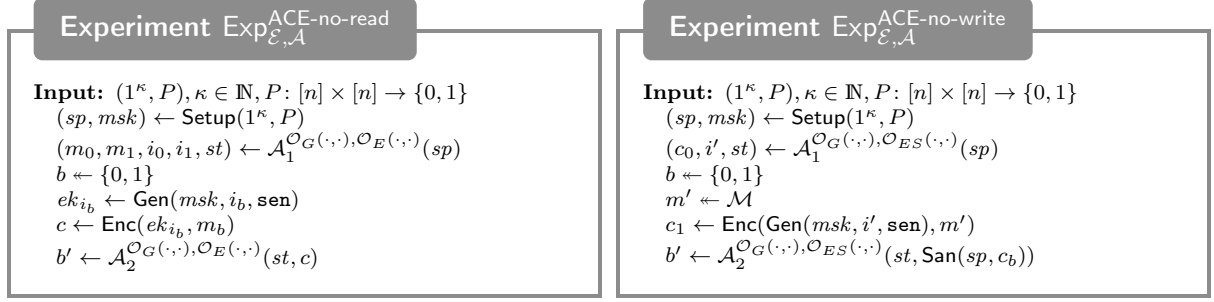


Figure 1: The no-read and no-write experiments for an ACE scheme \mathcal{E} and an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. The oracles are defined as $\mathcal{O}_G(\cdot, \cdot) := \text{Gen}(msk, \cdot, \cdot)$, $\mathcal{O}_E(\cdot, \cdot) := \text{Enc}(\text{Gen}(msk, \cdot, \text{sen}), \cdot)$, and $\mathcal{O}_{ES}(\cdot, \cdot) := \text{San}(sp, \text{Enc}(\text{Gen}(msk, \cdot, \text{sen}), \cdot))$.

role of the sender. The security game allows an adversary to interact with a key-generation oracle (to obtain encryption and decryption keys for selected roles), and an encryption oracle to obtain encryptions of chosen messages for roles for which the adversary does not possess the encryption key. This attack model reflects that an adversary cannot obtain useful information by observing the ciphertexts that are sent to the sanitizer. To exclude trivial attacks, it is not considered a privacy breach if the adversary knows a decryption key that allows to decrypt the challenge ciphertext according to the policy. Similarly, it is not considered an anonymity breach if the encrypted messages are different. We next state the definition of the no-read rule.¹

Definition 2.3. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an ACE scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-read}}$ in Figure 1 and let J be the set of all j such that \mathcal{A}_1 or \mathcal{A}_2 issued the query (j, rec) to the oracle \mathcal{O}_G . The *payload-privacy advantage* and the *sender-anonymity advantage* of \mathcal{A} are defined as

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-read, priv}} &:= 2 \cdot \Pr[b' = b \wedge |m_0| = |m_1| \wedge \forall j \in J P(i_0, j) = P(i_1, j) = 0] - 1, \\ \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-read, anon}} &:= 2 \cdot \Pr[b' = b \wedge m_0 = m_1 \wedge \forall j \in J P(i_0, j) = P(i_1, j)] - 1, \end{aligned}$$

respectively, where the probabilities are over the randomness of all algorithms in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-read}}$. The scheme \mathcal{E} satisfies the *payload-privacy no-read rule* and the *sender-anonymity no-read rule* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-read, priv}}$ and $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-read, anon}}$ are negligible for all efficient \mathcal{A} , respectively. If it satisfies both, it is said to satisfy the *no-read rule*.

The no-write rule of ACE is the core property to capture access control. In a nutshell, if the adversary only possesses encryption keys for roles i and decryption keys for roles j with $P(i, j) = 0$, then he should not be able to create a ciphertext from which, after being sanitized, he can retrieve any information. Technically, in the corresponding security game, the adversary is given a key-generation oracle as above, and in addition an oracle to obtain *sanitized* ciphertexts for selected messages and roles. This attack model corresponds to a setting where an adversary only sees the outputs of a sanitizer, but not its inputs, and in particular no unsanitized ciphertexts generated for roles for which he does not possess the encryption key. The adversary wins if he manages to distinguish the sanitized version of a ciphertext of his choice from a sanitized

¹For anonymity, we adopt here the definition of [DHO16], which is stronger than the one used by Fuchsbauer et al. [FGKO17] since there, anonymity is not guaranteed against parties who can decrypt.

version of a freshly generated encryption of a random message, and if he does not obtain the encryption key for any role i and the decryption key of any role j for which $P(i, j) = 1$, as this would trivially allow him to distinguish.

Definition 2.4. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an ACE scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-write}}$ in Figure 1, let I_1 be the set of all i such that \mathcal{A}_1 issued the query (i, sen) to \mathcal{O}_G , and let J be the set of all j such that \mathcal{A}_1 or \mathcal{A}_2 issued the query (j, rec) to \mathcal{O}_G . We define the *no-write advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-write}} := 2 \cdot \Pr[b' = b \wedge i' \in I_1 \wedge \forall i \in I_1 \forall j \in J P(i, j) = 0 \wedge \text{San}(sp, c_0) \neq \perp] - 1,$$

where the probability is over the randomness of all algorithms in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-write}}$. The scheme \mathcal{E} satisfies the *no-write rule* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-write}}$ is negligible for all efficient \mathcal{A} .

Remark. Our definition follows the one by Fuchsbauer et al. [FGKO17] by requiring $\text{San}(sp, c_0) \neq \perp$ in the winning condition for the no-write rule, which was not required in the original definition by Damgård et al. [DHO16]. Schemes can be made secure with respect to the original definition by letting the algorithm San create a fresh ciphertext for a random message when given an invalid ciphertext.

The condition $i' \in I_1$ together with $\forall i \in I_1 \forall j \in J P(i, j) = 0$ ensures that \mathcal{A} does not have a key to decrypt c_1 , which would trivially allow to distinguish. Requiring that \mathcal{A} obtains a key for i' however excludes adversaries that obtain no key at all. The original definitions [DHO16] therefore include a special role 0 with $P(0, j) = 0$ for all j . One can then assume without loss of generality that anyone obtains a key for this role. Since assuming the existence of such a role appears to be a technicality that is only needed for the no-write rule, we do not make this assumption and present new security definitions in Section 4.2 that do not rely on such a role.

3 Ciphertext-Revealing Attacks Against Existing Schemes

3.1 Generic Description of Attack

We describe a fundamental practical issue of schemes which meet the above no-read and no-write definitions and show why the guarantees expected from an ACE scheme need to be strengthened. We show that schemes fulfilling the definition can suffer from what we call a malleability attack, which effectively bypasses the given policy and allows communication that is forbidden by the policy. The attack does not abuse any peculiarities of existing models and in fact only requires that the semi-honest sanitizer shares its inputs and outputs with colluding parties, which is arguably possible when the sanitizer is outsourced. In particular, security against such a sanitizer is desirable from a practical point of view.

We first give a high-level explanation of the attack, formalize it as a second step, and show that several existing schemes are vulnerable. Assume there are three parties, Alice, Bob, and Charlie, each having a different role assigned. We denote by A , B , and C the associated roles. In our example, Alice and Charlie are always honest. Alice is allowed to communicate with Bob and Charlie. Bob is dishonest and forbidden to send messages to Charlie (and to Alice). The attack now proceeds as follows: When Alice sends her first message, Bob requests the corresponding ciphertext and the sanitized ciphertext from the semi-honest sanitizer. He then decrypts the sanitized ciphertext and thus receives the message Alice has sent. With the knowledge of this message, as we show below, he is able to create a valid ciphertext for a chosen message m' ,

which will be correctly sanitized and later decrypted by Charlie, hence allowing unrestricted communication from Bob to Charlie. Details follow.

Consider the policy defined by

$$P(i, j) := \begin{cases} 1, & i = A, \\ 0, & \text{otherwise.} \end{cases}$$

For the sake of presentation, we assume that the ACE scheme \mathcal{E} under consideration enjoys perfect correctness. Also, we assume that the setup-phase has completed and the three parties thus possess the encryption and decryption keys, ek_i and dk_i , respectively. Now, imagine that the ACE scheme admits an efficient function $\mathbf{maul}_{\mathcal{E}}$ with the following property (later we show how to implement such a function for some existing schemes): For all messages m and m' , any role i , and sanitizer parameters sp in the range of \mathbf{Setup} , and for any fixed randomness r ,

$$\mathbf{maul}_{\mathcal{E}}(\mathbf{Enc}(ek_i, m; r), sp, m, m') = \mathbf{Enc}(ek_i, m'; r). \quad (1)$$

If such a malleability function exists, the communication policy can be bypassed as follows:

1. Alice encrypts a message $c \leftarrow \mathbf{Enc}(ek_A, m)$ and the sanitizer computes $c' \leftarrow \mathbf{San}(sp, c)$ and gives c and c' to Bob.
2. Bob computes $m \leftarrow \mathbf{Dec}(dk_B, c')$ and creates a new ciphertext $\hat{c} \leftarrow \mathbf{maul}_{\mathcal{E}}(c, sp, m, m')$ and sends it to the sanitizer.
3. The ciphertext is sanitized $\hat{c}' \leftarrow \mathbf{San}(sp, \hat{c})$ and subsequently sent to Charlie. By the (perfect) correctness of the assumed ACE scheme and by our assumption on $\mathbf{maul}_{\mathcal{E}}$, \hat{c}' is a valid ciphertext (under the encryption key of Alice) and Charlie is able to decrypt $m' \leftarrow \mathbf{Dec}(dk_C, \hat{c}')$, effectively receiving Bob's message m' .

In the following sections, we show that several existing ACE schemes \mathcal{E} admit an efficient function $\mathbf{maul}_{\mathcal{E}}$. More specifically, we consider the ‘‘linear’’ scheme by Damg ard et al. [DHO16] based on ElGamal and the ElGamal-based scheme by Fuchsbauer et al. [FGKO17].

3.2 DHO Scheme Based on ElGamal

We briefly recall the ElGamal based ACE scheme for a single identity. The sanitizer parameters of the scheme contain the description of a finite cyclic group $G = \langle g \rangle$ and its group order q , and additionally an element $h = g^x$ for a uniform random $x \in \mathbb{Z}_q$. The encryption key for A is a random value $ek \in \mathbb{Z}_q$, and the decryption key is $-x$. The algorithm \mathbf{Enc} on input an encryption key ek_i and a message $m \in \mathcal{M}$, samples $r_1, r_2 \in \mathbb{Z}_q$ uniformly at random and outputs the ciphertext

$$c = (c_0, c_1, c_2, c_3) := (g^{r_1}, h^{r_1} g^{ek_i}, g^{r_2}, m \cdot h^{r_2}).$$

We can define the function $\mathbf{maul}_{\text{DHO}}$ as

$$\mathbf{maul}_{\text{DHO}}((c_0, c_1, c_2, c_3), sp, m, m') := (c_0, c_1, c_2, m' \cdot m^{-1} \cdot c_3).$$

Since the group order q is part of sp , this function is efficiently computable. For $c_3 = m \cdot h^{r_2}$, we thus get a new fourth component $c'_3 = m' \cdot h^{r_2}$ and equation (1) is satisfied.

The malleability for more than one identity (and in particular in our scenario described above) follows since the scheme for several identities is composed of independent instances of the basic single-identity scheme.

3.3 FGKO Scheme Based on ElGamal

Description of the scheme. In that scheme, the sanitizer parameters consist of the description of a finite cyclic group $G = \langle g \rangle$ including the group order q and a generator g , a verification key vk^{Sig} of a signature scheme Sig , and a common-reference string crs^{NIZK} of a NIZK proof system NIZK for the language $L := \{x \mid \exists w (x, w) \in R\}$, where R is defined as follows: for $x = (vk^{\text{Sig}}, c_0, c_1, c_2, c_3)$ and a witness $w = (g^x, \sigma^{\text{Sig}}, m, r, s)$, $R(x, w) = 1$ if and only if

$$\text{Sig.Ver}(vk^{\text{Sig}}, g^x, \sigma^{\text{Sig}}) = 1 \wedge (c_0, c_1, c_2, c_3) = (g^r, g^{x \cdot r}, g^s, m \cdot g^{x \cdot s}).$$

The encryption and decryption keys are given by $ek := (g^x, \sigma^{\text{Sig}})$, $dk := x$ for a uniformly chosen $x \leftarrow \mathbb{Z}_q$, where σ^{Sig} is a signature on g^x . To encrypt a message m , first choose $r \leftarrow \mathbb{Z}_q^*$ and $s \leftarrow \mathbb{Z}_q$ uniformly at random and compute $(c_0, c_1, c_2, c_3) := (g^r, g^{x \cdot r}, g^s, m \cdot g^{x \cdot s})$. Then run $\pi^{\text{NIZK}} \leftarrow \text{NIZK.Prove}(crs^{\text{NIZK}}, (vk^{\text{Sig}}, c_0, c_1, c_2, c_3), (g^x, \sigma^{\text{Sig}}, m, r, s))$ and output the ciphertext $c := (c_0, c_1, c_2, c_3, \pi)$.

Potential malleability. We define the function $\text{maul}_{\text{FGKO}}$ as

$$\text{maul}_{\text{FGKO}}((c_0, c_1, c_2, c_3, \pi), sp, m, m') := (c_0, c_1, c_2, m' \cdot m^{-1} \cdot c_3, \pi).$$

This function satisfies [equation \(1\)](#) if, for example, the non-interactive zero-knowledge proof is independent of the last component c_3 . We show that such a NIZK proof system exists without violating the properties assumed by Fuchsbauer et al. [\[FGKO17\]](#). To this end, let NIZK' be a NIZK proof system for the language $L' := \{x \mid \exists w (x, w) \in R'\}$, where the relation R' is defined as follows: for $x = (vk^{\text{Sig}}, c_0, c_1, c_2)$ and $w = (g^x, \sigma^{\text{Sig}}, r, s)$, $(x, w) \in R'$ if and only if

$$\text{Sig.Ver}(vk^{\text{Sig}}, g^x, \sigma^{\text{Sig}}) = 1 \wedge (c_0, c_1, c_2) = (g^r, g^{x \cdot r}, g^s).$$

Given NIZK' , we construct a NIZK proof system NIZK for the original language L as follows:

$$\text{NIZK.Gen}(1^\kappa) := \text{NIZK}'.\text{Gen}(1^\kappa),$$

$$\text{NIZK.Prove}(crs^{\text{NIZK}}, (vk^{\text{Sig}}, c_0, c_1, c_2, c_3), (g^x, \sigma^{\text{Sig}}, m, r, s)) :=$$

$$\text{NIZK}'.\text{Prove}(crs^{\text{NIZK}}, (vk^{\text{Sig}}, c_0, c_1, c_2), (g^x, \sigma^{\text{Sig}}, r, s)),$$

$$\text{NIZK.Ver}(crs^{\text{NIZK}}, (vk^{\text{Sig}}, c_0, c_1, c_2, c_3), \pi^{\text{NIZK}}) := \text{NIZK}'.\text{Ver}(crs^{\text{NIZK}}, (vk^{\text{Sig}}, c_0, c_1, c_2), \pi^{\text{NIZK}}).$$

Correctness and zero-knowledge of NIZK follow straightforwardly from the underlying scheme NIZK' . For knowledge-extraction, assume that NIZK' is capable of extracting a valid witness $(g^x, \sigma^{\text{Sig}}, r, s)$ given a valid proof for the statement $(vk^{\text{Sig}}, c_0, c_1, c_2)$. Given a statement $(vk^{\text{Sig}}, c_0, c_1, c_2, c_3)$ in the original language L , we can obtain a valid message encoded in c_3 by computing $m := c_3 \cdot (g^{x \cdot s})^{-1}$, and thus also a witness $(g^x, \sigma^{\text{Sig}}, m, r, s)$ for the given statement. Finally, for soundness, note that if $(vk^{\text{Sig}}, c_0, c_1, c_2) \in L'$, this implies that any group element $c_3 \in G$ is a valid last component, i.e., $(vk^{\text{Sig}}, c_0, c_1, c_2, c_3) \in L$ for any $c_3 \in G$, since there exists the message $m := c_3 \cdot (g^{x \cdot s})^{-1}$, and thus a valid witness $w = (g^x, \sigma^{\text{Sig}}, m, r, s)$.

For the constructed scheme NIZK and the function $\text{maul}_{\text{FGKO}}$, [equation \(1\)](#) clearly holds. Hence, the FGKO scheme can be instantiated such that the malleability attack works. It could potentially be excluded by requiring stronger properties from the NIZK scheme.

4 A Stronger Notion of ACE

In this section, we introduce our new security definitions, which exclude the issues we have discovered.

4.1 ACE with Modification Detection

To be resilient against the ciphertext-revealing attacks described in [Section 3](#), the sanitizer should ideally only sanitize fresh encryptions and block ciphertexts that are either replays or obtained by modifying previous ciphertexts. Therefore, we introduce an additional algorithm for detecting modified ciphertexts. If the sanitizer receives a ciphertext that is detected to be a modification of a previously received one, this ciphertext is blocked. Since such ciphertexts will not be stored in the repository and consequently not be decrypted, we define chosen-ciphertext security with respect to a decryption oracle that does not return a decryption if the received ciphertext is detected to be a modification of the challenge ciphertext. Our definitions can therefore be seen as a variant of publicly-detectable replayable-CCA security, which was introduced by Canetti et al. [CKN03] for public key encryption. Before defining the security, we define the syntax of ACE schemes with this additional algorithm.

Definition 4.1. An *access control encryption with modification detection scheme* is an ACE scheme \mathcal{E} together with a PPT algorithm DMod that on input sanitizer parameters sp and two ciphertexts $c, \tilde{c} \in \mathcal{C}$, outputs a bit b (where $b = 1$ means that \tilde{c} was obtained via modifying c).

Except for [Section 4.3](#), where we show that our new definitions imply the existing ones, we will from now on only consider ACE schemes with modification detection and thus often refer to them simply as ACE schemes.

The algorithm DMod should output 1 if \tilde{c} is an adversarial modification of c , and 0 otherwise. We have the following intuitive requirements on DMod :

1. All ciphertexts \tilde{c} an adversary can produce given ciphertexts c_1, \dots, c_l and no encryption key, are either invalid (i.e., sanitize to \perp) or we have $\text{DMod}(sp, c_i, \tilde{c}) = 1$ for some $i \in \{1, \dots, n\}$.
2. Given encryption and decryption keys, an adversary is unable to produce a ciphertext c such that a ciphertext produced by Enc for a message of the adversary's choice is detected to be a modification of c . In particular, independent encryptions of messages collide only with negligible probability.

The first requirement is captured by role-respecting security as defined in [Definition 4.5](#), the second one by non-detection of fresh encryptions defined in [Definition 4.4](#).

Remark. Canetti et al. (translated to our setting) also require that if $\text{DMod}(sp, c, \tilde{c}) = 1$, then c and \tilde{c} decrypt to the same message [CKN03]. For our purpose, this is not needed. This means that we do not want to detect replays in the sense that the same message is replayed, but more generally, whether the given ciphertext was obtain via some modification of another ciphertext.

4.2 New Security Definitions

We formalize chosen-ciphertext attacks by giving the adversary access to an oracle \mathcal{O}_{SD} that first sanitizes a given ciphertext and then decrypts the result. One could also consider *chosen-sanitized-ciphertext attacks* by providing the adversary access to an oracle \mathcal{O}_D that only decrypts.

This is potentially stronger since the adversary can emulate the oracle \mathcal{O}_{SD} by first sanitizing the ciphertexts and then giving the result to \mathcal{O}_D , but given \mathcal{O}_{SD} , it is not necessarily possible to emulate \mathcal{O}_D . Since in the application, users can only send ciphertexts to the sanitizer but not directly write ciphertexts to the repository such that they are decrypted without being sanitized, the weaker notion is sufficient.

In principle, the adversary has in all definitions access to \mathcal{O}_{SD} , as well as to an encryption oracle and a key-generation oracle. To simplify the definitions, we omit the encryption or decryption oracles if the winning condition places no restriction on the encryption or decryption keys obtained from the key-generation oracle, respectively.

Privacy and anonymity. We now define (payload) privacy and sender-anonymity. The former guarantees that encryptions of different messages under the same encryption key cannot be distinguished as long as the adversary has no decryption key that allows to decrypt. We also require this for messages of different length, i.e., schemes satisfying our definition do not leak the length of the encrypted message, which means that the message space has to be bounded. Anonymity guarantees that encryptions of the same message under different keys cannot be distinguished. We distinguish a weak and a strong variant of anonymity, where the weak one provides no guarantees if the adversary can decrypt the ciphertext, and the strong one guarantees that even if the adversary has decryption keys, nothing is leaked about the sender role beyond which of the adversary's decryption keys can be used to decrypt.

Definition 4.2. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$, be an ACE with modification detection scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$ in Figure 2 and let J be the set of all j such that \mathcal{A}_1 or \mathcal{A}_2 issued the query (j, rec) to the oracle \mathcal{O}_G . We define the *privacy under chosen-ciphertext attacks advantage* and the *sender-anonymity under chosen-ciphertext attacks advantages* of \mathcal{A} as

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-PRV-CCA}} &:= 2 \cdot \Pr[b' = b \wedge i_0 = i_1 \wedge \forall j \in J P(i_0, j) = 0] - 1, \\ \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-wANON-CCA}} &:= 2 \cdot \Pr[b' = b \wedge m_0 = m_1 \wedge \forall j \in J P(i_0, j) = P(i_1, j) = 0] - 1, \\ \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sANON-CCA}} &:= 2 \cdot \Pr[b' = b \wedge m_0 = m_1 \wedge \forall j \in J P(i_0, j) = P(i_1, j)] - 1, \end{aligned}$$

respectively, where all probabilities are in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$. The scheme \mathcal{E} is called *private under chosen-ciphertext attacks (PRV-CCA secure)*, *weakly sender-anonymous under chosen-ciphertext attacks (wANON-CCA secure)*, and *strongly sender-anonymous under chosen-ciphertext attacks (sANON-CCA secure)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-PRV-CCA}}$, $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-wANON-CCA}}$, and $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sANON-CCA}}$ are negligible for all efficient \mathcal{A} , respectively.

Remark. Weak anonymity corresponds to the anonymity notion considered by Fuchsbauer et al. [FGKO17] and strong anonymity to the one considered by Damgård et al. [DHO16]. We state both definitions because weak anonymity is easier to achieve but strong anonymity might be required by some applications. If anonymity is only required against the sanitizer or if all messages are anyway signed by the sender, weak anonymity is sufficient. Strong anonymity is required in settings where senders also want to retain as much anonymity as possible against legitimate receivers.

Sanitization security. We next define sanitization security, which excludes that dishonest parties can communicate via the ciphertexts. We formalize this by requiring that the output

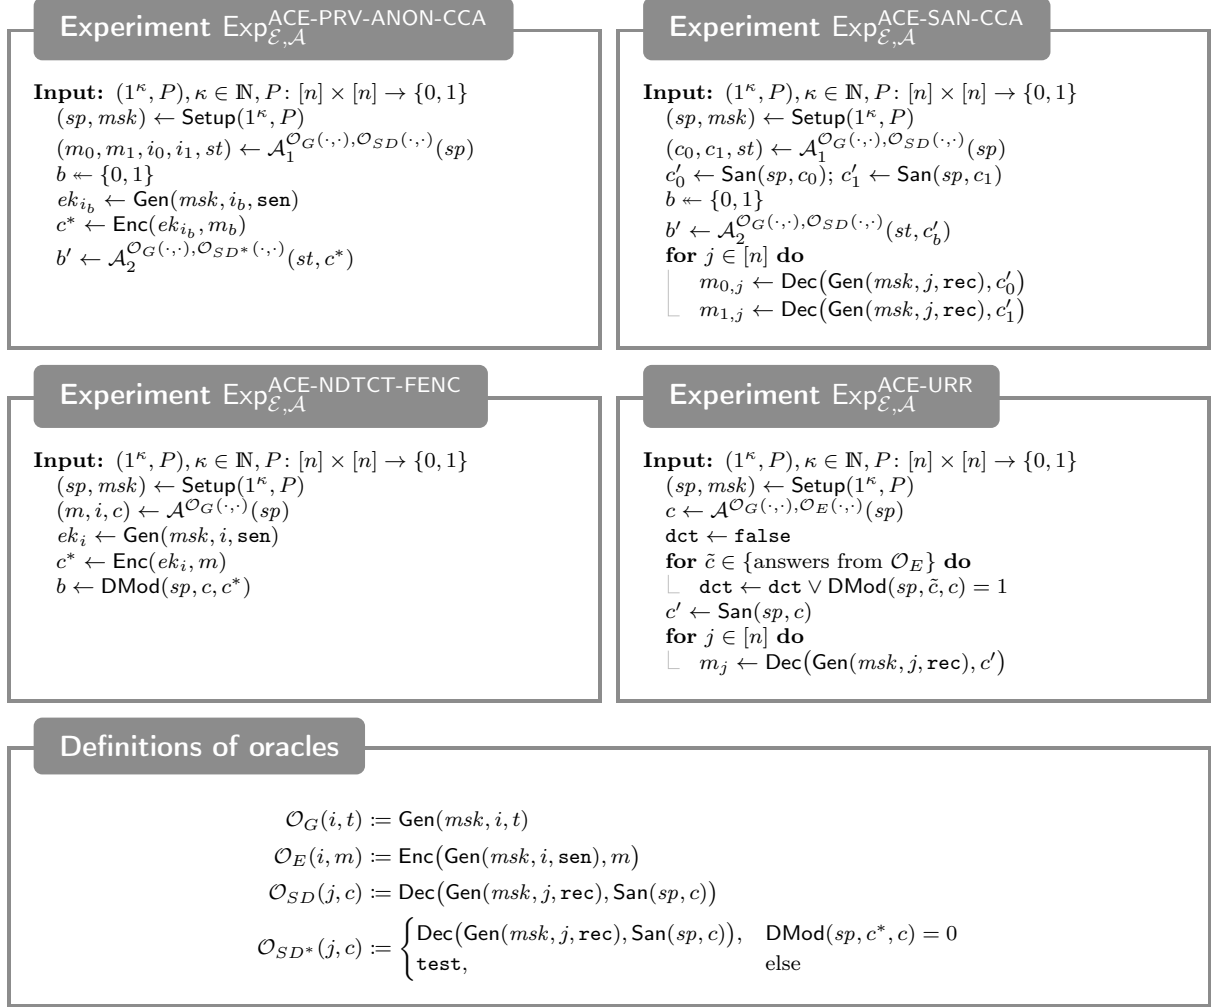


Figure 2: Security experiments for an ACE with modification detection scheme \mathcal{E} and an adversary \mathcal{A} , where $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the first two experiments.

of the sanitizer for two different ciphertexts cannot be distinguished, as long as both sanitized ciphertexts are not \perp and the adversary has no decryption key that decrypts one of the ciphertexts. This provides no security guarantees if the adversary can decrypt the ciphertexts, which does not seem to be an issue since in this case, the parties can anyway directly communicate via the messages. However, we additionally consider a stronger variant, where the adversary is allowed to possess a decryption key that decrypts the ciphertexts, as long as they both decrypt to the same message. This stronger variant excludes subliminal channels, i.e., even if the involved parties are allowed to communicate by the policy, they cannot exchange information via ciphertexts beyond the encrypted message.

Since the adversary provides the two ciphertexts that are sanitized, we do not know to which roles they correspond; they could even be particularly crafted without belonging to an existing role. Hence, we cannot state the requirement (in the weak variant) that the adversary should not be able to decrypt by only considering the policy and the obtained decryption keys, as in the

no-write rule in [Definition 2.4](#). Instead, we require that the decryption algorithm returns \perp for all decryption keys the adversary possesses. For this to provide the intended security, we need that the decrypt algorithm returns \perp whenever the receiver role corresponding to the used key is not supposed to read the message. This is guaranteed by role-respecting security which is defined later.

Definition 4.3. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-SAN-CCA}}$ in [Figure 2](#) and let J be the set of all j such that \mathcal{A}_1 or \mathcal{A}_2 issued the query (j, rec) to the oracle \mathcal{O}_G . We define the *sanitization under chosen-ciphertext attacks advantage* and the *strong sanitization under chosen-ciphertext attacks advantage* of \mathcal{A} as

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-SAN-CCA}} &:= 2 \cdot \Pr[b' = b \wedge c'_0 \neq \perp \neq c'_1 \wedge \forall j \in J m_{0,j} = m_{1,j} = \perp] - 1, \\ \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sSAN-CCA}} &:= 2 \cdot \Pr[b' = b \wedge c'_0 \neq \perp \neq c'_1 \wedge \forall j \in J m_{0,j} = m_{1,j}] - 1, \end{aligned}$$

respectively, where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-SAN-CCA}}$. The scheme \mathcal{E} is called *sanitization under chosen-ciphertext attacks secure (SAN-CCA secure)* and *strongly sanitization under chosen-ciphertext attacks secure (sSAN-CCA secure)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-SAN-CCA}}$ and $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-sSAN-CCA}}$ are negligible for all efficient \mathcal{A} , respectively.

Non-detection of fresh encryptions. In the intended way of using a scheme satisfying our notions, the sanitizer only adds sanitized ciphertexts to the repository if the given ciphertext is not detected to be a modification of a previously received ciphertext. This means that if an adversary can find a ciphertext c such that another ciphertext c^* that is later honestly generated is detected as a modification of c , the delivery of the message at that later point can be prevented by sending the ciphertext c to the sanitizer earlier. We exclude this by the following definition, which can be seen as an extended correctness requirement.

Definition 4.4. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let \mathcal{A} be a probabilistic algorithm. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}}$ in [Figure 2](#). We define the *non-detection of fresh encryptions advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}} := \Pr[b = 1],$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}}$. The scheme \mathcal{E} is said to have *non-detecting fresh encryptions (NDTCT-FENC)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}}$ is negligible for all efficient \mathcal{A} .

Role-respecting and uniform-decryption security. We finally define role-respecting and uniform-decryption security. The former means that an adversary cannot produce a ciphertext for which the pattern of roles that can decrypt does not correspond to a role for which the adversary has an encryption key. For example, if the adversary has only an encryption key for the role i such that roles j_0 and j_1 are the only roles j with $P(i, j) = 1$, all ciphertexts produced by the adversary are either invalid (i.e., sanitized to \perp or detected as a modification) or decrypt to a message different from \perp precisely under the decryption keys for j_0 and j_1 . On the one hand, this means that receivers who are not allowed to receive the message get \perp and hence

know that the message is not for them.² On the other hand, it also guarantees that the adversary cannot prevent receivers with role j_1 from receiving a message that is sent to receivers with role j_0 . Furthermore, uniform decryption guarantees for all ciphertexts c output by an adversary that if c decrypts to a message different from \perp for different decryption keys, it always decrypts to the same message. In the example above, this means that j_0 and j_1 not only both receive *some* message, but they both receive *the same* one.

Definition 4.5. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let \mathcal{A} be a probabilistic algorithm. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-URR}}$ in Figure 2 and let I and J be the sets of all i and j such that \mathcal{A} issued the query (i, sen) and (j, rec) to the oracle \mathcal{O}_G , respectively. We define the *role-respecting advantage* and the *uniform-decryption advantage* of \mathcal{A} as

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-RR}} &:= \Pr[c' \neq \perp \wedge \text{dct} = \text{false} \wedge \neg(\exists i \in I \forall j \in J (m_j \neq \perp \leftrightarrow P(i, j) = 1))], \\ \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-UDEC}} &:= \Pr[\exists j, j' \in J m_j \neq \perp \neq m_{j'} \wedge m_j \neq m_{j'}], \end{aligned}$$

respectively, where the probabilities are over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-URR}}$. The scheme \mathcal{E} is *role-respecting (RR secure)* and *uniform-decryption (UDEC) secure* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-RR}}$ and $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-UDEC}}$ are negligible for all efficient \mathcal{A} , respectively.

Remark. Note that in Definition 4.5, we only check the decryptions for receiver roles for which \mathcal{A} has requested the corresponding decryption key. This means that an adversary in addition to producing a ciphertext that causes an inconsistency, also has to find a receiver role for which this inconsistency manifests. If the total number of roles n is small (say polynomial in the security parameter), \mathcal{A} can simply query \mathcal{O}_G on all receiver keys, but for large n this condition is nontrivial. For example, we consider a scheme secure if an adversary can efficiently produce a ciphertext such that there is a receiver role that can decrypt it even though the policy does not allow it, as long as this receiver role is hard to find. The rationale is that in this case, the inconsistency cannot be exploited and will only be observed with negligible probability in an execution of the protocol.

4.3 Relation to the Original Security Notions

In this section, we discuss how our notions relate to the original security definitions (see Section 2.4). First note that we assume the scheme has an additional algorithm DMod . As explained in Section 4.1, the intended usage of such a scheme is that the sanitizer discards ciphertexts that are detected to be a modification of a previous ciphertext. This means that if dishonest parties want to communicate even though disallowed by the policy (i.e., they want to break the no-write rule), the sender must produce a ciphertext that is not detected as a modification of a previous ciphertext. With this in mind, it is natural to adjust the no-write rule such that an adversary only wins if the ciphertext he outputs is not detected to be a modification of a ciphertext generated by the oracle \mathcal{O}_{ES} (before sanitizing it).

²Detectability (Definition 2.2) provides this guarantee for honest encryptions, role-respecting security extends this to maliciously generated ciphertexts. Note, however, that detectability is not implied by role-respecting security: If an adversary has encryption keys for two roles i and i' , role-respecting security does not exclude that encrypting some message (depending on i') with the key for role i can be decrypted with keys for roles that are allowed to receive from i' .

Definition 4.6. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. The experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$ is identical to $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-write}}$ in Figure 1 except that after \mathcal{A}_1 returns (c_0, i', st) , it is checked whether the oracle \mathcal{O}_{ES} has generated some \tilde{c} and returned its sanitization such that $\text{DMod}(sp, \tilde{c}, c_0) = 1$. If this is the case, set $\text{dct} \leftarrow \text{true}$, else $\text{dct} \leftarrow \text{false}$. Let I_1 be the set of all i such that \mathcal{A}_1 issued the query (i, sen) to \mathcal{O}_G , and let J be the set of all j such that \mathcal{A}_1 or \mathcal{A}_2 issued the query (j, rec) to \mathcal{O}_G . We define the *no-write with modification detection advantage* of \mathcal{A} as

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}} := & 2 \cdot \Pr[b' = b \wedge \text{dct} = \text{false} \wedge i' \in I_1 \\ & \wedge \forall i \in I_1 \forall j \in J P(i, j) = 0 \wedge \text{San}(sp, c_0) \neq \perp] - 1, \end{aligned}$$

where the probability is over the randomness of all algorithms in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$. The scheme \mathcal{E} satisfies the *no-write with modification detection rule* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$ is negligible for all efficient \mathcal{A} .

We show that our new security definitions from Section 4.2 imply the no-read rule and the no-write with modification detection rule. We have to assume that the policy P allows for all i that one can efficiently find some j with $P(i, j) = 1$. This seems to be the case for all practically relevant policies, though. The results are summarized in the following theorem.

