On the Impossibility of Sender-Deniable Public Key Encryption

Dana Dachman-Soled[⋆]

Microsoft Research New England

Abstract. The primitive of deniable encryption was first introduced by Canetti et al. (CRYPTO, 1997). Deniable encryption is a regular public key encryption scheme with the added feature that after running the protocol honestly and transmitting a message m, both Sender and Receiver may produce random coins showing that the transmitted ciphertext was an encryption of any message m' in the message space. Deniable encryption is a key tool for constructing incoercible protocols, since it allows a party to send one message and later provide apparent evidence to a coercer that a different message was sent. In addition, deniable encryption may be used to obtain *adaptively*-secure multiparty computation (MPC) protocols and is secure under *selective-opening* attacks. Different flavors such as sender-deniable and receiver-deniable encryption, where only the Sender or Receiver can produce fake random coins, have been considered.

Recently, several open questions regarding the feasibility of deniable encryption have been resolved (c.f. (O'Neill et al., CRYPTO, 2011), (Bendlin et al., ASIACRYPT, 2011)). A fundamental remaining open question is whether it is possible to construct sender-deniable Encryption Schemes with super-polynomial security, where an adversary has negligible advantage in distinguishing real and fake openings.

The primitive of simulatable public key encryption (PKE), introduced by Damgård and Nielsen (CRYPTO, 2000), is a public key encryption scheme with additional properties that allow oblivious sampling of public keys and ciphertexts. It is one of the low-level primitives used to construct adaptively-secure MPC protocols and was used by O'Neill et al. in their construction of bi-deniable encryption in the multi-distributional model (CRYPTO, 2011). Moreover, the original construction of sender-deniable encryption with polynomial security given by Canetti et al. can be instantiated with simulatable PKE. Thus, a natural question to ask is whether it is possible to construct sender-deniable encryption with *super-polynomial security* from simulatable PKE.

In this work, we investigate the possibility of constructing sender-deniable public key encryption from the primitive of simulatable PKE in a black-box manner. We show that, in fact, there is no black-box construction of sender-deniable encryption with super-polynomial security from simulatable PKE. This indicates that the original construction of sender-deniable public key encryption given by Canetti et al. is in some sense optimal, since improving on it will require the use of non-black-box techniques, stronger underlying assumptions or interaction.

Key words: sender-deniable encryption, simulatable PKE, black-box separation

[⋆] Email: dadachma@microsoft.com

1 Introduction

Deniable encryption was first introduced by Canetti et al. [3]. In its strongest form, called bi-deniable encryption, this primitive allows a Sender and Receiver to communicate via a public key encryption scheme (sending some message m) and then later allows both parties to produce apparent evidence (i.e. secret key and random coins) that the ciphertext sent/received was actually an encryption of any message m' in the message space. Deniable encryption is useful for designing protocols that resist coercion (c.f. [5]) as well as for designing *adaptively*-secure protocols. Moreover, deniable encryption is is secure under *selective-opening* attacks. As a concrete example, consider a voting scheme where parties encrypt their votes using the voting authority's public key and send the ciphertext to the voting authority over a public channel. The voting authority is then trusted to decrypt and tally the votes[1]. In the voting scheme described, voters can carry away a *receipt*, the ciphertext sent to the authority along with the random coins used to encrypt, which can later be used to prove to a third party that a particular vote was cast. Although obtaining a receipt may seem desirable, it also means that voters or the voting authority can later be coerced by some third party to reveal the vote cast by a particular ciphertext. Thus, such a voting scheme is highly susceptible to coercion. However, using a bi-deniable encryption scheme instead of a regular public key encryption scheme allows both the voters and the authority to claim that a specific ciphertext corresponds to a vote for a particular candidate regardless of the actual effective vote. One may also consider weaker versions of bi-deniable encryption such as sender-deniable encryption and receiver-deniable encryption, where only the Sender (resp. Receiver) can produce fake coins.

Constructing deniable encryption schemes seems difficult due to two conflicting goals: Parties must be able to communicate effectively with each other, but if coerced, both parties must be able to produce seemingly correctly distributed randomness and/or secret keys consistent with *any* message m in the message space. Now it seems that surely deniability must interfere with effective communication since the Receiver cannot tell which message m was the intended message and the Sender cannot be assured that his intended message m was received. Indeed, it was shown by [2] that (non-interactive) receiver-deniable encryption (with negligible distinguishing advantage), and thus (non-interactive) bi-deniable encryption is impossible to achieve. However, the existence of a sender-deniable encryption scheme, although paradoxical for the same reasons as above, has not been ruled out. Moreover, sender-deniable encryption schemes can still be effective in preventing coercion in cases such as the voting scheme described above: In the case of voting schemes, individual voters are highly susceptible to coercion, but succeeding in coercing the voting authority to reveal its secret key is unlikely.

There are, in fact, known constructions of deniable encryption [3] with *non-negligible* distinguishing advantage, where an adversary may distinguish real and fake openings of ciphertexts with probability $1/\mathrm{poly}$ for some polynomial. Alternatively, we say that such schemes have *polynomial security*. These constructions get around the paradox above by having the Sender and Receiver share some secret information that is never revealed to the coercing party. To date, over 15 years since the primitive was first introduced, the best known construction of sender-deniable encryption is the one found in the original work of [3] and achieves only polynomial security for some poly, where poly depends directly on the ciphertext size. This leaves open the following important question pointed out by [2]:

> Is it possible to construct sender-deniable public-key encryption with better than polynomial security?

[1] Alternatively, the voting authority may be required to give a zero-knowledge proof that the final tally is consistent with the transmitted ciphertexts.

Relationship to Adaptive Security and Simulatable Public Key Encryption

There is a strong link between deniable encryption and another primitive known as *non-committing encryption* [4]. The main difference between the two is that a Non-Committing Encryption scheme consists of two sets of Key Generation and Encryption algorithms–one for honest players and one for the simulator. Moreover, only honest parties need to communicate effectively, while only the simulator needs to equivocate ciphertexts. Both deniable encryption and non-committing encryption can be used to achieve *adaptively* secure multiparty computation and both are secure under *selective opening* attacks. One of the standard low-level assumptions used to construct non-committing encryption is a primitive known as *simulatable public key encryption (PKE)* introduced by Damgård and Nielsen[8][2]. O'Neill et al. showed how to use simuatable PKE to construct bi-deniable encryption in the multi-distributional model [22]. Moreover, although in their paper they do not explicitly use simulatable PKE, it is not hard to see that the same construction given by [3] can be instantiated with simulatable PKE instead of trapdoor permutations.

Thus, a natural and imperative direction to explore is whether it is possible to construct sender-deniable encryption with super-polynomial security from simulatable PKE.

Our Results

We consider the possibility of constructing non-interactive sender-deniable encryption, known as *sender-deniable public key encryption*, with super-polynomial security in a *black-box manner* from simulatable PKE. We provide a negative answer to the above question by showing the following:

Theorem 1 (Main Theorem, Informal). *There is no (fully) black-box reduction of sender-deniable public key encryption with super-polynomial security to simulatable PKE.*

In particular, we show that every black-box construction of a sender-deniable public key encryption scheme from simulatable PKE with query complexity $m = m(n)$ cannot achieve security better than $O(m^4(n))$. Our results indicate that the scheme of [3] is in some sense optimal, since improving on it will require the use of non-black-box techniques, stronger underlying assumptions or interaction.

Black-Box Separations

Impagliazzo and Rudich [17] were the first to develop a technique to rule out the existence of an important class of reductions between primitives known as black-box reductions. Indeed, most known reductions between cryptographic primitives are black-box (see the works of [25, 14, 24, 15, 13, 18, 16, 21, 20] for a small sampling). Intuitively, black-box reductions are reductions where the primitive is treated as an oracle or a "black-box". There are actually several flavors of black-box reductions (fully black-box, semi black-box and weakly black-box [23]). In our work, we only deal with fully black-box reductions, and so we will focus on this notion here. Informally, a fully black-box reduction from a primitive $\mathcal{Q}$ to a primitive $\mathcal{P}$ is a pair of *oracle* $\mathcal{PPT}$ Turing machines (G, S) such that the following two properties hold:

Correctness: For every implementation f of primitive $\mathcal{P}$, $g = G^f$ implements $\mathcal{Q}$.

[2] In fact, an even weaker primitive called *trapdoor-simulatable PKE* [6] is sufficient for non-committing encryption.

Security: For every implementation f of primitive $\mathcal{P}$, and every adversary A, if A breaks G^f (as an implementation of $\mathcal{Q}$) then $S^{A,f}$ breaks f. (Thus, if f is "secure", then so is G^f.)

We remark that an *implementation* of a primitive is any specific scheme that meets the requirements of that primitive (e.g., an implementation of a public-key encryption scheme provides samplability of key pairs, encryption with the public-key, and decryption with the private key). Correctness thus states that when G is given oracle access to any valid implementation of $\mathcal{P}$, the result is a valid implementation of $\mathcal{Q}$. Furthermore, security states that any adversary breaking G^f yields an adversary breaking f. The reduction here is *fully* black-box in the sense that the adversary S breaking f uses A in a black-box manner.

Techniques

Following the paradigm introduced by [17], we define an oracle O and consider constructions of simulatable PKE and sender-deniable public key encryption relative to this oracle. The oracle O that we use is similar to the by now standard oracle first introduced by [11]. This oracle implements a trapdoor function with the important property that it is difficult to obliviously sample from the range of the function. Namely, it is hard to find an image in the range of the function without first sampling the corresponding preimage.

Relative to the oracle O, we show the following:

- There exists a simulatable PKE scheme, $\mathcal{E}_{\mathsf{Sim}}$ secure against all (computationally unbounded) adversaries $\mathcal{A}$ making at most polynomial number of queries.
- For every implementation $\mathcal{E}$ of a sender-deniable public key encryption scheme relative to O, there exists an adversary $\mathcal{A}$ making at most polynomial number of queries such that $\mathcal{A}$ breaks $\mathcal{E}$.

The above implies that there is no fully black-box construction of sender-deniable public key encryption from simulatable PKE. To see why this is so, assume that there exists a fully black-box construction of sender-deniable public key encryption from simulatable PKE. Then there must be a construction of sender-deniable public key encryption relative to O (since $\mathcal{E}_{\mathsf{Sim}}$ exists relative to O). Moreover, there must exist a poly-time reduction S which gets black-box access to $\mathcal{A}$ and $\mathcal{E}_{\mathsf{Sim}}$ and breaks the security of $\mathcal{E}_{\mathsf{Sim}}$. Thus, there also exists a poly-time oracle machine $\tilde{S}$, such that $\tilde{S}^{\mathcal{A},\mathsf{O}}$ breaks the security of $\mathcal{E}_{\mathsf{Sim}}$. But, since $\mathcal{S}$ is polynomial-time and since $\mathcal{A}$ makes at most polynomial number of oracle queries, we have that $\tilde{S}^{\mathcal{A},\mathsf{O}}$ also makes at most polynomial number of oracle queries and thus by the security of $\mathcal{E}_{\mathsf{Sim}}$, $\tilde{S}^{\mathcal{A},\mathsf{O}}$ cannot break $\mathcal{E}_{\mathsf{Sim}}$.

Now, recall that a sender-deniable public key encryption scheme is a public key encryption scheme with an additional algorithm, Fake, which takes an honestly generated Sender's view $\mathsf{View}_{\mathsf{S}_0}$ encrypting a bit b and returns a fake view, $\mathsf{View}_{\mathsf{S}_1} = \mathsf{Fake}(\mathsf{View}_{\mathsf{S}_0})$, encrypting the bit $1 - b$. A simple but key observation is the following: If the distributions over the corresponding views, $\mathsf{View}_{\mathsf{S}_0}$ and $\mathsf{View}_{\mathsf{S}_1}$ are indistinguishable, then one should be able to now compute $\mathsf{View}_{\mathsf{S}_2} = \mathsf{Fake}(\mathsf{View}_{\mathsf{S}_1})$ obtaining a fake view encrypting the bit b and such that the distributions over the views, $\mathsf{View}_{\mathsf{S}_1}$ and $\mathsf{View}_{\mathsf{S}_2}$ are again indistinguishable. More generally, in any sender-deniable public key encryption scheme with negligible distinguishing advantage, one must be able to run Fake iteratively on the output of the previous Fake invocation for any (unbounded) polynomial number of times. Otherwise, if there is a fixed polynomial upper bound $p(n)$ on the number of times that Fake can be applied to a fresh ciphertext (before outputting $\perp$), then we can distinguish $\mathsf{View}_{\mathsf{S}_0}$ from $\mathsf{View}_{\mathsf{S}_{p(n)}} = \perp = \mathsf{Fake}^{p(n)}(\mathsf{View}_{\mathsf{S}_0})$ (where by $\mathsf{Fake}^{p(n)}$ we denote the composition of Fake, $p(n)$ times). So by a hybrid argument there must be some i such that $\mathsf{Fake}^{i}(\mathsf{View}_{\mathsf{S}_0})$, $\mathsf{Fake}^{i+1}(\mathsf{View}_{\mathsf{S}_0})$ can be distinguished with probability $1/p(n)$. Finally, this means that real and fake openings $\mathsf{View}_{\mathsf{S}_0}$ and $\mathsf{View}_{\mathsf{S}_1}$

can be distinguished, contradicting the security of the sender-deniable public key encryption scheme[3]. Thus, in order to prove the lower bound it is sufficient to show that relative to our oracle, Fake can be repeatedly applied only a fixed polynomial number of times before failure. We note that somewhat similar arguments were used in [2]

To gain some intuition for why this is the case, it is instructive to recall the [3] construction[4]. Let $\{F_{pk}\}$ be a family of trapdoor functions with pseudorandom range such that given the secret key sk of F_{pk}, one can distinguish between elements y in the range of F_{pk} and random elements, but given only pk, random elements in the range of F_{pk} are indistinguishable from random strings. In [3], the secret key of the sender-deniable public key encryption scheme is the secret key sk of the trapdoor function F. The public key pk is the public key of F. Each ciphertext consists of m number of strings $s_1, \ldots, s_m$. To encrypt a 1, choose an a set of indeces $I \subseteq [m]$ of odd cardinality; otherwise choose a set $I \subseteq [m]$ of even cardinality. Compute m strings in the following way: For the i-th string, if $i \in I$, choose a random x_i and compute $y_i = f(x_i)$. If $i \notin I$, choose y_i to be a random string. The Sender sends these m strings to the Receiver. The Receiver then checks which of the m strings $y_1, \ldots, y_m$ are valid images. If an odd number of strings are valid, output 1. Otherwise, output 0. It is not hard to see that the Fake algorithm works by having the Sender claim that a pseudorandom string is really random (but note that the Sender cannot claim the reverse).