Theorem 4.7. *Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let $\mathcal{E}' = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be the corresponding ACE scheme. If \mathcal{E} is correct and PRV-CCA, sANON-CCA, SAN-CCA, and RR secure, then it satisfies the the no-write with modification detection rule for policies P such that for all i , one can efficiently find some j with $P(i, j) = 1$, and \mathcal{E}' satisfies the no-read rule. More precisely, for all adversaries \mathcal{A} , \mathcal{A}' , and \mathcal{A}'' , there exist adversaries \mathcal{A}_{PRV} and $\mathcal{A}_{\text{WANON}}$ (both roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$), an adversary $\mathcal{A}'_{\text{SANON}}$ (roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}', \mathcal{A}'}$), and adversaries $\mathcal{A}''_{\text{SAN}}$, $\mathcal{A}''_{\text{RR}}$, and $\mathcal{A}''_{\text{CORR}}$ (all roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}, \mathcal{A}''}^{\text{ACE-MD-no-write}}$) such that*

$$\begin{aligned} \text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,priv}} &\leq \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{PRV}}}^{\text{ACE-PRV-CCA}} + \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}}, \\ \text{Adv}_{\mathcal{E}', \mathcal{A}'}^{\text{ACE-no-read,anon}} &= \text{Adv}_{\mathcal{E}, \mathcal{A}'_{\text{SANON}}}^{\text{ACE-sANON-CCA}}, \\ \text{Adv}_{\mathcal{E}, \mathcal{A}''}^{\text{ACE-MD-no-write}} &\leq \text{Adv}_{\mathcal{E}, \mathcal{A}''_{\text{SAN}}}^{\text{ACE-SAN-CCA}} + 4 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}''_{\text{RR}}}^{\text{ACE-RR}} + 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}''_{\text{CORR}}}^{\text{ACE-CORR}}. \end{aligned}$$

We here sketch the proof idea, a detailed proof of the theorem is provided in Appendix B. To prove the claim about the payload-privacy no-read rule, consider the hybrid experiment H that is identical to $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$ except that after \mathcal{A}_1 returns (m_0, m_1, i_0, i_1, st) , i_1 is replaced by i_0 . If \mathcal{A} wins the no-read privacy game, $P(i_0, j) = P(i_1, j) = 0$ for all j for which \mathcal{A} obtained a decryption key. Hence, in this case $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$ and H are indistinguishable by weak sender-anonymity. If \mathcal{A} wins in H , one can construct an adversary against PRV-CCA security by running \mathcal{A} , returning (m_0, m_1, i_0, i_0, st) when \mathcal{A}_1 returns (m_0, m_1, i_0, i_1, st) , and returning the same guess as \mathcal{A}_2 . Note that \mathcal{A} has access to an encryption oracle \mathcal{O}_E in $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$, which is not available in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$. However, since the winning conditions do not restrict the encryption keys obtained from \mathcal{O}_G , the oracle \mathcal{O}_E can be emulated by obtaining the encryption key and then encrypting the message.

Relating the sender-anonymity no-read rule to sANON-CCA security is a straightforward reduction.

To prove the claim about the no-write rule, assume \mathcal{A}'' wins the corresponding game. If \mathcal{A}'' does not obtain a decryption key that decrypts c_0 or c_1 to a message different from \perp , this adversary can be used to break SAN-CCA security as follows: when \mathcal{A}''_1 returns (c_0, i', st) , output c_0 and the encryption of a uniformly chosen message for sender role i' as c_1 ; finally output the same guess b' as \mathcal{A}''_2 . Correctness ensures that c_1 does not sanitize to \perp ,³ so the winning condition of the SAN-CCA game is satisfied. If \mathcal{A}'' does obtain a decryption key that decrypts c_0 or c_1 to a message different from \perp , one can construct an adversary against role-respecting security.

Relation to original no-write rule. Perhaps surprisingly, one can also show that our new notions imply the original no-write rule if DMod is symmetric in the sense that $\Pr[\text{DMod}(sp, c_0, c_1) = 1] = \Pr[\text{DMod}(sp, c_1, c_0) = 1]$ (which is the case for all schemes considered in this paper). The proof idea is to construct adversaries against correctness, and sanitization and role-respecting security as above. Now, the role-respecting game is not won if the adversary \mathcal{A} returns a ciphertext c_0 that is detected to be a modification of a ciphertext generated by \mathcal{O}_{ES} . We show that in this case, we can break sSAN-CCA security. Note that \mathcal{O}_{ES} only gives \mathcal{A} the sanitized ciphertexts. The proof idea is as follows. \mathcal{A}_1 makes several queries to \mathcal{O}_{ES} . For a uniformly chosen one, encrypt the message twice, give the resulting ciphertexts \tilde{c}_0, \tilde{c}_1 to the sSAN-CCA challenger, and give the obtained sanitized ciphertext \tilde{c}'_b to \mathcal{A}_1 . For all other queries, encrypt and sanitize the message normally. When \mathcal{A}_1 returns a ciphertext c_0 , check whether c_0 is detected to be a modification of \tilde{c}_0 or \tilde{c}_1 . Since the ciphertext \tilde{c}_{1-b} is information-theoretically hidden from \mathcal{A} , it can be considered to be a fresh encryption. By our assumption, the probability that c_0 is detected to be a modification of \tilde{c}_{1-b} is equal to the probability that \tilde{c}_{1-b} is detected to be a modification of c_0 , which contradicts non-detection of fresh encryptions. Hence, by checking which of the two ciphertexts is detected, one can guess b and thus break sSAN-CCA security. Note that \mathcal{A} is allowed to obtain a decryption key that decrypts \tilde{c}'_b , which is why we need *strong* sanitization security.

Theorem 4.8. *Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme such that $\Pr[\text{DMod}(sp, c_0, c_1) = 1] = \Pr[\text{DMod}(sp, c_1, c_0) = 1]$ for all sp returned by Setup and all ciphertexts $c_0, c_1 \in \mathcal{C}$. Further let $\mathcal{E}' = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be the corresponding ACE scheme. If \mathcal{E} is correct, detectable, has NDTCT-FENC, and is sSAN-CCA and RR secure, then \mathcal{E}' satisfies the no-write rule for policies P such that for all i , one can efficiently find some j with $P(i, j) = 1$. More precisely, for all adversaries \mathcal{A} that make at most q_{ES} queries to the oracle \mathcal{O}_{ES} and at most q_{dk} queries of the form (\cdot, rec) to \mathcal{O}_G , there exist adversaries $\mathcal{A}_{\text{SAN}}, \mathcal{A}_{\text{RR}}, \mathcal{A}_{\text{sSAN}}, \mathcal{A}_{\text{NDTCT}}, \mathcal{A}_{\text{CORR}}$, and $\mathcal{A}_{\text{dtct}}$ (all roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-write}}$) such that*

$$\begin{aligned} \text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-write}} &\leq \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{SAN}}}^{\text{ACE-SAN-CCA}} + 4 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{RR}}}^{\text{ACE-RR}} + 2q_{ES} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-sSAN-CCA}} \\ &\quad + 4q_{ES} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{NDTCT}}}^{\text{ACE-NDTCT-FENC}} + (8q_{ES}q_{dk} + 2) \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{CORR}}}^{\text{ACE-CORR}} + 8q_{ES}q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{dtct}}}^{\text{ACE-DTCT}}. \end{aligned}$$

See [Appendix B](#) for a detailed proof.

³This is the only place where we need that one can efficiently find j with $P(i', j) = 1$ since the adversary in the correctness game has to provide such j .

5 Enhanced Sanitizable Public-Key Encryption

5.1 Definitions

As a stepping stone toward ACE schemes satisfying our new security definitions, we introduce *enhanced sanitizable public-key encryption*. Sanitizable public-key encryption has been considered by Damgård et al. [DHO16] and Fuchsbauer et al. [FGKO17] as a relaxation of universal re-encryption [GJS04] and rerandomizable encryption [Gro04; PR07]. It allows to *sanitize* a ciphertext to obtain a *sanitized ciphertext* that cannot be linked to the original ciphertext except that it decrypts to the correct message. In contrast to rerandomizable encryption, sanitized ciphertexts can have a different syntax than ciphertexts, i.e., it is not required that a sanitized ciphertext is indistinguishable from a fresh encryption. We introduce an enhanced variant with a different syntax and stronger security guarantees.

Definition 5.1. An *enhanced sanitizable public-key encryption (sPKE) scheme* consists of the following five PPT algorithms:

Setup: The algorithm **Setup** on input a security parameter 1^κ , outputs *sanitizer parameters* sp , and a *master secret key* msk . We implicitly assume that all parameters and keys include the finite *message space* \mathcal{M} and the *ciphertext spaces* $\mathcal{C}, \mathcal{C}'$.

Key generation: The algorithm **Gen** on input a master secret key msk , outputs an *encryption key* ek and a *decryption key* dk .

Encryption: The algorithm **Enc** on input an encryption key ek and a message $m \in \mathcal{M}$, outputs a ciphertext $c \in \mathcal{C}$.

Sanitization: The algorithm **San** on input sanitizer parameters sp and a ciphertext $c \in \mathcal{C}$, outputs a *sanitized ciphertext* $c' \in \mathcal{C}' \cup \{\perp\}$.

Decryption: The algorithm **Dec** on input a decryption key dk and a sanitized ciphertext $c' \in \mathcal{C}'$, outputs a message $m \in \mathcal{M} \cup \{\perp\}$; on input dk and \perp , it outputs \perp .

For correctness, we require for all (sp, msk) in the range of **Setup**, all (ek, dk) in the range of **Gen**(msk), and all $m \in \mathcal{M}$ that

$$\text{Dec}(dk, \text{San}(sp, \text{Enc}(ek, m))) = m$$

with probability 1.

We require robustness in the sense that no ciphertext decrypts to a message different from \perp for two different decryption keys (except with negligible probability). This is similar to detectability for ACE schemes, but we allow the adversary to directly output a ciphertext, instead of a message, which is then honestly encrypted. Our notion therefore closely resembles *unrestricted strong robustness (USROB)*, introduced by Farshim et al. [FLPQ13] for public-key encryption, which also allows the adversary to choose a ciphertext and, in contrast to strong robustness by Abdalla et al. [ABN10], gives the adversary access to decryption keys.

Definition 5.2. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an sPKE scheme. For a probabilistic algorithm \mathcal{A} , we define the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-USROB}}$ that executes $(sp, msk) \leftarrow \text{Setup}(1^\kappa)$ and $(c, i_0, i_1) \leftarrow \mathcal{A}^{\mathcal{O}_G(\cdot)}(sp)$, where the oracle \mathcal{O}_G on input **getNew**, outputs a fresh key pair

$(ek, dk) \leftarrow \text{Gen}(msk)$. Let q be the number of oracle queries and let for $i \in \{1, \dots, q\}$, (ek_i, dk_i) be the i -th answer from \mathcal{O}_G . We define the (*unrestricted strong*) *robustness advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-USROB}} := \Pr[1 \leq i_0, i_1 \leq q \wedge i_0 \neq i_1 \\ \wedge \text{Dec}(dk_{i_0}, \text{San}(sp, c)) \neq \perp \neq \text{Dec}(dk_{i_1}, \text{San}(sp, c))],$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-USROB}}$ and the random coins of San and Dec (both executed independently twice). The scheme \mathcal{E} is (*unrestricted strongly*) *robust (USROB secure)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-USROB}}$ is negligible for all efficient \mathcal{A} .

We next define IND-CCA security analogously to the definition for ordinary public-key encryption. In contrast to the usual definition, we do not require the adversary to output two messages of equal length, which implies that schemes satisfying our definition do not leak the length of the encrypted message.

Definition 5.3. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an sPKE scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IND-CCA}}$ in Figure 3 and let $C_{\mathcal{A}_2}$ be the set of all ciphertexts that \mathcal{A}_2 queried to the oracle \mathcal{O}_{SD} . We define the *ciphertext indistinguishability under chosen-ciphertext attacks advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IND-CCA}} := 2 \cdot \Pr[b' = b \wedge c^* \notin C_{\mathcal{A}_2}] - 1,$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IND-CCA}}$. The scheme \mathcal{E} has *indistinguishable ciphertexts under chosen-ciphertext attacks (is IND-CCA secure)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IND-CCA}}$ is negligible for all efficient \mathcal{A} .

We also need that it is hard to predict a ciphertext generated by Enc from a message of the adversary's choice given encryption and decryption keys. Note that this is not implied by IND-CCA security since the adversary here obtains the decryption key.

Definition 5.4. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an sPKE scheme and let \mathcal{A} be a probabilistic algorithm. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-UPD-CTXT}}$ in Figure 3. We define the *ciphertext unpredictability advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-UPD-CTXT}} := \Pr[c = c^*],$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-UPD-CTXT}}$. The scheme \mathcal{E} has *unpredictable ciphertexts (is UPD-CTXT secure)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-UPD-CTXT}}$ is negligible for all efficient \mathcal{A} .

We further define anonymity or indistinguishability of keys following Bellare et al. [BBDP01].

Definition 5.5. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an sPKE scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IK-CCA}}$ in Figure 3 and let $C_{\mathcal{A}_2}$ be the set of all ciphertexts that \mathcal{A}_2 queried to the oracle \mathcal{O}_{SD_0} or \mathcal{O}_{SD_1} . We define the *indistinguishability of keys under chosen-ciphertext attacks advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IK-CCA}} := 2 \cdot \Pr[b' = b \wedge c^* \notin C_{\mathcal{A}_2}] - 1,$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IK-CCA}}$. The scheme \mathcal{E} has *indistinguishable keys under chosen-ciphertext attacks (is IK-CCA secure)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IK-CCA}}$ is negligible for all efficient \mathcal{A} .

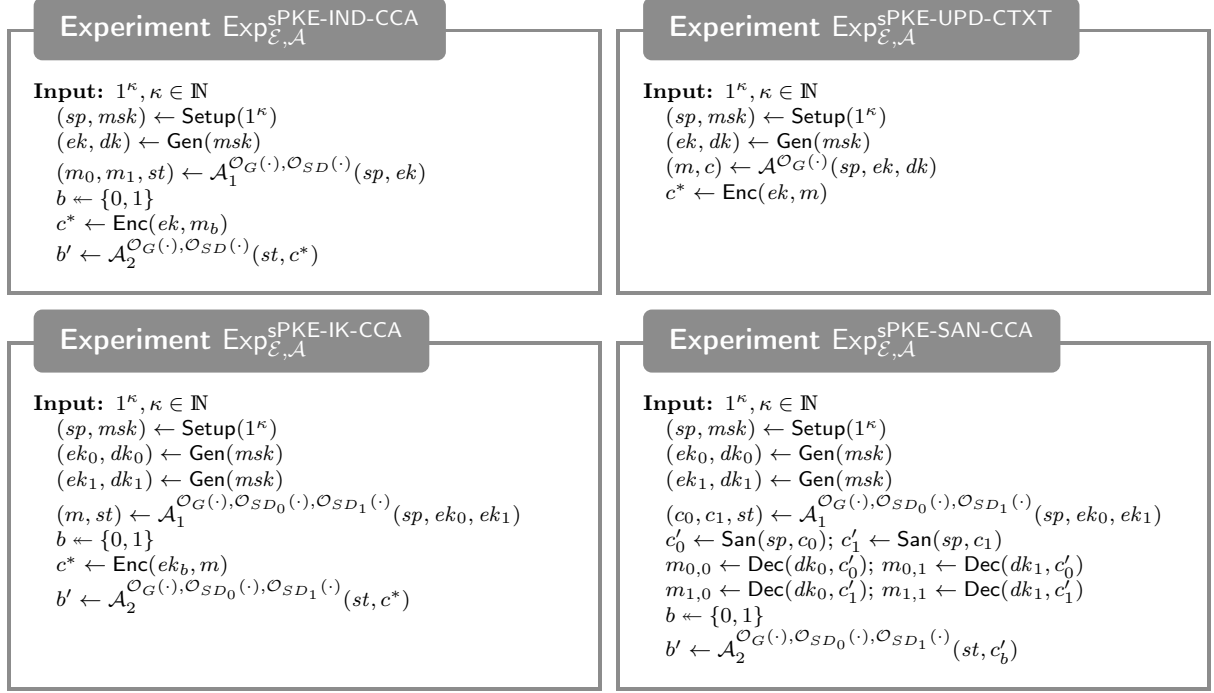


Figure 3: Security experiments for an sPKE scheme \mathcal{E} and an adversary \mathcal{A} , where $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the experiments $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IND-CCA}}$, $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IK-CCA}}$, and $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-SAN-CCA}}$. The oracle \mathcal{O}_{SD} is defined as $\mathcal{O}_{SD}(c) = \text{Dec}(dk, \text{San}(sp, c))$ and the oracle \mathcal{O}_{SD_j} as $\mathcal{O}_{SD_j}(c) = \text{Dec}(dk_j, \text{San}(sp, c))$. Moreover, the oracle \mathcal{O}_G on input `getNew`, outputs a fresh key pair $(ek, dk) \leftarrow \text{Gen}(msk)$.

Sanitization security formalizes that given certain public keys and a sanitized ciphertext, it is hard to tell which of two adversarially chosen ciphertexts was actually sanitized. To exclude trivial attacks, we require that both ciphertexts are encryptions relative to the two challenge public keys ek_0 and ek_1 . Otherwise, the adversary could use the oracle \mathcal{O}_G to obtain a fresh key-pair (ek, dk) and encrypt two different messages under ek . It could then decrypt the challenge ciphertext using dk and win the game.

Definition 5.6. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an sPKE scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-SAN-CCA}}$ in Figure 3. We define the *sanitization under chosen-ciphertext attacks advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-SAN-CCA}} := 2 \cdot \Pr[b' = b \wedge \exists j, j' \in \{0, 1\} m_{0,j} \neq \perp \neq m_{1,j'}] - 1,$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-IK-CCA}}$. The scheme \mathcal{E} is *sanitization under chosen-ciphertext attacks (SAN-CCA) secure* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-SAN-CCA}}$ is negligible for all efficient \mathcal{A} .

We finally define the probability that two independent executions of the key-generation algorithm produce the same encryption key. This probability has to be small for all IND-CCA-secure schemes because an attacker can otherwise obtain a new key pair from \mathcal{O}_G and if the obtained encryption key matches the one with which the challenge ciphertext is generated, the attacker can decrypt and win the IND-CCA game. We anyway explicitly define this probability to simplify our reductions later.

Definition 5.7. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be an sPKE scheme. We define the *encryption-key collision probability* $\text{Col}_{\mathcal{E}}^{\text{ek}}$ as

$$\text{Col}_{\mathcal{E}}^{\text{ek}} := \Pr^{(sp, msk) \leftarrow \text{Setup}(1^\kappa); (ek_0, dk_0) \leftarrow \text{Gen}(msk); (ek_1, dk_1) \leftarrow \text{Gen}(msk)} [ek_0 = ek_1].$$

5.2 Constructing an sPKE Scheme

We next construct an sPKE scheme satisfying our security definitions. Our construction resembles the weakly sanitizable PKE scheme by Fuchsbauer et al. [FGKO17]. We use a variant of ElGamal encryption and obtain security against chosen-ciphertext attacks using the technique of Naor and Yung [NY90], i.e., encrypting the message under two independent keys and proving in zero-knowledge that the ciphertexts are encryptions of the same message, which was shown to achieve full IND-CCA security if the zero-knowledge proof is one-time simulation sound by Sahai [Sah99].

Let PKE be a (IND-CPA secure) public-key encryption scheme, let Sig be a (EUF-CMA-secure) signature scheme, and let NIZK be a (one-time simulation sound) NIZK proof system for the language $L := \{x \mid \exists w (x, w) \in R\}$, where the relation R is defined as follows: for $x = (g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma)$ and $w = (m, g^a, g^b, r_1, s_1, r_2, s_2, \sigma, r)$, we have $(x, w) \in R$ if and only if

$$\begin{aligned} c_1 &= (g^{r_1}, g^{a \cdot r_1}, g^{s_1}, g^{a \cdot s_1} \cdot m) \wedge c_2 = (g^{r_2}, g^{b \cdot r_2}, g^{s_2}, g^{b \cdot s_2} \cdot m) \\ &\wedge \text{Sig.Ver}(vk^{\text{Sig}}, (g^a, g^b), \sigma) = 1 \wedge c_\sigma = \text{PKE.Enc}(ek^{\text{PKE}}, (g^a, g^b, \sigma); r). \end{aligned}$$

We define an sPKE scheme as follows:

Setup: The setup algorithm sPKE.Setup first generates

$$\begin{aligned} (ek^{\text{PKE}}, dk^{\text{PKE}}) &\leftarrow \text{PKE.Gen}(1^\kappa), \\ (vk^{\text{Sig}}, sk^{\text{Sig}}) &\leftarrow \text{Sig.Gen}(1^\kappa), \\ crs &\leftarrow \text{NIZK.Gen}(1^\kappa). \end{aligned}$$

Let $G = \langle g \rangle$ be a cyclic group with prime order p generated by g , with $p \geq 2^\kappa$, and let $\mathcal{M} \subseteq G$ such that $|\mathcal{M}|/p \leq 2^{-\kappa}$. The sanitizer parameters sp^{sPKE} contain $ek^{\text{PKE}}, vk^{\text{Sig}}, crs$, and a description of G , including g and p . The master secret key msk^{sPKE} consists of $ek^{\text{PKE}}, vk^{\text{Sig}}, sk^{\text{Sig}}, crs$, and a description of G , including g and p .

Key generation: The algorithm sPKE.Gen on input msk^{sPKE} , samples two elements $dk_1, dk_2 \leftarrow \mathbb{Z}_p^*$ and computes $ek_1 \leftarrow g^{dk_1}, ek_2 \leftarrow g^{dk_2}$, as well as $\sigma \leftarrow \text{Sig.Sign}(sk^{\text{Sig}}, (ek_1, ek_2))$. Finally, it outputs $ek^{\text{sPKE}} := (g, p, crs, ek^{\text{PKE}}, vk^{\text{Sig}}, ek_1, ek_2, \sigma)$ and $dk^{\text{sPKE}} := (dk_1, dk_2)$.

Encryption: The algorithm sPKE.Enc on input an encryption key $ek^{\text{sPKE}} = (g, p, crs, ek^{\text{PKE}}, vk^{\text{Sig}}, ek_1, ek_2, \sigma)$ and a message $m \in \mathcal{M}$, samples randomness r , chooses $r_1, s_1, r_2, s_2 \leftarrow \mathbb{Z}_p^*$ uniformly at random, and computes

$$\begin{aligned} c_1 &\leftarrow (g^{r_1}, ek_1^{r_1}, g^{s_1}, ek_1^{s_1} \cdot m), \\ c_2 &\leftarrow (g^{r_2}, ek_2^{r_2}, g^{s_2}, ek_2^{s_2} \cdot m), \\ c_\sigma &\leftarrow \text{PKE.Enc}(ek^{\text{PKE}}, (ek_1, ek_2, \sigma); r). \end{aligned}$$

It then generates $\pi \leftarrow \text{NIZK.Prove}(crs, x := (g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma), w := (m, ek_1, ek_2, r_1, s_1, r_2, s_2, \sigma, r))$. It finally outputs the ciphertext $c := (c_1, c_2, c_\sigma, \pi)$.

Sanitization: The algorithm sPKE.San on input sanitizer parameters sp^{sPKE} and a ciphertext $c = (c_1, c_2, c_\sigma, \pi)$, first verifies the NIZK proof by evaluating $\text{NIZK.Ver}(crs, x := (g, ek^{\text{sPKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma), \pi)$. It then parses $(c_{1,1}, c_{1,2}, c_{1,3}, c_{1,4}) \leftarrow c_1$. If the verification succeeds and $c_{1,1} \neq 1 \neq c_{1,2}$, then it chooses a random $t \leftarrow \mathbb{Z}_p^*$ and outputs the sanitized ciphertext

$$c' := ((c_{1,1})^t \cdot c_{1,3}, (c_{1,2})^t \cdot c_{1,4}).$$

If the verification fails or if $c_{1,1} = 1$ or $c_{1,2} = 1$, it outputs \perp .

Decryption: The algorithm sPKE.Dec on input a decryption key $dk^{\text{sPKE}} = (dk_1, dk_2)$ and a sanitized ciphertext $c' = (c'_1, c'_2)$, computes the message $m \leftarrow c'_2 \cdot ((c'_1)^{dk_1})^{-1}$. It outputs m if $m \in \mathcal{M}$, and otherwise it outputs \perp . On input dk^{sPKE} and \perp , it outputs \perp .

We first prove correctness and other straightforward properties of the scheme.

Proposition 5.8. *If Sig is correct and NIZK has perfect completeness, the scheme sPKE from above is correct, robust, has unpredictable ciphertexts, and negligible encryption-key collision probability.*

Proof. To prove correctness, let $(sp^{\text{sPKE}}, msk^{\text{sPKE}})$ in the range of sPKE.Setup , $(ek^{\text{sPKE}}, dk^{\text{sPKE}})$ in the range of $\text{sPKE.Gen}(msk^{\text{sPKE}})$, and let $m \in \mathcal{M}$. By correctness of Sig and completeness of NIZK, the NIZK verification in sPKE.San in the correctness experiment succeeds with probability 1. Moreover, since g generates G and $r_1, dk_1 \in \mathbb{Z}_p^*$, we have $c_{1,1} = g^{r_1} \neq 1$ and $c_{1,2} = ek_1^{r_1} = g^{dk_1 \cdot r_1} \neq 1$. Hence,

$$\begin{aligned} c' &= \text{sPKE.San}(sp^{\text{sPKE}}, \text{sPKE.Enc}(ek^{\text{sPKE}}, m)) = ((c_{1,1})^t \cdot c_{1,3}, (c_{1,2})^t \cdot c_{1,4}) \\ &= (g^{r_1 \cdot t + s_1}, ek_1^{r_1 \cdot t + s_1} \cdot m). \end{aligned}$$

and

$$\text{sPKE.Dec}(dk^{\text{sPKE}}, c') = ek_1^{r_1 \cdot t + s_1} \cdot m \cdot \left((g^{r_1 \cdot t + s_1})^{dk_1} \right)^{-1} = g^{dk_1(r_1 \cdot t + s_1)} \cdot m \cdot \left(g^{dk_1(r_1 \cdot t + s_1)} \right)^{-1} = m.$$

This shows that sPKE is correct.

For ciphertext unpredictability, note that each ciphertext contains g^{r_1} , g^{s_1} , g^{r_2} , and g^{s_2} for uniformly chosen $r_1, s_1, r_2, s_2 \in \mathbb{Z}_p^*$. Each of these elements can only be guessed with probability $1/|\mathbb{Z}_p^*| = 1/(p-1)$, where $p \geq 2^\kappa$. We can therefore conclude that for any \mathcal{A} ,

$$\text{Adv}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-UPD-CTXT}} \leq \frac{1}{(p-1)^4} \leq \frac{1}{(2^\kappa - 1)^4}.$$

Similarly, since the encryption keys contain the pairs $(ek_1 = g^{dk_1}, ek_2 = g^{dk_2})$ for uniformly chosen $dk_1, dk_2 \in \mathbb{Z}_p^*$, we have

$$\text{Col}_{\text{sPKE}}^{\text{ek}} \leq \frac{1}{(p-1)^2} \leq \frac{1}{(2^\kappa - 1)^2}.$$

We finally prove robustness. To this end, let \mathcal{A} be a probabilistic algorithm that makes at most q queries to \mathcal{O}_G and consider $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{sPKE-USROB}}$. Further let ek_i^{sPKE} and $dk_i^{\text{sPKE}} = (dk_{i,1}, dk_{i,2})$ be the keys returned from \mathcal{O}_G for the i -th query and let (c, i_0, i_1) be the output of \mathcal{A} , where $c := (c_1, c_2, c_\sigma, \pi)$ and $c_1 = (g^a, g^b, g^c, g^d)$. Assume that $i_0 \neq i_1$ and that c passes sanitization,

since \mathcal{A} cannot win otherwise. This implies $a \neq 0 \neq b$ and sanitizing and decrypting the ciphertext with the two decryption keys yield $m_0 = g^{bt_0+d-dk_{i_0,1}(at_0+c)}$ and $m_1 = g^{bt_1+d-dk_{i_1,1}(at_1+c)}$, respectively, where $t_0, t_1 \in \mathbb{Z}_p^*$ are chosen uniformly during sanitization. We then have that \mathcal{A} wins if $m_0, m_1 \in \mathcal{M}$. Assume that $m_0 \in \mathcal{M}$. Then,

$$\begin{aligned} m_1 &= m_0 \cdot g^{-b \cdot t_0 - d + dk_{i_0,1}(a \cdot t_0 + c)} \cdot g^{b \cdot t_1 + d - dk_{i_1,1}(a \cdot t_1 + c)} \\ &= m_0 \cdot g^{(dk_{i_0,1} - dk_{i_1,1})c} \cdot g^{(a \cdot dk_{i_0,1} - b)t_0} \cdot g^{(b - a \cdot dk_{i_1,1})t_1}. \end{aligned}$$

Note that if $dk_{i_0,1} \neq dk_{i_1,1}$ and $a \neq 0 \neq b$, then $a \cdot dk_{i_0,1} - b$ and $b - a \cdot dk_{i_1,1}$ cannot both be 0. Hence, in this case, m_1 is a uniformly random element in the group G . The probability that $m_1 \in \mathcal{M}$ is therefore $|\mathcal{M}|/|G| \leq 2^{-\kappa}$. Since \mathcal{A} obtains at most q decryption keys and the $dk_{i,1}$ are uniform elements in \mathbb{Z}_p^* , the probability that $dk_{i_0,1} = dk_{i_1,1}$ is bounded by $q^2 \cdot 1/(p-1) \leq q^2 \cdot 1/(2^\kappa - 1)$. We can therefore conclude that

$$\text{Adv}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-USROB}} \leq 2^{-\kappa} + \frac{q^2}{2^\kappa - 1} \leq \frac{q^2 + 1}{2^\kappa - 1}. \quad \square$$

The main result of this section is the security of the scheme, summarized in the following theorem.

Theorem 5.9. *If the DDH assumption holds in the group G , PKE is IND-CPA secure, Sig is EUF-CMA secure, and if NIZK is zero-knowledge, computationally sound, and one-time simulation sound, then the scheme sPKE from above is IND-CCA secure, IK-CCA secure, and SAN-CCA secure.*

On a high level, our proof proceeds as follows. It is rather straightforward to show that our variant of ElGamal encryption satisfies the CPA versions of the three properties. The proof of CCA security follows the proof by Sahai for public-key encryption [Sah99]: Since the NIZK ensures that both ciphertext components are encryptions of the same message, it does not matter which component is decrypted. In a reduction, where we assume an adversary \mathcal{A} against the CCA variants of the desired properties, and we want to break the corresponding CPA variants, we only get one public key and no decryption oracle from the challenger. In order to emulate the view toward \mathcal{A} , the reduction chooses an additional public key and a CRS for the NIZK scheme. Since the reduction thus knows one of the secret keys, it can emulate a decryption oracle. To generate a challenge ciphertext, the reduction obtains one challenge ciphertexts from its CPA challenger, and encrypts another, arbitrary message to get a second ciphertext. The reduction uses the NIZK simulator to obtain an accepting proof that is indistinguishable from a “real proof”, even if the underlying statement is not true. A crucial point here is that the NIZK scheme has to be *one-time simulation sound* (see Definition A.11). This ensures that even if the adversary sees one simulated (accepting) proof of a wrong statement, it is not capable of producing accepting proofs of wrong statements, except by reproducing the exact proof obtained within the challenge, but which \mathcal{A} is not allowed to ask to the decryption oracle by the CCA definition. The fundamental result of Sahai [Sah99] is that the above strategy successfully simulates a complete CCA attack toward \mathcal{A} .

An additional obstacle we have is that to preserve anonymity, the NIZK needs to be verified without knowing which encryption keys were used. On the other hand, the reduction only works if the two used keys “match”, since otherwise, the emulated decryption oracle would use an incorrect key to decrypt. To prevent an adversary from mixing different key pairs for encryptions, the key-generation process signs valid key pairs, and the NIZK ensures that a signed pair was

used. Due to anonymity, this signature cannot be directly contained in the ciphertexts. Instead, it is part of the witness. To prove that if a ciphertext is accepted, the used key pair was indeed signed by the key-generation process, we show that if \mathcal{A} manages to produce a ciphertext that is accepted but the keys were not signed, we can break EUF-CMA security of the signature scheme. In this reduction, we have to provide a forgery. Hence, the reduction needs to extract the signature and the used encryption keys from the ciphertext. This could be achieved by assuming that the NIZK is extractable. Extractability and simulation-soundness at the same time is, however, a quite strong assumption. Instead, we add an encryption of the signature and the key pair under a separate PKE scheme to the ciphertexts. The reduction can then generate the keys for this PKE scheme itself and perform extraction by decrypting that ciphertext.

Since the proofs for IND-CCA and IK-CCA security closely follow the proof by Sahai [Sah99], we here prove SAN-CCA security and defer the other proofs to [Appendix C](#).

Lemma 5.10. *Let sPKE be the scheme from above and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms such that \mathcal{A}_1 and \mathcal{A}_2 together make at most q_G queries to \mathcal{O}_G and at most q_{SD} queries to \mathcal{O}_{SD_0} and \mathcal{O}_{SD_1} combined. Then, there exist adversaries \mathcal{A}_{DDH} , \mathcal{A}_{snd} , and \mathcal{A}_{Sig} (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-SAN-CCA}}$) such that*

$$\begin{aligned} \text{Adv}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-SAN-CCA}} \leq & 8 \cdot \text{Adv}_{g, \mathcal{A}_{DDH}}^{\text{DDH}} + (24q_{SD} + 48) \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{snd}}^{\text{NIZK-snd}} \\ & + 24 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{Sig}}^{\text{Sig-EUF-CMA}} + \frac{52q_G^2 + 192q_G + 196}{2^\kappa - 1}. \end{aligned}$$

Proof. Let W_{san} be the event that \mathcal{A} wins the sanitization game, i.e.,

$$W_{\text{san}} := [b' = b \wedge \exists j, j' \in \{0, 1\} m_{0,j} \neq \perp \neq m_{1,j'}].$$

We define hybrid experiments H_0 to H_2 as follows:

- $H_0 := \text{Exp}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-SAN-CCA}}$ is the sanitization experiment.
- H_1 is identical to H_0 , except that if $c'_0 \neq \perp$, then c'_0 is replaced by two uniformly random group elements (g^b, g^c) .
- H_2 is identical to H_1 , except that if $c'_1 \neq \perp$, then c'_1 is replaced by two uniformly random group elements (g^b, g^c) .

In H_2 , if $c'_0 \neq \perp$ and $c'_1 \neq \perp$, the view of \mathcal{A} is independent of the bit b . Hence, \mathcal{A} cannot guess b with probability more than $1/2$ in this case. On the other hand, if $c'_0 = \perp$ or $c'_1 = \perp$, then $m_{0,0} = m_{0,1} = \perp$ or $m_{1,0} = m_{1,1} = \perp$, respectively, since \perp decrypts to \perp . By definition of the sanitization advantage, \mathcal{A} cannot win in this case. Thus,

$$\Pr^{H_2}[W_{\text{san}}] \leq \frac{1}{2}. \quad (2)$$

To conclude the proof, we show that the probability of W_{san} in H_0 differs only negligibly from its probability in H_2 . To this end, we first prove that three bad events occur only with negligible probability in any of the hybrids.

Claim 1. *Let $i \in \{0, 1, 2\}$ and consider the experiment H_i . Further let B_1 be the event that \mathcal{A} outputs as c_0 or c_1 or queries at least one of its decryption oracles with a valid but improper*

ciphertext $(c_1, c_2, c_\sigma, \pi)$, i.e., $(g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma) \notin L$, but where π is an accepting proof, i.e., $\text{NIZK.Ver}(crs, x := (g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma), \pi) = 1$. Then, there exists an adversary $\mathcal{A}_{\text{snd}}^i$ such that

$$\Pr^{H_i}[B_1] \leq (q_{SD} + 2) \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^i}^{\text{NIZK-snd}}.$$

Proof of claim. On input crs from the soundness challenger, $\mathcal{A}_{\text{snd}}^i$ uses this CRS, generates all needed keys itself, and emulates an execution of H_i toward \mathcal{A} . It initially chooses $q_0 \leftarrow \{-1, 0, 1, \dots, q_{SD}\}$ uniformly at random. If $q_0 > 0$ and when \mathcal{A} submits the q_0 -th query to a decryption oracle, $\mathcal{A}_{\text{snd}}^i$ outputs the corresponding statement and proof to the challenger. If $q_0 \leq 0$ and when \mathcal{A} outputs (c_0, c_1, st) , then $\mathcal{A}_{\text{snd}}^i$ submits the statement and proof from c_{q_0+1} to the challenger. If B_1 occurs, then for some q_0 , $\mathcal{A}_{\text{snd}}^i$ outputs an accepting proof for an incorrect statement. Hence, the claim follows. \diamond

Claim 2. Let $i \in \{0, 1, 2\}$ and consider the experiment H_i . Further let B_2 be the event that \mathcal{A} outputs as c_0 or c_1 or queries at least one of its decryption oracles with a valid and proper ciphertext $(c_1, c_2, c_\sigma, \pi)$, i.e., $(g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma) \in L$ and π is accepting, but where c_σ is the encryption of a triple (ek_1, ek_2, σ) , such that the pair (ek_1, ek_2) has never been output by the experiment or the oracle \mathcal{O}_G . Then, there exists an adversary $\mathcal{A}_{\text{Sig}}^i$ such that

$$\Pr^{H_i}[B_2] \leq \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^i}^{\text{Sig-EUF-CMA}}.$$

Proof of claim. On input a signature verification key vk^{Sig} , $\mathcal{A}_{\text{Sig}}^i$ generates all keys except for vk^{Sig} and sk^{Sig} , and emulates an execution of H_i . To generate the encryption keys ek_0^{PKE} and ek_1^{PKE} and to answer queries to \mathcal{O}_G , $\mathcal{A}_{\text{Sig}}^i$ obtains the needed signature using the signing oracle of $\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^i}^{\text{Sig-EUF-CMA}}$. The rest of H_i is straightforward to emulate since $\mathcal{A}_{\text{Sig}}^i$ possesses all keys except for sk^{Sig} . Whenever \mathcal{A} returns or submits a ciphertext $(c_1, c_2, c_\sigma, \pi)$ to one of the decryption oracles, $\mathcal{A}_{\text{Sig}}^i$ decrypts c_σ to obtain a pair (ek'_1, ek'_2) and a signature σ' . If it has never queried (ek'_1, ek'_2) to its signing oracle and if the signature is valid, then it outputs $((ek'_1, ek'_2), \sigma')$ as its forgery. Note that if B_2 occurs, $\mathcal{A}_{\text{Sig}}^i$ obtains a forgery, so the claim follows. \diamond

Claim 3. Let $i \in \{0, 1, 2\}$ and consider the experiment H_i . Further let B_3 be the event that H_i generates two different encryption keys $ek^{\text{PKE}} = (g, p, crs, ek^{\text{PKE}}, vk^{\text{Sig}}, ek_1, ek_2, \sigma)$ and $(ek^{\text{PKE}})' = (g, p, crs, ek^{\text{PKE}}, vk^{\text{Sig}}, ek'_1, ek'_2, \sigma')$ such that $ek_1 = ek'_1$ or $ek_2 = ek'_2$. Then,

$$\Pr^{H_i}[B_3] \leq \frac{2(q_G + 2)^2}{2^\kappa - 1}.$$

Proof of claim. The experiment H_i initially generates two encryption keys and then one for each query to \mathcal{O}_G . Hence, there are at most $(q_G + 2)^2$ such pairs. For each of these pairs, the probability that one of the two components collides is at most $2 \cdot (1/|\mathbb{Z}_p^*|) = 2/(p - 1)$. Using $p \geq 2^\kappa$ and the union bound implies the claim. \diamond

We now bound the difference of the probabilities of W_{san} in different hybrids. To this end, let $B := B_1 \cup B_2 \cup B_3$.

Claim 4. For all $i \in \{0, 1\}$, there exists an adversary $\mathcal{A}_{\text{DDH}}^i$ such that

$$\Pr^{H_i}[W_{\text{san}}] - \Pr^{H_{i+1}}[W_{\text{san}}] \leq 2 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{DDH}} + 2 \cdot \Pr^{H_i}[B] + 4 \cdot \Pr^{H_{i+1}}[B] + \frac{q_G^2 + 1}{2^\kappa - 1}.$$

Proof of claim. Let $i \in \{0, 1\}$ and let G_0 and G_1 be the events that c_i output by \mathcal{A} is an encryption under ek_0^{SPKE} and ek_1^{SPKE} , respectively. If B , G_0 , and G_1 all do not occur, then c_i is either invalid or a valid encryption under a key different from ek_0^{SPKE} and ek_1^{SPKE} . Since W_{san} can only occur if the ciphertext decrypts to a message different from \perp under one of these keys, this only happens if robustness is violated. Using the result on robustness derived in the proof of [Proposition 5.8](#), this implies

$$\Pr^{H_i}[W_{\text{san}} \cap \neg B \cap \neg G_1 \cap \neg G_2] \leq \Pr^{H_i}[W_{\text{san}} \mid \neg B \cap \neg G_1 \cap \neg G_2] \leq \frac{q_G^2 + 1}{2^\kappa - 1}.$$

We also have

$$\begin{aligned} & \Pr^{H_i}[W_{\text{san}}] - \Pr^{H_{i+1}}[W_{\text{san}}] \\ &= \Pr^{H_i}[W_{\text{san}} \cap \neg B \cap (G_1 \cup G_2)] + \Pr^{H_i}[W_{\text{san}} \cap (B \cup \neg(G_1 \cup G_2))] \\ & \quad - \Pr^{H_{i+1}}[W_{\text{san}} \cap \neg B \cap (G_1 \cup G_2)] - \Pr^{H_{i+1}}[W_{\text{san}} \cap (B \cup \neg(G_1 \cup G_2))] \\ &\leq \Pr^{H_i}[W_{\text{san}} \cap \neg B \cap (G_1 \cup G_2)] - \Pr^{H_{i+1}}[W_{\text{san}} \cap \neg B \cap (G_1 \cup G_2)] \\ & \quad + \Pr^{H_i}[B] + \Pr^{H_i}[W_{\text{san}} \cap \neg(G_1 \cup G_2)], \end{aligned}$$

and

$$\Pr^{H_i}[W_{\text{san}} \cap \neg(G_1 \cup G_2)] \leq \Pr^{H_i}[W_{\text{san}} \cap \neg B \cap \neg(G_1 \cup G_2)] + \Pr^{H_i}[B].$$

This implies

$$\begin{aligned} \Pr^{H_i}[W_{\text{san}}] - \Pr^{H_{i+1}}[W_{\text{san}}] &\leq \Pr^{H_i}[W_{\text{san}} \cap \neg B \cap (G_1 \cup G_2)] \\ & \quad - \Pr^{H_{i+1}}[W_{\text{san}} \cap \neg B \cap (G_1 \cup G_2)] + 2 \cdot \Pr^{H_i}[B] + \frac{q_G^2 + 1}{2^\kappa - 1}. \quad (3) \end{aligned}$$

We now define the adversary $\mathcal{A}_{\text{DDH}}^i$. On input (X, Y, T) , $\mathcal{A}_{\text{DDH}}^i$ chooses $j \leftarrow \{0, 1\}$ uniformly at random and sets $ek_{j,1} \leftarrow X$. All remaining keys, including $ek_{j,2}$, are generated as in H_i , and \mathcal{A} is invoked on $(sp^{\text{SPKE}}, ek_0^{\text{SPKE}} = (ek_{0,1}, ek_{0,2}), ek_1^{\text{SPKE}} = (ek_{1,1}, ek_{1,2}))$. The adversary $\mathcal{A}_{\text{DDH}}^i$ then emulates an execution of H_i . Since it has all keys except for the decryption key $dk_{j,1}$, only the emulation of the decryption oracle \mathcal{O}_{SD_j} is nontrivial. To answer queries to this oracle, $\mathcal{A}_{\text{DDH}}^i$ sanitizes and decrypts the second ciphertext component instead of the first one using $dk_{j,2}$. When \mathcal{A} outputs (c_0, c_1, st) , both ciphertexts are sanitized and decrypted as in the emulation of the decryption oracles, except that $m_{i,j}$ is not set to \perp during decryption if $m_{i,j} \notin \mathcal{M}$. If $c'_i \neq \perp$, it is replaced by $c'_i \leftarrow (Y, T \cdot m_{i,j})$. Moreover, $\mathcal{A}_{\text{DDH}}^i$ decrypts $c_{i,\sigma}$ and checks whether it contains the encryption keys corresponding to ek_j^{SPKE} . If this is not the case, it terminates and returns 0. Otherwise, it continues with the emulation. Finally, when \mathcal{A} terminates, $\mathcal{A}_{\text{DDH}}^i$ outputs $d = 1$ if W_{san} occurs, and $d = 0$ otherwise.

Note that B not occurring implies that $\mathcal{A}_{\text{DDH}}^i$ emulates the decryption oracle perfectly since in this case, all submitted valid ciphertexts contain two encryptions of the same message under a signed key pair. Moreover, due to $\neg B_3$, the first encryption key matches the first key of the oracle if and only if the second keys match. If they match, decryption with either key yields the correct message with probability 1. Otherwise, the message (before potentially being set to \perp) is a uniform group element for both keys, as shown in the robustness proof of [Proposition 5.8](#).

Furthermore, if (X, Y, T) are three independent uniform group elements, c'_i gets replaced by two uniformly random group elements if $c'_i \neq \perp$, as in H_{i+1} . On the other hand, if $\neg B$ and

G_j occur and if $c'_i \neq \perp$, then $c_i = (c_{i,1}, c_{i,2}, c_{i,\sigma}, \pi_i)$ is a valid encryption of $m_{i,j}$ under ek_j^{SPKE} . Hence, there exist $r_1, s_1 \in \mathbb{Z}_p^*$ such that

$$c_{i,1} = (g^{r_1}, (ek_{j,1})^{r_1}, g^{s_1}, (ek_{j,1})^{s_1} \cdot m_{i,j}) = (g^{r_1}, X^{r_1}, g^{s_1}, X^{s_1} \cdot m_{i,j}).$$

In H_i , this ciphertext is sanitized to $c'_i = (g^{r_1 \cdot t + s_1}, X^{r_1 \cdot t + s_1} \cdot m_{i,j})$ for $t \leftarrow \mathbb{Z}_p^*$. If we further have $X = g^a$, $Y = g^b$, and $T = g^{ab}$, then this corresponds to $c'_i = (g^{r_1 \cdot t + s_1}, g^{a \cdot (r_1 \cdot t + s_1)} \cdot m_{i,j})$, which is equally distributed as the sanitization $(Y, T \cdot m_{i,j})$ generated by $\mathcal{A}_{\text{DDH}}^i$. Since we also have that the probability of $\neg B \cap G_j$ is equal in $\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{real}}$ and H_i , as well as in $\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}$ and H_{i+1} , we can conclude

$$\begin{aligned} \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{real}}} [d = 1 \cap \neg B \cap G_j] &= \Pr^{H_i} [W_{\text{san}} \cap \neg B \cap G_j], \\ \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap \neg B \cap G_j] &= \Pr^{H_{i+1}} [W_{\text{san}} \cap \neg B \cap G_j]. \end{aligned}$$

Hence,

$$\begin{aligned} \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{DDH}} &= \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{real}}} [d = 1] - \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1] \\ &= \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{real}}} [d = 1 \cap \neg B \cap G_j] + \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{real}}} [d = 1 \cap (B \cup \neg G_j)] \\ &\quad - \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap \neg B \cap G_j] - \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap (B \cup \neg G_j)] \\ &\geq \Pr^{H_i} [W_{\text{san}} \cap \neg B \cap G_j] - \Pr^{H_{i+1}} [W_{\text{san}} \cap \neg B \cap G_j] \\ &\quad - \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap B] - \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap \neg G_j]. \end{aligned}$$

If $\neg B$ occurs, then $c_{i,\sigma}$ contains the correct encryption keys and thus, if also $\neg G_j$ occurs, $\mathcal{A}_{\text{DDH}}^i$ always returns 0. This implies $\Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap \neg G_j \cap \neg B] = 0$, and therefore

$$\begin{aligned} \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap \neg G_j] &= \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap \neg G_j \cap \neg B] + \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [d = 1 \cap \neg G_j \cap B] \\ &\leq \Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [B]. \end{aligned}$$

Using $\Pr^{\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{rand}}} [B] = \Pr^{H_{i+1}} [B]$, we obtain

$$\text{Adv}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{DDH}} \geq \Pr^{H_i} [W_{\text{san}} \cap \neg B \cap G_j] - \Pr^{H_{i+1}} [W_{\text{san}} \cap \neg B \cap G_j] - 2 \cdot \Pr^{H_{i+1}} [B].$$

Combining this with [equation \(3\)](#) and the fact that given $G_1 \cup G_2$ and $\neg B$, G_j occurs with probability 1/2 (independently of W_{san}), we can conclude

$$\begin{aligned} \Pr^{H_i} [W_{\text{san}}] - \Pr^{H_{i+1}} [W_{\text{san}}] &\leq 2 \cdot \Pr^{H_i} [W_{\text{san}} \cap \neg B \cap G_j] - 2 \cdot \Pr^{H_{i+1}} [W_{\text{san}} \cap \neg B \cap G_j] \\ &\quad + 2 \cdot \Pr^{H_i} [B] + \frac{q_G^2 + 1}{2^\kappa - 1} \\ &\leq 2 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{DDH}} + 4 \cdot \Pr^{H_{i+1}} [B] + 2 \cdot \Pr^{H_i} [B] + \frac{q_G^2 + 1}{2^\kappa - 1}. \quad \diamond \end{aligned}$$

Using Claim 4 and equation (2), we obtain

$$\begin{aligned}
& \text{Adv}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-SAN-CCA}} \\
&= 2 \cdot \Pr^{H_0}[W_{\text{san}}] - 1 \\
&= 2 \cdot (\Pr^{H_0}[W_{\text{san}}] - \Pr^{H_1}[W_{\text{san}}] + \Pr^{H_1}[W_{\text{san}}] - \Pr^{H_2}[W_{\text{san}}] + \Pr^{H_2}[W_{\text{san}}]) - 1 \\
&\leq 4 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^0}^{\text{DDH}} + 4 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}} + 4 \cdot \Pr^{H_0}[B] + 12 \cdot \Pr^{H_1}[B] + 8 \cdot \Pr^{H_2}[B] + 4 \cdot \frac{q_G^2 + 1}{2^\kappa - 1}.
\end{aligned}$$

Claims 1 to 3 further imply

$$\Pr^{H_i}[B] \leq (q_{SD} + 2) \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^i}^{\text{NIZK-snd}} + \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^i}^{\text{Sig-EUF-CMA}} + \frac{2(q_G + 2)^2}{2^\kappa - 1}.$$

Hence,

$$\begin{aligned}
\text{Adv}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-SAN-CCA}} &\leq 4 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^0}^{\text{DDH}} + 4 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}} + (4q_{SD} + 8) \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^0}^{\text{NIZK-snd}} \\
&\quad + (12q_{SD} + 24) \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^1}^{\text{NIZK-snd}} + (8q_{SD} + 16) \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^2}^{\text{NIZK-snd}} \\
&\quad + 4 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^0}^{\text{Sig-EUF-CMA}} + 12 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^1}^{\text{Sig-EUF-CMA}} + 8 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^2}^{\text{Sig-EUF-CMA}} \\
&\quad + \frac{48(q_G + 2)^2 + 4q_G^2 + 4}{2^\kappa - 1}.
\end{aligned}$$

We define the adversary \mathcal{A}_{DDH} as running $\mathcal{A}_{\text{DDH}}^0$ and $\mathcal{A}_{\text{DDH}}^1$ with probability $\frac{1}{2}$ each, the adversary \mathcal{A}_{snd} as running $\mathcal{A}_{\text{snd}}^0$ with probability $\frac{4q_{SD}+8}{24q_{SD}+48}$, $\mathcal{A}_{\text{snd}}^1$ with probability $\frac{12q_{SD}+24}{24q_{SD}+48}$, and $\mathcal{A}_{\text{snd}}^2$ with probability $\frac{8q_{SD}+16}{24q_{SD}+48}$, and the adversary \mathcal{A}_{Sig} as running $\mathcal{A}_{\text{Sig}}^0$ with probability $\frac{4}{24}$, $\mathcal{A}_{\text{Sig}}^1$ with probability $\frac{12}{24}$, and $\mathcal{A}_{\text{Sig}}^2$ with probability $\frac{8}{24}$. Using the result above, we finally conclude

$$\begin{aligned}
\text{Adv}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-SAN-CCA}} &\leq 8 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}}^{\text{DDH}} + (24q_{SD} + 48) \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}}^{\text{NIZK-snd}} \\
&\quad + 24 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + \frac{52q_G^2 + 192q_G + 196}{2^\kappa - 1}. \quad \square
\end{aligned}$$

6 Construction of an ACE Scheme

6.1 Construction for Equality

Following Fuchsbaauer et al. [FGKO17], we first construct an ACE scheme for the equality policy, i.e., $P(i, j) = 1 \Leftrightarrow i = j$, and then use such a scheme in another construction for richer policies. We base our construction on an sPKE scheme, which already has many important properties needed for a secure ACE scheme. A syntactical difference between sPKE and ACE schemes is that the key generation of the former on every invocation produces a fresh key pair, while the latter schemes allow the generation of keys for a given role. To bind key pairs to some role $i \in [n]$, we use the output of a pseudorandom function on input i as the randomness for the sPKE key generation. For role-respecting security, we have to ensure that an adversary can only produce ciphertexts for keys obtained from the key generation oracle. This is achieved by signing all keys with a signing key generated at setup. To prevent malleability attacks as the ones described in Section 3, the encryption algorithm additionally signs all ciphertexts with a separate signing key that is tied to the encryption key. To maintain anonymity, the signatures are not part of the

ciphertext but the encrypters prove in zero-knowledge that they know such signatures. Finally, the modification detection simply checks whether the ciphertexts (without the NIZK proofs) are equal. Intuitively, this is sufficient since we assume the underlying sPKE scheme to be CCA secure, which implies that it is not possible to meaningfully modify a given ciphertext. Hence, a ciphertext is either equal to an existing one (and thus detected by the algorithm) or a fresh encryption.

Our construction. Let sPKE be a sanitizable public-key encryption scheme, let Sig be a signature scheme, and let F be a PRF. Further let NIZK be a NIZK proof of knowledge system for the language $L := \{x \mid \exists w (x, w) \in R\}$, where the relation R is defined as follows: for $x = (vk_i^{\text{Sig}}, \tilde{c})$ and $w = (ek_i^{\text{sPKE}}, m, r, vk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, \sigma_c^{\text{Sig}})$, $(x, w) \in R$ if and only if

$$\begin{aligned} \tilde{c} = \text{sPKE.Enc}(ek_i^{\text{sPKE}}, m; r) \wedge \text{Sig.Ver}(vk_i^{\text{Sig}}, [ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}], \sigma_i^{\text{Sig}}) = 1 \\ \wedge \text{Sig.Ver}(vk_i^{\text{Sig}}, \tilde{c}, \sigma_c^{\text{Sig}}) = 1. \end{aligned}$$

We define an ACE with modification detection scheme ACE as follows:

Setup: On input a security parameter 1^κ and a policy $P: [n] \times [n] \rightarrow \{0, 1\}$ with $P(i, j) = 1 \Leftrightarrow i = j$, the algorithm ACE.Setup picks a random PRF key K for a PRF F , and runs

$$\begin{aligned} (sp^{\text{sPKE}}, msk^{\text{sPKE}}) &\leftarrow \text{sPKE.Setup}(1^\kappa), \\ (vk_i^{\text{Sig}}, sk_i^{\text{Sig}}) &\leftarrow \text{Sig.Gen}(1^\kappa), \\ crs^{\text{NIZK}} &\leftarrow \text{NIZK.Gen}(1^\kappa). \end{aligned}$$

It outputs the master secret key $msk^{\text{ACE}} := (K, msk^{\text{sPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, crs^{\text{NIZK}})$ and the sanitizer parameters $sp^{\text{ACE}} := (sp^{\text{sPKE}}, vk_i^{\text{Sig}}, crs^{\text{NIZK}})$.

Key generation: The algorithm ACE.Gen on input a master secret key $msk^{\text{ACE}} = (K, msk^{\text{sPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, crs^{\text{NIZK}})$, a role $i \in [n]$, and a type $t \in \{\text{sen}, \text{rec}\}$, computes

$$(ek_i^{\text{sPKE}}, dk_i^{\text{sPKE}}) \leftarrow \text{sPKE.Gen}(msk^{\text{sPKE}}; F_K([i, 0])).$$

If $t = \text{sen}$, it further computes

$$\begin{aligned} (vk_i^{\text{Sig}}, sk_i^{\text{Sig}}) &\leftarrow \text{Sig.Gen}(1^\kappa; F_K([i, 1])), \\ \sigma_i^{\text{Sig}} &\leftarrow \text{Sig.Sign}(sk_i^{\text{Sig}}, [ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}]; F_K([i, 2])). \end{aligned}$$

If $t = \text{sen}$, it outputs the encryption key $ek_i^{\text{ACE}} := (vk_i^{\text{Sig}}, ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, crs^{\text{NIZK}})$; if $t = \text{rec}$, it outputs the decryption key $dk_i^{\text{ACE}} := dk_i^{\text{sPKE}}$.

Encryption: On input an encryption key $ek_i^{\text{ACE}} = (vk_i^{\text{Sig}}, ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, crs^{\text{NIZK}})$ and a message $m \in \mathcal{M}^{\text{ACE}}$, the algorithm ACE.Enc samples randomness r and computes

$$\begin{aligned} \tilde{c} &\leftarrow \text{sPKE.Enc}(ek_i^{\text{sPKE}}, m; r), \\ \sigma_c^{\text{Sig}} &\leftarrow \text{Sig.Sign}(sk_i^{\text{Sig}}, \tilde{c}), \\ \pi^{\text{NIZK}} &\leftarrow \text{NIZK.Prove}(crs^{\text{NIZK}}, x := (vk_i^{\text{Sig}}, \tilde{c}), w := (ek_i^{\text{sPKE}}, m, r, vk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, \sigma_c^{\text{Sig}})). \end{aligned}$$

It outputs the ciphertext $c := (\tilde{c}, \pi^{\text{NIZK}})$.

Sanitization: On input sanitizer parameters $sp^{\text{ACE}} = (sp^{\text{sPKE}}, vk^{\text{Sig}}, crs^{\text{NIZK}})$ and a ciphertext $c = (\tilde{c}, \pi^{\text{NIZK}})$, ACE.San outputs the sanitized ciphertext $c' \leftarrow \text{sPKE.San}(sp^{\text{sPKE}}, \tilde{c})$ if $\text{NIZK.Ver}(crs^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}) = 1$; otherwise, it outputs \perp .

Decryption: The algorithm ACE.Dec on input a decryption key dk_j^{ACE} and a sanitized ciphertext c' , outputs the message $m \leftarrow \text{sPKE.Dec}(dk_j^{\text{ACE}}, c')$.

Modification detection: The algorithm ACE.DMod on input sp^{ACE} , $c_1 = (\tilde{c}_1, \pi_1^{\text{NIZK}})$, and $c_2 = (\tilde{c}_2, \pi_2^{\text{NIZK}})$, outputs 1 if $\tilde{c}_1 = \tilde{c}_2$, and 0 otherwise.

We first show that our scheme is correct and strongly detectable.

Proposition 6.1. *Let ACE be the scheme from above. Then, ACE is perfectly correct, i.e., $\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-CORR}} = 0$ for all \mathcal{A} . Moreover, if F is pseudorandom and sPKE is unrestricted strongly robust, then ACE is strongly detectable.*

Proof. Perfect correctness follows from the perfect correctness of the sPKE and signature schemes and the perfect completeness of the NIZK proof system.

To prove strong detectability, let \mathcal{A} be a probabilistic algorithm. We assume without loss of generality that \mathcal{A} returns (m, r, i, j) with $P(i, j) = 0$ since doing otherwise can only reduce the advantage. Let $H_0 := \text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-sDTCT}}$, let H_1 be as H_0 where F_K is replaced by a truly uniform function U , and let W be the event that \mathcal{A} wins the strong detectability game, i.e.,

$$W := [\text{ACE.Dec}(dk_j^{\text{ACE}}, \text{ACE.San}(sp^{\text{ACE}}, \text{ACE.Enc}(ek_i^{\text{ACE}}, m; r))) \neq \perp].$$

We first show that the difference in the winning probability in H_0 and H_1 is bounded by the PRF advantage.

Claim 1. *There exists a probabilistic algorithm $\mathcal{A}_{\text{PRF}}^{\mathcal{O}(\cdot)}$ such that*

$$\Pr^{H_0}[W] - \Pr^{H_1}[W] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}}.$$

Proof of claim. Consider $\mathcal{A}_{\text{PRF}}^{\mathcal{O}(\cdot)}$ that emulates an execution of H_0 , where all invocations of $F_K(\cdot)$ are replaced by a call to the oracle $\mathcal{O}(\cdot)$. When \mathcal{A} wins, \mathcal{A}_{PRF} outputs 1, and 0 otherwise. In case $\mathcal{O}(\cdot)$ corresponds to $F_K(\cdot)$, \mathcal{A}_{PRF} perfectly emulates H_0 , if it corresponds to $U(\cdot)$, it perfectly emulates H_1 . Hence,

$$\Pr^{H_0}[W] - \Pr^{H_1}[W] = \Pr[\mathcal{A}_{\text{PRF}}^{F_K(\cdot)}(1^\kappa) = 1] - \Pr[\mathcal{A}_{\text{PRF}}^{U(\cdot)}(1^\kappa) = 1] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}}. \quad \diamond$$

We now construct a winner \mathcal{A}_{rob} for the robustness game for sPKE. The algorithm \mathcal{A}_{rob} on input sp^{sPKE} emulates an execution of H_1 . To answer queries of \mathcal{A} to the key-generation oracle, \mathcal{A}_{rob} uses the oracle \mathcal{O}_G to obtain encryption and decryption keys for sPKE; the required signature keys are generated internally. For each query (i, t) , \mathcal{A}_{rob} remembers the generated keys ek_i^{ACE} and dk_i^{ACE} , and returns the same keys for subsequent queries with the same i . When \mathcal{A} returns (m, r, i, j) , \mathcal{A}_{rob} first checks whether i and j have been queried by \mathcal{A} to the key-generation oracle. If not, \mathcal{A}_{rob} now generates these keys as above. Let \tilde{r} be the randomness used by $\text{ACE.Enc}(ek_i^{\text{ACE}}, m; r)$ for the algorithm sPKE.Enc. Then, \mathcal{A}_{rob} computes $c \leftarrow \text{sPKE.Enc}(ek_i^{\text{sPKE}}, m; \tilde{r})$, and returns (c, i_0, i_1) , such that the i_0 -th query and the i_1 -th query to the key-generation oracle were for the roles i and j , respectively. Since P is the equality

predicate, $P(i, j) = 0$ is equivalent to $i_0 \neq i_1$. We further have by the perfect correctness of sPKE that $\text{sPKE.Dec}(dk_i^{\text{sPKE}}, \text{sPKE.San}(sp^{\text{sPKE}}, c)) \neq \perp$. Hence, \mathcal{A}_{rob} wins the robustness game if and only if \mathcal{A} wins the strong detectability game in H_1 . Using [Claim 1](#), we can therefore conclude

$$\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-sDTCT}} = \Pr^{H_0}[W] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \Pr^{H_1}[W] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}. \quad \square$$

In the following, we prove the security of our scheme, which is summarized by the theorem below.

Theorem 6.2. *If F is pseudorandom, NIZK is zero-knowledge and extractable, Sig is EUF-CMA secure, and sPKE is IND-CCA, IK-CCA, SAN-CCA, USROB, and UPD-CTXT secure and has negligible encryption-key collision probability, then the scheme ACE from above is PRV-CCA, sANON-CCA, SAN-CCA, UDEC, and RR secure, and has NDTCT-FENC.*

We first show that our scheme satisfies the privacy definition from [Definition 4.2](#) if the underlying sanitizable public-key encryption scheme is IND-CCA secure, the PRF is secure, and the NIZK is zero-knowledge.

Lemma 6.3. *Let ACE be the scheme from above, let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an attacker on the privacy such that \mathcal{A}_1 makes at most q_S queries of the form (\cdot, sen) to the oracle \mathcal{O}_G , and at most q_D queries to \mathcal{O}_{SD} . Then, there exist probabilistic algorithms \mathcal{A}_{PRF} , \mathcal{A}_{ZK} , and $\mathcal{A}_{\text{sPKE}}$ (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$) such that*

$$\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-PRV-CCA}} \leq 2 \cdot \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + (q_S + q_D + 1) \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-IND-CCA}}.$$

Proof. We assume without loss of generality that \mathcal{A} ensures $i_0 = i_1$ and $P(i_0, j) = 0$ for all $j \in J$, since doing otherwise can only decrease the advantage. Let $H_0 := \text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$ and H_1 be as H_0 where F_K is replaced by a truly uniform random function U . The following can be proven as [Claim 1](#) in the proof of [Proposition 6.1](#).

Claim 1. *There exists a probabilistic algorithm $\mathcal{A}_{\text{PRF}}^{\mathcal{O}(\cdot)}$ such that*

$$\Pr^{H_0}[b' = b] - \Pr^{H_1}[b' = b] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}}.$$

Now let H_2 be as H_1 , where we replace $crs^{\text{NIZK}} \leftarrow \text{NIZK.Gen}(1^\kappa)$ by $(crs^{\text{NIZK}}, \tau^{\text{NIZK}}) \leftarrow S_1^{\text{NIZK}}(1^\kappa)$ in ACE.Setup, and for the generation of the challenge ciphertext c^* , we replace $\pi^{\text{NIZK}} \leftarrow \text{NIZK.Prove}(crs^{\text{NIZK}}, x, w)$ in ACE.Enc by $\pi^{\text{NIZK}} \leftarrow S_2^{\text{NIZK}}(crs^{\text{NIZK}}, \tau^{\text{NIZK}}, x)$.

Claim 2. *There exists a probabilistic algorithm $\mathcal{A}_{\text{ZK}}^{\mathcal{O}(\cdot, \cdot)}$ such that*

$$\Pr^{H_1}[b' = b] - \Pr^{H_2}[b' = b] = \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}}.$$

Proof of claim. The algorithm $\mathcal{A}_{\text{ZK}}^{\mathcal{O}(\cdot, \cdot)}$ on input crs^{NIZK} proceeds as follows. It emulates an execution of H_1 , where in ACE.Setup, crs^{NIZK} is used instead of generating it, and for the generation of c^* , $\text{NIZK.Prove}(crs^{\text{NIZK}}, x, w)$ in ACE.Enc is replaced by the oracle query (x, w) . Finally, $\mathcal{A}_{\text{ZK}}^{\mathcal{O}(\cdot, \cdot)}$ outputs $\tilde{b} = 1$ if \mathcal{A}_2 returns $b' = b$, and $\tilde{b} = 0$ otherwise. Note that if crs^{NIZK} is generated by NIZK.Gen and $\mathcal{O}(\cdot, \cdot)$ corresponds to $\text{NIZK.Prove}(crs^{\text{NIZK}}, \cdot, \cdot)$, $\mathcal{A}_{\text{ZK}}^{\mathcal{O}(\cdot, \cdot)}$ perfectly emulates H_1 . Moreover, if crs^{NIZK} is generated together with τ^{NIZK} by S_1^{NIZK} and $\mathcal{O}(x, w)$ returns $S_2^{\text{NIZK}}(crs^{\text{NIZK}}, \tau^{\text{NIZK}}, x)$, $\mathcal{A}_{\text{ZK}}^{\mathcal{O}(\cdot, \cdot)}$ perfectly emulates H_2 . Thus, the claim follows. \diamond

We finally show how to transform any winner \mathcal{A} for H_2 to a winner $\mathcal{A}_{\text{sPKE}}$ for the IND-CCA game for the scheme sPKE . The strategy of our reduction is to guess which oracle queries of \mathcal{A}_1 are for the role i_0 , use the key from the sPKE -scheme for these queries, and generate all other keys as H_2 . Details follow. On input $(sp^{\text{sPKE}}, ek^{\text{sPKE}})$, $\mathcal{A}_{\text{sPKE}}$ initializes $i_{q_0} \leftarrow \perp$, $k_q \leftarrow 1$, chooses $q_0 \leftarrow \{0, \dots, q_S + q_D\}$ uniformly at random, runs $(vk^{\text{Sig}}, sk^{\text{Sig}}) \leftarrow \text{Sig.Gen}(1^\kappa)$, and $(crs^{\text{NIZK}}, \tau^{\text{NIZK}}) \leftarrow S_1^{\text{NIZK}}(1^\kappa)$, and gives $sp^{\text{ACE}} := (sp^{\text{sPKE}}, vk^{\text{Sig}}, crs^{\text{NIZK}})$ to \mathcal{A}_1 . It emulates the oracles for \mathcal{A}_1 as follows.

$\mathcal{O}_G(\cdot, \cdot)$: On query (i, sen) , if $k_q \neq q_0$ and $i \neq i_{q_0}$, then generate an encryption key $ek_i^{\text{ACE}} := (vk_i^{\text{Sig}}, ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, crs^{\text{NIZK}})$ as H_2 does, where $(ek_i^{\text{sPKE}}, dk_i^{\text{sPKE}})$ is obtained via \mathcal{O}_G and remembered for future queries. If $k_q = q_0$ or $i = i_{q_0}$, replace ek_i^{sPKE} by ek_i^{sPKE} and set $i_{q_0} \leftarrow i$. In both cases, set $k_q \leftarrow k_q + 1$ at the end. On query (j, rec) , obtain a decryption key via \mathcal{O}_G .