Clearly, the Fake algorithm described above can be run iteratively at most m times for a given ciphertext, since the Sender claims to have made one less query each time Fake is run and there are at most m queries total. Unfortunately, our analysis is more complicated since we must also consider candidate schemes where the Fake algorithm might *add* queries to the outputted view. It may seem at first glance that it is impossible for Fake to add new queries to the Sender's view that were not in the original view since it would seem to require inverting a random image y without access to the corresponding secret key. However, this is not necessarily the case (see Appendix B for a toy example where this occurs).

Thus, we must show that even for candidate schemes whose Fake algorithms may both remove and add queries, Fake can be repeatedly applied only a fixed polynomial number of times before failure. Intuitively, the reason we can handle such schemes is that it is infeasible to add an unbounded number of new queries to the fake view, since many queries must be removed from the previous view for each new query that is added. In order to show that this intuition indeed holds, we leverage the fact that in our oracle, with overwhelming probability, random strings are not valid images of the trapdoor function. Much of the technical part of the proof is in showing that the above intuition holds for *all* possible constructions of sender-deniable public key encryption schemes relative to our oracle.

Technical Overview of Proof. The high-level approach of the proof will be to consider the distribution $\mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$, where $m(n)$ is the maximum number of queries made by Sender and Receiver, and a draw from $\mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ is obtained in the following way:

- Draw an oracle O and original views, $\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{R}}$, for Sender and Receiver from the correct distributions.
- For $1 \leq i \leq 10m^2(n)$, set $\mathsf{View}_{\mathsf{S}_i} = \mathsf{Fake}^{\mathsf{O}}(\mathsf{View}_{\mathsf{S}_{i-1}})$.
- Output $\mathsf{View}_{\mathsf{S}_0}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}$

In our analysis, we will look at the properties of sequences of fake openings $\mathsf{View}_{\mathsf{S}_0}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}$ drawn from this distribution. Note that for any deniable public key encryption scheme it should (at the very least) be

[3] Simply run Fake iteratively i number of times on $\mathsf{View}_{\mathsf{S}_0}$ and then use the distinguisher above.

[4] We simplify their construction here somewhat.

the case that w.v.h.p. for every consecutive $i, i+1$, $\mathsf{View}_{\mathsf{S}_i}$ and $\mathsf{View}_{\mathsf{S}_{i+1}}$ are valid encryptions of bits b_i and $b_{i+1} = 1 - b_i$, respectively. Furthermore, we show that if a public key encryption scheme has the deniability property then with high probability a sequence drawn from $\mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ will have several additional properties. However, we will also argue that it is impossible for a sequence of fake openings of length $10m^2(n)$ to satisfy all of the required properties simultaneously. Thus, a sequence drawn from $\mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ will with high probability not satisfy at least one of the required properties. This leads to contradiction and so we conclude that the encryption scheme is not deniable.

In what follows, we give a slightly innacurate but intuitive overview of what these properties are and the techniques we use to prove that with high probability a sequence of fake openings will possess these properties.

First, note that a fake opening is simply a view $\mathsf{View}_{\mathsf{S}_i}$ of the Sender which consists of a transcript, W (i.e. a public key, PK, and ciphertext c), and a set of queries $Q(\mathsf{S}_i)$ made by the Sender. We also consider the set $Q(E)_i$ which, intuitively, is a set of queries that includes all queries the honest Sender (with view $\mathsf{View}_{\mathsf{S}_i}$) believes may have been made by both him and the Receiver. The set of queries in $Q(E)_i$ can be found by running an algorithm that is very similar to the Eve algorithm of [1], which finds intersection queries based only on the transcript (and does not depend on the Sender's view, as in our case). During the execution of the Eve algorithm, Eve will find pairs (pk^*, y^*) such that it is likely the Sender queried $F(pk^*, x) = y^*$ for some x. If Eve identifies a such a pair (pk^*, y^*) and, indeed, a corresponding $F(pk^*, x^*) = y^*$ is found in $\mathsf{View}_{\mathsf{S}_i}$, then the query is placed in Q_i^{made} and we think of the query as having been "added". If Eve identifies a such a pair (pk^*, y^*), however, and no corresponding $F(pk^*, x^*) = y^*$ is found in $\mathsf{View}_{\mathsf{S}_i}$, then the query is placed in Q_i^{skipped} and we think of the query as having been "removed".

Now for each fake opening $\mathsf{View}_{\mathsf{S}_i}$ we consider two corresponding types of queries "A" type queries and "B" type queries. Intuitively, "A" type queries are those queries that were originally in $\mathsf{View}_{\mathsf{S}_0}$ and have either not been removed in some Q_j^{skipped} set (for $j \leq i$), or were removed and then added again in some Q_k^{made} set (for $j < k \leq i$). Intuitively, "B" type queries are new queries that do not appear in the original view $\mathsf{View}_{\mathsf{S}_0}$, were added in some Q_j^{made} set (for $j \leq i$) and have not subsequently been removed in a Q_k^{skipped} set (for $j < k \leq i$). Thus, each view $\mathsf{View}_{\mathsf{S}_i}$ is associated with a set, A^i, of "A" type queries and a set, B^i, of "B" type queries.

We will show that with high probability a draw of fake openings $\mathsf{View}_{\mathsf{S}_0}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}$ and corresponding sequence $(A^0, B^0), \ldots, (A^{10m^2(n)}, B^{10m^2(n)})$ will satisfy the following:

- $(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}})$ are valid openings.
- $A^i \subseteq A^{i-1}$ for $1 \leq i \leq 10m^2(n)$
- $(A^{i-1}, B^{i-1}) \neq (A^i, B^i)$ for $1 \leq i \leq 10m^2(n)$

To show this, we apply the following insight: If a certain event is unlikely to occur with respect to the *first* fake view $\mathsf{View}_{\mathsf{S}_1}$, then it should also be unlikely to occur in all subsequent fake views $\mathsf{View}_{\mathsf{S}_i}$ for $2 \leq i \leq 10m^2(n)$ (since otherwise we can build a distinguisher that distinguishes $\mathsf{View}_{\mathsf{S}_0}$ and $\mathsf{View}_{\mathsf{S}_1}$). This insight is formalized with the definition of *iterative indistinguishability* (see Definition 4). Now, to show that the properties should hold w.h.p. for the *first* fake view, we use techniques similar to those introduced by [17, 11, 1, 7, 19]. These properties (and two other properties that we don't discuss here) are stated formally in Definition 5 and Lemma 4 and the formal proofs are found in Section 4.4 and Section 5.1. Some additional proof intuition is given in Section 4.4 and Section 5.

The next property of fake openings says that with high probability a draw of fake openings $\mathsf{View}_{\mathsf{S}_0}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}$ with a corresponding sequence $(A^0, B^0), \ldots, (A^{10m^2(n)}, B^{10m^2(n)})$ satisfies the following: If the same set

A^* appears consecutively β times in subsequence $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$ and all consecutive B^{j+k}'s are different, then $\beta \leq 10m(n)$. To show this, we require new techniques that leverage the fact that a fixed string y^*, where $F(pk, x) = y^*$ has not yet been queried, is very unlikely to be in the image of F, where the probability is taken over the choice of oracle O, conditioned on the queries already made to the oracle. Our argument is information-theoretic and hinges on showing that a subsequence of pairs $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$ as above has a short description relative to an oracle. This property is stated formally in Lemma 5 and the formal proof is found in Section 6. Some additional proof intuition is given in Section 5.

Related Work

In their seminal paper, Canetti et al. [3] introduce the primitive of deniable encryption and present constructions. However, for the strongest form of deniable encryption which assumes that the same key generation and encryption algorithms are always used, [3] achieve only sender-deniable and receiver-deniable schemes with polynomial security. [3] also rule out the existence of a specific type of sender-deniable encryption scheme with negligible distinguishing advantage (or super-polynomial security) called *separable* schemes. Our impossibility result is incomparable to theirs since ours rules out a larger class of reductions (black-box reductions), but only rules out reductions to the specific primitive of simulatable PKE.

O'Neill et al. [22] recently constructed a bi-deniable encryption scheme in the multi-distributional model, in which the parties run alternative key-generation and encryption algorithms for equivocable communication, but claim under coercion to have run the prescribed algorithms. This weaker model was also initially considered by [3]. Although useful in some settings, the multi-distributional model does not achieve the strongest form of deniability which we consider in this work. We note that it is essential for our impossibility result that the *same* encryption algorithm is run for both real and equivocable communication, which is why our result does not contradict the work of [22].

Recently, Dürmuth and Freeman announced a fully-deniable (receiver/sender)-deniable interactive cryptosystem with negligible security [9]. However their result was later showed to be incorrect by Peikert and Waters (see [10] for details). The protocol constructed by [9] was both interactive and utilized the fact that for the trapdoor function used, a random element in the range could be sampled obliviously. We note that in our analysis it is essential both that the schemes we consider are non-interactive and that the trapdoor function implemented by our oracle does not allow oblivious sampling of the range. Thus, an interesting open question is whether removing these two restrictions can help achieve fully-deniable encryption schemes.

Finally, [2] recently showed, using an information-theoretic argument, that (non-interactive) receiver-deniable encryption with negligible distinguishing advantage do not exist, unconditionally. We note, however, that the work of [2] does not address the case of sender-deniable encryption and it does not seem that their techniques may be applied to our case.

Organization

We define the cryptographic primitives relating to our result in Section 2. In Section 3 we define our oracle and in Section 4 we define some additional useful notations, distributions, algorithms and corresponding properties which will be used in the main result. Finally, in Section 5 we prove our main theorem.

2 Definitions

Definition 1 (Sender-Deniable Public Key Encryption). *A* sender-deniable (bit) public key encryption scheme *is a tuple of algorithms* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *defined as follows:*

- *The key-generation, encryption and decryption algorithms* $\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}$ *are defined as usual for public-key encryption.*
- *The* sender faking algorithm $\mathsf{Fake}(\mathrm{PK}, r_S, b)$*, given a public key* PK*, original coins* r_S *and bit* b *of* Enc*, outputs faked random coins* r_S^* *for* Enc *and the bit* $1 - b$*.*

We require the following properties:

Correctness. $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *forms a correct public-key encryption scheme.*

Deniability. For $b \in \{0, 1\}$*, we require that the following two probability ensembles are computationally indistinguishable:*

- $\{(\mathrm{PK}, c, r_S) | \mathrm{PK} \leftarrow \mathsf{Gen}(1^n; r_G), c \leftarrow \mathsf{Enc}(\mathrm{PK}, b; r_S)\}_n$
- $\{(\mathrm{PK}, c, r_S^*) | \mathrm{PK} \leftarrow \mathsf{Gen}(1^n; r_G), c \leftarrow \mathsf{Enc}(\mathrm{PK}, 1 - b; r_S), r_s^* \leftarrow \mathsf{Fake}(\mathrm{PK}, r_S, b)\}_n$

It follows from the definition that a sender-deniable public key encryption scheme is also semantically secure.

Remark 1. In this work, we also consider constructions of deniable public key encryption schemes that do not achieve negligible distinguishing advantage. We say that a deniable encryption scheme has security $p(n)$ for some polynomial $p(\cdot)$ if correctness holds and every probabilistic polynomial time adversary $\mathcal{A}$ distinguishes the following two probability ensembles with advantage at most $1/p(n)$:

- $\{(\mathrm{PK}, c, r_S) | \mathrm{PK} \leftarrow \mathsf{Gen}(1^n; r_G), c \leftarrow \mathsf{Enc}(\mathrm{PK}, b; r_S)\}_n$
- $\{(\mathrm{PK}, c, r_S^*) | \mathrm{PK} \leftarrow \mathsf{Gen}(1^n; r_G), c \leftarrow \mathsf{Enc}(\mathrm{PK}, b; r_S), r_s^* \leftarrow \mathsf{Fake}(\mathrm{PK}, r_S, b)\}_n$.

We note that in this case semantic security does not follow from deniability and is an additional requirement.

Definition 2 (Simulatable PKE). *A* ℓ-bit simulatable encryption scheme *consists of an encryption scheme* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *augmented with* $(\mathsf{oGen}, \mathsf{oRndEnc}, \mathsf{rGen}, \mathsf{rRndEnc})$*. Here,* oGen *and* $\mathsf{oRndEnc}$ *are the oblivious sampling algorithms for public keys and ciphertexts, and* rGen *and* $\mathsf{rRndEnc}$ *are the respective inverting algorithms,* rGen *(resp.* $\mathsf{rRndEnc}$*) takes* r_{G} *(resp.* $(\mathrm{PK}, r_{\mathrm{E}}, m)$*) as the trapdoor information. We require that, for all messages* $m \in \{0, 1\}^\ell$*, the following distributions are computationally indistinguishable:*

$$\{\mathsf{rGen}(r_{\mathrm{G}}), \mathsf{rRndEnc}(\mathrm{PK}, r_{\mathrm{E}}, m), \mathrm{PK}, c \mid (\mathrm{PK}, \mathrm{SK}) = \mathsf{Gen}(1^k; r_{\mathrm{G}}), c = \mathsf{Enc}_{\mathrm{PK}}(m; r_{\mathrm{E}})\}$$
$$\textit{and } \{\hat{r}_{\mathrm{G}}, \hat{r}_{\mathrm{E}}, \hat{\mathrm{PK}}, \hat{c} \mid (\hat{\mathrm{PK}}, \perp) = \mathsf{oGen}(1^k; \hat{r}_{\mathrm{G}}), \hat{c} = \mathsf{oRndEnc}_{\hat{\mathrm{PK}}}(1^k; \hat{r}_{\mathrm{E}})\}$$

It follows from the definition that a simulatable encryption scheme is also semantically secure.