$\mathcal{O}_{SD}(\cdot, \cdot)$: On query $(j, c = (\tilde{c}, \pi^{\text{NIZK}}))$, if $k_q \neq q_0$ and $j \neq i_{q_0}$, run $c' \leftarrow \text{ACE.San}(sp^{\text{ACE}}, c)$, generate a decryption key dk_j^{ACE} as above, decrypt c' using dk_j^{ACE} , and return the resulting message. If $k_q = q_0$ or $j = i_{q_0}$, set $i_{q_0} \leftarrow j$ and use the oracle \mathcal{O}_{SD} of the IND-CCA experiment to obtain a decryption m of \tilde{c} . If $\text{NIZK.Ver}(crs^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}) = 1$, return m , otherwise, return \perp . In all cases, set $k_q \leftarrow k_q + 1$ at the end.

When \mathcal{A}_1 returns (m_0, m_1, i_0, i_1, st) , output (m_0, m_1) to the challenger of the IND-CCA experiment to obtain a challenge ciphertext \tilde{c}^* . Then run $\pi^{\text{NIZK}} \leftarrow S_2^{\text{NIZK}}(crs^{\text{NIZK}}, \tau^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}^*))$, and give st and the ciphertext $c^* := (\tilde{c}^*, \pi^{\text{NIZK}})$ to \mathcal{A}_2 . Emulate the oracles for \mathcal{A}_2 as follows.

$\mathcal{O}_G(\cdot, \cdot)$: On query (i, sen) , if $i \neq i_0$, then generate an encryption key $ek_i^{\text{ACE}} := (vk_i^{\text{Sig}}, ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, crs^{\text{NIZK}})$ as H_2 does, where $(ek_i^{\text{sPKE}}, dk_i^{\text{sPKE}})$ is obtained via \mathcal{O}_G and remembered for future queries. If $i = i_0$, replace ek_i^{sPKE} by ek_i^{sPKE} . On query (j, rec) , obtain a decryption key from \mathcal{O}_G .

$\mathcal{O}_{SD^*}(\cdot, \cdot)$: On query $(j, c = (\tilde{c}, \pi^{\text{NIZK}}))$, run $\text{ACE.DMod}(sp^{\text{ACE}}, c^*, c)$. If the output is 1, return **test**. Otherwise, if $j \neq i_0$, run $c' \leftarrow \text{ACE.San}(sp^{\text{ACE}}, c)$, generate a decryption key dk_j^{ACE} as above, decrypt c' using dk_j^{ACE} , and return the resulting message. If $j = i_0$, use the oracle \mathcal{O}_{SD} of the IND-CCA experiment to obtain a decryption m of \tilde{c} . If $\text{NIZK.Ver}(crs^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}) = 1$, return m , otherwise, return \perp .

Note that we never query the decryption oracle of the IND-CCA experiment on \tilde{c}^* because we return **test** whenever this would be necessary. Denote by Q the event that either $i_{q_0} = i_0$, or $q_0 = 0$ and \mathcal{A}_1 does not make the query (i_0, sen) to \mathcal{O}_G and no queries for role i_0 to \mathcal{O}_{SD} . When \mathcal{A}_2 returns a bit b' and Q holds, $\mathcal{A}_{\text{sPKE}}$ returns the same bit $b'' \leftarrow b'$, if $\neg Q$, $\mathcal{A}_{\text{sPKE}}$ returns a uniform bit $b'' \leftarrow \{0, 1\}$.

Let \tilde{b} be the bit chosen by the IND-CCA challenger. Note that by our assumption on \mathcal{A} , $i_0 = i_1$ and \mathcal{A} does not query (i_0, rec) to \mathcal{O}_G , i.e., $i_0 \notin J$, since $P(i_0, i_0) = 1$. Hence, if Q occurs, the view of \mathcal{A} is identical to the one in H_2 with $b = \tilde{b}$. This implies

$$\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-IND-CCA}}} [b'' = \tilde{b} \mid Q] = \Pr^{H_2} [b' = b],$$

and therefore

$$\begin{aligned}
\Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [b'' = \tilde{b}] &= \Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [b'' = \tilde{b} \mid Q] \cdot \Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [Q] \\
&\quad + \Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [b'' = \tilde{b} \mid \neg Q] \cdot \Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [\neg Q] \\
&= \Pr^{H_2} [b' = b] \cdot \Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [Q] + \frac{1}{2} \Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [\neg Q].
\end{aligned}$$

Using that the probability of Q is $1/(q_S + q_D + 1)$, this yields

$$\begin{aligned}
&\Pr^{H_2} [b' = b] \\
&= \frac{1}{\Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [Q]} \cdot \left(\Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [b'' = \tilde{b}] - \frac{1}{2} \cdot \left(1 - \Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [Q] \right) \right) \\
&= (q_S + q_D + 1) \cdot \left(\Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [b'' = \tilde{b}] - \frac{1}{2} \right) + \frac{1}{2}.
\end{aligned}$$

Combining this with [Claims 1](#) and [2](#), we can conclude

$$\begin{aligned}
&\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-PRV-CCA}} \\
&= 2 \cdot \Pr^{H_0} [b' = b] - 1 \\
&= 2 \cdot \left(\Pr^{H_0} [b' = b] - \Pr^{H_1} [b' = b] + \Pr^{H_1} [b' = b] - \Pr^{H_2} [b' = b] + \Pr^{H_2} [b' = b] \right) - 1 \\
&= 2 \cdot \left[\text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + (q_S + q_D + 1) \left(\Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}} [b'' = \tilde{b}] - \frac{1}{2} \right) + \frac{1}{2} \right] - 1 \\
&= 2 \cdot \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + (q_S + q_D + 1) \cdot \text{Adv}_{\mathcal{S}_{\text{PKE}}, \mathcal{A}_{\text{PKE}}}^{\text{SPKE-IND-CCA}}. \quad \square
\end{aligned}$$

The proofs of the other properties use similar techniques and can be found in [Appendix D](#).

6.2 Lifting Equality to Disjunction of Equalities

We finally show how an ACE scheme for equality, as the one from [Section 6.1](#), can be used to construct a scheme for the policy $P_{\text{DEq}}: \mathcal{D}^\ell \times \mathcal{D}^\ell \rightarrow \{0, 1\}$ with

$$P_{\text{DEq}}(\mathbf{x} = (x_1, \dots, x_\ell), \mathbf{y} = (y_1, \dots, y_\ell)) = 1 \iff \bigvee_{i=1}^{\ell} x_i = y_i,$$

where \mathcal{D} is some finite set and $\ell \in \mathbb{N}$.⁴ This policy can for example be used to implement the no read-up and now write-down principle ($P(i, j) = 1 \Leftrightarrow i \leq j$) from the Bell–LaPadula model [\[BL73\]](#) via an appropriate encoding of the roles [\[FGKO17\]](#).

The intuition of our construction is as follows. A key for a role $\mathbf{x} = (x_1, \dots, x_\ell)$ contains one key of the ACE scheme for equality for each component x_i of the role vector. To encrypt a message, this message is encrypted with each of these keys. To decrypt, one tries to decrypt each ciphertext component with the corresponding key. If at least one component of the sender and receiver roles match (i.e., if the policy is satisfied), one of the decryptions is successful. So

⁴In this section, we denote roles by \mathbf{x} and \mathbf{y} instead of i and j . To be compatible with our definitions that consider policies $[n] \times [n] \rightarrow \{0, 1\}$, one needs to identify elements of \mathcal{D}^ℓ with numbers in $[n]$. We will ignore this technicality to simplify the presentation.

far, the construction is identical to the one by Fuchsbauer et al. [FGKO17]. That construction is, however, not role-respecting, since a dishonest sender with keys for more than one role can arbitrarily mix the components of the keys for the encryption. Moreover, the construction does not guarantee uniform decryption, because different messages can be encrypted in different components. We fix these issues using the same techniques we used in our construction of the scheme for equality, i.e., we add a signature of the key vector to the encryption keys, sign the ciphertexts, and require a zero-knowledge proof of knowledge that a valid key combination was used to encrypt the same message for each component and that all signatures are valid.

Our construction. Let $\text{ACE}_=$ be an ACE with modification detection scheme for the equality predicate on $\mathcal{D} \times [\ell]$, let Sig be a signature scheme, let F be a PRF, and let NIZK be a NIZK proof of knowledge system for the language $L := \{x \mid \exists w (x, w) \in R\}$, where the relation R is defined as follows: for $x = (vk_{\mathbf{x}}^{\text{Sig}}, c_1, \dots, c_\ell)$ and $w = (ek_{(x_1,1)}^{\text{ACE}_=}, \dots, ek_{(x_\ell,\ell)}^{\text{ACE}_=}, m, r_1, \dots, r_\ell, vk_{\mathbf{x}}^{\text{Sig}}, \sigma_{\mathbf{x}}^{\text{Sig}}, \sigma_c^{\text{Sig}})$, $(x, w) \in R$ if and only if

$$\bigwedge_{i=1}^{\ell} c_i = \text{ACE}_=. \text{Enc}(ek_{(x_i,i)}^{\text{ACE}_=}, m; r_i) \wedge \text{Sig.Ver}(vk_{\mathbf{x}}^{\text{Sig}}, [c_1, \dots, c_\ell], \sigma_c^{\text{Sig}}) = 1 \\ \wedge \text{Sig.Ver}(vk_{\mathbf{x}}^{\text{Sig}}, [ek_{(x_1,1)}^{\text{ACE}_=}, \dots, ek_{(x_\ell,\ell)}^{\text{ACE}_=}, vk_{\mathbf{x}}^{\text{Sig}}], \sigma_{\mathbf{x}}^{\text{Sig}}) = 1.$$

We define an ACE scheme ACE_{DEq} as follows:

Setup: On input a security parameter 1^κ and the policy P_{DEq} , the algorithm $\text{ACE}_{\text{DEq}}.\text{Setup}$ picks a random key K for F and runs

$$(msk^{\text{ACE}_=}, sp^{\text{ACE}_=}) \leftarrow \text{ACE}_=. \text{Setup}(1^\kappa), \\ (vk_{\mathbf{x}}^{\text{Sig}}, sk_{\mathbf{x}}^{\text{Sig}}) \leftarrow \text{Sig.Gen}(1^\kappa), \\ crs^{\text{NIZK}} \leftarrow \text{NIZK.Gen}(1^\kappa).$$

It outputs the master secret key $msk^{\text{ACE}_{\text{DEq}}} := (K, msk^{\text{ACE}_=}, vk_{\mathbf{x}}^{\text{Sig}}, sk_{\mathbf{x}}^{\text{Sig}}, crs^{\text{NIZK}})$ and the sanitizer parameters $sp^{\text{ACE}_{\text{DEq}}} := (sp^{\text{ACE}_=}, vk_{\mathbf{x}}^{\text{Sig}}, crs^{\text{NIZK}})$.

Key generation: The algorithm $\text{ACE}_{\text{DEq}}.\text{Gen}$ on input a master secret key $msk^{\text{ACE}_{\text{DEq}}} = (K, msk^{\text{ACE}_=}, vk_{\mathbf{x}}^{\text{Sig}}, sk_{\mathbf{x}}^{\text{Sig}}, crs^{\text{NIZK}})$, a role $\mathbf{x} \in \mathcal{D}^\ell$, and the type \mathbf{sen} , generates

$$ek_{(x_i,i)}^{\text{ACE}_=} \leftarrow \text{ACE}_=. \text{Gen}(msk^{\text{ACE}_=}, (x_i, i), \mathbf{sen}) \quad (\text{for } i \in [\ell]), \\ (vk_{\mathbf{x}}^{\text{Sig}}, sk_{\mathbf{x}}^{\text{Sig}}) \leftarrow \text{Sig.Gen}(1^\kappa; F_K([\mathbf{x}, 0])), \\ \sigma_{\mathbf{x}}^{\text{Sig}} \leftarrow \text{Sig.Sign}(sk_{\mathbf{x}}^{\text{Sig}}, [ek_{(x_1,1)}^{\text{ACE}_=}, \dots, ek_{(x_\ell,\ell)}^{\text{ACE}_=}, vk_{\mathbf{x}}^{\text{Sig}}]; F_K([\mathbf{x}, 1])),$$

and outputs the encryption key $ek_{\mathbf{x}}^{\text{ACE}_{\text{DEq}}} := (vk_{\mathbf{x}}^{\text{Sig}}, ek_{(x_1,1)}^{\text{ACE}_=}, \dots, ek_{(x_\ell,\ell)}^{\text{ACE}_=}, vk_{\mathbf{x}}^{\text{Sig}}, sk_{\mathbf{x}}^{\text{Sig}}, \sigma_{\mathbf{x}}^{\text{Sig}}, crs^{\text{NIZK}})$; on input $msk^{\text{ACE}_{\text{DEq}}}$, a role $\mathbf{y} \in \mathcal{D}^\ell$, and the type \mathbf{rec} , it generates for $i \in [\ell]$,

$$dk_{(y_i,i)}^{\text{ACE}_=} \leftarrow \text{ACE}_=. \text{Gen}(msk^{\text{ACE}_=}, (y_i, i), \mathbf{rec}),$$

and outputs the decryption key $dk_{\mathbf{y}}^{\text{ACE}_{\text{DEq}}} := (dk_{(y_1,1)}^{\text{ACE}_=}, \dots, dk_{(y_\ell,\ell)}^{\text{ACE}_=})$.

Encryption: On input an encryption key $ek_{\mathbf{x}}^{\text{ACE}_{\text{DEq}}} = (vk^{\text{Sig}}, ek_{(x_1,1)}^{\text{ACE}_{=}}, \dots, ek_{(x_\ell,\ell)}^{\text{ACE}_{=}}, vk_{\mathbf{x}}^{\text{Sig}}, sk_{\mathbf{x}}^{\text{Sig}}, \sigma_{\mathbf{x}}^{\text{Sig}}, crs^{\text{NIZK}})$ and a message $m \in \mathcal{M}^{\text{ACE}_{\text{DEq}}}$, the algorithm $\text{ACE}_{\text{DEq}}.\text{Enc}$ samples randomness r_1, \dots, r_ℓ and computes

$$\begin{aligned} c_i &\leftarrow \text{ACE}_{=}. \text{Enc}(ek_{(x_i,i)}^{\text{ACE}_{=}}, m; r_i) \quad (\text{for } i \in [\ell]), \\ \sigma_c^{\text{Sig}} &\leftarrow \text{Sig}.\text{Sign}(sk_{\mathbf{x}}^{\text{Sig}}, [c_1, \dots, c_\ell]), \\ \pi^{\text{NIZK}} &\leftarrow \text{NIZK}.\text{Prove}(crs^{\text{NIZK}}, x := (vk^{\text{Sig}}, c_1, \dots, c_\ell), \\ &\quad w := (ek_{(x_1,1)}^{\text{ACE}_{=}}, \dots, ek_{(x_\ell,\ell)}^{\text{ACE}_{=}}, m, r_1, \dots, r_\ell, vk_{\mathbf{x}}^{\text{Sig}}, \sigma_{\mathbf{x}}^{\text{Sig}}, \sigma_c^{\text{Sig}})). \end{aligned}$$

It outputs the ciphertext $c := (c_1, \dots, c_\ell, \pi^{\text{NIZK}})$.

Sanitization: On input sanitizer parameters $sp^{\text{ACE}_{\text{DEq}}} = (sp^{\text{ACE}_{=}}, vk^{\text{Sig}}, crs^{\text{NIZK}})$ and a ciphertext $c = (c_1, \dots, c_\ell, \pi^{\text{NIZK}})$, the algorithm $\text{ACE}_{\text{DEq}}.\text{San}$ checks whether $\text{NIZK}.\text{Ver}(crs^{\text{NIZK}}, x := (vk^{\text{Sig}}, c_1, \dots, c_\ell), \pi^{\text{NIZK}}) = 1$. If this is the case, it runs $c'_i \leftarrow \text{ACE}_{=}. \text{San}(c_i)$ for $i \in [\ell]$. If $c'_i \neq \perp$ for all $i \in [\ell]$, it outputs the sanitized ciphertext $c' := (c'_1, \dots, c'_\ell)$. If the verification fails or any of the sanitized ciphertexts is \perp , it outputs \perp .

Decryption: On input a decryption key $dk_{\mathbf{y}}^{\text{ACE}_{\text{DEq}}} = (dk_{(y_1,1)}^{\text{ACE}_{=}}, \dots, dk_{(y_\ell,\ell)}^{\text{ACE}_{=}})$ and a sanitized ciphertext $c' := (c'_1, \dots, c'_\ell)$, the algorithm $\text{ACE}_{\text{DEq}}.\text{Dec}$ computes for $i \in [\ell]$ the message $m_i \leftarrow \text{ACE}_{=}. \text{Dec}(dk_{(y_i,i)}^{\text{ACE}_{=}}, c'_i)$. If $m_i \neq \perp$ for some $i \in [\ell]$, $\text{ACE}_{\text{DEq}}.\text{Dec}$ outputs the first such m_i ; otherwise it outputs \perp .

Modification detection: On input sanitizer parameters $sp^{\text{ACE}_{\text{DEq}}} := (sp^{\text{ACE}_{=}}, vk^{\text{Sig}}, crs^{\text{NIZK}})$ and two ciphertexts $c = (c_1, \dots, c_\ell, \pi^{\text{NIZK}})$ and $\tilde{c} := (\tilde{c}_1, \dots, \tilde{c}_\ell, \tilde{\pi}^{\text{NIZK}})$, the algorithm $\text{ACE}_{\text{DEq}}.\text{DMod}$ checks for $i \in [\ell]$ whether $\text{ACE}_{=}. \text{DMod}(sp^{\text{ACE}_{=}}, c_i, \tilde{c}_i) = 1$. If this is the case for some $i \in [\ell]$, it outputs 1; otherwise, it outputs 0.

Weak and strong anonymity. As we show below, our scheme enjoys weak anonymity. It is easy to see that it does not have strong anonymity: Given a decryption key for the role (1, 2), one can decrypt ciphertexts encrypted under a key for the roles (1, 1) and (2, 2). One does, however, also learn which of the two components decrypted successfully. If it is the first one, the sender role must be (1, 1), if it is the second one, the sender role must be (2, 2). For similar reasons, we do not achieve strong sanitization security.

A similar construction can be used to achieve strong anonymity for less expressive policies: If a sender role still corresponds to a vector $(x_1, \dots, x_\ell) \in \mathcal{D}^\ell$ but a receiver role only to one component $(j, y) \in [\ell] \times \mathcal{D}$, one can consider the policy that allows to receive if $x_j = y$. Now, we do not need several components for the decryption key and the problem sketched above disappears.

Proposition 6.4. *If $\text{ACE}_{=}$ is correct and detectable, then the scheme ACE_{DEq} from above is correct and detectable. If $\text{ACE}_{=}$ is strongly detectable, then ACE_{DEq} is also strongly detectable. More precisely, for all probabilistic algorithms \mathcal{A} , there exist probabilistic algorithms $\mathcal{A}_{\text{corr}}$, $\mathcal{A}_{\text{dtct}}$, $\mathcal{A}'_{\text{dtct}}$, and $\mathcal{A}_{\text{sdtct}}$ such that*

$$\begin{aligned} \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-CORR}} &\leq \text{Adv}_{\text{ACE}_{=}, \mathcal{A}_{\text{corr}}}^{\text{ACE-CORR}} + (\ell - 1) \cdot \text{Adv}_{\text{ACE}_{=}, \mathcal{A}_{\text{dtct}}}^{\text{ACE-DTCT}}, \\ \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-DTCT}} &\leq \ell \cdot \text{Adv}_{\text{ACE}_{=}, \mathcal{A}'_{\text{dtct}}}^{\text{ACE-DTCT}}, \\ \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-sDTCT}} &\leq \ell \cdot \text{Adv}_{\text{ACE}_{=}, \mathcal{A}_{\text{sdtct}}}^{\text{ACE-sDTCT}}. \end{aligned}$$

Proof. We first prove correctness. Let \mathcal{A} be a probabilistic algorithm and let $(m, \mathbf{x}, \mathbf{y})$ with $P_{\text{DEq}}(\mathbf{x}, \mathbf{y}) = 1$ be the output of \mathcal{A} in an execution of $\text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-CORR}}$. Correctness of the signature scheme and completeness of the NIZK imply that the verification in the sanitizer algorithm succeeds with probability 1. Note that $P_{\text{DEq}}(\mathbf{x}, \mathbf{y}) = 1$ implies that there exists $i \in [\ell]$ with $x_i = y_i$. Let i_0 be the first such i . Then, \mathcal{A} only wins the correctness game if either $\text{ACE}_{=}.\text{Dec}(dk_{(y_{i_0}, i_0)}^{\text{ACE}_{=}}, c'_{i_0}) \neq m$, or $\text{ACE}_{=}.\text{Dec}(dk_{(y_i, i)}^{\text{ACE}_{=}}, c'_i) \neq \perp$ for some $i < i_0$. The probability of the former event is bounded by $\text{Adv}_{\text{ACE}_{=}, \mathcal{A}_{\text{corr}}}^{\text{ACE-CORR}}$ where $\mathcal{A}_{\text{corr}}$ emulates this experiment and returns $(m, (x_{i_0}, i_0), (y_{i_0}, i_0))$. For the latter event, note that there are at most $\ell - 1$ such i , so the probability that $\text{ACE}_{=}.\text{Dec}$ returns a message different from \perp for any of them can be bounded by $(\ell - 1) \cdot \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}_{\text{dtct}}}^{\text{ACE-DTCT}}$ for the adversary $\mathcal{A}_{\text{dtct}}$ that emulates the experiment and returns $(m, (x_i, i), (y_i, i))$ for a uniformly chosen $i < i_0$.

For detectability, the adversary $\mathcal{A}'_{\text{dtct}}$ emulates an execution of $\text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-DTCT}}$ and when \mathcal{A} returns $(m, \mathbf{x}, \mathbf{y})$, $\mathcal{A}'_{\text{dtct}}$ outputs $(m, (x_i, i), (y_i, i))$ for a uniformly chosen $i \in \{1, \dots, \ell\}$. Note that \mathcal{A} only wins if $P_{\text{DEq}}(\mathbf{x}, \mathbf{y}) = 0$, which implies that $x_i \neq y_i$ for all $i \in [\ell]$. In this case, \mathcal{A} wins if any of the ciphertext components decrypt to something different from \perp . Thus, $\mathcal{A}'_{\text{dtct}}$ also wins if the component i was guesses correctly, which happens with probability $1/\ell$. The proof for strong detectability is similar, while $\mathcal{A}_{\text{sdctct}}$ additionally outputs the randomness used for encrypting the chosen component when the randomness output by \mathcal{A} is used to generate the whole ciphertext. \square

The following theorem summarizes the security properties we prove for our scheme.

Theorem 6.5. *If F is pseudorandom, NIZK is zero-knowledge and extractable, Sig is EUF-CMA secure, and $\text{ACE}_{=}$ is perfectly correct, strongly detectable, has NDTCT-FENC, and is PRV-CCA, wANON-CCA, SAN-CCA, RR, and UDEC secure, then the scheme ACE_{DEq} from above has NDTCT-FENC and is PRV-CCA, wANON-CCA, SAN-CCA, RR, and UDEC secure.*

We prove this theorem in a sequence of lemmata proving the individual properties. We begin by showing that privacy and weak anonymity of the scheme follow from the corresponding properties of the underlying scheme for equality and the zero-knowledge property of the NIZK. Note that security of the PRF is not needed for this step since it is only used for the signatures, which are irrelevant here.

Lemma 6.6. *Let ACE_{DEq} , be the scheme from above, let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a probabilistic algorithm. Then, there exist probabilistic algorithms \mathcal{A}_{ZK} , \mathcal{A}_{ACE} , \mathcal{A}'_{ZK} , and $\mathcal{A}'_{\text{ACE}}$ (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$) such that*

$$\begin{aligned} \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-PRV-CCA}} &\leq 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + \ell \cdot \text{Adv}_{\text{ACE}_{=}, \mathcal{A}_{\text{ACE}}}^{\text{ACE-PRV-CCA}}, \\ \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-wANON-CCA}} &\leq 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}'_{\text{ZK}}}^{\text{NIZK-ZK}} + \ell \cdot \text{Adv}_{\text{ACE}_{=}, \mathcal{A}'_{\text{ACE}}}^{\text{ACE-wANON-CCA}}. \end{aligned}$$

Proof. We only prove the statement about the privacy advantage. The proof for weak anonymity is completely analogous. We assume without loss of generality that \mathcal{A} ensures $\mathbf{x}^0 = \mathbf{x}^1$ and $P(\mathbf{x}^0, \mathbf{y}) = 0$ for all $\mathbf{y} \in J$, since doing otherwise can only decrease the privacy advantage. Let $H_0 := \text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$ and let H_1 be as H_0 where we replace $\text{crs}^{\text{NIZK}} \leftarrow \text{NIZK.Gen}(1^\kappa)$ by $(\text{crs}^{\text{NIZK}}, \tau^{\text{NIZK}}) \leftarrow S_1^{\text{NIZK}}(1^\kappa)$ in $\text{ACE}_{\text{DEq}}.\text{Setup}$, and for the generation of the challenge ciphertext c^* , we replace $\pi^{\text{NIZK}} \leftarrow \text{NIZK.Prove}(\text{crs}^{\text{NIZK}}, x, w)$ in $\text{ACE}_{\text{DEq}}.\text{Enc}$ by

$\pi^{\text{NIZK}} \leftarrow S_2^{\text{NIZK}}(crs^{\text{NIZK}}, \tau^{\text{NIZK}}, x)$. It can be shown as in the proof of [Lemma 6.3](#) that there exists a probabilistic algorithm \mathcal{A}_{ZK} such that

$$\Pr^{H_0}[b' = b] - \Pr^{H_1}[b' = b] = \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}}. \quad (4)$$

For $k \in \{0, \dots, \ell\}$, we define $H_{2,k}$ as follows. It is identical to H_1 except that after \mathcal{A} returns $(m_0, m_1, \mathbf{x}^0, \mathbf{x}^1, st)$, we replace the ciphertext components in c^* by

$$\begin{aligned} c_i &\leftarrow \text{ACE}_=. \text{Enc}(ek_{(x_i^0, i)}^{\text{ACE}_=}, m_0; r_i) \quad (\text{for } 1 \leq i \leq k), \\ c_i &\leftarrow \text{ACE}_=. \text{Enc}(ek_{(x_i^1, i)}^{\text{ACE}_=}, m_1; r_i) \quad (\text{for } k < i \leq \ell). \end{aligned}$$

Note that $H_{2,0}$ corresponds to H_1 with $b = 1$ and $H_{2,\ell}$ corresponds to H_1 with $b = 0$. Now consider the adversary \mathcal{A}_{ACE} that on input sp chooses $k_0 \leftarrow \{1, \dots, \ell\}$ uniformly at random and emulates an execution of H_1 . It emulates the oracle \mathcal{O}_G by obtaining all the required sub-keys from its own oracle \mathcal{O}_G . To emulate the oracle \mathcal{O}_{SD} , it first checks the NIZK proof as $\text{ACE}_{\text{DEq}}.\text{San}$ and if the verification succeeds, it uses its oracle \mathcal{O}_{SD} to sanitize and decrypt all ciphertext components. As $\text{ACE}_{\text{DEq}}.\text{Dec}$, it outputs the first message different from \perp , or \perp if no such message exists.

When \mathcal{A} returns $(m_0, m_1, \mathbf{x}^0, \mathbf{x}^1, st)$, \mathcal{A}_{ACE} generates the challenge ciphertext c^* by encrypting m_0 under the key $ek_{(x_i^0, i)}^{\text{ACE}_=}$ to obtain c_i for $1 \leq i < k_0$, and by encrypting m_1 under the key $ek_{(x_i^1, i)}^{\text{ACE}_=}$ for $k_0 < i \leq \ell$, where these keys can be obtained from \mathcal{O}_G without changing the advantage. For the k_0 -th component, it returns $(m_0, m_1, x_{k_0}^0, x_{k_0}^1)$ to the challenger and uses the obtained challenge ciphertext as c_{k_0} . It then proceeds with the emulation of H_1 . It emulates the oracle \mathcal{O}_G as above and the oracle \mathcal{O}_{SD^*} as \mathcal{O}_{SD} with the difference that if its own oracle returns **test** for any of the components, it returns **test** as well. Finally, when \mathcal{A}_2 returns b' , \mathcal{A}_{ACE} returns $b'' \leftarrow b'$. Note that if $b = 0$ or $b = 1$, \mathcal{A}_{ACE} perfectly emulates an execution of H_{2,k_0} or H_{2,k_0-1} , respectively. Further note that since \mathcal{A} by assumption does not query \mathcal{O}_G on a decryption key for any \mathbf{y} with $P(\mathbf{x}^0, \mathbf{y}) = 1$, \mathcal{A}_{ACE} also does not ask for a decryption that could decrypt the challenge ciphertext. Hence, \mathcal{A}_{ACE} wins if $b'' = b$ and we have

$$\begin{aligned} \text{Adv}_{\text{ACE}_=, \mathcal{A}_{\text{ACE}}}^{\text{ACE-PRV-CCA}} &= 2 \cdot \Pr^{\text{Exp}_{\text{ACE}_=, \mathcal{A}_{\text{ACE}}}^{\text{ACE-PRV-ANON-CCA}}}[b'' = b] - 1 \\ &= \Pr^{\text{Exp}_{\text{ACE}_=, \mathcal{A}_{\text{ACE}}}^{\text{ACE-PRV-ANON-CCA}}}[b'' = 1 \mid b = 1] - \Pr^{\text{Exp}_{\text{ACE}_=, \mathcal{A}_{\text{ACE}}}^{\text{ACE-PRV-ANON-CCA}}}[b'' = 1 \mid b = 0] \\ &= \sum_{k=1}^{\ell} \frac{1}{\ell} \Pr^{H_{2,k-1}}[b' = 1] - \sum_{k=1}^{\ell} \frac{1}{\ell} \Pr^{H_{2,k}}[b' = 1] \\ &= (\Pr^{H_{2,0}}[b' = 1] - \Pr^{H_{2,\ell}}[b' = 1]) / \ell \\ &= (\Pr^{H_1}[b' = 1 \mid b = 1] - \Pr^{H_1}[b' = 1 \mid b = 0]) / \ell. \end{aligned}$$

We therefore have that $2 \cdot \Pr^{H_1}[b' = b] - 1 = \ell \cdot \text{Adv}_{\text{ACE}_=, \mathcal{A}_{\text{ACE}}}^{\text{ACE-PRV-CCA}}$. Combining this with [equation \(4\)](#) concludes the proof. \square

Next, we sketch how to prove sanitization security.

Lemma 6.7. *If F is pseudorandom, NIZK is extractable, Sig is EUF-CMA secure, and $\text{ACE}_=$ is perfectly correct, strongly detectable, and SAN-CCA secure, then the scheme ACE_{DEq} from above is SAN-CCA secure.*

Proof sketch. The basic idea is to construct an adversary \mathcal{A}_{SAN} against the sanitization security of $\text{ACE}_=$ that chooses $k_0 \leftarrow \{1, \dots, \ell\}$ uniformly at random and emulates an execution of $\text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-SAN-CCA}}$. When \mathcal{A}_1 returns two ciphertexts c_0, c_1 , \mathcal{A}_{SAN} gives the sanitized ciphertext (c'_1, \dots, c'_ℓ) to \mathcal{A}_2 where $c'_i \leftarrow \text{ACE}_=. \text{San}(c_{0,i})$ for $1 \leq i < k_0$, $c'_i \leftarrow \text{ACE}_=. \text{San}(c_{1,i})$ for $k_0 < i \leq \ell$, and c'_{k_0} is obtained from the challenger by submitting (c_{0,k_0}, c_{1,k_0}) . When \mathcal{A}_2 returns the bit b' , \mathcal{A}_{SAN} returns the same bit b' . Note that \mathcal{A}_{SAN} wins if the bit is guessed correctly and if both returned ciphertexts sanitize properly and no decryption key has been obtained that decrypts any of the ciphertexts. If the last two conditions are not satisfied, then also \mathcal{A} does not win. For the hybrid argument to go through, however, we need to ensure that \mathcal{A}_{SAN} still wins with probability $1/2$ when \mathcal{A} violates one of these two conditions. To achieve this, \mathcal{A}_{SAN} needs to detect that this would next happen and in this case abort the emulation, return two valid ciphertexts (if not done already) and guess a uniform bit. To detect this event before it happens, extract witnesses from the ciphertexts returned by \mathcal{A}_1 . If the ciphertexts are valid, the extractions are successful, the signature scheme is EUF-CMA secure, and the PRF is pseudorandom, then the ciphertexts have with overwhelming probability been obtained by encrypting messages with encryption keys that \mathcal{A}_1 has obtained from the oracle \mathcal{O}_G . Hence, \mathcal{A}_{SAN} knows in this case for which roles the messages have been encrypted. When \mathcal{A}_2 asks for a decryption key, \mathcal{A}_{SAN} checks whether the policy allows this key to decrypt any of the two ciphertexts. Given perfect correctness and strong detectability, the decryptions yield \perp if and only if the policy does not allow decryption. Therefore, \mathcal{A}_{SAN} can detect when the bad event is about to happen and abort in this case. \square

Non-detection of fresh encryptions directly follows from the same property of the underlying ACE scheme.

Lemma 6.8. *Let ACE_{DEq} , be the scheme from above and let \mathcal{A} be an attacker on the non-detection of fresh encryptions. Then, there exists a probabilistic algorithm \mathcal{A}' (which is roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}}$) such that*

$$\text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}} \leq \ell \cdot \text{Adv}_{\text{ACE}_=, \mathcal{A}'}^{\text{ACE-NDTCT-FENC}}.$$

Proof. Let \mathcal{A}' emulate an execution of $\text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}}$, using \mathcal{O}_G to answer oracle queries from \mathcal{A} . When \mathcal{A} returns $(m, \mathbf{x}, c = (c_1, \dots, c_\ell, \pi^{\text{NIZK}}))$, \mathcal{A}' chooses $k \leftarrow \{1, \dots, \ell\}$ uniformly at random, and returns $(m, (x_k, k), c_k)$. If \mathcal{A} wins, a fresh encryption of m under \mathbf{x} is detected as a modification of c . Since encryption and modification detection are defined component-wise, this means that there exists a component k_0 such that a fresh encryption of m under (x_{k_0}, k_0) is detected to be a modification of c_{k_0} . Hence, \mathcal{A}' also wins if additionally $k = k_0$, which happens with probability $1/\ell$. \square

We finally prove role-respecting and uniform decryption security.

Lemma 6.9. *Let ACE_{DEq} , be the scheme from above and let \mathcal{A} be a probabilistic algorithm that makes at most at most q_E queries to the oracle \mathcal{O}_E . Then, there exist probabilistic algorithms \mathcal{A}_{PRF} , $\mathcal{A}_{\text{ZK}_1}$, $\mathcal{A}_{\text{ZK}_2}$, \mathcal{A}_{Sig} , and \mathcal{A}_{ACE} (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-URR}}$) such that*

$$\begin{aligned} \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-RR}} + \text{Adv}_{\text{ACE}_{\text{DEq}}, \mathcal{A}}^{\text{ACE-UDEC}} &\leq 2 \cdot \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2} \\ &\quad + 2(q_E + 1) \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + 2\ell \cdot \left(\text{Adv}_{\text{ACE}_=, \mathcal{A}_{\text{ACE}}}^{\text{ACE-RR}} + \text{Adv}_{\text{ACE}_=, \mathcal{A}_{\text{ACE}}}^{\text{ACE-UDEC}} \right). \end{aligned}$$

Proof sketch. As in the proof of Lemma D.5, we define hybrids $H_0 := \text{Exp}_{\text{ACE-DEq}, \mathcal{A}}^{\text{ACE-URR}}$, H_1 as H_0 where F_K is replaced by a uniform random function U , H_2 as H_1 where crs^{NIZK} is generated by E_1^{NIZK} , H_3 as H_2 where a witness $w = (ek_{(x_1,1)}^{\text{ACE=}}, \dots, ek_{(x_\ell, \ell)}^{\text{ACE=}}, m, r_1, \dots, r_\ell, vk_x^{\text{Sig}}, \sigma_x^{\text{Sig}}, \sigma_c^{\text{Sig}})$ is extracted from π^{NIZK} by E_2^{NIZK} after \mathcal{A} returned $c := (c_1, \dots, c_\ell, \pi^{\text{NIZK}})$. We can bound the probability that no valid witness is extracted even though π^{NIZK} is a valid proof by the knowledge extraction advantage of a suitable adversary, and the probability that a valid witness was extracted and the contained encryption key was not obtained via an oracle call by the signature forgery advantage of another adversary as in the proof of Lemma D.5. If these events do not occur, the ciphertext c is an encryption of the message m under a valid key that was returned by \mathcal{O}_G . Hence, \mathcal{A} can in this case only win the role-respecting game or the uniform decryption game if some ciphertext component violates one of these properties. We can construct an adversary \mathcal{A}_{ACE} that emulates the execution, guesses this component, and uses the corresponding ciphertext component to win the game for the underlying scheme for equality. \square

7 Conclusion and Directions for Future Work

In this paper, we have critically revisited existing notions for access control encryption, proposed stronger security definitions, and presented a new scheme that provably achieves our strong requirements. The need for stronger notions is not only a theoretical one as we have shown: In particular, we have described a practical attack based on the observation that a semi-honest sanitizer might leak an unsanitized ciphertext to a dishonest party.

An important question is whether all realistic attacks are excluded by our definitions. Furthermore, we would like to understand the fundamental limits of ACE. This includes investigating in which scenarios it can or cannot be used. To settle these questions, the authors are currently working on a theoretical model to capture the use case of ACE in a simulation-based framework. Another interesting research direction is to find more efficient schemes for useful policies.

A Standard Cryptographic Primitives and Games

A.1 Decisional Diffie-Hellman Assumption

Definition A.1. Let $G = \langle g \rangle$ be a cyclic group of prime-order q and let g be a generator. Let \mathcal{A} be a probabilistic algorithm that on input q, g , and three elements $X, Y, T \in G$ returns a bit d . Let $\text{DDH}_{g, \mathcal{A}}^{\text{real}}$ be the experiment where \mathcal{A} is given two random group elements $X = g^a, Y = g^b$, and the value $T = g^{ab}$. Let $\text{DDH}_{g, \mathcal{A}}^{\text{rand}}$ be the experiment where \mathcal{A} is given three random group elements $X = g^a, Y = g^b$, and $T = g^c$. We define the *decisional Diffie-Hellman (DDH) advantage* of \mathcal{A} as

$$\text{Adv}_{g, \mathcal{A}}^{\text{DDH}} := \Pr^{\text{DDH}_{g, \mathcal{A}}^{\text{real}}}[d = 1] - \Pr^{\text{DDH}_{g, \mathcal{A}}^{\text{rand}}}[d = 1].$$

The *decisional Diffie-Hellman (DDH) assumption* for the group G states that $\text{Adv}_{g, \mathcal{A}}^{\text{DDH}}$ is negligible for all efficient \mathcal{A} .

A.2 Pseudorandom Functions

Definition A.2. For $\kappa \in \mathbb{N}$, let $\mathcal{K}_\kappa, \mathcal{X}_\kappa$, and \mathcal{Y}_κ be finite sets and let $F_\kappa: \mathcal{K}_\kappa \times \mathcal{X}_\kappa \rightarrow \mathcal{Y}_\kappa$ be a function. For $K \in \mathcal{K}_\kappa$, we use the notation $F_K := F_\kappa(K, \cdot)$. Further let \mathcal{A} be a probabilistic

algorithm and consider the experiment in which \mathcal{A} outputs a bit after interacting with an oracle that either corresponds to F_K for a uniformly chosen $K \in \mathcal{K}_\kappa$, or to a uniformly chosen function $U: \mathcal{X}_\kappa \rightarrow \mathcal{Y}_\kappa$. We define the *pseudorandom function advantage* of \mathcal{A} as

$$\text{Adv}_{F,\mathcal{A}}^{\text{PRF}} := \Pr[\mathcal{A}^{F_K(\cdot)}(1^\kappa) = 1] - \Pr[\mathcal{A}^{U(\cdot)}(1^\kappa) = 1],$$

where the first probability is over the random coins of \mathcal{A} and the choice of K , and the second probability is over the random coins of \mathcal{A} and the choice of U . The function family F is called *pseudorandom* if $\text{Adv}_{F,\mathcal{A}}^{\text{PRF}}$ is negligible for all efficient \mathcal{A} .

A.3 Public-Key Encryption

Definition A.3. A *public-key encryption (PKE) scheme* consist of the following three PPT algorithms:

Key generation: The algorithm Gen on input a security parameter 1^κ , outputs a *public key* ek and a *private key* dk .

Encryption: The algorithm Enc on input a public key ek and a message $m \in \mathcal{M}$, outputs a *ciphertext* c .

Decryption: The algorithm Dec on input a private key dk and a ciphertext c , outputs a message $m \in \mathcal{M} \cup \{\perp\}$.

We require for all (ek, dk) in the range of Gen and all $m \in \mathcal{M}$ that

$$\text{Dec}(dk, \text{Enc}(ek, m)) = m$$

with probability 1.

Definition A.4. Let $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a PKE scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Consider the experiment $\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{PKE-IND-CPA}}$ in Figure 4. We define the *ciphertext indistinguishability under chosen-plaintext attacks advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{PKE-IND-CPA}} := 2 \cdot \Pr[b' = b \wedge |m_0| = |m_1|] - 1$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{PKE-IND-CPA}}$. The scheme \mathcal{E} has *indistinguishable ciphertexts under chosen-plaintext attacks (is IND-CPA secure)* if $\text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{PKE-IND-CPA}}$ is negligible for all efficient \mathcal{A} .

A.4 Digital Signature Schemes

Definition A.5. A (*digital*) *signature scheme* consist of the following three PPT algorithms:

Key generation: The algorithm Gen on input a security parameter 1^κ , outputs a *public key* vk and a *private key* sk .

Signing: The algorithm Sign on input a private key sk and a message $m \in \mathcal{M}$, outputs a *signature* σ .

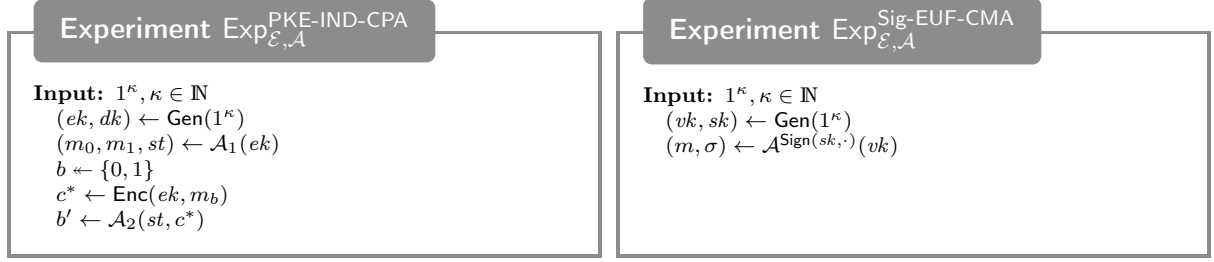


Figure 4: Experiments for the security definitions of public-key encryption and digital signature schemes.

Verification: The algorithm Ver is deterministic and on input a public key vk , a message m , and a signature σ , outputs a bit b (where $b = 1$ means “valid” and $b = 0$ means “invalid”).

We require for all (vk, sk) in the range of Gen and all $m \in \mathcal{M}$ that

$$\text{Ver}(vk, m, \text{Sign}(sk, m)) = 1$$

with probability 1.

Definition A.6. Let $\mathcal{E} = (\text{Gen}, \text{Sign}, \text{Ver})$ be a signature scheme and let \mathcal{A} be a probabilistic algorithm. Consider the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{Sig-EUF-CMA}}$ in Figure 4 and let Q be the set of queries \mathcal{A} issued to its oracle. We define the *existential unforgeability under adaptive chosen-message attacks advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{Sig-EUF-CMA}} := \Pr[\text{Ver}(vk, m, \sigma) = 1 \wedge m \notin Q],$$

where the probability is over the randomness in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{Sig-EUF-CMA}}$. The scheme \mathcal{E} is *existentially unforgeable under adaptive chosen-message attacks (EUF-CMA secure)* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{Sig-EUF-CMA}}$ is negligible for all efficient \mathcal{A} .

A.5 Non-Interactive Zero-Knowledge Proofs

We define non-interactive zero-knowledge proofs following Groth [Gro06].

Definition A.7. Let R be an efficiently computable binary relation and consider the *language* $L := \{x \mid \exists w (x, w) \in R\}$. A *non-interactive proof system* for L (or for R) consists of the following three PPT algorithms:

Key generation: The algorithm Gen on input a security parameter 1^κ , outputs a *common reference string* crs .

Proving: The algorithm Prove on input a common reference string crs , a *statement* x , and a *witness* w , outputs a *proof* π .

Verification: The algorithm Ver on input a common reference string crs , a statement x , and a proof π , outputs a bit b (where $b = 1$ means “accept” and $b = 0$ means “reject”).

We require *perfect completeness*, i.e., for all crs in the range of Gen and for all $(x, w) \in R$, we have

$$\text{Ver}(crs, x, \text{Prove}(crs, x, w)) = 1$$

with probability 1.

Definition A.8 (Soundness). Let $\mathcal{E} = (\text{Gen}, \text{Prove}, \text{Ver})$ be a non-interactive proof system for a language L and let \mathcal{A} be a probabilistic algorithm. We define the *soundness advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{NIZK-snd}} := \Pr^{crs \leftarrow \text{Gen}(1^\kappa); (x, \pi) \leftarrow \mathcal{A}(crs)} [x \notin L \wedge \text{Ver}(crs, x, \pi) = 1].$$

The scheme \mathcal{E} is *computationally sound* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{NIZK-snd}}$ is negligible for all efficient \mathcal{A} and *perfectly sound* if $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{NIZK-snd}} = 0$ for all \mathcal{A} .

Definition A.9 (Computational zero-knowledge). Let $\mathcal{E} = (\text{Gen}, \text{Prove}, \text{Ver})$ be a non-interactive proof system for a relation R and let $S = (S_1, S_2)$ be a pair of PPT algorithms, called *simulator*. Further let $S'(crs, \tau, x, w) = S_2(cr, \tau, x)$ for $(x, w) \in R$, and $S'(crs, \tau, x, w) = \text{failure}$ for $(x, w) \notin R$. We define the *zero-knowledge advantage* of a probabilistic algorithm \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, S, \mathcal{A}}^{\text{NIZK-ZK}} := \Pr^{crs \leftarrow \text{Gen}(1^\kappa)} [\mathcal{A}^{\text{Prove}(crs, \cdot, \cdot)}(crs) = 1] - \Pr^{(crs, \tau) \leftarrow S_1(1^\kappa)} [\mathcal{A}^{S'(crs, \tau, \cdot, \cdot)}(crs) = 1].$$

We call $(\text{Gen}, \text{Prove}, \text{Ver}, S_1, S_2)$ a *non-interactive zero-knowledge (NIZK) proof system* for R if $\text{Adv}_{\mathcal{E}, S, \mathcal{A}}^{\text{NIZK-ZK}}$ is negligible for all efficient \mathcal{A} ; it is called *single-theorem NIZK proof system* if $\text{Adv}_{\mathcal{E}, S, \mathcal{A}}^{\text{NIZK-ZK}}$ is negligible for all efficient \mathcal{A} that make at most one query to their oracle.

Definition A.10 (Knowledge extraction). Let $\mathcal{E} = (\text{Gen}, \text{Prove}, \text{Ver})$ be a non-interactive proof system for a relation R and let $E = (E_1, E_2)$ be a pair of PPT algorithms, called *knowledge extractor*. We define the *knowledge extraction advantages* of a probabilistic algorithm \mathcal{A} as

$$\begin{aligned} \text{Adv}_{\mathcal{E}, E, \mathcal{A}}^{\text{NIZK-ext}_1} &:= \Pr^{crs \leftarrow \text{Gen}(1^\kappa)} [\mathcal{A}(crs) = 1] - \Pr^{(crs, \xi) \leftarrow E_1(1^\kappa)} [\mathcal{A}(crs) = 1], \\ \text{Adv}_{\mathcal{E}, E, \mathcal{A}}^{\text{NIZK-ext}_2} &:= \Pr^{(crs, \xi) \leftarrow E_1(1^\kappa); (x, \pi) \leftarrow \mathcal{A}(crs); w \leftarrow E_2(cr, \xi, x, \pi)} [\text{Ver}(crs, x, \pi) = 1 \wedge (x, w) \notin R]. \end{aligned}$$

We call $(\text{Gen}, \text{Prove}, \text{Ver}, E_1, E_2)$ a *non-interactive proof of knowledge system* for R if $\text{Adv}_{\mathcal{E}, E, \mathcal{A}}^{\text{NIZK-ext}_1}$ and $\text{Adv}_{\mathcal{E}, E, \mathcal{A}}^{\text{NIZK-ext}_2}$ are negligible for all efficient \mathcal{A} .

Definition A.11 (Simulation soundness). Let $\mathcal{E} = (\text{Gen}, \text{Prove}, \text{Ver})$ be a non-interactive proof system for a language L , let $S = (S_1, S_2)$ be a pair of PPT algorithms, and let \mathcal{A} be a probabilistic algorithm. Consider the experiment $\text{Exp}_{\mathcal{E}, S, \mathcal{A}}^{\text{NIZK-sim-snd}}$ that executes $(crs, \tau) \leftarrow S_1(1^\kappa)$ and $(x, \pi) \leftarrow \mathcal{A}^{S_2(cr, \tau, \cdot)}(crs)$. Further let Q be the set of all (x', π') such that \mathcal{A} queried x' to its oracle and received π' as a response. We define the *simulation soundness advantage* of \mathcal{A} as

$$\text{Adv}_{\mathcal{E}, S, \mathcal{A}}^{\text{NIZK-sim-snd}} := \Pr[(x, \pi) \notin Q \wedge x \notin L \wedge \text{Ver}(crs, x, \pi) = 1].$$

We say $(\text{Gen}, \text{Prove}, \text{Ver}, S_1, S_2)$ is *simulation sound* if $\text{Adv}_{\mathcal{E}, S, \mathcal{A}}^{\text{NIZK-sim-snd}}$ is negligible for all efficient \mathcal{A} ; it is *one-time simulation sound* if $\text{Adv}_{\mathcal{E}, S, \mathcal{A}}^{\text{NIZK-sim-snd}}$ is negligible for all efficient \mathcal{A} that make at most one query to the oracle S_2 .

Note that in the above definition, \mathcal{A} is allowed to issue queries $x' \notin L$ to its oracle. This means that soundness is preserved even if an adversary sees simulated proofs of false statements.

B Proof of Relation to the Original Security Notions

We prove the three claims in [Theorem 4.7](#) as separate lemmata, starting with the payload-privacy no-read rule.

Lemma B.1. *Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let $\mathcal{E}' = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be the corresponding ACE scheme. Further let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Then, there exist adversaries \mathcal{A}_{PRV} and $\mathcal{A}_{\text{WANON}}$ (both roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-read}}$) such that*

$$\text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,priv}} \leq \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{PRV}}}^{\text{ACE-PRV-CCA}} + \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}}.$$

Proof. We assume without loss of generality that \mathcal{A} ensures $|m_0| = |m_1|$ and $P(i_0, j) = P(i_1, j) = 0$ for all $j \in J$, where J is the set of all j such that \mathcal{A}_1 or \mathcal{A}_2 issued the query (j, rec) to the oracle \mathcal{O}_G . Let H be identical to $\text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,priv}}$ except that after \mathcal{A}_1 returns (m_0, m_1, i_0, i_1, st) , i_1 is replaced by i_0 . We first show that the probability that b is guessed correctly in H and $\text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,priv}}$ differ only negligibly if the scheme satisfies weak anonymity. Note that if $b = 0$, the two experiments are identical, which implies

$$\Pr^{\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}} [b' = b \mid b = 0] = \Pr^H [b' = b \mid b = 0]. \quad (5)$$

Claim 1. *There exists an adversary $\mathcal{A}_{\text{WANON}}$ such that*

$$\Pr^{\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}} [b' = b \mid b = 1] - \Pr^H [b' = b \mid b = 1] = \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}}.$$

Proof of claim. We construct $\mathcal{A}_{\text{WANON}}$ as follows. On input sp , it emulates an execution of $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$, where the oracles for \mathcal{A} are emulated as follows.

$\mathcal{O}_G(\cdot, \cdot)$: Relay queries to the oracle \mathcal{O}_G of $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-PRV-ANON-CCA}}$.

$\mathcal{O}_E(\cdot, \cdot)$: On query (j, m) , query (j, sen) to the oracle \mathcal{O}_G to receive the encryption key ek_j . Then compute $c \leftarrow \text{Enc}(ek_j, m)$ and return c .

When \mathcal{A} outputs (m_0, m_1, i_0, i_1, st) , $\mathcal{A}_{\text{WANON}}$ gives (m_1, m_1, i_0, i_1) to the challenger to obtain a ciphertext c^* , which is given to \mathcal{A}_2 . When \mathcal{A}_2 returns b' , $\mathcal{A}_{\text{WANON}}$ returns the same bit b' . Note that if $b = 0$, $\mathcal{A}_{\text{WANON}}$ perfectly emulates H with $b = 1$, and if $b = 1$, $\mathcal{A}_{\text{WANON}}$ perfectly emulates $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$ with $b = 1$. Hence,

$$\begin{aligned} & \Pr^{\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}} [b' = b \mid b = 1] - \Pr^H [b' = b \mid b = 1] \\ &= \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = 1 \mid b = 1] - \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = 1 \mid b = 0] \\ &= 2 \cdot \left(\frac{1}{2} \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = b \mid b = 1] - \frac{1}{2} \left(1 - \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = b \mid b = 0] \right) \right) \\ &= 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = b] - 1. \end{aligned}$$

Note that $\mathcal{A}_{\text{WANON}}$ returns the same message m_1 twice and we have $P(i_0, j) = P(i_1, j) = 0$ for all $j \in J$ by the assumption on \mathcal{A} . This implies $\text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}} = 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = b] - 1$ and concludes the proof of the claim. \diamond

Combining Claim 1 and equation (5), we obtain

$$\begin{aligned}
\Pr^{\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}} [b' = b] &= \frac{1}{2} \cdot \Pr^{\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}} [b' = b \mid b = 0] + \frac{1}{2} \cdot \Pr^{\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}} [b' = b \mid b = 1] \\
&= \frac{1}{2} \cdot \Pr^H [b' = b \mid b = 0] + \frac{1}{2} \cdot \left(\text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}} + \Pr^H [b' = b \mid b = 1] \right) \\
&= \Pr^H [b' = b] + \frac{1}{2} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}}.
\end{aligned}$$

Hence,

$$\text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,priv}} = 2 \cdot \Pr^H [b' = b] - 1 + \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}}. \quad (6)$$

We now construct the adversary \mathcal{A}_{PRV} . When invoked on input sp , it starts an emulation of H by passing sp to \mathcal{A} . The oracles for \mathcal{A} are emulated as in the proof of Claim 1. When \mathcal{A}_1 returns (m_0, m_1, i_0, i_1, st) , $\mathcal{A}_{\text{WANON}}$ gives (m_0, m_1, i_0, i_1) to the challenger to obtain a ciphertext c^* , which is then given to \mathcal{A}_2 . When \mathcal{A}_2 returns b' , \mathcal{A}_{PRV} returns the same bit b' . Note that the view of \mathcal{A} in this emulation is identical to its view in H . Since \mathcal{A}_{PRV} returns the same role i_0 twice and $P(i_0, j) = 0$ for all $j \in J$ by the assumption on \mathcal{A} , we have

$$\text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{PRV}}}^{\text{ACE-PRV-CCA}} = 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{PRV}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = b] - 1 = 2 \cdot \Pr^H [b' = b] - 1.$$

Using equation (6), we conclude

$$\text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,priv}} = \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{PRV}}}^{\text{ACE-PRV-CCA}} + \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{WANON}}}^{\text{ACE-wANON-CCA}}. \quad \square$$

We next show that the sender-anonymity no-read rule is implied by strong sender anonymity.

Lemma B.2. *Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let $\mathcal{E}' = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be the corresponding ACE scheme. Further let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Then, there exists an adversary $\mathcal{A}_{\text{sANON}}$ (roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}', \mathcal{A}'}^{\text{ACE-no-read}}$) such that*

$$\text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,anon}} = \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{sANON}}}^{\text{ACE-sANON-CCA}}.$$

Proof. We construct $\mathcal{A}_{\text{sANON}}$ as follows. On input sp , it emulates an execution of $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$, where the oracles \mathcal{O}_G and \mathcal{O}_E for \mathcal{A} are emulated as in the proof of Lemma B.1. When \mathcal{A}_1 returns (m_0, m_1, i_0, i_1, st) , $\mathcal{A}_{\text{sANON}}$ gives (m_0, m_1, i_0, i_1) to the challenger to obtain the ciphertext c^* . Then, \mathcal{A}_2 is invoked on input (st, c^*) and the oracles are emulated as before. When \mathcal{A}_2 terminates with output b' , $\mathcal{A}_{\text{sANON}}$ returns the same bit b' . We observe that the view $\mathcal{A}_{\text{sANON}}$ emulates toward \mathcal{A} is identical to the view of \mathcal{A} in the experiment $\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}$. Thus,

$$\begin{aligned}
&\text{Adv}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read,anon}} \\
&= 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}', \mathcal{A}}^{\text{ACE-no-read}}} [b' = b \wedge m_0 = m_1 \wedge \forall j \in J P(i_0, j) = P(i_1, j)] - 1 \\
&= 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sANON}}}^{\text{ACE-PRV-ANON-CCA}}} [b' = b \wedge m_0 = m_1 \wedge \forall j \in J P(i_0, j) = P(i_1, j)] - 1 \\
&= \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{sANON}}}^{\text{ACE-sANON-CCA}}. \quad \square
\end{aligned}$$

To conclude the proof of Theorem 4.7, we prove the claim about the no-write with modification detection rule.

Lemma B.3. Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms. Then, there exist adversaries \mathcal{A}_{SAN} , \mathcal{A}_{RR} , and $\mathcal{A}_{\text{CORR}}$ (all roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$) such that for policies P where for all i , one can efficiently find some j with $P(i, j) = 1$,

$$\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}} \leq \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{SAN}}}^{\text{ACE-SAN-CCA}} + 4 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{RR}}}^{\text{ACE-RR}} + 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{CORR}}}^{\text{ACE-CORR}}.$$

Proof. We first construct the adversary \mathcal{A}_{SAN} . When invoked on input sp , it gives sp to \mathcal{A}_1 and emulates $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$. The oracles for \mathcal{A} are emulated as follows.

$\mathcal{O}_G(\cdot, \cdot)$: Relay queries to the oracle \mathcal{O}_G of $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{SAN}}}^{\text{ACE-SAN-CCA}}$.

$\mathcal{O}_{ES}(\cdot, \cdot)$: On query (j, m) , \mathcal{A}_{SAN} queries (j, sen) to its oracle \mathcal{O}_G to receive the encryption key ek_j .⁵ It then computes $c' \leftarrow \text{San}(sp, \text{Enc}(ek_j, m))$ and outputs c' to \mathcal{A} .

When \mathcal{A}_1 outputs (c_0, i', st) , \mathcal{A}_{SAN} chooses a uniformly random message $m \leftarrow \mathcal{M}$, queries (i', sen) to its oracle \mathcal{O}_G to receive the encryption key $ek_{i'}$, and computes $c_1 \leftarrow \text{Enc}(ek_{i'}, m)$. Then, \mathcal{A}_{SAN} gives (c_0, c_1) to the challenger to obtain a sanitized ciphertext c'_b . It then invokes \mathcal{A}_2 on input (st, c'_b) and emulates the oracles as above. When \mathcal{A}_2 outputs its guess b' , \mathcal{A}_{SAN} outputs the same bit b' as its own guess. Note that the view \mathcal{A}_{SAN} emulates toward \mathcal{A} is identical to the view of \mathcal{A} in the experiment $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$. Let W_{noW} and W_{san} be the events that \mathcal{A} wins in the no-write with modification detection experiment and \mathcal{A}_{SAN} wins the sanitization experiment, respectively, i.e.,

$$\begin{aligned} W_{\text{noW}} &:= [b' = b \wedge \text{dct} = \text{false} \wedge i' \in I_1 \\ &\quad \wedge \forall i \in I_1 \forall j \in J P(i, j) = 0 \wedge \text{San}(sp, c_0) \neq \perp], \\ W_{\text{san}} &:= [b' = b \wedge c'_0 \neq \perp \neq c'_1 \wedge \forall j \in J m_{0,j} = m_{1,j} = \perp]. \end{aligned}$$

Further consider the events

$$\begin{aligned} C &:= [\text{San}(sp, c_1) \neq \perp], \\ R &:= [\forall j \in J \text{Dec}(\text{Gen}(msk, j, \text{rec}), \text{San}(sp, c_0)) = \text{Dec}(\text{Gen}(msk, j, \text{rec}), \text{San}(sp, c_1)) = \perp]. \end{aligned}$$

We then have

$$\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{SAN}}}^{\text{ACE-SAN-CCA}}} [W_{\text{san}}] \geq \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}} \cap C \cap R]. \quad (7)$$

We next show that the events $\neg C$ and $\neg R$ only occur with negligible probability if the ACE scheme is correct and role-respecting, respectively.

Claim 1. *There exists an adversary $\mathcal{A}_{\text{CORR}}$ (roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$) such that*

$$\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [\neg C] \leq \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{CORR}}}^{\text{ACE-CORR}}.$$

Proof of claim. On input sp , the adversary $\mathcal{A}_{\text{CORR}}$ begins an emulation of $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$ as \mathcal{A}_{SAN} above. When \mathcal{A}_1 outputs (c_0, i', st) , $\mathcal{A}_{\text{CORR}}$ chooses a uniformly random message $m \leftarrow \mathcal{M}$ and finds j with $P(i', j) = 1$. It finally returns (m, i', j) . By definition of Dec , we have $\text{Dec}(\text{Gen}(msk, j, \text{rec}), \perp) = \perp$. Hence, if $\neg C$ occurs, then encrypting m for role i' and sanitizing and decrypting the result yields $\perp \neq m$. Therefore, $\mathcal{A}_{\text{CORR}}$ wins the correctness game in this case, which implies the claim. \diamond

⁵Looking ahead, we note that obtaining additional encryption keys is not problematic in the sanitization game, since the winning condition does not restrict the obtained encryption keys.

Claim 2. *There exists an adversary \mathcal{A}_{RR} (roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$) such that*

$$\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}} \cap C \cap \neg R] \leq 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{RR}}^{\text{ACE-RR}}.$$

Proof of claim. When invoked on input sp , \mathcal{A}_{RR} internally emulates an execution of \mathcal{A} on input sp and emulates the oracles as follows.

$\mathcal{O}_G(\cdot, \cdot)$: Relay queries to the oracle \mathcal{O}_G of $\text{Exp}_{\mathcal{E}, \mathcal{A}_{RR}}^{\text{ACE-URR}}$.

$\mathcal{O}_{ES}(\cdot, \cdot)$: On query (j, m) , query (j, m) to \mathcal{O}_E to receive the ciphertext c . Then, compute $c' \leftarrow \text{San}(sp, c)$ and return c' .

When \mathcal{A}_1 outputs (c_0, i', st) , \mathcal{A}_{RR} chooses a uniformly random message $m \leftarrow \mathcal{M}$, queries (i', sen) to its oracle \mathcal{O}_G to receive the encryption key $ek_{i'}$, and computes $c_1 \leftarrow \text{Enc}(ek_{i'}, m)$. Then, \mathcal{A}_{RR} chooses $c \leftarrow \{c_0, c_1\}$ uniformly at random and outputs c to the challenger. Let W_{noW} be the event that \mathcal{A}_{RR} wins the role-respecting game, i.e.,

$$W_{RR} := [c' \neq \perp \wedge \text{dct} = \text{false} \wedge \neg(\exists i \in I \forall j \in J (m_j \neq \perp \leftrightarrow P(i, j) = 1))].$$

Note that W_{noW} and C imply that $c' \neq \perp$, $\text{dct} = \text{false}$, and $\forall i \in I \forall j \in J P(i, j) = 0$. Hence, if we additionally have $\neg R$, at least one of the two possible choices for c yield $m_j \neq \perp$ for some $j \in J$, and thus

$$\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}} \cap C \cap \neg R] \leq 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{RR}}^{\text{ACE-URR}}} [W_{RR}] = 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{RR}}^{\text{ACE-RR}}. \quad \diamond$$

Combining [equation \(7\)](#) and [Claims 1](#) and [2](#), we obtain

$$\begin{aligned} \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}}] &\leq \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}} \cap C \cap R] \\ &\quad + \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}} \cap C \cap \neg R] + \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [\neg C] \\ &\leq \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{SAN}}^{\text{ACE-SAN-CCA}}} [W_{\text{san}}] + 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{RR}}^{\text{ACE-RR}} + \text{Adv}_{\mathcal{E}, \mathcal{A}_{CORR}}^{\text{ACE-CORR}}. \end{aligned}$$

We can thus conclude

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}} &= 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}}] - 1 \\ &\leq 2 \cdot \underbrace{\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{SAN}}^{\text{ACE-SAN-CCA}}} [W_{\text{san}}] - 1}_{= \text{Adv}_{\mathcal{E}, \mathcal{A}_{SAN}}^{\text{ACE-SAN-CCA}}} + 4 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{RR}}^{\text{ACE-RR}} + 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{CORR}}^{\text{ACE-CORR}}. \quad \square \end{aligned}$$

We finally prove [Theorem 4.8](#), which we first restate.

Theorem 4.8. *Let $\mathcal{E} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec}, \text{DMod})$ be an ACE with modification detection scheme such that $\Pr[\text{DMod}(sp, c_0, c_1) = 1] = \Pr[\text{DMod}(sp, c_1, c_0) = 1]$ for all sp returned by Setup and all ciphertexts $c_0, c_1 \in \mathcal{C}$. Further let $\mathcal{E}' = (\text{Setup}, \text{Gen}, \text{Enc}, \text{San}, \text{Dec})$ be the corresponding ACE scheme. If \mathcal{E} is correct, detectable, has NDTCT-FENC, and is sSAN-CCA and RR secure, then \mathcal{E}' satisfies the no-write rule for policies P such that for all i , one can efficiently find some j with $P(i, j) = 1$. More precisely, for all adversaries \mathcal{A} that make at most q_{ES} queries to the oracle \mathcal{O}_{ES} and at most q_{dk} queries of the form (\cdot, rec) to \mathcal{O}_G , there exist adversaries \mathcal{A}_{SAN} ,*

\mathcal{A}_{RR} , \mathcal{A}_{sSAN} , \mathcal{A}_{NDTCT} , \mathcal{A}_{CORR} , and \mathcal{A}_{dtct} (all roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-no-write}}$) such that

$$\begin{aligned} \text{Adv}_{\mathcal{E}',\mathcal{A}}^{\text{ACE-no-write}} &\leq \text{Adv}_{\mathcal{E},\mathcal{A}_{sSAN}}^{\text{ACE-SAN-CCA}} + 4 \cdot \text{Adv}_{\mathcal{E},\mathcal{A}_{RR}}^{\text{ACE-RR}} + 2q_{ES} \cdot \text{Adv}_{\mathcal{E},\mathcal{A}_{sSAN}}^{\text{ACE-sSAN-CCA}} \\ &\quad + 4q_{ES} \cdot \text{Adv}_{\mathcal{E},\mathcal{A}_{NDTCT}}^{\text{ACE-NDTCT-FENC}} + (8q_{ES}q_{dk} + 2) \cdot \text{Adv}_{\mathcal{E},\mathcal{A}_{CORR}}^{\text{ACE-CORR}} + 8q_{ES}q_{dk} \cdot \text{Adv}_{\mathcal{E},\mathcal{A}_{dtct}}^{\text{ACE-DTCT}}. \end{aligned}$$

Proof. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary and consider $\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}}$. Let W_{noW} and $W_{\text{MD-noW}}$ be the events that \mathcal{A} wins the no-write and no-write with modification detection game, respectively. Note that we have $W_{\text{MD-noW}} = W_{\text{noW}} \cap [\text{dct} = \text{false}]$. Hence,

$$\begin{aligned} \text{Adv}_{\mathcal{E}',\mathcal{A}}^{\text{ACE-no-write}} &= 2 \cdot \Pr^{\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}}] - 1 \\ &= 2 \left(\Pr^{\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{MD-noW}}] + \Pr^{\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}}} [W_{\text{noW}} \cap [\text{dct} = \text{true}]] \right) - 1 \\ &\leq \text{Adv}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}} + 2 \cdot \Pr^{\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}}} [\text{dct} = \text{true}]. \end{aligned}$$

Lemma B.3 implies that there exist adversaries \mathcal{A}_{sSAN} , \mathcal{A}_{RR} , and \mathcal{A}'_{CORR} (all roughly as efficient as emulating an execution of $\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}}$) such that

$$\begin{aligned} \text{Adv}_{\mathcal{E}',\mathcal{A}}^{\text{ACE-no-write}} &\leq \text{Adv}_{\mathcal{E},\mathcal{A}_{sSAN}}^{\text{ACE-SAN-CCA}} + 4 \cdot \text{Adv}_{\mathcal{E},\mathcal{A}_{RR}}^{\text{ACE-RR}} + 2 \cdot \text{Adv}_{\mathcal{E},\mathcal{A}'_{CORR}}^{\text{ACE-CORR}} \\ &\quad + 2 \cdot \Pr^{\text{Exp}_{\mathcal{E},\mathcal{A}}^{\text{ACE-MD-no-write}}} [\text{dct} = \text{true}]. \quad (8) \end{aligned}$$

To bound the probability of $[\text{dct} = \text{true}]$, we construct the adversary \mathcal{A}_{sSAN} . When invoked on input sp , it first chooses $q_0 \leftarrow \{1, \dots, q_{ES}\}$ uniformly at random, sets $k \leftarrow 1$ and internally emulates an execution of \mathcal{A} on input sp . Oracle queries by \mathcal{A} are answered as follows:

$\mathcal{O}_G(\cdot, \cdot)$: Relay queries to the oracle \mathcal{O}_G of $\text{Exp}_{\mathcal{E},\mathcal{A}_{sSAN}}^{\text{ACE-SAN-CCA}}$.