Definition 3 (Sender-Deniable Public Key Encryption from Simulatable PKE). *For oracle algorithms* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *we call* $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *a black-box construction of sender-deniable public key encryption based on simulatable PKE if the following properties hold:*

- **Implementation:** *The algorithms* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *get oracle access to simulatable PKE scheme* $\mathcal{E}_{\mathsf{Sim}}$ *and* $\mathcal{E}$ *is an implementation of sender-deniable public key encryption.*
- **Security:** *There is a polynomial-time oracle algorithm* S *with the following property. For any simulatable PKE* $\mathcal{E}_{\mathsf{Sim}} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{oGen}, \mathsf{oRndEnc}, \mathsf{rGen}, \mathsf{rRndEnc})$*, given as oracle, if* $\mathcal{A}$ *breaks the security of* $\mathcal{E}$ *then* $S^{\mathcal{E}_{\mathsf{Sim}}, \mathcal{A}}$ *breaks the security of* $\mathcal{E}_{\mathsf{Sim}}$*.*

Remark 2. If there exists a construction of simulatable PKE scheme $\mathcal{E}_{\mathsf{Sim}}$ relative to Oracle O, then for simplicity we consider an implementation of a sender-deniable public key encryption scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ relative to O, where algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ get oracle access to O (instead of just $\mathcal{E}_{\mathsf{Sim}}$). Notice that every sender-deniable public key encryption scheme from simulatable PKE is also a sender-deniable public key encryption scheme relative to O.

3 Oracle

The oracle O consists of three functions G, F, F^{-1} defined below for every security parameter n.

- $G : \{0,1\}^n \to \{0,1\}^{3n}$ is an injective function taking inputs sk of length n bits to outputs pk of length $3n$ bits.
- $F : \{0,1\}^{4n} \to \{0,1\}^{12n}$ is an injective function taking inputs pk, x of length $4n$ bits to outputs y of length $12n$ bits.
- $F^{-1} : \{0,1\}^{13n} \to \{0,1\}^n$ takes inputs of the form sk, y where $sk \in \{0,1\}^n$ and $y \in \{0,1\}^{12n}$. F^{-1} returns $x \in \{0,1\}^n$ if $G(sk) = pk$ and $F(pk, x) = y$ and $\perp$ otherwise.

Note that the oracle above behaves like a trapdoor function, where G is the key generation functionality, F evaluates the trapdoor function and F^{-1} is the inversion function. Additionally, note that we may easily construct a simulatable PKE scheme relative to this oracle.

We denote by Υ the uniform distribution over all possible oracles O.

Lemma 1. *There is a construction of a simulatable PKE scheme $\mathcal{E}_{\mathsf{Sim}}$ relative to oracle O, such that for every unbounded adversary $\mathcal{A}$, making a polynomial number of queries to O:*

$$\Pr_{\mathsf{O} \sim \Upsilon}[\mathcal{A}^{\mathsf{O}} \textit{ breaks } \mathcal{E}^{\mathsf{O}}_{\mathsf{Sim}}] \leq \mathrm{neg}(n).$$

The proof of the Lemma above is by now standard (c.f. [11, 12]) and so we omit it.

4 Preliminaries

In this section we introduce some useful notation, algorithms and properties of sender-deniable public key encryption schemes.

4.1 Sender/Receiver game

Given a deniable public key encryption scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$, we will consider the following natural two-message protocol $\langle \mathsf{S}, \mathsf{R} \rangle$ between a *Sender*, S, and *Receiver*, R (See Figure 1).

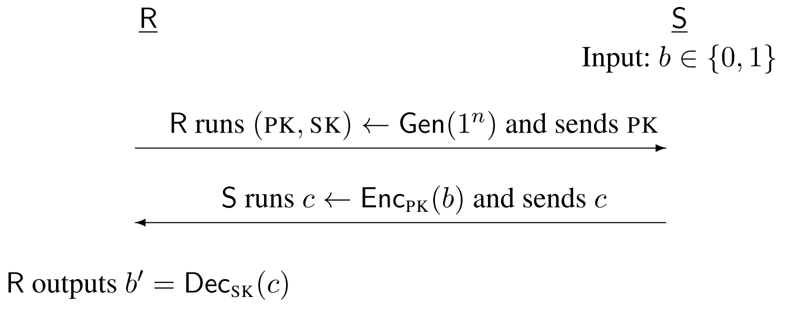
R
S
Input: $b \in \{0,1\}$
R runs $(\mathrm{PK}, \mathrm{SK}) \leftarrow \mathsf{Gen}(1^n)$ and sends PK
S runs $c \leftarrow \mathsf{Enc}_{\mathrm{PK}}(b)$ and sends c
R outputs $b' = \mathsf{Dec}_{\mathrm{SK}}(c)$

Fig. 1. Protocol $\langle \mathsf{S}, \mathsf{R} \rangle$

The view of the Receiver (resp. Sender) consists of the transcript W, random tape, r_{R} (resp. r_{S}) and queries made to the oracle along with the responses. The view of the Receiver, denoted by $\mathsf{View}_{\mathsf{R}} =$

$(\mathsf{View}_G, \mathsf{View}_D)$, consists of two parts where View_G includes queries and responses made during Gen and View_D includes queries and responses made during Dec. The view of the Sender, denoted by View_S includes queries and responses made during Enc. We denote the queries to O in View_R by $Q(\mathsf{R}) = Q(G) \cup Q(D)$. We denote the queries to O in View_S by $Q(\mathsf{S})$.

We assume without loss of generality that:

- No queries to F^{-1} are made during Gen.
- Each party queries $G(sk) = pk$ before querying $F^{-1}(sk, y)$.
- Either Fake returns a valid opening or returns $\perp$ and $\mathsf{Fake}(\perp) = \perp$.

Additionally, relative to our oracle O, we assume WLOG that Fake takes View_S and returns another View_S. More specifically, $\mathsf{View}_{\mathsf{S}_{i+1}} = \mathsf{Fake}^{\mathsf{O}}(\mathsf{View}_{\mathsf{S}_i})$. By $\mathsf{Fake}^{\mathsf{O},i}$ we denote composing Fake with itself i times.

4.2 Useful Distributions

Distribution $\mathcal{D}$: $\mathcal{D}$ is a distribution over tuples $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ resulting from an execution of $\langle \mathsf{S}, \mathsf{R} \rangle$. A draw from $\mathcal{D}$ is obtained as follows:

- Draw $\mathsf{O} \sim \varUpsilon$, $b \leftarrow \{0,1\}$, $\mathsf{r}_\mathsf{R}, \mathsf{r}_\mathsf{S} \leftarrow \{0,1\}^{p(n)}$, for some polynomial $p(\cdot)$ and execute $\langle \mathsf{S}, \mathsf{R} \rangle$ with $\mathsf{O}, \mathsf{r}_\mathsf{R}, \mathsf{r}_\mathsf{S}$ and input bit b.
- Output the views $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ resulting from the execution of $\langle \mathsf{S}, \mathsf{R} \rangle$ above.

Distribution $\mathcal{D}^i$: $\mathcal{D}^i$ is a distribution over tuples $(\mathsf{View}_{\mathsf{S}_i}, \mathsf{View}_\mathsf{R})$ as before, but here we begin to use the Fake algorithm. A draw from $\mathcal{D}^i$ is obtained as follows:

- Draw $\mathsf{O} \sim \varUpsilon$, $b \leftarrow \{0,1\}$, $\mathsf{r}_\mathsf{R}, \mathsf{r}_\mathsf{S} \leftarrow \{0,1\}^{p(n)}$. and execute $\langle \mathsf{S}, \mathsf{R} \rangle$ with $\mathsf{O}, \mathsf{r}_\mathsf{R}, \mathsf{r}_\mathsf{S}$ and input bit b.
- Let $\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_\mathsf{R}$ containing $\textsc{pk}, c, b, \mathsf{r}_\mathsf{S}$ be the resulting views from the execution of $\langle \mathsf{S}, \mathsf{R} \rangle$. Compute $\mathsf{View}_{\mathsf{S}_i} = \mathsf{Fake}^{\mathsf{O},i}(\mathsf{View}_\mathsf{S})$.
- Output O and the views $(\mathsf{View}_{\mathsf{S}_i}, \mathsf{View}_\mathsf{R})$.

For every fixed polynomial $p(\cdot)$, we additionally define the following distribution:

Distribution $\mathcal{D}^{p(n)}_{\mathsf{Fake}}$: $\mathcal{D}^{p(n)}_{\mathsf{Fake}}$ is a distribution over tuples $(\mathsf{O}, \mathsf{View}_\mathsf{S}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{p(n)}})$. A draw from $\mathcal{D}^{p(n)}_{\mathsf{Fake}}$ is obtained in the following way:

- Draw $(\mathsf{O}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_\mathsf{R}) \sim \mathcal{D}$.
- Output $\Big(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1} = \mathsf{Fake}^{\mathsf{O}}(\mathsf{View}_\mathsf{S}), \mathsf{View}_{\mathsf{S}_2} = \mathsf{Fake}^{\mathsf{O}}(\mathsf{View}_{\mathsf{S}_1}), \ldots, \mathsf{View}_{\mathsf{S}_{p(n)}} = \mathsf{Fake}^{\mathsf{O}}(\mathsf{View}_{\mathsf{S}_{p(n)-1}})\Big)$

4.3 Algorithms for Finding Likely Queries

As in [17, 1, 7, 11, 19], we will be concerned with finding *intersection queries*, or common information about the oracle shared by S and R. We note that in our setting there are two ways to get an intersection query:

- One party makes a query of the form $G(sk) = pk$, $F(pk, x) = y$, or $F^{-1}(sk, y)$ and the other party makes the same query.
- One of the parties queries both $G(sk)$, $F^{-1}(sk, y) = x$ and the other party queries $F(pk, x) = y$.

We now (informally) define the Eve algorithm: For a more formal specification, see Appendix A. Eve runs the following algorithm, using threshold $\varepsilon = \varepsilon_1 = 1/m^{16}$ during the first pass (before S sends its message) and using threshold $\varepsilon = \varepsilon_2 = 1/m^6$ during the second pass (after S sends its message).

(0) *Eve queries F on all possible inputs up to length $4\hat{n} = 4\log(10m^{34})$ and adds all queries and responses to E.*

(1) *As long as there exists a query q of the form $G(sk)$, $F(pk, x)$, or $F^{-1}(sk, y)$ that was previously made by* S *or* R *with probability at least ε (conditioned on* Eve*'s current knowledge, E), then ask q from the oracle and add q paired with its answer to E.*

(2) *As long as there exists a pair (pk^*, y^*) such that $G(sk) = pk^* \in Q(E)$, $F(pk^*, x) = y^* \notin Q(E)$ and with probability at least ε,* R *made a query of the form $F(pk^*, x) = y^*$ for some x (conditioned on* Eve*'s current knowledge, E), then query the oracle on $F^{-1}(sk, y^*)$. If $F^{-1}(sk, y^*)$ returns some value x, then add $F(pk^*, x) = y^*$ to E. If $F^{-1}(sk, y^*)$ returns $\perp$ then add $F^{-1}(sk, y^*) = \perp$ to E.*

(3) *As long as there exists a pair (pk^*, y^*) such that $F(pk^*, x) = y^* \notin Q(E)$ and with probability at least ε,* S *made a query of the form $F(pk^*, x) = y^*$ for some x (conditioned on* Eve*'s current knowledge, E), then if $F(pk^*, x) = y^* \in Q(\mathsf{S})$, add q paired with its answer to E and add (pk^*, y^*) to Q^{made}. Otherwise, add the information "query $F(pk^*, x) = y^*$ not made by* S*" to E and add (pk^*, y^*) to Q^{skipped}.*

We denote by $Q(E)_G$ the Eve queries made after the first message is sent from R to S and denote by $Q(E)_{\mathsf{S}}$ the Eve queries made after the second message is sent from S to R.

The following Lemma appeared in [7], but there was proven with respect to a random oracle.

Lemma 2. *Let $\langle \mathsf{S}, \mathsf{R} \rangle$ be a protocol as specified above in which the Sender and Receiver ask at most $2m$ queries each from the oracle* O*. Then there is a universal constant c such that on input parameter ε:*

- (cm/ε)**-Efficiency:** *Eve is deterministic and, over the randomness of the oracle and* S *and* R*'s private randomness, the expected number of Eve queries from the oracle* O *is at most cm/ε_1.*
- $(c\sqrt{m\varepsilon})$**-Security:** *Let W be the transcript of messages sent between* R *and* S *so far, and let E be the additional information that Eve has learned till the* end *of the i'th round. We denote by $Q(E)$ the oracle query/answer pairs that Eve has asked. Let $\mathcal{D}(W, E)$ be the* joint *distribution over the views* $(\mathsf{View}_{\mathsf{S}}, \mathsf{View}_{\mathsf{R}})$ *of* S *and* R only *conditioned on (W, E). By $\mathcal{D}_{\mathsf{R}}(\cdot, \cdot)$ and $\mathcal{D}_{\mathsf{S}}(\cdot, \cdot)$ we refer to the projections of $\mathcal{D}(W, E)$ over its first or second components.*
  With probability at least $1 - c\sqrt{m\varepsilon}$ over the randomness of S, R*, and the random oracle* O *the following holds at* all *moments during the protocol when Eve is done with her learning phase in that round: There are* independent *distributions $\mathcal{S}(W, E), \mathcal{R}(W, E)$ such that:*
  1. *The statistical distance between $\mathcal{S}(W, E) \times \mathcal{R}(W, E)$ and $\mathcal{D}(W, E)$ is at most $\Delta(\mathcal{S}(W, E) \times \mathcal{R}(W, E), \mathcal{D}(W, E)) \leq c\sqrt{m\varepsilon}$.*
  2. *For every oracle query $q \notin Q(E)$ it holds that $\Pr_{(\mathsf{View}_{\mathsf{S}} \sim \mathcal{S}(W,E), \mathsf{View}_{\mathsf{R}} \sim \mathcal{R}(W,E))}[q \in Q(\mathsf{S}) \cup Q(\mathsf{R})] \leq \varepsilon$.*
- **Robustness.** *The learning algorithm is robust to the input parameter ε in the following sense. If the parameter ε changes in the interval $\varepsilon \in [\varepsilon_1, \varepsilon_2]$ arbitrarily during the learner's execution (even inside a learning phase of a specific round), it still preserves $O(cm/\varepsilon_1)$-efficiency and $(c\sqrt{m\varepsilon_2})$-security.*

See Appendix A for the proof of Lemma 2 which is based on the proofs found in [1, 7, 19].

Remark 3. Note that the Eve algorithm as described above requires knowledge of $\mathsf{View}_{\mathsf{S}}$ but not of $\mathsf{View}_{\mathsf{R}}$. Thus, Eve can only be simulated by a party who has knowledge of $\mathsf{View}_{\mathsf{S}}$. This is a key difference between our results and the results of [11]. Note that we can actually implement oblivious transfer relative to our oracle, since although it is hard to sample valid public keys without knowing the corresponding secret key, a party can call $F(pk, \cdot)$ with any string pk and receive a value y indistinguishable from a "valid" image. In contrast, [11] show that oblivious transfer does not exist relative to their oracle. The fact that only S can simulate Eve but not R is the reason that our results do not contradict those of [11].

Remark 4. Note that since the expected number of Eve queries is at most cm/ε, we may consider a modified algorithm Eve' which simulates Eve but aborts if Eve makes more than cm/ε^2 number of queries. By Markov's inequality, this occurs with probability at most $O(\varepsilon)$ and so executions of Eve and Eve' are identical with probability $1 - O(\varepsilon)$. Thus, all properties stated above for Eve hold also for Eve'. In the following, we assume that we run Eve', making at most $N = O(m^{33}) = \mathrm{poly}(n)$ number of queries, to generate the sets $E, Q(E)$. We additionally assume that $N \leq 2^{\hat{n}}/1600m^2$.

4.4 Properties of Fake openings

Definition 4 (Iterative Indistinguishability). *Let* $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *be an implementation of a sender-deniable public key encryption scheme relative to oracle* O*. We say that* $\mathcal{E}$ *is* iteratively indistinguishable up to $p(n)$*, where* $p(\cdot)$ *is some polynomial, if for every* i *where* $1 \leq i \leq p(n)$*, and every adversary* $\mathcal{A}$ *making at most a polynomial number of oracle queries we have:*

$$\Pr_{\mathsf{View}_\mathsf{S} \sim \mathcal{D}_\mathsf{S}}[\mathcal{A}^\mathsf{O}(\mathsf{View}_\mathsf{S}) \text{ outputs } 1] - \Pr_{\mathsf{View}_{\mathsf{S}_i} \sim \mathcal{D}^i_\mathsf{S}}[\mathcal{A}^\mathsf{O}(\mathsf{View}_{\mathsf{S}_i}) \text{ outputs } 1] \leq i/80p(n).$$

In what follows, we split the queries found in a given view $\mathsf{View}_{\mathsf{S}_i}$ into two types: "A" type queries and "B" type queries. Informally, "A" type queries are queries that were also made in the original $\mathsf{View}_{\mathsf{S}_0} = \mathsf{View}_\mathsf{S}$. "B" type queries are new queries that were added which do not appear in $\mathsf{View}_{\mathsf{S}_0}$. Details follow.

For a given draw $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{p(n)}}) \sim \mathcal{D}^{p(n)}_{\mathsf{Fake}}$, we consider a run of the Eve' algorithm with $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0})$ yielding sets $Q(E), Q^{\mathsf{made}}, Q^{\mathsf{skipped}}$ and a run of the Eve' algorithm with $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_i})$ for each $1 \leq i \leq p(n)$ yielding sets $Q(E)_i, Q^{\mathsf{made}}_i, Q^{\mathsf{skipped}}_i$.

We define the sets A^0, B^0 corresponding to $(\mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0})$ as follows: $A^0 = Q(\mathsf{S}_0)$, $B^0 = \emptyset$. For $i \geq 1$, we define the sets A^i, B^i corresponding to $(\mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_i})$ as follows [5]:

$$A^i = \left(A^{i-1} \setminus Q^{\mathsf{skipped}}_i\right) \cup \left(Q^{\mathsf{made}}_i \cap Q(\mathsf{S}_0)\right),$$

$$B^i = \left(B^{i-1} \setminus Q^{\mathsf{skipped}}_i\right) \cup \left(Q^{\mathsf{made}}_i \setminus Q(\mathsf{S}_0)\right).$$

Note that every draw $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{p(n)}}) \sim \mathcal{D}^{p(n)}_{\mathsf{Fake}}$, is associated with a unique sequence $(A^0, B^0), (A^1, B^1), \ldots, (A^{p(n)}, B^{p(n)})$.

Definition 5 (Well-formed Sequences). *Let* $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *be an implementation of a sender-deniable public key encryption scheme relative to oracle* O*. We say that an opening* $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{p(n)}}) \sim \mathcal{D}^{p(n)}_{\mathsf{Fake}}$ *is* well-formed *if it has the following properties:*

(1) $(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{p(n)}})$ *are valid openings.*

[5] By the notation below, we mean to remove from A^{i-1} all queries of the form $F(pk, x) = y$ such that the pair $(pk, y) \in Q^{\mathsf{skipped}}$. The same holds for the following definitions.

(2) $\left(Q(G) \cap \bigcup_{i=1}^{p(n)} Q(\mathsf{S}_i)\right) \setminus Q(E)_G = \emptyset$.
(3) $A^i \subseteq A^{i-1}$ *for* $1 \leq i \leq p(n)$.
(4) *For every query of the form* $F(pk, x) = y$ *that appears in* $Q(E)_i$ *for some* $1 \leq i \leq p(n)$*, the pair* (pk, y) *does not appear in* Q_j^{skipped} *for all* $1 \leq j \leq i$.

Claim 2 *Let* $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *be an implementation of a sender-deniable public key encryption scheme relative to oracle* O *and let* $m = m(n)$ *be the maximum number of queries made by* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. *If* $\mathcal{E}$ *is iteratively indistinguishable up to* $10m^2(n)$ *then* $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{Views}_{\mathsf{S}_0}, \mathsf{Views}_{\mathsf{S}_1}, \ldots, \mathsf{Views}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ *is well-formed with probability* $9/10$.

Proof Intuition for Claim 2. It is clear that Property (1) must hold with high probability. The fact that Property (2) holds with high probability follows directly from Lemma 2 (see below for details).

We give some intuition here for why Property (4) holds with high probability; the argument for why Property (3) holds with high probability is similar. The formal proof of the claim appears below.

Consider only the set Q_0^{skipped}. We will argue that it is unlikely for there to be some $(pk, y) \in Q_0^{\mathsf{skipped}}$ such that there exists $x : F(pk, x) = y$ relative to O. This means that it is unlikely that any subsequent $\mathsf{Views}_{\mathsf{S}_i}, Q(E)_{\mathsf{S}_i}$ contains a query of the form $F(pk, x) = y$ for some $(pk, y) \in Q_0^{\mathsf{skipped}}$. Once we have shown this for Q_0^{skipped}, we may argue that the same holds for every Q_i^{skipped} due to iterative indistinguishability.

Now, assume there is some pair $(pk, y) \in Q_0^{\mathsf{skipped}}$ such that R did not previously query $F(pk, x) = y$ during Gen (we deal with the remaining cases in the formal proof). Then, since $F(pk, x) = y$ also does not appear in $\mathsf{Views}_{\mathsf{S}_0}$ (otherwise (pk, y) would not be in Q_0^{skipped}) we may run the protocol $\langle \mathsf{S}, \mathsf{R} \rangle$ without querying $F(pk, x) = y$. Thus, when the pair (pk, y) is encountered during the run of the Eve' algorithm, then the probability (over choice of oracle) that y is a valid image for $F(pk, *)$ is very small. Thus, since the size of Q_0^{skipped} is bounded, the probability that for some pair $(pk, y) \in Q_0^{\mathsf{skipped}}$ (where R did not query $F(pk, x) = y$), y is a valid image, is also small.

Proof of Claim 2 We show that each of the four properties holds with probability at least $39/40$. Thus, by a union bound, all four properties hold with probability at least $9/10$.

The fact that Property (1) holds with probability $39/40$ follows immediately from iterative indistinguishability.

The fact that Property (2) holds with probability at least $39/40$ follows immediately from the security of the Eve' algorithm (See Lemma 2) and the fact that $\left|\bigcup_{i=1}^{10m^2(n)} Q(\mathsf{S}_i)\right| \leq 10m^2(n) \cdot m(n)$.

Next, we turn to Properties (3) and (4). Note that if for some i, $A^i \setminus A^{i-1} \neq \emptyset$. Then there is some query q in Q_i^{made} such that q is in Q_j^{skipped} for some $j < i$. Thus, we may simultaneously prove that Properties (3), (4) hold with probability $39/40$ by showing that: With probability $39/40$, for every query q that appears in $Q(E)_{\mathsf{S}_i}$, q does not appear in Q_j^{skipped} for $j < i$.

We show that the probability that there is some query q of the form $F(pk^*, x^*) = y^*$ in $\bigcup_{i=1}^{10m^2(n)} Q(E)_{\mathsf{S}_i}$ such that $(pk^*, y^*) \in Q_0^{\mathsf{skipped}}$ is at most $1/800m^2(n)$. By iterative indistinguishability, this implies that for every $1 \leq j \leq 10m^2(n)$, the probability that there is some query q of the form $F(pk^*, x^*) = y^*$ in $\bigcup_{i=1}^{m^2(n)} Q(E)_{\mathsf{S}_i}$ such that $(pk^*, y^*) \in Q_j^{\mathsf{skipped}}$ is at most $1/400m^2(n)$. By a union bound, this implies that with probability $39/40$, for every query q that appears in $Q(E)_{\mathsf{S}_i}$, q does not appear in Q_j^{skipped} for $j < i$.

Let $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_\mathsf{S}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}^{10m^2(n)}_{\mathsf{Fake}}$ and consider the corresponding $(\mathsf{View}_\mathsf{S}, Q(E)_G, Q(E)_\mathsf{S})$. Assume towards contradiction that with non-negligible probability (over draws) there is some query q of the form $F(pk^*, x^*) = y^*$ in $\bigcup_{i=1}^{10m^2(n)} Q(E)_{\mathsf{S}_i}$ such that $(pk^*, y^*) \in Q_0^{\mathsf{skipped}}$. We first consider the case that $q \in Q(G)$. If $q \in Q(G)$ then $q \in \left(Q(G) \cap \bigcup_{i=1}^{10m^2(n)} Q(E)_{\mathsf{S}_i}\right) \setminus Q(E)_G$. But since $\left|\bigcup_{i=1}^{10m^2(n)} Q(E)_{\mathsf{S}_i}\right| \leq m^{13}$, we have by Lemma 2 that with probability at least $1 - 1/m^3$, $\left(Q(G) \cap \bigcup_{i=1}^{10m^2(n)} Q(E)_{\mathsf{S}_i}\right) = \emptyset$. Since we may assume (WLOG) that $m \geq 1600$, we have that will all but probability $1/1600m^2$, $\left(Q(G) \cap \bigcup_{i=1}^{m^2(n)} Q(E)_{\mathsf{S}_i}\right) = \emptyset$.

Alternatively, if $q \notin Q(G)$ then $q \notin Q(G) \cup Q(\mathsf{S})$. Then, we can emulate an interaction of R and S and then run the Eve algorithm. When, during the execution of the Eve algorithm, a pair of the form (pk, y) is added to Q_0^{skipped} then at that point, neither party has queried $F(pk, x) = y$. Thus, when each such query is encountered, with probability $1 - 1/2^{\hat{n}}$ over choice of O, the pair (pk, y) is invalid for O. Since there are at most $O(m^{15}) \leq N$ queries in Q_0^{skipped}, and since $N \leq 2^{\hat{n}}/1600m^2$ we have that with all but probability $1/1600m^2$, all of the pairs (pk, y) in Q_0^{skipped} are invalid. So, for $1 \leq i \leq 10m^2(n)$, $Q(E)_{\mathsf{S}_i}$ cannot contain a query of the form $F(pk, x) = y$ for $(pk, y) \in Q_0^{\mathsf{skipped}}$ with all but probability $1/1600m^2$.

Combining the two cases, we have that the probability that there is some query q of the form $F(pk^*, x^*) = y^*$ in $\bigcup_{i=1}^{10m^2(n)} Q(E)_{\mathsf{S}_i}$ such that $(pk^*, y^*) \in Q_0^{\mathsf{skipped}}$ is at most $1/1600m^2 + 1/1600m^2 = 1/800m^2(n)$ and so the lemma is proved.