$\mathcal{O}_{ES}(\cdot, \cdot)$: On query (i, m) , if $k \neq q_0$, \mathcal{A}_{sSAN} queries (i, sen) to its oracle \mathcal{O}_G to receive the encryption key ek_i . It then computes $c' \leftarrow \text{San}(sp, \text{Enc}(ek_i, m))$ and outputs c' to \mathcal{A} . Finally, it sets $k \leftarrow k + 1$.

If $k = q_0$, then \mathcal{A}_{sSAN} queries (i, sen) to its oracle \mathcal{O}_G to receive the encryption key ek_i . It then creates two independent encryptions of m by computing $\tilde{c}_0 \leftarrow \text{Enc}(ek_i, m)$ and $\tilde{c}_1 \leftarrow \text{Enc}(ek_i, m)$, sets $i_{q_0} \leftarrow i$, $m_{q_0} \leftarrow m$, $k \leftarrow k + 1$, and gives \tilde{c}_0, \tilde{c}_1 to the challenger to obtain \tilde{c}'_b .

If \mathcal{A}_1 terminates before $k = q_0$ is reached, \mathcal{A}_{sSAN} gives two fresh encryptions of some fixed message m_{q_0} for a fixed role i_{q_0} to the challenger and then returns a uniform bit $b' \leftarrow \{0, 1\}$. Otherwise, when \mathcal{A}_1 returns i' and c_0 , \mathcal{A}_{sSAN} evaluates $d_0 \leftarrow \text{DMod}(sp, \tilde{c}_0, c_0)$ and $d_1 \leftarrow \text{DMod}(sp, \tilde{c}_1, c_0)$. If $d_0 = d_1$, then \mathcal{A}_{sSAN} also returns a uniform bit; if $d_b = 1$ for exactly one $b' \in \{0, 1\}$, \mathcal{A}_{sSAN} returns b' .

Let Q be the event that $d_0 = 1$ or $d_1 = 1$ and let D be the event that $d_{1-b} = 1$. Note that if Q and $\neg D$ occur, \mathcal{A}_{sSAN} returns the correct bit $b' = b$. Moreover, if Q does not occur, \mathcal{A}_{sSAN} returns a uniform bit. Hence,

$$\begin{aligned} \Pr[b' = b] &= \Pr[[b' = b] \cap Q \cap \neg D] + \Pr[[b' = b] \cap \neg(Q \cap \neg D)] \\ &\geq \Pr[Q \cap \neg D] + \Pr[[b' = b] \cap \neg Q] \\ &= \Pr[Q \cap \neg D] + \Pr[b' = b \mid \neg Q] \cdot \Pr[\neg Q] \\ &= \Pr[Q \cap \neg D] + \frac{1}{2} \cdot (1 - \Pr[Q]), \end{aligned}$$

where all probabilities are in $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}$. This implies

$$\Pr[Q] = \Pr[Q \cap D] + \Pr[Q \cap \neg D] \leq \Pr[D] + \Pr[b' = b] - \frac{1}{2} + \frac{1}{2} \Pr[Q]$$

and therefore

$$\Pr[Q] \leq 2 \cdot \Pr[D] + 2 \cdot \Pr[b' = b] - 1.$$

Note that if $[\text{dct} = \text{true}]$ occurs in $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}$, then Q occurs in $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}$ with probability at least $1/q_{ES}$. We can therefore conclude that

$$\begin{aligned} \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}}[\text{dct} = \text{true}] &\leq q_{ES} \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[Q] \\ &\leq 2q_{ES} \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[D] + q_{ES} \left(2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[b' = b] - 1 \right). \end{aligned}$$

Let W_{sSAN} be the event that $\mathcal{A}_{\text{sSAN}}$ wins the strong sanitization game and consider the events

$$\begin{aligned} C &:= [\tilde{c}'_0 \neq \perp \neq \tilde{c}'_1 \wedge \forall j \in J (P(i_{q_0}, j) = 1 \rightarrow m_{0,j} = m_{1,j} = m_{q_0})], \\ R &:= [\forall j \in J (P(i_{q_0}, j) = 0 \rightarrow m_{0,j} = m_{1,j} = \perp)]. \end{aligned}$$

We then have that $[b' = b]$, C , and R together imply W_{sSAN} . Thus,

$$\begin{aligned} \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[b' = b] &= \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[[b' = b] \cap C \cap R] + \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[[b' = b] \cap \neg(C \cap R)] \\ &\leq \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[W_{\text{sSAN}}] + \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[\neg C \cup \neg R]. \end{aligned}$$

Together with the previous result, this yields

$$\begin{aligned} &\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}}[\text{dct} = \text{true}] \\ &\leq 2q_{ES} \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[D] + q_{ES} \left[2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[W_{\text{sSAN}}] - 1 + 2 \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[\neg C \cup \neg R] \right] \\ &= 2q_{ES} \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[D] + q_{ES} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-sSAN-CCA}} + 2q_{ES} \cdot \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[\neg C \cup \neg R]. \end{aligned}$$

Now consider $\mathcal{A}_{\text{NDTCT}}$ that emulates $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}$ and outputs (m_{q_0}, i_{q_0}, c_0) . Note that the view of \mathcal{A}_1 in the emulation is independent of \tilde{c}_{1-b} . One can therefore assume that \tilde{c}_{1-b} is generated after \mathcal{A}_1 outputs c_0 , as c^* in $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{NDTCT}}}^{\text{ACE-NDTCT-FENC}}$. By assumption, we have $\Pr[\text{DMod}(sp, \tilde{c}_{1-b}, c_0) = 1] = \Pr[\text{DMod}(sp, c_0, \tilde{c}_{1-b}) = 1]$, and therefore

$$\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[D] = \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{NDTCT}}}^{\text{ACE-NDTCT-FENC}}.$$

Further consider $\mathcal{A}_{\text{CORR}}''$ that emulates $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}$ and if there exists $j \in J$ with $P(i_{q_0}, j) = 1$, it chooses $j \leftarrow J$ uniformly at random and outputs (m_{q_0}, i_{q_0}, j) . If such $j \in J$ does not exist, $\mathcal{A}_{\text{CORR}}''$ finds $j \in [n]$ with $P(i_{q_0}, j) = 1$ and then outputs (m_{q_0}, i_{q_0}, j) . Note that $m_{0,j} = m_{1,j} = m_{q_0}$ implies $\tilde{c}'_0 \neq \perp \neq \tilde{c}'_1$ since \perp decrypts to \perp . Hence, if such $j \in J$ exists and $\neg C$ occurs, $\mathcal{A}_{\text{CORR}}''$ wins the correctness game with probability at least $1/(2|J|) \geq 1/(2q_{dk})$; and if no such $j \in J$ exists and $\neg C$ occurs, $\mathcal{A}_{\text{CORR}}''$ wins with probability at least $1/2$. The factor $1/2$ is due to the fact that the message is encrypted twice in $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}$ but only once in the correctness experiment. Overall, we get

$$\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[\neg C] \leq 2q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{CORR}}''}^{\text{ACE-CORR}}.$$

Finally consider $\mathcal{A}_{\text{dtct}}$ that emulates $\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}$, chooses $j \leftarrow J$ uniformly at random and outputs (m_{q_0}, i_{q_0}, j) . If $\neg R$ occurs, $\mathcal{A}_{\text{dtct}}$ wins the detectability game with probability at least $1/(2q_{dk})$. Hence,

$$\Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}}}[\neg R] \leq 2q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{dtct}}}^{\text{ACE-DTCT}}.$$

Combining our results, we obtain

$$\begin{aligned} \Pr^{\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-MD-no-write}}}[\text{dct} = \text{true}] &\leq 2q_{ES} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{NDTCT}}}^{\text{ACE-NDTCT-FENC}} + q_{ES} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-sSAN-CCA}} \\ &\quad + 4q_{ES}q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}'_{\text{CORR}}}^{\text{ACE-CORR}} + 4q_{ES}q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{dtct}}}^{\text{ACE-DTCT}}. \end{aligned}$$

Together with [equation \(8\)](#), this yields

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{ACE-no-write}} &\leq \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-SAN-CCA}} + 4 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{RR}}}^{\text{ACE-RR}} + 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}'_{\text{CORR}}}^{\text{ACE-CORR}} + 4q_{ES} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{NDTCT}}}^{\text{ACE-NDTCT-FENC}} \\ &\quad + 2q_{ES} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{sSAN}}}^{\text{ACE-sSAN-CCA}} + 8q_{ES}q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}'_{\text{CORR}}}^{\text{ACE-CORR}} + 8q_{ES}q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{dtct}}}^{\text{ACE-DTCT}}. \end{aligned}$$

For the adversary $\mathcal{A}_{\text{CORR}}$ that runs $\mathcal{A}'_{\text{CORR}}$ with probability $\frac{2}{8q_{ES}q_{dk}+2}$ and $\mathcal{A}''_{\text{CORR}}$ with probability $\frac{8q_{ES}q_{dk}}{8q_{ES}q_{dk}+2}$, we have $(8q_{ES}q_{dk} + 2) \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}_{\text{CORR}}}^{\text{ACE-CORR}} = 2 \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}'_{\text{CORR}}}^{\text{ACE-CORR}} + 8q_{ES}q_{dk} \cdot \text{Adv}_{\mathcal{E}, \mathcal{A}''_{\text{CORR}}}^{\text{ACE-CORR}}$ and the claim of the theorem follows. \square

C Proofs of Privacy and Anonymity of the sPKE Scheme

To complete the proof of [Theorem 5.9](#), we first show that our sPKE scheme is IND-CCA secure. The proof follows Lindell's proof for the construction of an IND-CCA secure public-key encryption scheme from a IND-CPA secure one [[Lin06](#)].

Lemma C.1. *Let sPKE be the scheme from [Section 5.2](#) and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms such that \mathcal{A}_1 and \mathcal{A}_2 together make at most q_G queries to \mathcal{O}_G and at most q_{SD} queries to \mathcal{O}_{SD} . Then, there exist adversaries \mathcal{A}_{DDH} , \mathcal{A}_{ZK} , \mathcal{A}_{snd} , and \mathcal{A}_{Sig} (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}$) such that*

$$\begin{aligned} \text{Adv}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}} &\leq 4 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}}^{\text{DDH}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + 2q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}}^{\text{NIZK-sim-snd}} \\ &\quad + 2 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + \frac{4(q_G + 1)^2 + 8}{2^\kappa - 1}. \end{aligned}$$

Proof. We assume without loss of generality that \mathcal{A}_2 does not query the challenge ciphertext c^* to its decryption oracle \mathcal{O}_{SD} since doing so can only decrease the advantage. For $b_1, b_2 \in \{0, 1\}$, we define the hybrid experiment H_{b_1, b_2} as follows: Let H_{b_1, b_2} be as $\text{Exp}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}$, but where the common reference string crs is obtained via $(crs, \tau) \leftarrow S_1(1^\kappa)$ (instead of an invocation of NIZK.Gen). When \mathcal{A}_1 outputs (m_0, m_1, st) , H_{b_1, b_2} computes c_1 as the encryption of m_{b_1} under ek_1 , c_2 as the encryption of m_{b_2} under ek_2 , and c_σ as in the real experiment (namely as the encryption of the two ElGamal public keys and the accompanying signature). It then simulates the proof π using S_2 and invokes \mathcal{A}_2 on input st and $c^* := (c_1, c_2, c_\sigma, \pi)$.

Claim 1. *There exist adversaries \mathcal{A}'_{ZK} and $\mathcal{A}''_{\text{ZK}}$ such that*

$$\begin{aligned} \Pr^{H_{0,0}}[b' = 1] - \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 1 \mid b = 0] &= \text{Adv}_{\text{NIZK}, \mathcal{A}'_{\text{ZK}}}^{\text{NIZK-ZK}}, \\ \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 1 \mid b = 1] - \Pr^{H_{1,1}}[b' = 1] &= \text{Adv}_{\text{NIZK}, \mathcal{A}''_{\text{ZK}}}^{\text{NIZK-ZK}}. \end{aligned}$$

Proof of claim. We only prove the first part of the claim, the second one can be shown analogously. The adversary \mathcal{A}'_{ZK} on input crs , emulates toward \mathcal{A} the experiment $H_{0,0}$. To this end, it generates all required keys. When generating the challenge ciphertext $c^* = (c_1, c_2, c_\sigma, \pi)$, it obtains π via the proof oracle. Note that in $H_{0,0}$, this ciphertext is a valid encryption of m_0 , so the statement is correct and the proof oracle consequently returns a valid proof. When \mathcal{A} returns a bit b' , \mathcal{A}'_{ZK} returns $1 - b'$. Observe that if the CRS and the proofs are real, then this emulation is equivalent to the experiment $\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{SPKE-IND-CCA}}$ when $b = 0$, and if the CRS and the proofs are simulated, then it is equivalent to $H_{0,0}$. Hence,

$$\begin{aligned} \text{Adv}_{\text{NIZK}, \mathcal{A}'_{\text{ZK}}}^{\text{NIZK-ZK}} &= \Pr^{crs \leftarrow \text{Gen}(1^\kappa)} [\mathcal{A}^{\text{Prove}(crs, \cdot, \cdot)}(crs) = 1] - \Pr^{(crs, \tau) \leftarrow S_1(1^\kappa)} [\mathcal{A}^{S'(crs, \tau, \cdot)}(crs) = 1] \\ &= 1 - \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{SPKE-IND-CCA}}} [b' = 1 \mid b = 0] - (1 - \Pr^{H_{0,0}} [b' = 1]) \\ &= \Pr^{H_{0,0}} [b' = 1] - \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{SPKE-IND-CCA}}} [b' = 1 \mid b = 0]. \quad \diamond \end{aligned}$$

Analogous to the proof of [Lemma 5.10](#), we next define three bad events. Let B_1 be the event that \mathcal{A} queries its decryption oracle with a valid but improper ciphertext $(c_1, c_2, c_\sigma, \pi)$, i.e., $(g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma) \notin L$, but where π is an accepting proof, i.e., $\text{NIZK.Ver}(crs, x := (g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma), \pi) = 1$. As in the proof of [Lemma 5.10](#), one can show that there exists an adversary $\mathcal{A}_{\text{snd}}^{b_1, b_2}$ such that

$$\Pr^{H_{b_1, b_2}} [B_1] \leq q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^{b_1, b_2}}^{\text{NIZK-sim-snd}},$$

except that we here need (one-time) simulation soundness since the proof in the challenge ciphertext is simulated.

Further let B_2 be the event that \mathcal{A} queries its decryption oracle with a valid and proper ciphertext $(c_1, c_2, c_\sigma, \pi)$, i.e., $(g, ek^{\text{PKE}}, vk^{\text{Sig}}, c_1, c_2, c_\sigma) \in L$ and π is accepting, but where c_σ is the encryption of a triple (ek_1, ek_2, σ) , such that the pair (ek_1, ek_2) has never been output by the experiment or the oracle \mathcal{O}_G . Again as in the proof of [Lemma 5.10](#), it can be shown that there exists an adversary $\mathcal{A}_{\text{Sig}}^{b_1, b_2}$ such that

$$\Pr^{H_{b_1, b_2}} [B_2] \leq \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^{b_1, b_2}}^{\text{Sig-EUF-CMA}}.$$

Finally, let B_3 be the event that H_{b_1, b_2} generates two different encryption keys $ek^{\text{SPKE}} = (g, p, crs, ek^{\text{PKE}}, vk^{\text{Sig}}, ek_1, ek_2, \sigma)$ and $(ek^{\text{SPKE}})' = (g, p, crs, ek^{\text{PKE}}, vk^{\text{Sig}}, ek'_1, ek'_2, \sigma')$ such that $ek_1 = ek'_1$ or $ek_2 = ek'_2$. Then,

$$\Pr^{H_{b_1, b_2}} [B_3] \leq \frac{2(q_G + 1)^2}{2^\kappa - 1},$$

which can be shown as in the proof of [Lemma 5.10](#). For $B := B_1 \cup B_2 \cup B_3$, we therefore have

$$\Pr^{H_{b_1, b_2}} [B] \leq q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^{b_1, b_2}}^{\text{NIZK-sim-snd}} + \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^{b_1, b_2}}^{\text{Sig-EUF-CMA}} + \frac{2(q_G + 1)^2}{2^\kappa - 1}. \quad (9)$$

Claim 2. *There exist adversaries $\mathcal{A}'_{\text{DDH}}$ and $\mathcal{A}''_{\text{DDH}}$ such that*

$$\begin{aligned} \Pr^{H_{1,0}} [b' = 1] - \Pr^{H_{0,0}} [b' = 1] &\leq 2 \cdot \text{Adv}_{g, \mathcal{A}'_{\text{DDH}}}^{\text{DDH}} + \Pr^{H_{0,0}} [B] + \Pr^{H_{1,0}} [B] + 2^{2-\kappa}, \\ \Pr^{H_{1,1}} [b' = 1] - \Pr^{H_{1,0}} [b' = 1] &\leq 2 \cdot \text{Adv}_{g, \mathcal{A}''_{\text{DDH}}}^{\text{DDH}} + 2^{2-\kappa}. \end{aligned}$$

Proof of claim. We define the adversary $\mathcal{A}'_{\text{DDH}}$ as follows. On input a triple (X, Y, T) , it sets $ek_1 \leftarrow X$. It further generates all the remaining keys of the experiment (and thus lacks only the decryption key dk_1), samples $b \leftarrow \{0, 1\}$, and emulates $H_{b,0}$ toward \mathcal{A} . The oracle \mathcal{O}_{SD} is emulated by decrypting the second ciphertext component instead of the first one using dk_2 . When \mathcal{A}_1 returns (m_0, m_1, st) , $\mathcal{A}'_{\text{DDH}}$ samples $r \leftarrow \mathbb{Z}_p^*$ and sets

$$c_1 \leftarrow (g^r, X^r, Y, T \cdot m_b).$$

It further computes c_2 as an ElGamal encryption of m_0 , encrypts both keys and their signature to obtain c_σ , and simulates the NIZK proof π using S_2 . It continues the emulation by giving $c^* := (c_1, c_2, c_\sigma, \pi)$ to \mathcal{A}_2 . When \mathcal{A}_2 outputs its decision bit b' , $\mathcal{A}'_{\text{DDH}}$ outputs $d = 1$ if $b' = b$, and $d = 0$ otherwise.

First note that if (X, Y, T) are three uniform group elements, c_1 is independent of the bit b , and thus

$$\Pr^{\text{DDH}^{\text{rand}}_{g, \mathcal{A}'_{\text{DDH}}}}[d = 1] = \frac{1}{2}.$$

On the other hand, if (X, Y, T) is a DDH triple, we have for a uniform $s \in \mathbb{Z}_p$,

$$c_1 = (g^r, X^r, Y, T \cdot m_b) = (g^r, ek_1^r, g^s, ek_1^s \cdot m_b),$$

which corresponds to a proper ElGamal encryption of m_b if $X \neq 1$ and $Y \neq 1$. Further note that, as in the proof of [Lemma 5.10](#), \mathcal{O}_{SD} is emulated perfectly if B does not occur. We therefore have

$$\Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[d = 1 \cap \neg B \mid X \neq 1 \neq Y] = \Pr^{H_{b,0}}[b' = b \cap \neg B].$$

This implies

$$\begin{aligned} \text{Adv}_{g, \mathcal{A}'_{\text{DDH}}}^{\text{DDH}} &= \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[d = 1] - \Pr^{\text{DDH}^{\text{rand}}_{g, \mathcal{A}'_{\text{DDH}}}}[d = 1] \\ &\geq \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[d = 1 \cap \neg B \mid X \neq 1 \neq Y] \cdot \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[X \neq 1 \neq Y] - \frac{1}{2} \\ &= \Pr^{H_{b,0}}[b' = b \cap \neg B] \cdot \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[X \neq 1 \neq Y] - \frac{1}{2}. \end{aligned}$$

Using the union bound and $|G| = p \geq 2^\kappa$, we further have

$$\begin{aligned} \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[X \neq 1 \neq Y] &= 1 - \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[X = 1 \vee Y = 1] \\ &\geq 1 - \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[X = 1] - \Pr^{\text{DDH}^{\text{real}}_{g, \mathcal{A}'_{\text{DDH}}}}[Y = 1] \\ &\geq 1 - 2 \cdot 2^{-\kappa}. \end{aligned}$$

Hence,

$$\text{Adv}_{g, \mathcal{A}'_{\text{DDH}}}^{\text{DDH}} \geq \Pr^{H_{b,0}}[b' = b \cap \neg B] - 2^{1-\kappa} - \frac{1}{2}.$$

Since

$$\Pr^{H_{b,0}}[b' = b] \leq \Pr^{H_{b,0}}[(b' = b \cap \neg B) \cup B] \leq \Pr^{H_{b,0}}[b' = b \cap \neg B] + \Pr^{H_{b,0}}[B],$$

we obtain

$$\begin{aligned}
\text{Adv}_{g, \mathcal{A}'_{\text{DDH}}}^{\text{DDH}} &\geq \Pr^{H_{b,0}}[b' = b] - \Pr^{H_{b,0}}[B] - 2^{1-\kappa} - \frac{1}{2} \\
&= \frac{1}{2} \Pr^{H_{0,0}}[b' = 0] + \frac{1}{2} \Pr^{H_{1,0}}[b' = 1] - \frac{1}{2} \Pr^{H_{0,0}}[B] - \frac{1}{2} \Pr^{H_{1,0}}[B] - 2^{1-\kappa} - \frac{1}{2} \\
&= \frac{1}{2} \Pr^{H_{1,0}}[b' = 1] - \frac{1}{2} \Pr^{H_{0,0}}[b' = 1] - \frac{1}{2} \Pr^{H_{0,0}}[B] - \frac{1}{2} \Pr^{H_{1,0}}[B] - 2^{1-\kappa}.
\end{aligned}$$

Rearranging the inequality concludes the proof of the first part of the claim.

The second part of the claim can be proven analogously, where $\mathcal{A}''_{\text{DDH}}$ sets $ek_2 \leftarrow X$ instead of $ek_1 \leftarrow X$. Since it therefore has dk_1 , which is the key used by the decryption algorithm, the decryption oracle can be emulated perfectly, even if B occurs. \diamond

Using [Claims 1](#) and [2](#), we get

$$\begin{aligned}
\text{Adv}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}} &\leq 2 \cdot \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = b] - 1 \\
&= 2 \left(\frac{1}{2} \cdot \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 0 \mid b = 0] + \frac{1}{2} \cdot \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 1 \mid b = 1] \right) - 1 \\
&= \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 1 \mid b = 1] - \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 1 \mid b = 0] \\
&= \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 1 \mid b = 1] - \Pr^{H_{1,1}}[b' = 1] \\
&\quad + \Pr^{H_{1,1}}[b' = 1] - \Pr^{H_{1,0}}[b' = 1] + \Pr^{H_{1,0}}[b' = 1] - \Pr^{H_{0,0}}[b' = 1] \\
&\quad + \Pr^{H_{0,0}}[b' = 1] - \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}}[b' = 1 \mid b = 0] \\
&\leq \text{Adv}_{\text{NIZK}, \mathcal{A}'_{\text{ZK}}}^{\text{NIZK-ZK}} + \text{Adv}_{\text{NIZK}, \mathcal{A}''_{\text{ZK}}}^{\text{NIZK-ZK}} + 2 \cdot \text{Adv}_{g, \mathcal{A}'_{\text{DDH}}}^{\text{DDH}} + 2 \cdot \text{Adv}_{g, \mathcal{A}''_{\text{DDH}}}^{\text{DDH}} \\
&\quad + \Pr^{H_{0,0}}[B] + \Pr^{H_{1,0}}[B] + 2^{3-\kappa}.
\end{aligned}$$

Let \mathcal{A}_{ZK} be the adversary that runs \mathcal{A}'_{ZK} and $\mathcal{A}''_{\text{ZK}}$ with probability 1/2 each, let \mathcal{A}_{DDH} run $\mathcal{A}'_{\text{DDH}}$ and $\mathcal{A}''_{\text{DDH}}$ with probability 1/2 each, let \mathcal{A}_{snd} run $\mathcal{A}_{\text{snd}}^{0,0}$ and $\mathcal{A}_{\text{snd}}^{1,0}$ with probability 1/2 each, and let \mathcal{A}_{Sig} run $\mathcal{A}_{\text{Sig}}^{0,0}$ and $\mathcal{A}_{\text{Sig}}^{1,0}$ with probability 1/2 each. Combing the result above with [equation \(9\)](#), we can then conclude

$$\begin{aligned}
\text{Adv}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}} &\leq 4 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}}^{\text{DDH}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + 2q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}}^{\text{NIZK-sim-snd}} \\
&\quad + 2 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + \frac{4(q_G + 1)^2 + 8}{2^\kappa - 1}. \quad \square
\end{aligned}$$

We finally show that sPKE is IK-CCA secure.

Lemma C.2. *Let sPKE be the scheme from [Section 5.2](#) and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a pair of probabilistic algorithms such that \mathcal{A}_1 and \mathcal{A}_2 together make at most q_G queries to \mathcal{O}_G and at most q_{SD} queries to \mathcal{O}_{SD_0} and \mathcal{O}_{SD_1} combined. Then, there exist adversaries \mathcal{A}_{DDH} , \mathcal{A}_{ZK} , \mathcal{A}_{snd} , \mathcal{A}_{PKE} , and \mathcal{A}_{Sig} (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}}$) such that*

$$\begin{aligned}
\text{Adv}_{\text{SPKE}, \mathcal{A}}^{\text{sPKE-IND-CCA}} &\leq 8 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}}^{\text{DDH}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + 8q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}}^{\text{NIZK-sim-snd}} \\
&\quad + 2 \cdot \text{Adv}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}} + 8 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + \frac{16(q_G + 2)^2 + 32}{2^\kappa - 1}.
\end{aligned}$$

Proof. We assume without loss of generality that \mathcal{A}_2 does not query the challenge ciphertext c^* to any of its decryption oracles \mathcal{O}_{SD_0} or \mathcal{O}_{SD_1} , since doing so can only decrease the advantage. We define hybrid experiments H_0 to H_5 as follows:

- H_0 is identical to $\text{Exp}_{\mathcal{S}_{\text{PKE},\mathcal{A}}}^{\text{SPKE-IK-CCA}}$, except that the common reference string crs is obtained via $(crs, \tau) \leftarrow S_1(1^\kappa)$ (instead of an invocation of NIZK.Gen), and the proof π in the challenge ciphertext c^* is simulated using S_2 .
- H_1 is identical to H_0 , but when \mathcal{A}_1 outputs (m, st) , the hybrid computes c_σ not as an encryption of $(ek_{b,1}, ek_{b,2}, \sigma_b)$, but as the encryption of 0^ℓ , where ℓ is the length of the encoding of $(ek_{b,1}, ek_{b,2}, \sigma_b)$ (where the encoding needs to be chosen such that this length is equal for all keys).
- H_2 is identical to H_1 , except that for the generation of the challenge ciphertext c^* , the key $ek_{0,1}$ is replaced by $g^{d_{0,1}}$ for a freshly sampled $d_{0,1} \leftarrow \mathbb{Z}_p^*$.
- H_3 is identical to H_2 , except that for the generation of the challenge ciphertext c^* , the key $ek_{0,2}$ is replaced by $g^{d_{0,2}}$ for a freshly sampled $d_{0,2} \leftarrow \mathbb{Z}_p^*$.
- H_4 is identical to H_3 , except that for the generation of the challenge ciphertext c^* , the key $ek_{1,1}$ is replaced by $g^{d_{1,1}}$ for a freshly sampled $d_{1,1} \leftarrow \mathbb{Z}_p^*$.
- H_5 is identical to H_4 , except that for the generation of the challenge ciphertext c^* , the key $ek_{1,2}$ is replaced by $g^{d_{1,2}}$ for a freshly sampled $d_{1,2} \leftarrow \mathbb{Z}_p^*$.

Note that the view of \mathcal{A} in H_5 is independent from the bit b , which implies

$$\Pr^{H_5}[b' = b] = \frac{1}{2}. \quad (10)$$

It can be shown as in the proof of [Lemma C.1](#) that there exist an adversary \mathcal{A}_{ZK} such that

$$\Pr^{\text{Exp}_{\mathcal{S}_{\text{PKE},\mathcal{A}}}^{\text{SPKE-IK-CCA}}}[b' = b] - \Pr^{H_0}[b' = b] = \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}}. \quad (11)$$

Claim 1. *There exists an adversary \mathcal{A}_{PKE} such that*

$$\Pr^{H_0}[b' = b] - \Pr^{H_1}[b' = b] = \text{Adv}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}.$$

Proof of claim. When \mathcal{A}_{PKE} obtains a public key ek from the CPA challenger, it generates all remaining keys itself and emulates H_0 (or H_1) toward \mathcal{A} . Note that \mathcal{A}_{PKE} never needs to decrypt any of the ciphertexts c_σ in the experiment and thus, the missing decryption key is not needed for the emulation. When \mathcal{A}_1 outputs (m, st) , \mathcal{A}_{PKE} gives $(0^\ell, (ek_{b,1}, ek_{b,2}, \sigma_b))$ to its CPA challenger to obtain a ciphertext c_σ , where ℓ is the length of the encoding of $(ek_{b,1}, ek_{b,2}, \sigma_b)$. The rest is done as in H_0 . When \mathcal{A}_2 returns a bit b' , \mathcal{A}_{PKE} returns $b'' = 1$ if $b' = b$, and $b'' = 0$ if $b' \neq b$.

Note that if the CPA challenger chooses the bit $b_{\text{CPA}} = 0$, c_σ is an encryption of 0^ℓ , as in H_1 , and if $b_{\text{CPA}} = 1$, c_σ is as in H_0 . Hence,

$$\begin{aligned} \Pr^{\text{Exp}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}}[b'' = 1 \mid b_{\text{CPA}} = 0] &= \Pr^{H_1}[b' = b], \\ \Pr^{\text{Exp}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}}[b'' = 1 \mid b_{\text{CPA}} = 1] &= \Pr^{H_0}[b' = b]. \end{aligned}$$

We can therefore conclude

$$\begin{aligned}
\text{Adv}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}} &= 2 \cdot \Pr^{\text{Exp}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}} [b'' = b_{\text{CPA}}] - 1 \\
&= \Pr^{\text{Exp}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}} [b'' = 0 \mid b_{\text{CPA}} = 0] + \Pr^{\text{Exp}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}} [b'' = 1 \mid b_{\text{CPA}} = 1] - 1 \\
&= \Pr^{\text{Exp}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}} [b'' = 1 \mid b_{\text{CPA}} = 1] - \Pr^{\text{Exp}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}}} [b'' = 1 \mid b_{\text{CPA}} = 0] \\
&= \Pr^{H_0} [b' = b] - \Pr^{H_1} [b' = b]. \quad \diamond
\end{aligned}$$

We define the event B analogous to the events in the proofs of [Lemmata 5.10](#) and [C.1](#). As there, one can show for $i \in \{0, \dots, 5\}$ that there exist adversaries $\mathcal{A}_{\text{snd}}^i$ and $\mathcal{A}_{\text{Sig}}^i$ such that

$$\Pr^{H_i} [B] \leq q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}^i}^{\text{NIZK-sim-snd}} + \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}^i}^{\text{Sig-EUF-CMA}} + \frac{2(q_G + 2)^2}{2^\kappa - 1}. \quad (12)$$

Claim 2. *There exist adversaries $\mathcal{A}_{\text{DDH}}^1, \dots, \mathcal{A}_{\text{DDH}}^4$ such that for $i \in \{1, 3\}$*

$$\Pr^{H_i} [b' = b] - \Pr^{H_{i+1}} [b' = b] \leq \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{DDH}} + \Pr^{H_i} [B] + \Pr^{H_{i+1}} [B] + 2^{2-\kappa},$$

and for $i \in \{2, 4\}$

$$\Pr^{H_i} [b' = b] - \Pr^{H_{i+1}} [b' = b] \leq \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{DDH}} + 2^{2-\kappa}.$$

Proof of claim. On input (X, Y, T) , $\mathcal{A}_{\text{DDH}}^1$ sets $ek_{0,1} \leftarrow X$, $\mathcal{A}_{\text{DDH}}^2$ sets $ek_{0,2} \leftarrow X$, $\mathcal{A}_{\text{DDH}}^3$ sets $ek_{1,1} \leftarrow X$, and $\mathcal{A}_{\text{DDH}}^4$ sets $ek_{1,2} \leftarrow X$. All adversaries generate the remaining keys themselves and emulate H_i (or H_{i+1}) toward \mathcal{A} . To emulate the decryption oracles, $\mathcal{A}_{\text{DDH}}^1$ and $\mathcal{A}_{\text{DDH}}^3$ decrypt the second ciphertext component instead of the first one; $\mathcal{A}_{\text{DDH}}^2$ and $\mathcal{A}_{\text{DDH}}^4$ can emulate all oracles perfectly. As in the proof of [Lemma 5.10](#), \mathcal{O}_{SD} is also emulated perfectly by $\mathcal{A}_{\text{DDH}}^1$ and $\mathcal{A}_{\text{DDH}}^3$ if the event B does not occur. When \mathcal{A}_1 returns (m, st) and if $b = 0$, then $\mathcal{A}_{\text{DDH}}^1$ samples $r \leftarrow \mathbb{Z}_p^*$, sets

$$c_1 \leftarrow (Y^r, T^r, Y, T \cdot m),$$

and generates the remaining ciphertext components as in the real experiment. The other adversaries generate the ciphertext components analogously. When \mathcal{A}_2 returns a bit b' , then $\mathcal{A}_{\text{DDH}}^i$ returns $d = 1$ if $b' = b$ and $d = 0$ if $b' \neq b$ for all $i \in \{1, \dots, 4\}$.

Consider the case $i = 1$ and note that if (X, Y, T) is a DDH triple, we have $Y = g^s$ and $T = X^s$ for a uniform $s \in \mathbb{Z}_p$, and thus, if $b = 0$,

$$c_1 = (Y^r, X^{s \cdot r}, Y, X^s \cdot m) = (g^{s \cdot r}, (ek_{0,1})^{s \cdot r}, g^s, (ek_{0,1})^s \cdot m).$$

If $X \neq 1 \neq Y$, this corresponds to an encryption of m under $ek_{0,1}$, as in H_1 . On the other hand, if (X, Y, T) are three uniform group elements and $X \neq 1 \neq Y$, then there are (uniformly distributed) $s, d_{0,1} \in \mathbb{Z}_p^*$ such that $Y = g^s$ and $T = g^{s \cdot d_{0,1}}$. Hence, we have in this case

$$c_1 = (g^{s \cdot r}, g^{s \cdot d_{0,1} \cdot r}, g^s, g^{s \cdot d_{0,1}} \cdot m) = (g^{s \cdot r}, (g^{d_{0,1}})^{s \cdot r}, g^s, (g^{d_{0,1}})^s \cdot m),$$

which corresponds to an encryption under the fresh key $g^{d_{0,1}}$, as in H_2 . Further note that if $b = 1$, the emulation, H_1 , and H_2 are all equivalent. We therefore have

$$\begin{aligned}
\Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{real}}} [d = 1 \cap \neg B \mid X \neq 1 \neq Y] &= \Pr^{H_1} [b' = b \cap \neg B], \\
\Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [d = 1 \cap \neg B \mid X \neq 1 \neq Y] &= \Pr^{H_2} [b' = b \cap \neg B].
\end{aligned}$$

As in the proof of Lemma C.1, we obtain

$$\Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{real}}} [d = 1] \geq \Pr^{H_1} [b' = b] - \Pr^{H_1} [B] - 2^{1-\kappa}.$$

Moreover,

$$\begin{aligned} \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [d = 1] &= \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [d = 1 \cap \neg B \mid X \neq 1 \neq Y] \cdot \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [X \neq 1 \neq Y] \\ &\quad + \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [d = 1 \cap (B \cup [X = 1 \vee Y = 1])] \\ &\leq \Pr^{H_2} [b' = b \cap \neg B] + \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [B] + \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [X = 1 \vee Y = 1] \\ &\leq \Pr^{H_2} [b' = b] + \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [B] + 2^{1-\kappa}. \end{aligned}$$

Since the probability of B in $\text{DDH}_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{rand}}$ is equal to its probability in H_2 , we conclude

$$\begin{aligned} \text{Adv}_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}} &= \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{real}}} [d = 1] - \Pr_{g, \mathcal{A}_{\text{DDH}}^1}^{\text{DDH}^{\text{rand}}} [d = 1] \\ &\geq \Pr^{H_1} [b' = b] - \Pr^{H_2} [b' = b] - \Pr^{H_1} [B] - \Pr^{H_2} [B] - 2^{2-\kappa}. \end{aligned}$$

The proofs for $i \in \{2, 3, 4\}$ are analogous, where for $i \in \{2, 4\}$, the occurrence of B does not matter since the decryption oracle can always be emulated perfectly. \diamond

Using equation (11), Claims 1 and 2, and equation (10), we obtain

$$\begin{aligned} \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{SPKE-IK-CCA}}} [b' = b] &= \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{SPKE-IK-CCA}}} [b' = b] - \Pr^{H_0} [b' = b] + \Pr^{H_0} [b' = b] - \Pr^{H_1} [b' = b] \\ &\quad + \sum_{i=1}^4 \left(\Pr^{H_i} [b' = b] - \Pr^{H_{i+1}} [b' = b] \right) + \Pr^{H_5} [b' = b] \\ &\leq \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + \text{Adv}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}} + \sum_{i=1}^4 \left(\text{Adv}_{g, \mathcal{A}_{\text{DDH}}^i}^{\text{DDH}} + \Pr^{H_i} [B] \right) + 2^{4-\kappa} + \frac{1}{2}. \end{aligned}$$

For the adversary \mathcal{A}_{snd} that runs $\mathcal{A}_{\text{snd}}^1, \dots, \mathcal{A}_{\text{snd}}^4$ with probability 1/4 each, and the adversary \mathcal{A}_{sig} that runs $\mathcal{A}_{\text{sig}}^1, \dots, \mathcal{A}_{\text{sig}}^4$ with probability 1/4 each, we obtain using equation (12),

$$\sum_{i=1}^4 \Pr^{H_i} [B] \leq 4q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}}^{\text{NIZK-sim-snd}} + 4 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{sig}}}^{\text{Sig-EUF-CMA}} + \frac{8(q_G + 2)^2}{2^\kappa - 1}.$$

Further defining \mathcal{A}_{DDH} as running $\mathcal{A}_{\text{DDH}}^1, \dots, \mathcal{A}_{\text{DDH}}^4$ with probability 1/4 each yields

$$\begin{aligned} \text{Adv}_{\text{SPKE}, \mathcal{A}}^{\text{SPKE-IK-CCA}} &\leq 2 \cdot \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}}^{\text{SPKE-IK-CCA}}} [b' = b] - 1 \\ &\leq 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + 2 \cdot \text{Adv}_{\text{PKE}, \mathcal{A}_{\text{PKE}}}^{\text{PKE-IND-CPA}} + 8 \cdot \text{Adv}_{g, \mathcal{A}_{\text{DDH}}}^{\text{DDH}} + 8q_{SD} \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{snd}}}^{\text{NIZK-sim-snd}} \\ &\quad + 8 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{sig}}}^{\text{Sig-EUF-CMA}} + \frac{16(q_G + 2)^2}{2^\kappa - 1} + 2^{5-\kappa}. \end{aligned}$$

Observing that $2^{5-\kappa} \leq \frac{32}{2^\kappa - 1}$ concludes the proof. \square

D Remaining Security Proofs of the ACE Scheme for Equality

After privacy, which was shown in Lemma 6.3, we prove anonymity, which can be shown similarly. We provide a proof for strong anonymity. Note, however, that for the equality policy, strong anonymity does not provide more guarantees than weak anonymity because anyone who can decrypt directly learns that the sender role is equal to the receiver role.

Lemma D.1. *Let ACE be the scheme from above, let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an attacker on the anonymity such that \mathcal{A}_1 makes at most q_S queries of the form (\cdot, sen) to the oracle \mathcal{O}_G , and at most q_D queries to \mathcal{O}_{SD} . Then, there exist probabilistic algorithms \mathcal{A}_{PRF} , \mathcal{A}_{ZK} , and $\mathcal{A}_{\text{sPKE}}$ (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$) such that*

$$\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-sANON-CCA}} \leq 2 \cdot \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + (q_S + q_D + 1)^2 \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-IK-CCA}}.$$

Proof. We assume without loss of generality that \mathcal{A} ensures $m_0 = m_1$ and $P(i_0, j) = P(i_1, j)$ for all $j \in J$, since doing otherwise can only decrease the advantage. Since we have $P(i, j) = 1 \Leftrightarrow i = j$, the latter condition implies that if $i_0 \in J$ or $i_1 \in J$, then $i_0 = i_1$. In case $i_0 = i_1$ and $m_0 = m_1$, \mathcal{A} cannot have positive advantage. Hence, we can further assume without loss of generality that $i_0 \notin J$ and $i_1 \notin J$. As in the proof of Lemma 6.3, let $H_0 := \text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-PRV-ANON-CCA}}$, let H_1 be as H_0 where F_K is replaced by a truly uniform random function U , and let H_2 be as H_1 , where $\text{crs}^{\text{NIZK}} \leftarrow \text{NIZK.Gen}(1^\kappa)$ in ACE.Setup is replaced by $(\text{crs}^{\text{NIZK}}, \tau^{\text{NIZK}}) \leftarrow S_1^{\text{NIZK}}(1^\kappa)$ and for the generation of the challenge ciphertext c^* , $\pi^{\text{NIZK}} \leftarrow \text{NIZK.Prove}(\text{crs}^{\text{NIZK}}, x, w)$ in ACE.Enc is replaced by $\pi^{\text{NIZK}} \leftarrow S_2^{\text{NIZK}}(\text{crs}^{\text{NIZK}}, \tau^{\text{NIZK}}, x)$. An identical proof as the one in the proof of Lemma 6.3 shows that there exist \mathcal{A}_{PRF} and \mathcal{A}_{ZK} such that

$$\Pr^{H_0}[b' = b] - \Pr^{H_2}[b' = b] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}}.$$

We now transform \mathcal{A} to a winner $\mathcal{A}_{\text{sPKE}}$ for the anonymity game for the scheme sPKE. The reduction is similar to the one in the proof of Lemma 6.3, but $\mathcal{A}_{\text{sPKE}}$ has to guess both i_0 and i_1 , which is why we loose the quadratic factor $(q_S + q_D + 1)^2$. On input $(sp^{\text{sPKE}}, ek_0^{\text{sPKE}}, ek_1^{\text{sPKE}})$, $\mathcal{A}_{\text{sPKE}}$ initializes $i_{q_0}, i_{q_1} \leftarrow \perp$, $k_q \leftarrow 1$, chooses $q_0, q_1 \leftarrow \{0, \dots, q_S + q_D\}$ uniformly at random, runs $(vk^{\text{Sig}}, sk^{\text{Sig}}) \leftarrow \text{Sig.Gen}(1^\kappa)$, and $(\text{crs}^{\text{NIZK}}, \tau^{\text{NIZK}}) \leftarrow S_1^{\text{NIZK}}(1^\kappa)$, and gives $sp^{\text{ACE}} := (sp^{\text{sPKE}}, vk^{\text{Sig}}, \text{crs}^{\text{NIZK}})$ to \mathcal{A}_1 . It emulates the oracles for \mathcal{A}_1 as follows.

$\mathcal{O}_G(\cdot, \cdot)$: On query (i, sen) , if $k_q \notin \{q_0, q_1\}$ and $i \notin \{i_{q_0}, i_{q_1}\}$, then generate an encryption key $ek_i^{\text{ACE}} := (vk^{\text{Sig}}, ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, \text{crs}^{\text{NIZK}})$ as H_2 does, where $(ek_i^{\text{sPKE}}, dk_i^{\text{sPKE}})$ is obtained via \mathcal{O}_G and remembered for future queries. If $k_q = q_l$ or $i = i_{q_l}$ for some $l \in \{0, 1\}$, replace ek_i^{sPKE} by ek_i^{sPKE} (by ek_0^{sPKE} if $q_0 = q_1$) and set $i_{q_l} \leftarrow i$. In both cases, set $k_q \leftarrow k_q + 1$ at the end. On query (j, rec) , obtain a decryption key from \mathcal{O}_G and remember it for later.

$\mathcal{O}_{SD}(\cdot, \cdot)$: On query $(j, c = (\tilde{c}, \pi^{\text{NIZK}}))$, if $k_q \notin \{q_0, q_1\}$ and $j \notin \{i_{q_0}, i_{q_1}\}$, then execute $c' \leftarrow \text{ACE.San}(sp^{\text{ACE}}, c)$, generate a decryption key dk_j^{ACE} as above, decrypt c' using dk_j^{ACE} , and return the resulting message. If $k_q = q_l$ or $j = i_{q_l}$ for some $l \in \{0, 1\}$, set $i_{q_l} \leftarrow j$ and use the oracle \mathcal{O}_{SD_l} of the IK-CCA experiment to obtain a decryption m of \tilde{c} . If $\text{NIZK.Ver}(\text{crs}^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}) = 1$, return m , otherwise, return \perp . In all cases, set $k_q \leftarrow k_q + 1$ at the end.

When \mathcal{A}_1 returns (m_0, m_1, i_0, i_1, st) , $\mathcal{A}_{\text{SPKE}}$ outputs m_0 to the challenger of the anonymity experiment to obtain a challenge ciphertext \tilde{c}^* . It then runs $S_2^{\text{NIZK}}(crs^{\text{NIZK}}, \tau^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}^*))$, and gives st and the ciphertext $c^* := (\tilde{c}^*, \pi^{\text{NIZK}})$ to \mathcal{A}_2 . It emulates the oracles for \mathcal{A}_2 as follows:

$\mathcal{O}_G(\cdot, \cdot)$: On query (i, sen) , if $i \notin \{i_0, i_1\}$, then generate an encryption key $ek_i^{\text{ACE}} := (vk_i^{\text{Sig}}, ek_i^{\text{SPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, crs^{\text{NIZK}})$ as H_2 does, where $(ek_i^{\text{SPKE}}, dk_i^{\text{SPKE}})$ is obtained via \mathcal{O}_G and remembered for future queries. If $i = i_{q_l}$ for some $l \in \{0, 1\}$, replace ek_i^{SPKE} by ek_l^{SPKE} . On query (j, rec) , obtain a decryption key as before.

$\mathcal{O}_{SD^*}(\cdot, \cdot)$: On query $(j, c = (\tilde{c}, \pi^{\text{NIZK}}))$, run $\text{ACE.DMod}(sp^{\text{ACE}}, c^*, c)$. If the output is 1, return **test**. Otherwise, if $j \notin \{i_0, i_1\}$, run $c' \leftarrow \text{ACE.San}(sp^{\text{ACE}}, c)$, generate a decryption key dk_j^{ACE} as above, decrypt c' using dk_j^{ACE} , and return the resulting message. If $j = i_{q_l}$ for some $l \in \{0, 1\}$, use the oracle \mathcal{O}_{SD_l} of the IK-CCA experiment to obtain a decryption m of \tilde{c} . If $\text{NIZK.Ver}(crs^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}) = 1$, return m , otherwise, return \perp .

Note that $\mathcal{A}_{\text{SPKE}}$ never queries any of the decryption oracles of the IK-CCA experiment on \tilde{c}^* because we return **test** whenever this would be necessary. Denote by Q the event that for all $l \in \{0, 1\}$ we have either $i_{q_l} = i_l$, or $q_l = 0$ and \mathcal{A}_1 does not make the query (i_l, sen) to \mathcal{O}_G and no queries for role i_l to \mathcal{O}_{SD} . When \mathcal{A}_2 returns a bit b' and Q holds, $\mathcal{A}_{\text{SPKE}}$ returns the same bit $b'' \leftarrow b'$, if $\neg Q$, $\mathcal{A}_{\text{SPKE}}$ returns a uniform bit $b'' \leftarrow \{0, 1\}$.

Let \tilde{b} be the bit chosen by the IK-CCA experiment. Note that if Q occurs, the view of \mathcal{A} is identical to the one in H_2 with $b = \tilde{b}$. This implies

$$\Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}_{\text{SPKE}}}^{\text{SPKE-IK-CCA}}} [b'' = \tilde{b} \mid Q] = \Pr^{H_2} [b' = b].$$

Using that the probability of Q is $1/(q_S + q_D + 1)^2$, it follows as in the proof of [Lemma 6.3](#) that

$$\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-SANON-CCA}} = 2 \cdot \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}}}^{\text{NIZK-ZK}} + (q_S + q_D + 1)^2 \cdot \text{Adv}_{\text{SPKE}, \mathcal{A}_{\text{SPKE}}}^{\text{SPKE-IK-CCA}}. \quad \square$$

We next prove sanitization security of our scheme.

Lemma D.2. *Let ACE be the scheme from above, and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an attacker on the sanitization security such that \mathcal{A}_1 makes at most q_{S_1} queries of the form (\cdot, sen) and at most q_{R_1} queries of the form (\cdot, rec) to the oracle \mathcal{O}_G , and at most q_{D_1} queries to \mathcal{O}_{SD} , and \mathcal{A}_2 makes at most q_{R_2} queries of the form (\cdot, rec) to the oracle \mathcal{O}_G . Then, there exist probabilistic algorithms \mathcal{A}_{PRF} , $\mathcal{A}_{\text{ZK}_1}$, $\mathcal{A}_{\text{ZK}_2}$, \mathcal{A}_{Sig} , $\mathcal{A}_{\text{SPKE}}$, and \mathcal{A}_{rob} (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-SAN-CCA}}$) such that*

$$\begin{aligned} \text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-SAN-CCA}} &\leq 2 \cdot \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1} + 4 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2} + 4 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} \\ &\quad + (q_{S_1} + q_{R_1} + q_{D_1})^2 \cdot \text{Adv}_{\text{SPKE}, \mathcal{A}_{\text{SPKE}}}^{\text{SPKE-SAN-CCA}} + 4(q_{R_1} + q_{R_2}) \cdot \text{Adv}_{\text{SPKE}, \mathcal{A}_{\text{rob}}}^{\text{SPKE-USROB}}. \end{aligned}$$

Proof. Let $H_0 := \text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-SAN-CCA}}$, let H_1 be as H_0 where F_K is replaced by a truly uniform random function U , and let H_2 be as H_1 , where $crs^{\text{NIZK}} \leftarrow \text{NIZK.Gen}(1^\kappa)$ in ACE.Setup is replaced by $(crs^{\text{NIZK}}, \xi^{\text{NIZK}}) \leftarrow E_1^{\text{NIZK}}(1^\kappa)$. Let W_{ACE} denote the event that \mathcal{A} wins, i.e.,

$$W_{\text{ACE}} := [b' = b \wedge c'_0 \neq \perp \neq c'_1 \wedge \forall j \in J m_{0,j} = m_{1,j} = \perp].$$

Similarly as in the proof of [Lemma 6.3](#), it can be shown that there exist \mathcal{A}_{PRF} and $\mathcal{A}_{\text{ZK}_1}$ such that

$$\Pr^{H_0} [W_{\text{ACE}}] - \Pr^{H_2} [W_{\text{ACE}}] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1}. \quad (13)$$

Let H_3 be identical to H_2 except that after \mathcal{A}_1 returns $(c_0 = (\tilde{c}_0, \pi_0^{\text{NIZK}}), c_1 = (\tilde{c}_1, \pi_1^{\text{NIZK}}), st)$, H_3 executes for $\tilde{b} \in \{0, 1\}$

$$w_{\tilde{b}} := (ek_{i_{\tilde{b}}}^{\text{sPKE}}, m_{\tilde{b}}, r_{\tilde{b}}, vk_{i_{\tilde{b}}}^{\text{Sig}}, \sigma_{i_{\tilde{b}}}^{\text{Sig}}, \sigma_{c_{\tilde{b}}}^{\text{Sig}}) \leftarrow E_2^{\text{NIZK}}(crs^{\text{NIZK}}, \xi^{\text{NIZK}}, x_{\tilde{b}} := (vk^{\text{Sig}}, \tilde{c}_{\tilde{b}}), \pi_{\tilde{b}}^{\text{NIZK}}).$$

We clearly have

$$\Pr^{H_3}[W_{\text{ACE}}] = \Pr^{H_2}[W_{\text{ACE}}]. \quad (14)$$

Let $V_{\tilde{b}} := [\text{NIZK.Ver}(crs^{\text{NIZK}}, x_{\tilde{b}}, \pi_{\tilde{b}}^{\text{NIZK}}) = 1]$ and let B_E be the event that (at least) one of the extractions fail, i.e.,

$$B_E := [(V_0 \wedge (x_0, w_0) \notin R) \vee (V_1 \wedge (x_1, w_1) \notin R)].$$

If B_E occurs, the knowledge extraction of NIZK is broken. To prove this, we define $\mathcal{A}_{\text{ZK}_2}$ as follows. On input crs^{NIZK} , it emulates an execution of H_3 , where in ACE.Setup , crs^{NIZK} is used instead of generating it. When \mathcal{A}_1 returns (c_0, c_1, st) , $\mathcal{A}_{\text{ZK}_2}$ flips a coin $\tilde{b} \leftarrow \{0, 1\}$ and returns $(x_{\tilde{b}}, \pi_{\tilde{b}}^{\text{NIZK}})$. If the \tilde{b} 's extraction fails, $\mathcal{A}_{\text{ZK}_2}$ wins the extraction game. Hence,

$$\Pr^{H_3}[B_E] \leq 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2}. \quad (15)$$

For $\tilde{b} \in \{0, 1\}$, let $B_{S, \tilde{b}}$ be the event that $(x_{\tilde{b}}, w_{\tilde{b}}) \in R$ and $ek_{i_{\tilde{b}}}^{\text{sPKE}}$ is not contained in an answer from \mathcal{O}_G to \mathcal{A}_1 , and let B_S be the union of $B_{S,0}$ and $B_{S,1}$. We next show that if B_S occurs, the adversary found a forgery for the signature scheme.

Claim 1. *There exists a probabilistic algorithm \mathcal{A}_{Sig} such that*

$$\Pr^{H_3}[B_S] \leq 2 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}}. \quad (16)$$

Proof of claim. On input vk^{Sig} , \mathcal{A}_{Sig} emulate an execution of H_3 , where vk^{Sig} is used in msk^{ACE} and sp^{ACE} . Queries (i, sen) by \mathcal{A}_1 to the oracle \mathcal{O}_G are answered by executing ACE.Gen (with F_K replaced by U) where σ_i^{Sig} is generated using the signing oracle of $\text{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{Sig-EUF-CMA}}$. After extracting w_0 and w_1 , \mathcal{A}_{Sig} flips a coin $\tilde{b} \leftarrow \{0, 1\}$ and returns $([ek_{i_{\tilde{b}}}^{\text{sPKE}}, vk_{i_{\tilde{b}}}^{\text{Sig}}], \sigma_{i_{\tilde{b}}}^{\text{Sig}})$. If $B_{S, \tilde{b}}$ occurs, $[ek_{i_{\tilde{b}}}^{\text{sPKE}}, vk_{i_{\tilde{b}}}^{\text{Sig}}]$ was not queried to the signing oracle and $(x_{\tilde{b}}, w_{\tilde{b}}) \in R$. The latter implies that $\sigma_{i_{\tilde{b}}}^{\text{Sig}}$ is a valid signature and hence \mathcal{A}_{Sig} successfully forged a signature. We conclude

$$\Pr^{H_3}[B_S] \leq 2 \cdot \left(\frac{1}{2} \Pr^{H_3}[B_{S,0}] + \frac{1}{2} \Pr^{H_3}[B_{S,1}] \right) = 2 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}}. \quad \diamond$$

Let H_4 be identical to H_3 with the difference that we replace for $k \in \{0, 1\}$ and $j \in J$, $m_{k,j} \leftarrow \text{ACE.Dec}(\text{ACE.Gen}(msk, j, \text{rec}), c'_k)$ by

$$m_{k,j} \leftarrow \begin{cases} m_k, & ek_j^{\text{sPKE}} = ek_{i_k}^{\text{sPKE}} \text{ for } (ek_j^{\text{sPKE}}, dk_j^{\text{sPKE}}) = \text{sPKE.Gen}(msk^{\text{sPKE}}; U([j, 0])), \\ \perp, & \text{else,} \end{cases} \quad (17)$$

where $ek_{i_k}^{\text{sPKE}}$ are the extracted keys. Note that if V_k , $\neg B_E$, and $\neg B_S$ occur, we have $c'_k = \text{San}(sp^{\text{sPKE}}, \tilde{c}_k)$, $\tilde{c}_k = \text{sPKE.Enc}(ek_{i_k}^{\text{sPKE}}, m_k; r_k)$, and $ek_{i_k}^{\text{sPKE}}$ was generated by \mathcal{O}_G . Hence, for $j \in J$ with $ek_j^{\text{sPKE}} = ek_{i_k}^{\text{sPKE}}$, we have by the correctness of the sPKE scheme that

$\text{ACE.Dec}(\text{ACE.Gen}(msk, j, \text{rec}), c'_k) = m_k$, i.e., $m_{k,j} = m_k$ in both H_3 and H_4 . For other $j \in J$, decryption only yields a message different from \perp if robustness of the sPKE scheme is violated. Since $|J| \leq q_{R_1} + q_{R_2}$, this implies for $V := V_0 \cap V_1$,

$$\Pr^{H_3}[W_{\text{ACE}} \mid V \cap \neg B_E \cap \neg B_S] - \Pr^{H_4}[W_{\text{ACE}} \mid V \cap \neg B_E \cap \neg B_S] \leq 2(q_{R_1} + q_{R_2}) \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}, \quad (18)$$

where \mathcal{A}_{rob} emulates the experiment and outputs \tilde{c}_k for a uniformly chosen $k \in \{0, 1\}$, i such that the i -th query to the key-generation oracle yields $ek_{i_k}^{\text{sPKE}}$, and a uniformly chosen j .⁶

We finally construct an adversary $\mathcal{A}_{\text{sPKE}}$ against the sanitization security of sPKE. On input $(sp^{\text{sPKE}}, ek_0^{\text{sPKE}}, ek_1^{\text{sPKE}})$, $\mathcal{A}_{\text{sPKE}}$ initializes $i_{q_0}, i_{q_1} \leftarrow \perp$, $k_q \leftarrow 1$, chooses distinct $q_0, q_1 \leftarrow \{1, \dots, q_{S_1} + q_{R_1} + q_{D_1}\}$ uniformly at random, executes $(vk^{\text{Sig}}, sk^{\text{Sig}}) \leftarrow \text{Sig.Gen}(1^\kappa)$, and $(crs^{\text{NIZK}}, \xi^{\text{NIZK}}) \leftarrow E_1^{\text{NIZK}}(1^\kappa)$, and gives $sp^{\text{ACE}} := (sp^{\text{sPKE}}, vk^{\text{Sig}}, crs^{\text{NIZK}})$ to \mathcal{A}_1 . It emulates the oracles for \mathcal{A}_1 as follows.

$\mathcal{O}_G(\cdot, \cdot)$: On query (i, sen) , if $k_q \notin \{q_0, q_1\}$ and $i \notin \{i_{q_0}, i_{q_1}\}$, generate an encryption key $(vk_i^{\text{Sig}}, ek_i^{\text{sPKE}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, crs^{\text{NIZK}})$ as H_4 , where $(ek_i^{\text{sPKE}}, dk_i^{\text{sPKE}})$ is obtained via \mathcal{O}_G and remembered for future queries. If $k_q = q_l$ or $i = i_{q_l}$ for some $l \in \{0, 1\}$, replace ek_i^{sPKE} by ek_l^{sPKE} and set $i_{q_l} \leftarrow i$. In both cases, set $k_q \leftarrow k_q + 1$ at the end.

On query (j, rec) , if $k_q \notin \{q_0, q_1\}$ and $j \notin \{i_{q_0}, i_{q_1}\}$, obtain a decryption key from \mathcal{O}_G , remember it, and set $k_q \leftarrow k_q + 1$. If $k_q = q_l$ or $j = i_{q_l}$ for some $l \in \{0, 1\}$, then return \perp and set $k_q \leftarrow k_q + 1$.

$\mathcal{O}_{SD}(\cdot, \cdot)$: On query $(j, c = (\tilde{c}, \pi^{\text{NIZK}}))$, if $k_q \notin \{q_0, q_1\}$ and $j \notin \{i_{q_0}, i_{q_1}\}$, then execute $c' \leftarrow \text{ACE.San}(sp^{\text{ACE}}, c)$, generate a decryption key dk_j^{ACE} as above, decrypt c' using dk_j^{ACE} , and return the resulting message. If $k_q = q_l$ or $j = i_{q_l}$ for some $l \in \{0, 1\}$, set $i_{q_l} \leftarrow j$, if $\text{NIZK.Ver}(crs^{\text{NIZK}}, x := (vk_i^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}) = 0$, return \perp , otherwise, use the oracle \mathcal{O}_{SD_l} of the sPKE-sanitization experiment to obtain a decryption of \tilde{c} and return it. In all cases, set $k_q \leftarrow k_q + 1$ at the end.

When \mathcal{A}_1 returns $(c_0 = (\tilde{c}_0, \pi_0^{\text{NIZK}}), c_1 = (\tilde{c}_1, \pi_1^{\text{NIZK}}), st)$, $\mathcal{A}_{\text{sPKE}}$ verifies the proofs π_0^{NIZK} and π_1^{NIZK} and extracts the witnesses to check the events V , B_E , and B_S . Denote by Q the event that $ek_{i_0}^{\text{sPKE}}, ek_{i_1}^{\text{sPKE}} \in \{ek_0^{\text{sPKE}}, ek_1^{\text{sPKE}}\}$, where $ek_{i_0}^{\text{sPKE}}, ek_{i_1}^{\text{sPKE}}$ are the extracted keys. Note that if V , $\neg B_E$, and $\neg B_S$ occur, both $ek_{i_0}^{\text{sPKE}}$ and $ek_{i_1}^{\text{sPKE}}$ have been returned by \mathcal{O}_G to \mathcal{A}_1 . This implies

$$\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}}[Q \mid V \cap \neg B_E \cap \neg B_S] \geq 1/(q_{S_1} + q_{R_1} + q_{D_1})^2. \quad (19)$$

If Q , V , $\neg B_E$, and $\neg B_S$ occur, $\mathcal{A}_{\text{sPKE}}$ returns $(\tilde{c}_0, \tilde{c}_1)$ to the challenger of the sPKE-sanitization experiment to obtain the sanitized ciphertext c'_b . It then gives (st, c'_b) to \mathcal{A}_2 and emulates the oracles as above. After \mathcal{A}_2 returned the bit b' , $\mathcal{A}_{\text{sPKE}}$ returns $b'' \leftarrow b'$. If $Q \cap V \cap \neg B_E \cap \neg B_S$ does not occur, $\mathcal{A}_{\text{sPKE}}$ runs $\bar{c} \leftarrow \text{sPKE.Enc}(ek_0^{\text{sPKE}}, \bar{m})$ for an arbitrary fixed message \bar{m} and returns $(c_0 := \bar{c}, c_1 := \bar{c})$ to the challenger. After receiving back a sanitized ciphertext c'_b , it returns a uniform bit $b'' \leftarrow \{0, 1\}$.

Let W_{sPKE} be the event that $\mathcal{A}_{\text{sPKE}}$ wins, i.e.,

$$W_{\text{sPKE}} := [b'' = \tilde{b} \wedge \exists j, j' \in \{0, 1\} m_{0,j}^{\text{sPKE}} \neq \perp \neq m_{1,j'}^{\text{sPKE}}],$$

⁶Note that robustness is only defined for encryption and decryption keys generated by sPKE.Gen. Hence, it is important to also condition on $\neg B_S$.

where the messages refer to the ones generated by $\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}$. Note that if $Q \cap V \cap \neg B_E \cap \neg B_S$ does not occur, we have $m_{0,0}^{\text{sPKE}} = m_{1,0}^{\text{sPKE}} = \bar{m} \neq \perp$ by the correctness of sPKE, and thus

$$\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}} \mid \neg(Q \cap V \cap \neg B_E \cap \neg B_S)] = \frac{1}{2}. \quad (20)$$

Next consider the case that $Q \cap V \cap \neg B_E \cap \neg B_S$ occurs. In this case, the view of \mathcal{A} is identical to the one in H_4 with $b = \tilde{b}$, as long as the emulated \mathcal{O}_G never returns \perp . Moreover, if \mathcal{A} wins, we have $m_{0,j}^{H_4} = m_{1,j}^{H_4} = \perp$ for all $j \in J^{H_4}$, where the messages here refer to the ones in H_4 , generated according to (17), and J^{H_4} is the set of all j such that \mathcal{A}_1 or \mathcal{A}_2 issued the query (j, rec) to the oracle \mathcal{O}_G . Therefore, \mathcal{O}_G is never gets a query for which it returns \perp in this case. The event $Q \cap V \cap \neg B_E$ implies that the ciphertexts are encryptions of some message under ek_0^{sPKE} or ek_1^{sPKE} . Correctness of sPKE now implies that $m_{0,0}^{\text{sPKE}} \neq \perp \neq m_{1,0}^{\text{sPKE}}$, i.e., the winning condition for $\mathcal{A}_{\text{sPKE}}$ is satisfied. We can conclude that

$$\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}} \mid Q \cap V \cap \neg B_E \cap \neg B_S] \geq \Pr^{H_4} [W_{\text{ACE}} \mid V \cap \neg B_E \cap \neg B_S]. \quad (21)$$

Let

$$p_G := \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [Q \cap V \cap \neg B_E \cap \neg B_S].$$

Putting our results together, we obtain

$$\begin{aligned} \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}}] &= \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}} \mid Q \cap V \cap \neg B_E \cap \neg B_S] \cdot p_G \\ &\quad + \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}} \mid \neg(Q \cap V \cap \neg B_E \cap \neg B_S)] \cdot (1 - p_G) \\ &\stackrel{(20)}{=} \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}} \mid Q \cap V \cap \neg B_E \cap \neg B_S] \cdot p_G + \frac{1}{2} (1 - p_G). \end{aligned}$$

This implies

$$\begin{aligned} \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}} \mid Q \cap V \cap \neg B_E \cap \neg B_S] &= \frac{1}{p_G} \left[\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}} [W_{\text{sPKE}}] - \frac{1}{2} (1 - p_G) \right] \\ &= \frac{1}{2p_G} \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}} + \frac{1}{2}. \end{aligned} \quad (22)$$

Furthermore,

$$\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-SAN-CCA}} = 2 \cdot \Pr^{H_0} [W_{\text{ACE}}] - 1 \stackrel{(13),(14)}{=} 2 \cdot \left(\text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1} + \Pr^{H_3} [W_{\text{ACE}}] \right) - 1.$$

Since B_E , $\neg B_E \cap B_S$, and $\neg B_E \cap \neg B_S$ partition the sample space, the law of total probability implies

$$\begin{aligned} \Pr^{H_3} [W_{\text{ACE}}] &= \Pr^{H_3} [W_{\text{ACE}} \cap B_E] + \Pr^{H_3} [W_{\text{ACE}} \cap \neg B_E \cap B_S] \\ &\quad + \Pr^{H_3} [W_{\text{ACE}} \cap \neg B_E \cap \neg B_S] \\ &\leq \Pr^{H_3} [B_E] + \Pr^{H_3} [B_S] + \Pr^{H_3} [W_{\text{ACE}} \cap \neg B_E \cap \neg B_S] \\ &\stackrel{(15),(16)}{\leq} 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2} + 2 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + \Pr^{H_3} [W_{\text{ACE}} \cap \neg B_E \cap \neg B_S]. \end{aligned}$$

Note that W_{ACE} implies $c'_0 \neq \perp \neq c'_1$ and thus also V because if the verification fails, ACE.San returns \perp . Hence,

$$\begin{aligned}
& \Pr^{H_3}[W_{\text{ACE}} \cap \neg B_E \cap \neg B_S] = \Pr^{H_3}[W_{\text{ACE}} \cap V \cap \neg B_E \cap \neg B_S] \\
& = \Pr^{H_3}[W_{\text{ACE}} \mid V \cap \neg B_E \cap \neg B_S] \cdot \Pr^{H_3}[V \cap \neg B_E \cap \neg B_S] \\
& \stackrel{(18)}{\leq} \left(\underbrace{\Pr^{H_4}[W_{\text{ACE}} \mid V \cap \neg B_E \cap \neg B_S]} + 2(q_{R_1} + q_{R_2}) \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}} \right) \cdot \Pr^{H_3}[V \cap \neg B_E \cap \neg B_S] \\
& \stackrel{(21)}{\leq} \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}}[W_{\text{sPKE}} \mid Q \cap V \cap \neg B_E \cap \neg B_S] \\
& \stackrel{(22)}{\leq} \left(\frac{1}{2p_G} \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}} + \frac{1}{2} + 2(q_{R_1} + q_{R_2}) \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}} \right) \cdot \Pr^{H_3}[V \cap \neg B_E \cap \neg B_S].
\end{aligned}$$

Since $\Pr^{H_3}[V \cap \neg B_E \cap \neg B_S] = \Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}}[V \cap \neg B_E \cap \neg B_S]$, we have

$$\begin{aligned}
\frac{\Pr^{H_3}[V \cap \neg B_E \cap \neg B_S]}{p_G} &= \frac{\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}}[V \cap \neg B_E \cap \neg B_S]}{\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}}[Q \cap V \cap \neg B_E \cap \neg B_S]} \\
&= \left(\Pr^{\text{Exp}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}}}[Q \mid V \cap \neg B_E \cap \neg B_S] \right)^{-1} \\
&\stackrel{(19)}{\leq} (q_{S_1} + q_{R_1} + q_{D_1})^2.
\end{aligned}$$

Therefore,

$$\Pr^{H_3}[W_{\text{ACE}} \cap \neg B_E \cap \neg B_S] \leq \frac{1}{2} (q_{S_1} + q_{R_1} + q_{D_1})^2 \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}} + \frac{1}{2} + 2(q_{R_1} + q_{R_2}) \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}.$$

This implies

$$\begin{aligned}
\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-SAN-CCA}} &\leq 2 \cdot \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1} + 4 \cdot \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2} + 4 \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} \\
&\quad + (q_{S_1} + q_{R_1} + q_{D_1})^2 \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-SAN-CCA}} + 4(q_{R_1} + q_{R_2}) \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}
\end{aligned}$$

and concludes the proof. \square

We next prove non-detection of fresh encryptions, which directly follows from ciphertext unpredictability of the underlying sPKE scheme.