5 Analysis

In this section, we prove our main theorem:

Theorem 3 (Main Theorem, Formal). *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ be a black-box construction of sender-deniable public key encryption from simulatable PKE and let $m = m(n)$ be the maximum number of queries made by $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. Then $\mathcal{E}$ has security at most $O(m^4)$.*

Towards proving Theorem 3 we first present the following Lemma:

Lemma 3. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ be an implementation of a sender-deniable public key encryption scheme relative to oracle O and let $m = m(n)$ be the maximum number of queries made by $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. Then $\mathcal{E}$ is* not *iteratively indistinguishable up to $10m^2 = 10m^2(n)$.*

Proving Lemma 3 will be the main technical part of the proof. First, we present the following corollary to Lemma 3 and use Corollary 1 to prove our main theorem:

Corollary 1. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ be an implementation of a sender-deniable public key encryption scheme relative to oracle O and let $m = m(n)$ be the maximum number of queries made by $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. Then there exists an adversary $\mathcal{A}$ making a polynomial number of oracle queries such that*

$$\Pr_{\mathsf{View}_\mathsf{S} \sim \mathcal{D}}[\mathcal{A}^{\mathsf{O}}(\mathsf{View}_\mathsf{S}) \textit{ outputs } 1] - \Pr_{\mathsf{View}_{\mathsf{S}_1} \sim \mathcal{D}^1}[\mathcal{A}^{\mathsf{O}}(\mathsf{View}_{\mathsf{S}_1}) \textit{ outputs } 1] \geq 1/8000m^4.$$

Proof (Proof of Corollary 1.). Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ be an implementation of a sender-deniable public key encryption scheme relative to oracle O. Then by Lemma 3, $\mathcal{E}$ is not iteratively indistinguishable up to $10m^2$. Now, by a hybrid argument we have that for some $1 \leq i \leq 10m^2$, there exists an adversary $\mathcal{A}'$ making at most a polynomial number of oracle queries such that

$$\Pr_{\mathsf{View}_{\mathsf{S}_i} \sim \mathcal{D}^i}[\mathcal{A}'^{\mathsf{O}}(\mathsf{View}_{\mathsf{S}}) \text{ outputs } 1] - \Pr_{\mathsf{View}_{\mathsf{S}_{i+1}} \sim \mathcal{D}^{i+1}}[\mathcal{A}'^{\mathsf{O}}(\mathsf{View}_{\mathsf{S}_{i+1}}) \text{ outputs } 1] \geq 1/8000m^4.$$

We now use the adversary $\mathcal{A}'$ to construct a distinguishing adversary $\mathcal{A}$ breaking the security of $\mathcal{E}$: $\mathcal{A}$ receives a view View drawn from either $\mathcal{D}$ or $\mathcal{D}^1$. $\mathcal{A}$ runs $\mathsf{Fake}^{\mathsf{O},i}(\mathsf{View})$ to obtain View^i. Then $\mathcal{A}$ hands View^i to $\mathcal{A}'$ and outputs whatever $\mathcal{A}'$ outputs. Now, if View was drawn from $\mathcal{D}$ then View^i is identically distributed as a draw from $\mathcal{D}^i$ and if View was drawn from $\mathcal{D}^1$ then View^i is identically distributed as a draw from $\mathcal{D}^{i+1}$. Thus, $\mathcal{A}$ distinguishes with the same (non-negligible) advantage as $\mathcal{A}'$.

Combined with Lemma 1, Corollary 1 implies our main theorem.

Proof (Proof of Main Theorem using Lemma 1 and Corollary 1.). Assume towards contradiction that there is some fully black-box reduction $(\mathcal{E}, S)$ of sender-deniable public key encryption with distinguishing advantage $o(1/m^4)$ to simulatable PKE, where S is a probabilitic polynomial time reduction. Then, since there exists a construction of simulatable PKE relative to oracle O, we have by Remark 2, $\mathcal{E}$ is also a sender-deniable public key encryption scheme relative to O. Now, Corollary 1 implies that with probability at least $1/16000m^4(n)$ over $\mathsf{O} \sim \Upsilon$, there exists an adversary $\mathcal{A}$ making at most a polynomial number of oracle queries such that $\mathcal{A}$ distinguishes with probability at least $1/16000m^4(n)$. Thus, with probability at least $1/16000m^4(n)$ over $\mathsf{O} \sim \Upsilon$, $\mathcal{A}$ breaks $\mathcal{E}$. However, since S makes at most a polynomial number of calls to $\mathcal{A}$, $S^{\mathcal{A}}$ also makes at most polynomial number of queries and so Lemma 1 implies that with probability $1 - \mathrm{neg}(n)$ over $\mathsf{O} \sim \Upsilon$, $S^{\mathcal{A}}$ does not break $\mathcal{E}_{\mathsf{Sim}}$. Thus, there must exist some fixed O such that $\mathcal{A}$ breaks $\mathcal{E}$ with distinguishing advantage $\Omega(1/m^4)$, but $S^{\mathsf{O},\mathcal{A}}$ does not break $\mathcal{E}_{\mathsf{Sim}}$, which means that the reduction $(\mathcal{E}, S)$ fails and so we arrive at contradiction.

We now turn to proving Lemma 3. We first define two events and then prove that they occur with small probability.

Event E_{rSets}: E_{rSets} is the event that a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}^{10m^2(n)}_{\mathsf{Fake}}$ has the property that $(A^i, B^i) = (A^{i+1}, B^{i+1})$ for some $0 \leq i \leq 10m^2(n) - 1$.

Event E_{rA}: E_{rA} is the event that a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}^{10m^2(n)}_{\mathsf{Fake}}$ has the property that for some A^* there are $\beta > 10m(n)$ number of consecutive pairs of the form $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$ such that $B^{j+i} \neq B^{j+i+1}$ for $0 \leq i \leq \beta - 2$.

Lemma 4. *Let* $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *be an implementation of a sender-deniable public key encryption scheme relative to* O *and let* $m = m(n)$ *be the maximum number of queries made by* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. *Let* $\mathcal{E}$ *be iteratively indistinguishable up to* $10m^2(n)$. *The probability that upon a draw* $(\mathsf{O}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}})$ $\mathcal{D}^{10m^2(n)}_{\mathsf{Fake}}$ *Event* E_{rSets} *occurs is at most* $1/2$.

Next, we give some intuition for the proof of Lemma 4. The detailed proof of Lemma 4 appears in Section 5.1.

Proof Intuition for Lemma 4. We show that if for two consecutive views $\mathsf{View}_{\mathsf{S}_i}, \mathsf{View}_{\mathsf{S}_{i+1}}$, we have that $(A^i, B^i) = (A^{i+1}, B^{i+1})$, then the set of "intersection queries" $Q(E)$ found by the Eve' algorithm when it is run on $\mathsf{View}_{\mathsf{S}_i}$ and $\mathsf{View}_{\mathsf{S}_{i+1}}$ are the same.

Now, intuitively, Lemma 2 tells us that conditioned on the transcript W and intersection queries $Q(E)$, the views of S and R are independent. Since both the transcript (which cannot be changed by the Fake algorithm) and the intersection queries $Q(E)$ are the same for the i-th and $i+1$-th opening, this means that the views of the receiver conditioned on $\mathsf{View}_{\mathsf{S}_i}$ and $\mathsf{View}_{\mathsf{S}_{i+1}}$ should be distributed nearly identically. But note that $\mathsf{View}_{\mathsf{S}_i}$ is supposed to be an encryption of a bit b, while $\mathsf{View}_{\mathsf{S}_i}$ is supposed to be an encryption of the bit $1-b$. Thus, by the correctness of the encryption scheme, the views of the receiver should be statistically far when conditioning on $\mathsf{View}_{\mathsf{S}_i}$ and $\mathsf{View}_{\mathsf{S}_{i+1}}$. This leads to a contradiction.

Lemma 5. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ be an implementation of a sender-deniable public key encryption scheme relative to O and let $m = m(n)$ be the maximum number of queries made by $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. Let $\mathcal{E}$ be iteratively indistinguishable up to $10m^2(n)$. The probability that upon a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}$ $\mathcal{D}^{10m^2(n)}_{\mathsf{Fake}}$ Event E_{rA} occurs is at most 1/5.*

Next, we give some intuition for the proof of Lemma 5. The detailed proof of Lemma 5 appears in Section 6.

Proof Intuition for Lemma 5. We show that given $\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{R}}$, oracle O, the set $A^* \subseteq Q(\mathsf{S}_0)$ plus some additional small amount of information we can reconstruct the entire sequence $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$. The following is an imprecise description of the reconstruction algorithm:

1. Execute the two-message protocol $\langle \mathsf{S}, \mathsf{R} \rangle$ with Receiver's view $\mathsf{View}_{\mathsf{R}}$ and Sender's view $\mathsf{View}_{\mathsf{S}_0}$.
2. Use the transcript W generated above and begin running the Eve' algorithm to reconstruct set $B^{j+\beta-1}$. Note that the only additional information necessary to reconstruct $B^{j+\beta-1}$ is upon encountering a pair (pk, y) whether to return $F^{-1}(sk, y) = x$ and add the query to $B^{j+\beta-1}$ or whether to add this query to $Q^{\mathsf{skipped}}_{j+\beta-1}$.
3. Continue to construct sets $B^{j+\beta-2}$ through B^j in the same way as above.

The additional information needed to reconstruct $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$ can be encoded by a list of α elements. More specifically, when encountering the pair (pk, y) as the ℓ-th query in the run of the Eve' algorithm reconstructing the set B^{j+i}, the algorithm checks whether the index ℓ appears on the list. If it does, the reconstruction algorithm adds $F^{-1}(sk, y) = x$ to B^{j+i}. Otherwise, it adds (pk, y) to $Q^{\mathsf{skipped}}_{j+i}$.

Now, since the Eve' algorithm is efficient and makes N queries (where $N \leq 2^{\hat{n}}/1600m^2$) to reconstruct each B set, we only need $\log N$ bits to encode each of the α elements of the list above. Thus, we need "additional information" of length at most $\alpha \cdot \log N$.

We use properties (2) and (4) of well-formed sequences (see Definition 5) to show that for almost all sequences, when a pair (pk, y) is encountered when running the Eve' algorithm to reconstruct set B^{j+i}, if the corresponding query ($F^{-1}(sk, y)$ or $F(pk, x) = y$) has already been made by the reconstruction algorithm, then (pk, y) is always added to B^{j+i}. Thus, we do not need to include such pairs in the list at all. This implies that since $B^{j+i} \neq B^{j+i+1}$ for all i, we must have $\alpha \geq \beta$. Moreover, the above implies that at the point when a pair (pk, y) is encountered as the ℓ-th Eve' query and the index ℓ appears on the list then it must be the case that the corresponding query $F(pk, x) = y$ has never been made by the reconstruction algorithm.

This means that at the point where we encounter each of these α queries on the list, the probability that an oracle O chosen *conditioned on the view of the reconstruction algorithm thus far* has the string y in its image is at most $1/2^{\hat{n}}$. Thus, the probability that an O chosen conditioned only on $\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R}$ has each of the α-many encountered strings $y_1, \ldots, y_\alpha$ in its image is at most $(1/2^{\hat{n}})^\alpha$.

Finally, taking a union bound over all sets $A^* \subseteq Q(\mathsf{S})$ and all sequences $\mathcal{S}$ we show that the probability that an oracle O chosen conditioned only on $\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_\mathsf{R}$ is consistent with *any* well-formed sequence corresponding to some set $A^* \subseteq Q(\mathsf{S}_0)$ and some and sequence $\mathcal{S}$ of length $\alpha \geq \beta$ is small.

First, we complete the proof of Lemma 3 using the above lemmas. Then, in Section 5.1 and Section 6 we present formal proofs of Lemmas 4 and 5, respectively

Proof (Proof of Lemma 3 using Lemmas 4 and 5). Assume towards contradiction that there is some implementation of a sender-deniable public key encryption scheme, $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, relative to oracle O that is iteratively indistinguishable up to $10m^2 = 10m^2(n)$. By Claim 2, we may assume that, with probability at least $9/10$, a draw $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ is well-formed. In particular, this implies that with probability at least $9/10$ over draws, Property (1) and (3) hold so we have that with probability $9/10$ over draws the openings $(\mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}})$ are all valid and $A^{i+1} \subseteq A^i$ for every $0 \leq i \leq 10m^2(n) - 1$. This implies that with probability $9/10$ over draws there must be some set A^* that appears at least $10m = 10m(n)$ times. Moreover, since Lemma 4 guarantees that event E_rSets occurs with probability at most $1/2$, we have that with probability at least $4/10 = 2/5$ over draws, there is some set A^* that appears at least $10m$ times consecutively and for this A^*, for all $0 \leq i \leq 10m - 2$, $B^{j+i} \neq B^{j+i+1}$. Now, by definition of Event E_rA, this means that with probability at least $4/10$ over draws $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$, we have that Event E_rA occurs. But by Lemma 5 we have that event E_rA occurs with probability at most $1/5$. Thus, we have arrived at contradiction and so the Lemma is proved.

5.1 Proof of Lemma 4

Recall the definition of Event E_rSets and the statement of Lemma 4:

Event E_rSets: E_rSets is the event that a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ has the property that $(A^i, B^i) = (A^{i+1}, B^{i+1})$ for some $0 \leq i \leq 10m^2(n) - 1$.

Lemma 6. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ be an implementation of a sender-deniable public key encryption scheme relative to O and let $m = m(n)$ be the maximum number of queries made by $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. Let $\mathcal{E}$ be iteratively indistinguishable up to $10m^2(n)$. The probability that upon a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}$ $\mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ Event E_rSets occurs is at most $1/2$.*

We begin with the following Claim:

Claim 4 *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ be a sender-deniable public key encryption scheme relative to O and let $m = m(n)$ be the maximum number of queries made by $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$. Let $\mathcal{E}$ be iteratively indistinguishable up to $10m^2(n)$. Let $(\mathsf{O}, \mathsf{View}_\mathsf{R}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$. For all $1 \leq i \leq 10m^2(n)$, if $(A^i, B^i) = (A^{i+1}, B^{i+1})$ then $E_i = E_{i+1}$.*

Proof. Assume towards contradiction that $Q(E)_i \neq Q(E)_{i-1}$. Then there must be some *first* point during the Eve' algorithm that either a query q is asked in i and not $i+1$ (or vice versa) or a query q receives a different response in $Q(E)_i$ and $Q(E)_{i+1}$.

If a query of the form $G(sk) = pk$ or $F(pk, x) = y$ is the first differing query. Then this query is always asked in both i and $i+1$ and receives the same response in both i and $i+1$. Thus, the first differing query cannot be of the form $G(sk)$ or $F(pk, x)$.

On the other hand, if a pair of the form (pk, y) (corresponding to a query of the form $F^{-1}(sk, y)$) is the first differing query and this query is included in $Q(E)_i$ (resp. $Q(E)_{i+1}$). Then this query must be in $A^i \cup B^i$ (resp. $A^{i+1} \cup B^{i+1}$) and hence also $A^{i+1} \cup B^{i+1}$ (resp. $A^i \cup B^i$), since the sets are equal. However, since this query is not in $Q(E)_{i+1}$ (resp. $Q(E)_i$), then it must be that the query is in $Q_{i+1}^{\mathsf{skipped}}$ (resp. Q_i^{skipped}). But then, by definition of the sets, this query cannot be in $A^{i+1} \cup B^{i+1}$ (resp. $A^i \cup B^i$). Thus, this query must also be asked in $Q(E)_{i+1}$ (resp. $Q(E)_i$) and, moreover, must receive the same response in both. Thus, the first differing query cannot be of the form $F^{-1}(sk, y)$ either and so we arrive at contradiction.

Let W be the transcript implicit in some fixed $\mathsf{View}^*_{\mathsf{S}}$. By $\mathcal{D}(\mathsf{View}^*_{\mathsf{S}}, E)$ (resp. $\mathcal{R}(\mathsf{View}^*_{\mathsf{S}}, E)$), we denote the distribution $\mathcal{D}$, conditioned on $\mathsf{View}_{\mathsf{S}} = \mathsf{View}^*_{\mathsf{S}}$ and E.

Next observe the following:

Claim 5 *Let* $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *be an implementation of a sender-deniable public key encryption scheme relative to oracle* O *that is iteratively indistinguishable up to* $10m^2(n)$*. Let* $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m}}$ $\mathcal{D}^{10m^2(n)}_{\mathsf{Fake}}$ *and let* b_i *be the bit encrypted by* $\mathsf{View}_{\mathsf{S}_i}$*. Then for every* $1 \leq i \leq 10m^2(n)$ *we have that with probability at least* $1 - 1/100m^2(n)$*:*

$$\Pr_{\mathsf{View}_{\mathsf{R}} \sim \mathcal{D}(\mathsf{View}_{\mathsf{S}_i}, E_i)}[\mathsf{View}_R \textit{ outputs } 1-b] \leq 1/100.$$

Proof. For the original view $\mathsf{View}_{\mathsf{S}_0}$, we have that with probability at least $1 - 1/200m^2$:

$$\Pr_{\mathsf{View}_{\mathsf{R}} \sim \mathcal{D}(\mathsf{View}_{\mathsf{S}_0}, E_0)}[\mathsf{View}_R \text{ outputs } 1-b] \leq \mathrm{neg}(n)$$

due to correctness. Thus, the Claim holds for $i \geq 1$ due to iterative indistinguishability of $\mathcal{E}$.

We are now ready to prove Lemma 4 using Claims 4 and 5. We will show that the probability over a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}^{10m^2(n)}_{\mathsf{Fake}}$ that $E_0 = E_1$ is at most $1/40m^2(n)$. By iterative indistinguishability, this means that for every fixed $1 \leq i \leq 10m^2(n) - 1$, the probability over a draw that $E_i = E_{i+1}$ is at most $1/20m^2(n)$. By a union bound we have that the probability over draws that for some $0 \leq i \leq m^2(n) - 1$ $E_i = E_{i+1}$, is at most $1/2$. Finally, by Claim 4 we have that $(A^i, B^i) = (A^{i+1}, B^{i+1})$ implies $E_i = E_{i+1}$. Thus, we have that Event E_{rSets} occurs with probability at most $1/2$, which proves the Lemma.

First, by Lemma 2, we have that with probability at least $1 - c\sqrt{m\varepsilon_2} = 1 - O(1/m^{2.5}(n))$, the two distributions $\mathcal{D}(W, E)$ and $\mathcal{S}(W, E) \times \mathcal{R}(W, E)$ have statistical distance at most $c\sqrt{m\varepsilon_2} = O(1/m^{2.5}(n))$.

Using Markov's inequality we have that with probability $1 - O(1/m^{2.5}(n))$ over W, E_0:

$$\Pr_{\mathsf{View}_{\mathsf{S}_0} \sim \mathcal{D}(W, E_0)}[\varDelta(\mathcal{R}(\mathsf{View}_{\mathsf{S}_0}, E_0), \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_0}, E_0)) \geq 1/5] \leq O(1/m^{2.5}(n)).$$

This implies that with probability at least $1 - O(1/m^{2.5}(n))$ over choice of $\mathsf{O}, \mathsf{View}_{\mathsf{S}_0} \sim \mathcal{D}_{\mathsf{S}}$ we have that

$$\varDelta(\mathcal{R}(\mathsf{View}_{\mathsf{S}_0}, E_0), \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_0}, E_0)) \leq 1/5.$$

Now due to iterative indistinguishability, we have that with probability at least $1 - 1/40m^2(n)$ over choice of $\mathsf{O}, \mathsf{View}_{\mathsf{S}_1} \sim \mathcal{D}^1_{\mathsf{S}}$ we have that

$$\Delta(\mathcal{R}(\mathsf{View}_{\mathsf{S}_1}, E_1), \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_1}, E_1)) \leq 1/5.$$

Let Y denote the event that upon a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{m^2(n)}}) \sim \mathcal{D}^{m^2(n)}_{\mathsf{Fake}}$ we have that for $j = 0, j = 1$: $\Delta(\mathcal{R}(\mathsf{View}_{\mathsf{S}_j}, E_j), \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_j}, E_j)) \leq 1/5$. Note that $\Pr_{(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1})}[\mathsf{Y}] \geq 1 - 1/400m^2(n)$. We now claim that $\Pr_{(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1})}[E_0 = E_1 | \mathsf{Y}] \leq 1/40m^2(n)$. Notice that this is sufficient to complete the proof since

$$\begin{aligned}\Pr_{(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1})}[E_0 = E_1] &\leq \Pr_{(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1})}[\overline{\mathsf{Y}}] + \Pr_{(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1})}[E_0 = E_1 | \mathsf{Y}] \\ &\leq 1/20m^2(n).\end{aligned}$$

where by $\overline{\mathsf{Y}}$ we denote the negation of Y.

Assume towards contradiction $\Pr_{(\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1})}[E_0 = E_1 | \mathsf{Y}] \geq 1/40m^2(n)$.
Let $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{m^2(n)}}) \sim \mathcal{D}^{m^2(n)}_{\mathsf{Fake}}$ be such that $E_0 = E_1$ and Y holds. Since $\mathcal{R}(W, E)$ is independent of $\mathsf{View}_{\mathsf{S}_0}$, we have $\mathcal{R}(\mathsf{View}_{\mathsf{S}_0}, E_0) \equiv \mathcal{R}(\mathsf{View}_{\mathsf{S}_1}, E_1) \equiv \mathcal{R}(W, E_1)$. Now, since Y holds we have that

$$\begin{aligned}\Delta(\mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_0}, E_0), \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_1}, E_1)) &\leq \Delta(\mathcal{R}(\mathsf{View}_{\mathsf{S}_0}, E_0), \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}}, E_0)) \\ &\quad + \Delta(\mathcal{R}(\mathsf{View}_{\mathsf{S}_1}, E_1), \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_1}, E_1)) \\ &\leq 2/5\end{aligned}$$

This implies that for $\mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, E_0, E_1$ such that $E_0 = E_1$ and Y holds, we have that for some $b \in \{0, 1\}$,

$$\left| \Pr_{\mathsf{View}_{\mathsf{R}} \sim \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_0}, E_0)}[\mathsf{View}_{\mathsf{R}} \text{ outputs } b] - \Pr_{\mathsf{View}_{\mathsf{R}} \sim \mathcal{D}_{\mathsf{R}}(\mathsf{View}_{\mathsf{S}_1}, E_1)}[\mathsf{View}_{\mathsf{R}} \text{ outputs } b] \right| \leq 2/5.$$

Note, however, that $\mathsf{View}_{\mathsf{S}_0}$ is supposed to be an encryption of a bit b while $\mathsf{View}_{\mathsf{S}_1}$ is an encryption of $1 - b$. Thus, for a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{m^2(n)}}) \sim \mathcal{D}^{m^2(n)}_{\mathsf{Fake}}$ such that $E_0 = E_1$ and Y holds, we have that $\Pr_{\mathsf{View}_{\mathsf{R}} \sim \mathcal{D}(\mathsf{View}_{\mathsf{S}_j}, E)}[\mathsf{View}_R \text{ outputs } 1 - b] \geq 1/100$ for either $j = 0$ or $j = 1$.

But now we have that for either $j = 0$ or $j = 1$:

$$\begin{aligned}\Pr_{\mathsf{View}_{\mathsf{S}_j}, E_j}\left[\Pr_{\mathsf{View}_{\mathsf{R}} \sim \mathcal{D}(\mathsf{View}_{\mathsf{S}_j}, E_j)}[\mathsf{View}_R \text{ outputs } 1 - b] \geq 1/100\right] &\geq 1/2 \cdot \Pr_{\mathsf{View}_{\mathsf{S}_j}, E_j}[E_0 = E_1 | \mathsf{Y}] \cdot \Pr[\mathsf{Y}] \\ &\geq 1/2 \cdot 1/40m^2(n) \cdot \Pr[\mathsf{Y}] \\ &> 1/100m^2(n)\end{aligned}$$

which is a contradiction to Claim 5, and so Lemma 4 is proved.

6 Proof of Lemma 5

Recall the definition of Event E_{rA} and the statement of Lemma 5:

Event E_{rA}*:* E_{rA} is the event that a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{S}}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ has the property that for some A^* there are $\beta > 10m$ number of consecutive pairs of the form $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$ such that $B^{j+i} \neq B^{j+i+1}$ for $0 \leq i \leq \beta - 2$.

Lemma 7. *Let* $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$ *be an implementation sender-deniable public key encryption scheme relative to* O *and let* $m = m(n)$ *be the maximum number of queries made by* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Fake})$*. Let* $\mathcal{E}$ *be iteratively indistinguishable up to* $10m^2(n)$*. The probability that upon a draw* $(\mathsf{O}, \mathsf{View}_{\mathsf{S}}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}})$ $\mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ *Event* E_{rA} *occurs is at most 1/5.*

Specifying a Sequence of Openings. Given $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$, a sequence of β consecutive pairs $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$ with the same set A^* can be specified by a sequence of tuples (k, ℓ) where k takes $\log \beta$ bits to record, and ℓ is of length $\log N$, where $N \leq 2^{\hat{n}}/1600m^2$, bits in the following way:

The tuples corresponding to the i-th set B^{j+i-1} are of the form $(i, *)$ For each i, starting with $i = \beta$ and going down to $i = 1$ we compute a sequence of tuples $(i, *), \ldots, (i, *)$ in the following way:

Emulate an execution of the Eve' algorithm: First, execute Step (0) exactly as in Eve'. For Steps (1), (2), (3), when the next query (the ℓ-th query thus far–not including Step (0)–in the current run of the Eve' algorithm) is of the form $G(sk)$, $F(pk, x)$ or $F^{-1}(sk, y)$, make the query to the oracle O. When the next query (the ℓ-th query) is an inversion query corresponding to the tuple (pk, y) and a query of the form $F(pk, x) = y$ does not yet appear in $Q(E)_i$ do the following:

- If the query $F(pk, x) = y$ is in A^* then add to E_i.
- If the query $F(pk, x) = y$ appeared in $Q(E)_t$ for $t > i$, then add to E_i
- If the query $F(pk, x) = y$ appears in $Q(\mathsf{S}_i)$, does not appear in $Q(G)$, and has not yet appeared in $Q(E)_t$ for $t \geq i$ then add a tuple of the form (i, ℓ) to the sequence and add $F(pk, x) = y$ to E_i.
- Otherwise, do not make the query and add (pk, y) to Q_i^{skipped}.

Note that all tuples (i, ℓ) added to the sequence must correspond to tuples (pk, y) where y has length at least $3\hat{n}$ (since all queries of input length less than or equal to $\hat{n}$ are automatically queried and added to $Q(E)_i$ in Step (0)).

We next describe how to convert a sequence of tuples of the form: $(k_1, \ell_1), \ldots, (k_\alpha, \ell_\alpha)$ back into a sequence of sets of the form: $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$.

Reconstruction Procedure. For *well-formed* openings $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{S}_1}, \ldots, \mathsf{View}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$, given $\mathsf{View}_G, \mathsf{View}_{\mathsf{S}}$, we have that: Given $\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}$, a set A^* and a sequence, $\mathcal{S} = (k_1, \ell_1), \ldots, (k_\alpha, \ell_\alpha)$ of length α, we can reconstruct $(A^*, B^j), \ldots, (A^*, B^{j+\beta-1})$ by reconstructing E_i for each i, starting from $i = \beta - 1$ and going down to $i = 0$. The relevant tuples for reconstructing the set E_i are those of the form $(i, *)$. For each i (starting with $i = \beta$ and going down to $i = 1$) do the following:

Emulate an execution of the Eve' algorithm: First, execute Step (0) exactly as in Eve'. For Steps (1), (2), (3), when the next query (the ℓ-th query thus far–not including Step (0)–in the current run of the Eve' algorithm) is of the form $F(pk, x)$, make the query to the oracle O. When next query (the ℓ-th query) is an inversion query corresponding to the tuple (pk, y) and a query of the form $F(pk, x) = y$ does not yet appear in $Q(E)_i$ do the following:

- If the query $F(pk, x) = y$ is in A^* then add to E_i.