Lemma D.3. *Let ACE be the scheme from above and let \mathcal{A} be an attacker on the non-detection of fresh encryptions that makes at most q queries to the oracle \mathcal{O}_G . Then, there exist probabilistic algorithms \mathcal{A}_{PRF} and $\mathcal{A}_{\text{sPKE}}$ (which are both roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}}$) such that*

$$\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}} \leq \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + (q + 1) \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{sPKE}}}^{\text{sPKE-UPD-CTXT}}.$$

Proof. Let $H_0 := \text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}}$ and H_1 be as H_0 where F_K is replaced by a truly uniform random function U . As in the proof of [Lemma 6.3](#), one can show that there exists \mathcal{A}_{PRF} such that

$$\Pr^{H_0}[b = 1] - \Pr^{H_1}[b = 1] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}}.$$

The adversary $\mathcal{A}_{\text{sPKE}}$ on input $(sp^{\text{sPKE}}, ek^{\text{sPKE}}, dk^{\text{sPKE}})$, sets $i_{q_0} \leftarrow \perp$, $k_q \leftarrow 1$, chooses $q_0 \leftarrow \{0, \dots, q\}$ uniformly at random, runs $(vk^{\text{Sig}}, sk^{\text{Sig}}) \leftarrow \text{Sig.Gen}(1^\kappa)$, $crs^{\text{NIZK}} \leftarrow \text{NIZK.Gen}(1^\kappa)$,

and gives $sp^{\text{ACE}} := (sp^{\text{SPKE}}, vk^{\text{Sig}}, crs^{\text{NIZK}})$ to \mathcal{A} . It emulates the oracle \mathcal{O}_G for \mathcal{A}_1 as follows. On query (i, t) , if $k_q \neq q_0$ and $i \neq i_{q_0}$, then generate an encryption key $ek_i^{\text{ACE}} := (vk^{\text{Sig}}, ek_i^{\text{SPKE}}, vk_i^{\text{Sig}}, sk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, crs^{\text{NIZK}})$ and a decryption key $dk_i^{\text{ACE}} := dk_i^{\text{SPKE}}$ as H_1 does, where $(ek_i^{\text{SPKE}}, dk_i^{\text{SPKE}})$ is obtained via \mathcal{O}_G and remembered for future queries. Return ek_i^{ACE} if $t = \text{sen}$, and dk_i^{ACE} if $t = \text{rec}$. If $k_q = q_0$ or $i = i_{q_0}$, replace ek_i^{SPKE} and dk_i^{SPKE} by ek^{SPKE} and dk^{SPKE} , respectively, and set $i_{q_0} \leftarrow i$. In both cases, set $k_q \leftarrow k_q + 1$ at the end. When \mathcal{A} returns $(m, i, c = (\tilde{c}, \pi^{\text{NIZK}}))$, $\mathcal{A}_{\text{SPKE}}$ returns (m, \tilde{c}) .

Let Q be the event that $i_{q_0} = i$, or $q_0 = 0$ and \mathcal{A} does not make the query (i, sen) or (i, rec) to \mathcal{O}_G . Note that the probability of Q is $1/(q+1)$ and since $b = \text{ACE.DMod}(sp^{\text{ACE}}, (\tilde{c}, \pi^{\text{NIZK}}), (\tilde{c}^*, \pi^{\text{NIZK}^*})) = 1$ if and only if $\tilde{c}^* = \tilde{c}$, we have

$$\Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}_{\text{SPKE}}}^{\text{SPKE-UPD-CTXT}}} [c = c^* \mid Q] = \Pr^{H_1} [b = 1].$$

Hence, we can conclude

$$\begin{aligned} \text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-NDTCT-FENC}} &= \Pr^{H_0} [b = 1] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \Pr^{H_1} [b = 1] \\ &= \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}_{\text{SPKE}}}^{\text{SPKE-UPD-CTXT}}} [c = c^* \mid Q] \\ &\leq \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + (q+1) \cdot \Pr^{\text{Exp}_{\text{SPKE}, \mathcal{A}_{\text{SPKE}}}^{\text{SPKE-UPD-CTXT}}} [c = c^*] \\ &= \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + (q+1) \cdot \text{Adv}_{\text{SPKE}, \mathcal{A}_{\text{SPKE}}}^{\text{SPKE-UPD-CTXT}}. \quad \square \end{aligned}$$

We finally prove the uniform decryption and role-respecting properties.

Lemma D.4. *Let ACE be the scheme from above and let \mathcal{A} be an attacker on the uniform-decryption security that makes at most q_R queries of the form (\cdot, rec) to the oracle \mathcal{O}_G . Then, there exist probabilistic algorithms \mathcal{A}_{PRF} , $\mathcal{A}_{\text{ZK}_1}$, $\mathcal{A}_{\text{ZK}_2}$, \mathcal{A}_{Sig} , and \mathcal{A}_{rob} (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-URR}}$) such that*

$$\text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-UDEC}} \leq \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2} + \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + q_R \cdot \text{Adv}_{\text{SPKE}, \mathcal{A}_{\text{rob}}}^{\text{SPKE-USROB}}.$$

Proof. Note that we can assume without loss of generality that \mathcal{A} does not use the oracle \mathcal{O}_E since obtaining encryption keys from \mathcal{O}_G does not decrease the advantage. Let $H_0 := \text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-URR}}$ and let W_{UDec} be the event that \mathcal{A} wins the uniform-decryption game:

$$W_{\text{UDec}} := [\exists j, j' \in J \ m_j \neq \perp \neq m_{j'} \wedge m_j \neq m_{j'}].$$

As in the proof of Lemma D.2, let H_1 be as H_0 with F_K replaced by a uniform random function U , let H_2 be as H_1 with crs^{NIZK} being generated by E_1^{NIZK} , and let H_3 be as H_2 , but after \mathcal{A} returns $c = (\tilde{c}, \pi^{\text{NIZK}})$, a witness

$$w := (ek_{i_w}^{\text{SPKE}}, m_w, r_w, vk_{i_w}^{\text{Sig}}, \sigma_{i_w}^{\text{Sig}}, \sigma_{c,w}^{\text{Sig}})$$

for the statement $x := (vk^{\text{Sig}}, \tilde{c})$ is extracted from the proof π^{NIZK} by E_2^{NIZK} . We define the events $V := [\text{NIZK.Ver}(crs^{\text{NIZK}}, x, \pi^{\text{NIZK}}) = 1]$, $B_E := [V \wedge (x, w) \notin R]$, and B_S as the event that $(x, w) \in R$ and $ek_{i_w}^{\text{SPKE}}$ is not contained in an answer from \mathcal{O}_G to \mathcal{A} . It can be shown as in the proof of Lemma D.2 that there exist \mathcal{A}_{PRF} , $\mathcal{A}_{\text{ZK}_1}$, $\mathcal{A}_{\text{ZK}_2}$, and \mathcal{A}_{Sig} such that

$$\begin{aligned} \Pr^{H_0} [W_{\text{UDec}}] - \Pr^{H_3} [W_{\text{UDec}}] &= \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1}, \\ \Pr^{H_3} [B_E] &\leq \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2}, \\ \Pr^{H_3} [B_S] &\leq \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}}, \end{aligned}$$

where the last inequality uses that \mathcal{A} does not query the oracle \mathcal{O}_E . Now let H_4 be as H_3 where for $j \in J$, $m_j \leftarrow \text{ACE.Dec}(\text{ACE.Gen}(msk, j, \text{rec}), c')$ is replaced by

$$m_j \leftarrow \begin{cases} m_w, & ek_j^{\text{sPKE}} = ek_{i_w}^{\text{sPKE}} \text{ for } (ek_j^{\text{sPKE}}, dk_j^{\text{sPKE}}) = \text{sPKE.Gen}(msk^{\text{sPKE}}; U([j, 0])), \\ \perp, & \text{else.} \end{cases}$$

One can show as in the proof of [Lemma D.2](#) that there exists a probabilistic algorithm \mathcal{A}_{rob} such that

$$\Pr^{H_3}[W_{\text{UDec}} \mid V \cap \neg B_E \cap \neg B_S] - \Pr^{H_4}[W_{\text{UDec}} \mid V \cap \neg B_E \cap \neg B_S] \leq q_R \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}.$$

Note that \mathcal{A} cannot win in H_4 since if $m_j \neq \perp \neq m_{j'}$, then $m_j = m_w = m_{j'}$. This implies that $\Pr^{H_3}[W_{\text{UDec}} \mid V \cap \neg B_E \cap \neg B_S] \leq q_R \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}$. Note that \mathcal{A} can only win in H_3 if V occurs since otherwise $c' = \perp$ and consequently $m_j = \perp$ for all $j \in J$. We therefore obtain

$$\begin{aligned} \Pr^{H_3}[W_{\text{UDec}}] &= \Pr^{H_3}[W_{\text{UDec}} \cap V \cap B_E] + \Pr^{H_3}[W_{\text{UDec}} \cap V \cap \neg B_E \cap B_S] \\ &\quad + \Pr^{H_3}[W_{\text{UDec}} \cap V \cap \neg B_E \cap \neg B_S] \\ &\leq \Pr^{H_3}[B_E] + \Pr^{H_3}[B_S] + \Pr^{H_3}[W_{\text{UDec}} \mid V \cap \neg B_E \cap \neg B_S] \\ &\leq \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2} + \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} + q_R \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}. \end{aligned}$$

Together with $\Pr^{H_0}[W_{\text{UDec}}] - \Pr^{H_3}[W_{\text{UDec}}] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1}$, this concludes the proof. \square

Lemma D.5. *Let ACE be the scheme from above and let \mathcal{A} be an attacker on the role-respecting security that makes at most q_S queries of the form (\cdot, sen) and at most q_R queries of the form (\cdot, rec) to the oracle \mathcal{O}_G , and at most q_E queries to the oracle \mathcal{O}_E . Then, there exist probabilistic algorithms \mathcal{A}_{PRF} , $\mathcal{A}_{\text{ZK}_1}$, $\mathcal{A}_{\text{ZK}_2}$, \mathcal{A}_{Sig} , and \mathcal{A}_{rob} (which are all roughly as efficient as emulating an execution of $\text{Exp}_{\text{ACE}, \mathcal{A}}^{\text{ACE-URR}}$) such that*

$$\begin{aligned} \text{Adv}_{\text{ACE}, \mathcal{A}}^{\text{ACE-RR}} &\leq \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2} + (q_E + 1) \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} \\ &\quad + q_R \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}} + (q_S + q_R + q_E)^2 \cdot \text{Col}_{\text{sPKE}}^{\text{ek}}. \end{aligned}$$

Proof. Let $H_0, \dots, H_4, V := [\text{NIZK.Ver}(crs^{\text{NIZK}}, x, \pi^{\text{NIZK}}) = 1]$, and $B_E := [V \wedge (x, w) \notin R]$ for the statement $x := (vk^{\text{Sig}}, \tilde{c})$ and the extracted witness $w := (ek_{i_w}^{\text{sPKE}}, m_w, r_w, vk_{i_w}^{\text{Sig}}, \sigma_{i_w}^{\text{Sig}}, \sigma_{c, w}^{\text{Sig}})$ be defined as in the proof of [Lemma D.4](#), and let W_{RR} be the event that \mathcal{A} wins the role-respecting game:

$$W_{\text{RR}} := [c' \neq \perp \wedge \text{dct} = \text{false} \wedge \neg(\exists i \in I \forall j \in J (m_j \neq \perp \leftrightarrow P(i, j) = 1))].$$

As in that proof, there exist \mathcal{A}_{PRF} , $\mathcal{A}_{\text{ZK}_1}$, and $\mathcal{A}_{\text{ZK}_2}$ such that

$$\Pr^{H_0}[W_{\text{RR}}] - \Pr^{H_3}[W_{\text{RR}}] = \text{Adv}_{F, \mathcal{A}_{\text{PRF}}}^{\text{PRF}} + \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_1}}^{\text{NIZK-ext}_1}, \quad (23)$$

and

$$\Pr^{H_3}[B_E] \leq \text{Adv}_{\text{NIZK}, \mathcal{A}_{\text{ZK}_2}}^{\text{NIZK-ext}_2}. \quad (24)$$

Let E_G be the event that the extracted key $ek_{i_w}^{\text{sPKE}}$ is contained in an answer from \mathcal{O}_G to \mathcal{A} . One can show similarly as in the proof of Lemma D.2 that there exists an algorithm \mathcal{A}_{rob} such that

$$\Pr^{H_3}[W_{\text{RR}} \cap V \cap \neg B_E \cap E_G] - \Pr^{H_4}[W_{\text{RR}} \cap V \cap \neg B_E \cap E_G] \leq q_R \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}}. \quad (25)$$

We first show that if V , $\neg B_E$, and E_G occur in H_4 , \mathcal{A} can only win if two encryption keys generated by sPKE.Gen are equal, which happens only with negligible probability.

Claim 1. *We have*

$$\Pr^{H_4}[W_{\text{RR}} \cap V \cap \neg B_E \cap E_G] \leq (q_S + q_R + q_E)^2 \cdot \text{Col}_{\text{sPKE}}^{\text{ek}}.$$

Proof of claim. If V , $\neg B_E$, and E_G occur, then there is an $i_0 \in I$ such that $ek_{i_0}^{\text{sPKE}} = ek_{i_w}^{\text{sPKE}}$ for $(ek_{i_0}^{\text{sPKE}}, dk_{i_0}^{\text{sPKE}}) = \text{sPKE.Gen}(msk^{\text{sPKE}}; U([i_0, 0]))$. Using $P(i, j) = 1 \Leftrightarrow i = j$, we have that \mathcal{A} only wins if there exists $j \in J \setminus \{i_0\}$ such that $m_j \neq \perp$ or if $i_0 \in J$ and $m_{i_0} = \perp$. Because in H_4 , m_j for $j \in J$ is equal to m_w if $ek_j^{\text{sPKE}} = ek_{i_w}^{\text{sPKE}}$ for $(ek_j^{\text{sPKE}}, dk_j^{\text{sPKE}}) = \text{sPKE.Gen}(msk^{\text{sPKE}}; U([j, 0]))$, and \perp otherwise, we have $m_{i_0} \neq \perp$ if $i_0 \in J$. Moreover, for $i_0 \neq j \in J$, we have $m_j = \perp$ unless $ek_j^{\text{sPKE}} = ek_{i_0}^{\text{sPKE}}$. This means that \mathcal{A} can only win if sPKE.Gen generates the same encryption key for the randomness values $U([i_0, 0])$ and $U([j, 0])$ for some $i_0 \neq j \in J$. Since at most $q_S + q_R + q_E$ key pairs are generated in the experiment, there are at most $(q_S + q_R + q_E)^2$ pairs of encryption keys that could collide. For each such pair, the collision probability is bounded by $\text{Col}_{\text{sPKE}}^{\text{ek}}$ because for $i \neq i'$, $U([i, 0])$ and $U([i', 0])$ are independent and uniformly distributed. Hence, the claim follows. \diamond

Now let E_E be the event that \mathcal{A} made a query (i, \cdot) to \mathcal{O}_E such that $ek_i^{\text{sPKE}} = ek_{i_w}^{\text{sPKE}}$ and $vk_i^{\text{Sig}} = vk_{i_w}^{\text{Sig}}$ for $(ek_i^{\text{sPKE}}, dk_i^{\text{sPKE}}) = \text{sPKE.Gen}(msk^{\text{sPKE}}; U([i, 0]))$ and $(vk_i^{\text{Sig}}, sk_i^{\text{Sig}}) = \text{Sig.Gen}(1^\kappa; U([i, 1]))$. We next show that if \mathcal{A} wins and $V \cap \neg B_E \cap \neg E_G \cap E_E$ occurs, \mathcal{A} forged a signature on \tilde{c} .

Claim 2. *There exists a probabilistic algorithm $\mathcal{A}_{\text{Sig}_1}$ such that*

$$\Pr^{H_3}[W_{\text{RR}} \cap V \cap \neg B_E \cap \neg E_G \cap E_E] \leq q_E \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}.$$

Proof of claim. On input vk^{Sig^*} , $\mathcal{A}_{\text{Sig}_1}$ initializes $i_{q_0} \leftarrow \perp$, $k_q \leftarrow 1$, chooses $q_0 \leftarrow \{1, \dots, q_E\}$ uniformly at random, generates $(sp^{\text{sPKE}}, msk^{\text{sPKE}}) \leftarrow \text{sPKE.Setup}(1^\kappa)$, $(vk^{\text{Sig}}, sk^{\text{Sig}}) \leftarrow \text{Sig.Gen}(1^\kappa)$, and $(crs^{\text{NIZK}}, \xi^{\text{NIZK}}) \leftarrow E_1^{\text{NIZK}}(1^\kappa)$ as H_3 , and gives $sp^{\text{ACE}} := (sp^{\text{sPKE}}, vk^{\text{Sig}}, crs^{\text{NIZK}})$ to \mathcal{A} . It emulates the oracles for \mathcal{A} as follows.

$\mathcal{O}_G(\cdot, \cdot)$: Generate the requested key exactly as H_3 does and return it.

$\mathcal{O}_E(\cdot, \cdot)$: On query (i, m) , if $k_q \neq q_0$ and $i \neq i_{q_0}$, generate an encryption key ek_i^{ACE} as H_3 , encrypt m using ek_i^{ACE} , and return the resulting ciphertext. If $k_q = q_0$ or $i = i_{q_0}$, set $i_{q_0} \leftarrow i$, execute $(ek_i^{\text{sPKE}}, dk_i^{\text{sPKE}}) \leftarrow \text{sPKE.Gen}(msk^{\text{sPKE}}; U([i, 0]))$, $\sigma_i^{\text{Sig}} \leftarrow \text{Sig.Sign}(sk^{\text{Sig}}, [ek_i^{\text{sPKE}}, vk_i^{\text{Sig}}]; U([i, 2]))$, and set $vk_i^{\text{Sig}} := vk_i^{\text{Sig}^*}$. Then, sample randomness r and compute $\tilde{c} \leftarrow \text{sPKE.Enc}(ek_i^{\text{sPKE}}, m; r)$, query the signing oracle on \tilde{c} to obtain a signature $\sigma_{\tilde{c}}^{\text{Sig}}$, and run

$$\pi^{\text{NIZK}} \leftarrow \text{NIZK.Prove}(crs^{\text{NIZK}}, x := (vk_i^{\text{Sig}}, \tilde{c}), w := (ek_i^{\text{sPKE}}, m, r, vk_i^{\text{Sig}}, \sigma_i^{\text{Sig}}, \sigma_{\tilde{c}}^{\text{Sig}}).$$

Finally, return the ciphertext $c := (\tilde{c}, \pi^{\text{NIZK}})$. In all cases, set $k_q \leftarrow k_q + 1$ at the end.

When \mathcal{A} returns $c = (\tilde{c}, \pi^{\text{NIZK}})$, $\mathcal{A}_{\text{Sig}_1}$ extracts a witness

$$w := (ek_{i_w}^{\text{sPKE}}, m_w, r_w, vk_{i_w}^{\text{Sig}}, \sigma_{i_w}^{\text{Sig}}, \sigma_{c,w}^{\text{Sig}}) \leftarrow E_2^{\text{NIZK}}(crs^{\text{NIZK}}, \xi^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}).$$

It finally returns the forgery attempt $(\tilde{c}, \sigma_{c,w}^{\text{Sig}})$.

Note that if \mathcal{A} wins the role-respecting game, $\text{ACE.DMod}(sp^{\text{ACE}}, \hat{c}, c) = 0$ for all \hat{c} that \mathcal{O}_E has returned. Since ACE.DMod checks for equality of sPKE ciphertexts, this means that $\mathcal{A}_{\text{Sig}_1}$ has not issued the query \tilde{c} to its signing oracle. Furthermore, if the extraction and verification succeed, $\sigma_{c,w}^{\text{Sig}}$ is a valid signature for \tilde{c} . Let Q be the event that $ek_{i_{q_0}}^{\text{sPKE}} = ek_{i_w}^{\text{sPKE}}$ and $vk_{i_{q_0}}^{\text{Sig}} = vk_{i_w}^{\text{Sig}}$. If Q and $V \cap \neg B_E \cap \neg E_G \cap E_E$ occur, \mathcal{A} has not requested $ek_{i_{q_0}}^{\text{ACE}}$ from \mathcal{O}_G and hence $\mathcal{A}_{\text{Sig}_1}$ perfectly emulates H_3 . This implies

$$\Pr^{\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}} [W_{\text{Sig}} \mid V \cap \neg B_E \cap \neg E_G \cap E_E \cap Q] \geq \Pr^{H_3} [W_{\text{RR}} \mid V \cap \neg B_E \cap \neg E_G \cap E_E],$$

where W_{Sig} denotes the event that $\mathcal{A}_{\text{Sig}_1}$ wins in the signature forgery game. We further have

$$\Pr^{\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}} [Q \mid V \cap \neg B_E \cap \neg E_G \cap E_E] = 1/q_E.$$

This implies for $p_G := \Pr^{\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}} [V \cap \neg B_E \cap \neg E_G \cap E_E \cap Q]$,

$$\begin{aligned} \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}} &= \Pr^{\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}} [W_{\text{Sig}}] \geq \Pr^{\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}} [W_{\text{Sig}} \mid V \cap \neg B_E \cap \neg E_G \cap E_E \cap Q] \cdot p_G \\ &\geq \Pr^{H_3} [W_{\text{RR}} \mid V \cap \neg B_E \cap \neg E_G \cap E_E] \cdot p_G \\ &= \Pr^{H_3} [W_{\text{RR}} \cap V \cap \neg B_E \cap \neg E_G \cap E_E] \cdot \frac{p_G}{\Pr^{H_3} [V \cap \neg B_E \cap \neg E_G \cap E_E]}. \end{aligned}$$

Since $[V \cap \neg B_E \cap \neg E_G \cap E_E]$ in H_3 has the same probability as in $\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}$, we have

$$\frac{p_G}{\Pr^{H_3} [V \cap \neg B_E \cap \neg E_G \cap E_E]} = \Pr^{\text{Exp}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}}} [Q \mid V \cap \neg B_E \cap \neg E_G \cap E_E] = \frac{1}{q_E},$$

which implies the claim. \diamond

Finally, we show that if \mathcal{A} wins and $V \cap \neg B_E \cap \neg E_G \cap \neg E_E$ occurs, \mathcal{A} forged a signature on $[ek_{i_w}^{\text{sPKE}}, vk_{i_w}^{\text{Sig}}]$.

Claim 3. *There exists a probabilistic algorithm $\mathcal{A}_{\text{Sig}_2}$ such that*

$$\Pr^{H_3} [W_{\text{RR}} \cap V \cap \neg B_E \cap \neg E_G \cap \neg E_E] \leq \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_2}}^{\text{Sig-EUF-CMA}}.$$

Proof of claim. The algorithm $\mathcal{A}_{\text{Sig}_2}$ on input vk^{Sig^*} runs $(sp^{\text{sPKE}}, msk^{\text{sPKE}}) \leftarrow \text{sPKE.Setup}(1^\kappa)$ and $(crs^{\text{NIZK}}, \xi^{\text{NIZK}}) \leftarrow E_1^{\text{NIZK}}(1^\kappa)$, and gives $sp^{\text{ACE}} := (sp^{\text{sPKE}}, vk^{\text{Sig}^*}, crs^{\text{NIZK}})$ to \mathcal{A} . It emulates the oracles for \mathcal{A} as follows.

$\mathcal{O}_G(\cdot, \cdot)$: Generate the requested key as H_3 , but obtain the signature σ_i^{Sig} via a query to the signing oracle. Remember the signature and when asked again for the same i , reuse σ_i^{Sig} instead of issuing another query. This ensures that the oracle behaves as the one in H_3 and returns the same key for repeated queries.

$\mathcal{O}_E(\cdot, \cdot)$: On query (i, m) , generate an encryption key as for a query (i, sen) to \mathcal{O}_G , encrypt m using that key, and return the resulting ciphertext.

When \mathcal{A} returns $c = (\tilde{c}, \pi^{\text{NIZK}})$, $\mathcal{A}_{\text{Sig}_1}$ extracts a witness

$$w := (ek_{i_w}^{\text{sPKE}}, m_w, r_w, vk_{i_w}^{\text{Sig}}, \sigma_{i_w}^{\text{Sig}}, \sigma_{c,w}^{\text{Sig}}) \leftarrow E_2^{\text{NIZK}}(crs^{\text{NIZK}}, \xi^{\text{NIZK}}, x := (vk^{\text{Sig}}, \tilde{c}), \pi^{\text{NIZK}}).$$

It finally returns the forgery attempt $([ek_{i_w}^{\text{sPKE}}, vk_{i_w}^{\text{Sig}}], \sigma_{i_w}^{\text{Sig}})$. Note that if $W_{\text{RR}} \cap V \cap \neg B_E \cap \neg E_G \cap \neg E_E$ occurs, $\sigma_{i_w}^{\text{Sig}}$ is a valid signature for $[ek_{i_w}^{\text{sPKE}}, vk_{i_w}^{\text{Sig}}]$ and $\mathcal{A}_{\text{Sig}_2}$ has not requested a signature for this value from the signing oracle. Therefore, $\mathcal{A}_{\text{Sig}_2}$ wins the forgery game and thus the probability of that event is bounded by $\text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_2}}^{\text{Sig-EUF-CMA}}$. \diamond

Combining [Claims 2](#) and [3](#), we obtain

$$\Pr^{H_3}[W_{\text{RR}} \cap V \cap \neg B_E \cap \neg E_G] \leq q_E \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}} + \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_2}}^{\text{Sig-EUF-CMA}}.$$

Let \mathcal{A}_{Sig} be the algorithm that runs $\mathcal{A}_{\text{Sig}_1}$ with probability $\frac{q_E}{q_E+1}$ and $\mathcal{A}_{\text{Sig}_2}$ with probability $\frac{1}{q_E+1}$. We then have

$$\begin{aligned} \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}} &= \frac{q_E}{q_E+1} \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_1}}^{\text{Sig-EUF-CMA}} + \frac{1}{q_E+1} \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}_2}}^{\text{Sig-EUF-CMA}} \\ &\geq \frac{1}{q_E+1} \cdot \Pr^{H_3}[W_{\text{RR}} \cap V \cap \neg B_E \cap \neg E_G]. \end{aligned} \quad (26)$$

Note that W_{RR} implies $c' \neq \perp$ and therefore V , i.e., the events W_{RR} and $W_{\text{RR}} \cap V$ are equal. Thus,

$$\begin{aligned} \Pr^{H_3}[W_{\text{RR}}] &= \Pr^{H_3}[W_{\text{RR}} \cap B_E] + \Pr^{H_3}[W_{\text{RR}} \cap V \cap \neg B_E \cap E_G] + \Pr^{H_3}[W_{\text{RR}} \cap V \cap \neg B_E \cap \neg E_G] \\ &\stackrel{(24),(25),(26)}{\leq} \text{Adv}_{\text{NIZK}, \mathcal{A}_{ZK_2}}^{\text{NIZK-ext}_2} + q_R \cdot \text{Adv}_{\text{sPKE}, \mathcal{A}_{\text{rob}}}^{\text{sPKE-USROB}} + \Pr^{H_4}[W_{\text{RR}} \cap V \cap \neg B_E \cap E_G] \\ &\quad + (q_E + 1) \cdot \text{Adv}_{\text{Sig}, \mathcal{A}_{\text{Sig}}}^{\text{Sig-EUF-CMA}}. \end{aligned}$$

Combined with [Claim 1](#) and [equation \(23\)](#), this concludes the proof. \square

References

- [ABN10] M. Abdalla, M. Bellare, and G. Neven, “Robust encryption,” in *Theory of Cryptography: 7th Theory of Cryptography Conference, TCC 2010*, D. Micciancio, Ed., Springer Berlin Heidelberg, 2010, pp. 480–497. DOI: [10.1007/978-3-642-11799-2_28](https://doi.org/10.1007/978-3-642-11799-2_28).
- [BL73] D. E. Bell and L. J. LaPadula, “Secure computer systems: Mathematical foundations,” MITRE, Tech. Rep. MTR-2547, 1973.
- [BBDP01] M. Bellare, A. Boldyreva, A. Desai, and D. Pointcheval, “Key-privacy in public-key encryption,” in *Advances in Cryptology — ASIACRYPT 2001: 7th International Conference on the Theory and Application of Cryptology and Information Security*, C. Boyd, Ed., Springer Berlin Heidelberg, 2001, pp. 566–582. DOI: [10.1007/3-540-45682-1_33](https://doi.org/10.1007/3-540-45682-1_33).

- [BSW11] D. Boneh, A. Sahai, and B. Waters, “Functional encryption: Definitions and challenges,” in *Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011*, Y. Ishai, Ed., Springer Berlin Heidelberg, 2011, pp. 253–273. DOI: [10.1007/978-3-642-19571-6_16](https://doi.org/10.1007/978-3-642-19571-6_16).
- [CKN03] R. Canetti, H. Krawczyk, and J. B. Nielsen, “Relaxing chosen-ciphertext security,” in *Advances in Cryptology - CRYPTO 2003: 23rd Annual International Cryptology Conference*, D. Boneh, Ed., Springer Berlin Heidelberg, 2003, pp. 565–582. DOI: [10.1007/978-3-540-45146-4_33](https://doi.org/10.1007/978-3-540-45146-4_33).
- [DHO16] I. Damgård, H. Haagh, and C. Orlandi, “Access control encryption: Enforcing information flow with cryptography,” in *Theory of Cryptography: 14th International Conference, TCC 2016-B*, M. Hirt and A. Smith, Eds., Springer Berlin Heidelberg, 2016, pp. 547–576. DOI: [10.1007/978-3-662-53644-5_21](https://doi.org/10.1007/978-3-662-53644-5_21).
- [Elg85] T. Elgamal, “A public key cryptosystem and a signature scheme based on discrete logarithms,” *IEEE Transactions on Information Theory*, vol. 31, no. 4, pp. 469–472, Jul. 1985. DOI: [10.1109/TIT.1985.1057074](https://doi.org/10.1109/TIT.1985.1057074).
- [FLPQ13] P. Farshim, B. Libert, K. G. Paterson, and E. A. Quaglia, “Robust encryption, revisited,” in *Public-Key Cryptography – PKC 2013: 16th International Conference on Practice and Theory in Public-Key Cryptography*, K. Kurosawa and G. Hanaoka, Eds., Springer Berlin Heidelberg, 2013, pp. 352–368. DOI: [10.1007/978-3-642-36362-7_22](https://doi.org/10.1007/978-3-642-36362-7_22).
- [FGKO17] G. Fuchsbauer, R. Gay, L. Kowalczyk, and C. Orlandi, “Access control encryption for equality, comparison, and more,” in *Public-Key Cryptography – PKC 2017: 20th IACR International Conference on Practice and Theory in Public-Key Cryptography*, S. Fehr, Ed., Springer Berlin Heidelberg, 2017, pp. 88–118. DOI: [10.1007/978-3-662-54388-7_4](https://doi.org/10.1007/978-3-662-54388-7_4).
- [GJJS04] P. Golle, M. Jakobsson, A. Juels, and P. Syverson, “Universal re-encryption for mixnets,” in *Topics in Cryptology – CT-RSA 2004: The Cryptographers’ Track at the RSA Conference 2004*, T. Okamoto, Ed., Springer Berlin Heidelberg, 2004, pp. 163–178. DOI: [10.1007/978-3-540-24660-2_14](https://doi.org/10.1007/978-3-540-24660-2_14).
- [GPSW06] V. Goyal, O. Pandey, A. Sahai, and B. Waters, “Attribute-based encryption for fine-grained access control of encrypted data,” in *Proceedings of the 13th ACM Conference on Computer and Communications Security*, ser. CCS ’06, ACM, 2006, pp. 89–98. DOI: [10.1145/1180405.1180418](https://doi.org/10.1145/1180405.1180418).
- [Gro04] J. Groth, “Rerandomizable and replayable adaptive chosen ciphertext attack secure cryptosystems,” in *Theory of Cryptography: First Theory of Cryptography Conference, TCC 2004*, M. Naor, Ed., Springer Berlin Heidelberg, 2004, pp. 152–170. DOI: [10.1007/978-3-540-24638-1_9](https://doi.org/10.1007/978-3-540-24638-1_9).
- [Gro06] —, “Simulation-sound NIZK proofs for a practical language and constant size group signatures,” in *Advances in Cryptology – ASIACRYPT 2006: 12th International Conference on the Theory and Application of Cryptology and Information Security*, X. Lai and K. Chen, Eds., Springer Berlin Heidelberg, 2006, pp. 444–459. DOI: [10.1007/11935230_29](https://doi.org/10.1007/11935230_29).

- [KW17] S. Kim and D. J. Wu, “Access control encryption for general policies from standard assumptions,” in *Advances in Cryptology—ASIACRYPT 2017*, to appear, Springer Berlin Heidelberg, 2017.
- [Lin06] Y. Lindell, “A simpler construction of CCA2-secure public-key encryption under general assumptions,” *Journal of Cryptology*, vol. 19, no. 3, pp. 359–377, 2006. DOI: [10.1007/s00145-005-0345-x](https://doi.org/10.1007/s00145-005-0345-x).
- [NY90] M. Naor and M. Yung, “Public-key cryptosystems provably secure against chosen ciphertext attacks,” in *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing*, ser. STOC ’90, ACM, 1990, pp. 427–437. DOI: [10.1145/100216.100273](https://doi.org/10.1145/100216.100273).
- [PR07] M. Prabhakaran and M. Rosulek, “Rerandomizable RCCA encryption,” in *Advances in Cryptology - CRYPTO 2007: 27th Annual International Cryptology Conference*, A. Menezes, Ed., Springer Berlin Heidelberg, 2007, pp. 517–534. DOI: [10.1007/978-3-540-74143-5_29](https://doi.org/10.1007/978-3-540-74143-5_29).
- [Sah99] A. Sahai, “Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security,” in *40th Annual Symposium on Foundations of Computer Science*, 1999, pp. 543–553. DOI: [10.1109/SFFCS.1999.814628](https://doi.org/10.1109/SFFCS.1999.814628).
- [SW05] A. Sahai and B. Waters, “Fuzzy identity-based encryption,” in *Advances in Cryptology - EUROCRYPT 2005: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, R. Cramer, Ed., Springer Berlin Heidelberg, 2005, pp. 457–473. DOI: [10.1007/11426639_27](https://doi.org/10.1007/11426639_27).
- [TZMT17] G. Tan, R. Zhang, H. Ma, and Y. Tao, “Access control encryption based on LWE,” in *Proceedings of the 4th ACM International Workshop on ASIA Public-Key Cryptography*, ser. APKC ’17, ACM, 2017, pp. 43–50. DOI: [10.1145/3055504.3055509](https://doi.org/10.1145/3055504.3055509).