- If the query $F(pk, x) = y$ appeared in $Q(E)_t$ for $t > i$, then add to E_i
- If a tuple of the form (i, ℓ) appears in the sequence, then return x such that $F(pk, x) = y$ (note that finding such x is not necessarily efficient) and add $F(pk, x) = y$ to $Q(E)_i$.
- Otherwise, do not make the query and add (pk, y) to Q_i^{skipped}

Note that we reconstruct the correct sequence $B^1, \dots, B^\beta$ for well-formed sequences since Properties (2) and (4) imply that every query q in $Q(G) \setminus Q(E)_G$ does not appear in $\bigcup_{i=1}^{10m^2(n)} Q(\mathsf{S}_i) \setminus Q(E)_G$ and for every query q of the form $F(pk, x) = y$ in $Q(E)_t$ such that the pair (pk, y) is also in $Q_i^{\mathsf{made}} \cup Q_i^{\mathsf{skipped}}$ for $i < t$, we have that $(pk, y) \in Q_i^{\mathsf{made}}$ and so q is always in $Q(E)_i$.

Remark 5. For simplicity, we described sequences $(k_1, \ell_1), \dots, (k_\alpha, \ell_\alpha)$ above as consisting of tuples where the value k_t in the t-th tuple indicates that we are considering the set B^{j+k_t-1} and is of length $\log \beta$. Note that given a sequence of this form we can transform to a sequence $(k_1, \ell_1), \dots, (k_\alpha, \ell_\alpha)$ where each k_t is a single bit indicating whether the corresponding set B^{j+k_t-1} is the same as the corresponding set of the previous tuple or not. Moreover, note that given a sequence where each k_t is a single bit, we can transform back to a sequence where each k_t is $\log(\beta)$ bits. Thus, we have that every opening $B^1, \dots, B^\beta$ with $\left|\bigcup_{i=1}^{\beta} B^i\right| = \alpha$ can be described by a sequence of length α in at most $\alpha \cdot (\log N + 1) \leq \alpha \cdot (\log 2N)$ bits. Additionally, note that for well-formed sequences, we must have $\alpha \geq \beta$.

Below, we show that the probability over draws $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{Views}_{\mathsf{S}_0}, \mathsf{Views}_{\mathsf{S}_1}, \dots, \mathsf{Views}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ that the opening is well-formed *and* event E_{rA} occurs is at most $1/10$. Since with probability $9/10$, a draw $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{Views}_{\mathsf{S}_0}, \mathsf{Views}_{\mathsf{S}_1}, \dots, \mathsf{Views}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ is well-formed, we have that the probability over draws $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{Views}_{\mathsf{S}_0}, \mathsf{Views}_{\mathsf{S}_1}, \dots, \mathsf{Views}_{\mathsf{S}_{10m^2(n)}}) \sim \mathcal{D}_{\mathsf{Fake}}^{10m^2(n)}$ that Event E_{rA} occurs is at most $1/10 + 1/10 \leq 1/5$ and thus the above is sufficient to prove the Lemma.

Consider an execution of the reconstruction procedure as given above. For every fixed $\Psi = (\mathsf{Views}_{\mathsf{S}_0}, \mathsf{View}_{\mathsf{R}}, A^*)$, and sequence $\mathcal{S}$, we consider partial views of the Reconstruction Procedure $\mathsf{View}_{\mathsf{Rec},j}^{\mathsf{O},\Psi}(\mathcal{S})$. Specifically, $\mathsf{View}_{\mathsf{Rec},j}^{\mathsf{O},\Psi}(\mathcal{S})$ contains the transcript, queries and responses of the Reconstruction Procedure at the point the j-th pair of sequence $\mathcal{S}$ is reached.

We say that a partial view $\mathsf{View}_{\mathsf{Rec},j}^{\mathsf{O},\Psi}(\mathcal{S})$ is *valid* if the following hold:

- The j-th pair (i_j, ℓ_j) corresponds to some pair (pk_j, y_j).
- A query of the form $F(pk_j, x_j) = y_j$ or $F^{-1}(sk_j, y_j)$ has not yet been made during the reconstruction procedure.
- For each $1 \leq v \leq j - 1$, where pair (i_v, ℓ_v) corresponds to pair (pk_v, y_v), the partial view contains a query of the form $F(pk_v, x) = y_v$.
- All previous partial views $\mathsf{View}_{\mathsf{Rec},v}^{\mathsf{O},\Psi}(\mathcal{S})$ for $v \leq j - 1$ are valid.

Finally, define $\mathsf{I}_{j,x_j}^{\mathsf{O},\Psi}(\mathcal{S})$ to be the event that during the execution of the Reconstruction Procedure with $\mathsf{O}, \mathsf{Views}_{\mathsf{S}}, \mathsf{View}_{\mathsf{R}}$, the set A and sequence $\mathcal{S}$, at the point the j-th pair is reached, we have that the corresponding pair (pk_j, y_j), is such that y_j is in the image of $F(pk_j, \cdot)$ with respect to oracle O.

The following fact follows straightforwardly from the definitions above:

Fact 6 *If* $\mathsf{View}_{\mathsf{Rec},\alpha+1}^{\mathsf{O},\Psi}(\mathcal{S})$ *is valid then we have that* $\mathsf{I}_{j,x_j}^{\mathsf{O},\Psi}(\mathcal{S})$ *occurred for* $1 \leq j \leq \alpha$ *and that* $\mathsf{View}_{\mathsf{Rec},j}^{\mathsf{O},\Psi}(\mathcal{S})$ *is valid for* $1 \leq j \leq \alpha$.

In the following, we proceed to show that the probability over draws $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}) \sim \mathcal{D}$ that there exists some corresponding well-formed opening where Event E_{rA} occurs is at most $1/10$.

Now, observe that for well-formed sequences $\mathcal{S}$, we must have that $\mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},\alpha+1}(\mathcal{S})$ is valid. Therefore, it is sufficient to upper bound the probability over draws $(\mathsf{O}, \mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}) \sim \mathcal{D}$ that there exists a set $A^* \subseteq Q(\mathsf{S}_0)$ and a sequence $\mathcal{S}$ of length $\alpha \geq \beta \geq 10m$ such that $\mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},\alpha+1}(\mathcal{S})$ is valid. In the following, we upperbound this quantity.

Claim 7 *For $\alpha \geq \beta \geq 10m \geq 10$, we have that*

$$\Pr_{\mathsf{O}\sim\Upsilon|Q(\mathsf{S}),Q(\mathsf{R})}[\exists A^*, \mathcal{S} : \mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},\alpha+1}(\mathcal{S}) \textit{ is valid}] \leq 1/10$$

Proof. Fix some $\Psi = (\mathsf{View}_{\mathsf{R}}, \mathsf{View}_{\mathsf{S}_0}, A^*)$ and some sequence $\mathcal{S} = \{(k_1, \ell_1), \ldots, (k_\alpha, \ell_\alpha)\}$ of length α.

Now, we must have that for every $1 \leq j \leq \alpha$:

$$\Pr_{\mathsf{O}\sim\Upsilon|Q(\mathsf{S}),Q(\mathsf{R})}[\mathsf{I}^{\mathsf{O},\Psi}_{j}(\mathcal{S}) \mid \mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},j}(\mathcal{S}) \text{ is valid}] \leq 1/2^{\hat{n}}. \tag{6.1}$$

This is the case since for every $1 \leq j \leq \alpha$ the length of y_j is at least $3\hat{n}$ bits and, moreover, if the partial view $\mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},j}(\mathcal{S})$ is valid then at the moment in the Reconstruction Procedure when pair (pk_j, y_j) is encountered, a query of the form $F(pk_j, *)$ or $F^{-1}(sk, y_j)$ (where $G(sk) = pk_j$) has never been made. In this case, conditioned on the view of the Reconstruction Procedure, the probability that y_j is in the image of $F(pk_j, \cdot)$ is at most $1/2^{\hat{n}}$.

Moreover, the above implies that:

$$\begin{aligned}(1/2^{\hat{n}})^\alpha &\geq \Pr_{\mathsf{O}\sim\Upsilon|Q(\mathsf{S}),Q(\mathsf{R})}[\mathsf{I}^{\mathsf{O},\Psi}_{1}(\mathcal{S}) \mid \mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},1}(\mathcal{S}) \text{ is valid}] \cdot \Pr_{\mathsf{O}\sim\Upsilon|Q(\mathsf{S}),Q(\mathsf{R})}[\mathsf{I}^{\mathsf{O},\Psi}_{2}(\mathcal{S}) \mid \mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},2}(\mathcal{S}) \text{ is valid}] \\ &\cdots \Pr_{\mathsf{O}\sim\Upsilon|Q(\mathsf{S}),Q(\mathsf{R})}[\mathsf{I}^{\mathsf{O},\Psi}_{\alpha}(\mathcal{S}) \mid \mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},\alpha}(\mathcal{S}) \text{ is valid}] \\ &\geq \Pr_{\mathsf{O}\sim\Upsilon|Q(\mathsf{S}),Q(\mathsf{R})}[\mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},\alpha+1}(\mathcal{S}) \text{ is valid}]\end{aligned}$$

where the first inequality follows from (6.1) and and the last inequality follows due to Fact 6

Now, by a union bound over the 2^m possible subsets $A^* \subseteq Q(\mathsf{S}_0)$, and the at most $(2N)^\alpha$ possible sequences $\mathcal{S}$ of length α:

$$\Pr_{\mathsf{O}\sim\Upsilon|Q(\mathsf{S}),Q(\mathsf{R})}[\exists A^*, \mathcal{S} : \mathsf{View}^{\mathsf{O},\Psi}_{\mathsf{Rec},\alpha+1}(\mathcal{S}) \text{ is valid}] \leq 2^m \cdot (2N)^\alpha \cdot \left(1/2^{\hat{n}}\right)^\alpha.$$

Additionally, we have that

$$\begin{aligned}2^m \cdot (2N)^\alpha \cdot \left(1/2^{\hat{n}}\right)^\alpha &\leq \left(\frac{4N}{2^{\hat{n}}}\right)^{10m} \\ &\leq \left(\frac{2^{\hat{n}}/400m^2}{2^{\hat{n}}}\right)^{10m} \\ &\leq \left(\frac{1}{400m^2}\right)^{10m} \\ &\leq 1/10\end{aligned}$$

when $10 \leq 10m < \beta \leq \alpha$.

References

1. Boaz Barak and Mohammad Mahmoody. Merkle puzzles are optimal - an O(n^2)-query attack on any key exchange from a random oracle. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 374–390. Springer, 2009.
2. Rikke Bendlin, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. Lower and upper bounds for deniable public-key encryption. In *ASIACRYPT*, pages 125–142, 2011.
3. Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. In *CRYPTO*, pages 90–104, 1997.
4. Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *STOC*, pages 639–648, 1996.
5. Ran Canetti and Rosario Gennaro. Incoercible multiparty computation (extended abstract). In *FOCS*, pages 504–513, 1996.
6. Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. Improved non-committing encryption with applications to adaptively secure protocols. In *ASIACRYPT*, pages 287–302, 2009.
7. Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin. On the black-box complexity of optimally-fair coin tossing. In *TCC*, pages 450–467, 2011.
8. Ivan Damgård and Jesper Buus Nielsen. Improved non-committing encryption schemes based on a general complexity assumption. In *CRYPTO*, pages 432–450, 2000.
9. Markus Dürmuth and David Mandell Freeman. Deniable encryption with negligible detection probability: An interactive construction. In *EUROCRYPT*, pages 610–626, 2011.
10. Markus Dürmuth and David Mandell Freeman. Deniable encryption with negligible detection probability: An interactive construction. *IACR Cryptology ePrint Archive*, 2011:66, 2011.
11. Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The relationship between public key encryption and oblivious transfer. In *FOCS*, pages 325–335, 2000.
12. Yael Gertner, Tal Malkin, and Steven Myers. Towards a separation of semantic and cca security for public key encryption. In *TCC*, pages 434–455, 2007.
13. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
14. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.
15. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
16. Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *FOCS*, pages 230–235, 1989.
17. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *STOC*, pages 44–61, 1989.
18. Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput.*, 17(2):373–386, 1988.
19. Hemanta Maji. *On Computational Intractability Assumptions in Cryptography*. PhD thesis, University of Illinois at Urbana-Champaign, Champaign, Illinois, 2011.
20. Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
21. Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *STOC*, pages 33–43, 1989.
22. Adam O'Neill, Chris Peikert, and Brent Waters. Bi-deniable public-key encryption. In *CRYPTO*, pages 525–542, 2011.
23. Omer Reingold, Luca Trevisan, and Salil Vadhan. Notions of reducibility between cryptographic primitives. In *TCC*, pages 1–20, 2004.
24. John Rompel. One-way functions are necessary and sufficient for secure signatures. In *STOC*, pages 387–394, 1990.
25. Andrew Chi-Chih Yao. Theory and applications of trapdoor functions. In *FOCS*, pages 80–91, 1982.

A The Eve Algorithm and its Properties

We first define the following events and distributions similarly to [1]. At any point during the protocol where Eve is done with her learning algorithm for that round, let W be the messages sent so far, let E be Eve's additional knowledge and let $Q(E)$ denote the set of queries made by Eve. The notation $\mathcal{D}(W, E)$ denotes the joint distribution over the Sender and Receiver's views so far conditioned on (W, E). The event $\mathsf{Good}(W, E)$ is defined over $\mathcal{D}(W, E)$ as follows. Let $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ be an instance of the space $\mathcal{D}(W, E)$. Then we say that $\mathsf{Good}(W, E)$ holds over $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ if there are no intersection queries in $Q(\mathsf{S})$, $Q(\mathsf{R})$

that are not in $Q(E)$. We define the distribution $\mathcal{D}_{\mathsf{Good}}(W,E)$ to be the distribution $\mathcal{D}(W,E)$ conditioned on $\mathsf{Good}(W,E)$.

We now specify the Eve algorithm formally: Eve runs the following with threshold $\varepsilon = \varepsilon_1$ before S sends its message and with threshold $\varepsilon = \varepsilon_2$ after S sends its message.

(1) *As long as there exists a query $q \notin Q(E)$ such that $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\leftarrow\mathcal{D}_{\mathsf{Good}}(W,E)}[q \in Q(\mathsf{S}) \cup Q(\mathsf{R})] \geq \varepsilon$ then ask q from the oracle and add q paired with its answer to E.*

(2) *As long as there exists a pair (pk^*, y^*) $G(sk) = pk^* \in Q(E)$, $F(pk^*,x) = y^* \notin Q(E)$ and $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\leftarrow\mathcal{D}_{\mathsf{Good}}(W,E)}[\exists x : F(pk^*,x) = y^* \in Q(\mathsf{R})] \geq \varepsilon$ then query the oracle on $F^{-1}(sk,y^*)$. If $F^{-1}(sk,y^*)$ returns some value x, then add $F(pk^*,x) = y^*$ to E. If $F^{-1}(sk,y^*)$ returns $\perp$ then add $F^{-1}(sk,y^*) = \perp$ to E.*

(3) *As long as there exists a pair (pk^*, y^*) such that $F(pk^*,x) = y^* \notin Q(E)$ and $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\leftarrow\mathcal{D}_{\mathsf{Good}}(W,E)}[\exists x : F(pk^*,x) = y^* \in Q(\mathsf{S})] \geq \varepsilon$ then if $F(pk^*,x) = y^* \in Q(\mathsf{S})$, add q paired with its answer to E. Otherwise, add the information "query $F(pk^*,x) = y^*$ not made by* S*" to E.*

The assymmetry of the Eve algorithm is due to the fact that the protocol has only one round. More specifically, suppose that when the Eve algorithm completes at the end of the first pass, there is some pair (pk^*, y^*) that is likely for R: $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\leftarrow\mathcal{D}_{\mathsf{Good}}(W,E)}[\exists x : F(pk^*,x) = y^* \in Q(\mathsf{R})] \geq \varepsilon$ [6]. Note that the query $F(pk^*,x) = y^*$ is only added to $Q(E)$ if $G(sk) = pk \in Q(E)$. We claim that if $G(sk) = pk^* \notin Q(E)$ then the probability that R queries $F(pk^*,*) = y^*$ and S queries $G(sk) = pk^*, F^{-1}(sk,y^*)$ is at most $\varepsilon + 1/2^n$. We distinguish between two cases:

$G(sk) \in \mathsf{View}_\mathsf{R}$**.** The probability over $(\mathsf{View}_\mathsf{R}, \mathsf{View}_\mathsf{S}) \sim \mathcal{D}(\textsc{pk}, Q(E)_G)$ that $G(sk) \in Q(\mathsf{R}) \cap Q(\mathsf{S})$ is at most ε (since otherwise, by definition of the Eve algorithm, $G(sk)$ would be in $Q(E)$).

$G(sk) \notin \mathsf{View}_\mathsf{R}$**.** In this case, the probability over $(\mathsf{View}_\mathsf{R}, \mathsf{View}_\mathsf{S}) \sim \mathcal{D}(\textsc{pk}, Q(E)_G)$ that $G(sk) \notin Q(\mathsf{R}), F(pk^*,*) = y^* \in Q(\mathsf{R})$ and $G(sk) = pk^* \in Q(\mathsf{S})$ is at most $1/2^n$ (since with probability $1 - 1/2^n$, pk^* is not a valid public key).

Thus, if $G(sk) = pk^* \notin Q(E)$ then an intersection query will be made with probability at most $\varepsilon + 1/2^n$.

Similarly to [19], we prove that $\mathcal{D}_{\mathsf{Good}}(W,E)$ is *nearly* a product distribution.

Claim 8 *There exists a set of views* Unlikely *such that* $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\sim\mathcal{D}}[\mathsf{Unlikely}] \leq O(m/2^n)$. *Moreover, for every W, E denoting* Eve*'s information up to just before the i-th query, if* $\Pr_{\mathcal{D}(W,E)}[\mathsf{Good}(W,E), \overline{\mathsf{Unlikely}}(W,E)] > 0$ *then there exists a distribution $\mathcal{S}(W,E)$ (resp. $\mathcal{R}(W,E)$) over the Senders (resp. Receiver's) view up to that point such that*

$$\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E) \approx_{1+O(q/2^n)} (\mathcal{S}(W,E) \times \mathcal{R}(W,E)) \mid \mathsf{Good}(W,E), \overline{\mathsf{Unlikely}}(W,E).$$

where we define the distribution $\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E)$ to be the distribution $\mathcal{D}(W,E)$ conditioned on $\mathsf{Good}(W,E)$ *and* $\overline{\mathsf{Unlikely}}(W,E)$.

Following [19], we use the notation $\approx_{1+O(q/2^n)}$ to represent that for every $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ in the support of $\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E)$, we have that the probability $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ is drawn from either distribution differs by at most a factor of $(1 + O(q/2^n)$.

Proof. The event $\mathsf{Unlikely}(W,E)$ is defined over $\mathcal{D}(W,E)$ as follows: Let $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ be an instance of the space $\mathcal{D}(W,E)$. Then we say that $\mathsf{Unlikely}(W,E)$ holds over $(\mathsf{View}_\mathsf{S}, \mathsf{View}_\mathsf{R})$ if some query of the form $F^{-1}(sk,y) \in Q(\mathsf{R}) \cup Q(\mathsf{S})$ such that $F(pk,x) = y \notin Q(E)$ returns a value $x \neq \perp$.

[6] Note that after the first pass $\mathcal{D}_{\mathsf{Good}}(W,E) \equiv \mathcal{D}(W,E)$

Note that $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\sim\mathcal{D}_{\mathsf{Good}}(W,E)}[\mathsf{Unlikely}] \leq O(m/2^n)$.

Now, for the second part of the lemma, we note that

$$\Pr[Q(\mathsf{S}), Q(\mathsf{R})|Q(E)] = \Pr[Q(\mathsf{S}) \mid Q(E)] \times \Pr[Q(\mathsf{R}) \mid Q(E), Q(\mathsf{S})]$$

We show that

$$\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\sim\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E)}[Q(\mathsf{R}) \mid Q(E), Q(\mathsf{S})] \approx_{1+O(q/2^n)} \Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\sim\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E)}[Q(\mathsf{R}) \mid Q(E)],$$

which is sufficient to prove the claim.

Consider the possible settings of $Q(\mathsf{R})$ and the possible settings of $Q(\mathsf{R}) \mid Q(\mathsf{S})$. Since $Q(\mathsf{R})$ and $Q(\mathsf{S})$ do not contain any intersection queries outside of $Q(E)$, the set of possible settings only differ only when $Q(\mathsf{R})$ contains some query q of the form $G(sk)$ or $F(pk, x)$ with image w such that a query of the form $G(*) = w$ or $F(*, *) = w$ appears in $Q(\mathsf{S})$. However, for every fixed setting of $Q(\mathsf{S})$, the probability that the above occurs when choosing $Q(\mathsf{R})$ (conditioned on $Q(E)$, Good, $\overline{\mathsf{Unlikely}}(W, E)$) is bounded by $2q \cdot 1/2^n$, which implies that $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\sim\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E)}[Q(\mathsf{R}) \mid Q(E), Q(\mathsf{S})] \approx_{1+O(q/2^n)} \Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\sim\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E)}[Q(\mathsf{R}) \mid Q(E)]$.

Thus, the distribution $\mathcal{S}(W, E)$ is defined as follows: For any view View_S, the probability of sampling View_S is proportional to $\Pr_{(\mathsf{View}_\mathsf{S},\mathsf{View}_\mathsf{R})\sim\mathcal{D}_{\mathsf{Good},\overline{\mathsf{Unlikely}}}(W,E)}[Q(\mathsf{S})|Q(E)]$. The distribution $\mathcal{R}(W, E)$ is defined analogously.

Given Claim 8, Lemma 2 follows from [1, 7, 19].

B Toy Scheme

Consider the following black-box construction of an encryption scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, from underlying simulatable PKE scheme $\mathcal{E}_\mathsf{Sim} = (\mathsf{Gen}_\mathsf{Sim}, \mathsf{Enc}_\mathsf{Sim}, \mathsf{Dec}_\mathsf{Sim}, \mathsf{oGen}, \mathsf{oRndEnc}, \mathsf{rGen}, \mathsf{rRndEnc})$. We assume WLOG that $\mathsf{Enc}_\mathsf{Sim}$ takes input messages of length n and outputs ciphertexts of length $p(n)$ for some polynomial $p(\cdot)$.

Key Generation $\mathsf{Gen}(1^n)$:

Run $\mathsf{Gen}_\mathsf{Sim}$ $p(n)+1$ times to produce public key, secret-key pairs $(\mathrm{PK}^1_\mathsf{Sim}, \mathrm{SK}^1_\mathsf{Sim}), \ldots, (\mathrm{PK}^{p(n)+1}_\mathsf{Sim}, \mathrm{SK}^{p(n)+1}_\mathsf{Sim})$.
Output $\mathrm{PK} = (\mathrm{PK}^1_\mathsf{Sim}, \ldots, \mathrm{PK}^{p(n)+1}_\mathsf{Sim})$ and $\mathrm{SK} = (\mathrm{SK}^1_\mathsf{Sim}, \ldots, \mathrm{SK}^{p(n)+1}_\mathsf{Sim})$.

Encryption $\mathsf{Enc}(\mathrm{PK}, m)$:

To encrypt a bit 0:

Compute $c_i = \mathsf{Enc}_\mathsf{Sim}(\mathrm{PK}^i_\mathsf{Sim}, 0^n)$ for each $1 \leq i \leq p(n)$. Ouput $c = (c_1, \ldots, c_{p(n)})$.

To encrypt a bit 1:

- Choose $r^* \leftarrow \{0,1\}^n$ at random and compute $c^* = \mathsf{Enc}_\mathsf{Sim}(\mathrm{PK}^{p(n)+1}_\mathsf{Sim}, r^*)$.
- For each of the $p(n)$ bits c^*_i of c^*, if $c^*_i = 0$ then compute $c_i = \mathsf{Enc}_\mathsf{Sim}(\mathrm{PK}^i_\mathsf{Sim}, 0^n)$. Otherwise, generate an oblivious ciphertext $c_i = \mathsf{oRndEnc}(1^n)$.
- Output $c = (c_1, \ldots, c_{p(n)})$.

Decryption $(\mathsf{Dec}(\mathrm{SK}, c)$:

- For $1 \leq i \leq p(n)$ compute $m_i = \mathsf{Dec}_\mathsf{Sim}(\mathrm{SK}^i_\mathsf{Sim}, c_i)$. If $m_i = 0^n$ then set $c^*_i = 0$. Otherwise, set $c^*_i = 1$.
- If the string $c^* = 0^{p(n)}$, output the bit 0.
- Otherwise, compute $r^* = \mathsf{Dec}_\mathsf{Sim}(\mathrm{SK}^{p(n)+1}_\mathsf{Sim}, c^*)$ and output the bit 1.

Notice that for an honest encryption of 0 of the form $c = (c_1, \ldots, c_{p(n)})$ we have that for each $1 \leq i \leq p(n)$ the Sender has computed $\mathsf{Enc}_{\mathsf{Sim}}(\text{PK}^i_{\mathsf{Sim}}, 0^n) = c_i$ and the Receiver has computed $\mathsf{Dec}_{\mathsf{Sim}}(\text{SK}^i_{\mathsf{Sim}}, c_i) = 0^n$. There are no other common queries.

We now present a Faking algorithm which will generate a fresh encryption of a new random message r^*, $c^* = \mathsf{Enc}_{\mathsf{Sim}}(\text{PK}^{p(n)+1}_{\mathsf{Sim}}, r^*)$ and will produce a view for the Sender where the Sender computes $\mathsf{Enc}_{\mathsf{Sim}}(\text{PK}^{p(n)+1}_{\mathsf{Sim}}, r) = c^*$ and (w.h.p. given this view) the Receiver computes $\mathsf{Dec}_{\mathsf{Sim}}(\text{SK}^{p(n)+1}_{\mathsf{Sim}}, c^*) = r^*$. Thus, the Faking algorithm has added a new common query to the views of the Sender and Receiver.

The Fake *algorithm:* Given an honest encryption of 0 of the form $c = (c_1, \ldots, c_{p(n)})$ along with the random coins $(r_1, \ldots, r_{p(n)})$ used to generate it, Fake does the following:

- Choose $r^* \leftarrow \{0, 1\}^n$ at random and compute $c^* = \mathsf{Enc}_{\mathsf{Sim}}(\text{PK}^{p(n)+1}_{\mathsf{Sim}}, r^*)$. Output the coins used to generate c^*.
- For each of the $p(n)$ bits c^*_i of c^*, if $c^*_i = 0$ then output the honest coins used to generate $c_i = \mathsf{Enc}_{\mathsf{Sim}}(\text{PK}^i_{\mathsf{Sim}}, 0^n)$. Otherwise, generate coins claiming that c_i was generated obliviously by outputting $\mathsf{rRndEnc}(\text{PK}^i_{\mathsf{Sim}}, r_i, 0^n)$.

Note that assuming the outputted coins for the Sender are honest, we have that with all but negligible probability both the Sender and Receiver have c^* as a new common query.