

Interactive Encryption and Message Authentication^{*}

Yevgeniy Dodis¹ and Dario Fiore^{2**}

¹ Department of Computer Science, New York University, USA
dodis@cs.nyu.edu

² IMDEA Software Institute, Madrid, Spain
dario.fiore@imdea.org

Abstract. Public-Key Encryption (PKE) and Message Authentication (PKMA, aka as digital signatures) are fundamental cryptographic primitives. Traditionally, both notions are defined as non-interactive (i.e., single-message). In this work, we initiate rigorous study of (possibly) *interactive* PKE and PKMA schemes. We obtain the following results demonstrating the power of interaction to resolve questions which are either open or impossible in the non-interactive setting.

Efficiency/Assumptions. One of the most well known open questions in the area of PKE is to build, in a “black-box way”, so called chosen ciphertext attack (CCA-) secure PKE from chosen plaintext attack (CPA-) secure PKE. In contrast, we show a simple 2-round CCA-secure PKE from any (non-interactive) CPA-secure PKE (in fact, these primitives turn out to be equivalent). Similarly, although non-interactive PKMA schemes can be inefficiently built from any one-way function, no efficient signature schemes are known from many popular number-theoretic assumptions, such as factoring, CDH or DDH. In contrast, we show an efficient 2-round PKMA from most popular assumptions, including factoring, CDH and DDH.

Advanced Properties. It is well known that no non-interactive signature (resp. encryption) scheme can be *deniable* (resp. *forward-secure*), since the signature (resp. ciphertext) can later “serve as an evidence of the sender’s consent” (resp. “be decrypted if the receiver’s key is compromised”). We also formalize a related notion of *replay-secure* (necessarily) interactive PKMA (resp. PKE) schemes, where the verifier (resp. encryptor) is assured that the “current” message can only be authenticated (resp. decrypted) by the secret key owner *now*, as opposed to some time in the past (resp. future). We observe that our 2-round PKMA scheme is both replay-secure and (passively) deniable, and our 2-round PKE scheme is both replay- and forward-secure.

Keywords: Public Key Encryption, Digital Signatures, Chosen Ciphertext Security, Man-in-the-Middle Attacks.

^{*} An extended abstract of this work appears in the proceedings of SCN 2014. This is the full version.

^{**} Work partially done while postdoc at NYU.

Table of Contents

Interactive Encryption and Message Authentication	1
<i>Yevgeniy Dodis and Dario Fiore</i>	
1 Introduction	1
1.1 Our Results	1
1.2 Preliminaries and Notation.	3
1.3 Organization of the paper	4
2 Defining Message Transmission Protocols	4
2.1 Interactive Chosen-Ciphertext-Secure Encryption.	6
2.2 Interactive Chosen-Message Secure Public Key Message Authentication.	7
3 Basic Constructions	7
3.1 iCCA Encryption from IND–CPA Encryption and iCMA PKMA	7
3.2 iCMA-secure PKMA from iCCA security	10
3.3 Secure Round Extension of Message Transmission Protocols	12
4 Advanced Security Properties from the Power of Interaction	13
4.1 Deniability	14
4.2 Forward Security	15
4.3 Replay Security	16
A Standard Cryptographic Primitives	22
A.1 (Non-Interactive) CCA-secure Public-Key Encryption	22
A.2 Digital Signatures	23
A.3 Message Authentication Codes	23
B DDN 3-round iCCA-Secure Encryption from IND–CPA Security	24
C A 1-bounded IND–CCA-Secure Encryption Scheme	24
D iCCA-Secure Interactive Encryption with Labels	25
E iCMA-secure PKMA from labeled iCCA-secure encryption.	26
F Postponed Proofs	27
F.1 Proof of Claim 1	27
F.2 Proof of Theorem 3 (iCCA security)	27
F.3 Proof of Theorem 3 (iCMA security)	29

1 Introduction

Digital signatures and public-key encryption (PKE) schemes are two of the most fundamental and well studied cryptographic primitives. Traditionally, both notions are defined as *non-interactive* (i.e., single message). Aside from obvious convenience — both the sender and receiver need not keep any state — such non-interactive “one-and-done” philosophy is also critical in many applications. Coupled with the fact that we now have many (often efficient) candidates for both signature and encryption schemes, it might appear that there is little value in extending the syntax of signature/encryption schemes to allow for (possibly) interactive realizations.

In this work we challenge this point of view, and initiate rigorous study of (possibly) *interactive signature and public-key encryption schemes*. For the former, we will actually use the term *Public-Key Message Authentication* (PKMA) scheme, as the term “signature” often comes with expectations of “non-repudiation”, which is orthogonal to the standard notion of “unforgeability” the moment interaction is allowed. First, although we agree that some applications might critically rely on the non-interactivity of PKE/PKMA schemes, we believe that many other applications, including *arguably the most basic one of sending/receiving private/authentic messages*, might not so be fixated on non-interactivity. For such applications, it appears natural to allow the sender and the receiver to interact, especially if they are involved in a conversation anyway.

Second, in this work we will show that, by allowing a single extra message (i.e., a 2-round protocol), we can “resolve” two arguably most important open problems in the area of non-interactive PKE/PKMA schemes:³ (a) “black-box” 2-round chosen-ciphertext attack (CCA-) secure PKE from any (non-interactive) chosen-plaintext attack (CPA-) secure PKE; (b) *efficient* 2-round strongly unforgeable PKMA scheme from a variety of simple assumptions, such as factoring and DDH.

Third, we point out several useful advanced properties of PKE/PKMA schemes which are *impossible to achieve without interaction*. While some of these properties (such as *deniable* PKMA [19,21,20,15]) were already extensively studied in the past, most others (such as interactive *forward-secure* PKE, and *replay-secure* PKE/PKMA) appear to be new.

RELATED WORK. Although our work is the first to offer a detailed and comprehensive study of interactive PKE/PKMA schemes, it is certainly not the first to consider these notions. The most related prior work in this regard is the famous “DDN-paper” on non-malleable cryptography [18,19]. This seminal work had many extremely important and influential results. Among them, it also considered non-malleable, interactive encryption and authentication, and briefly sketched⁴ elegant constructions for both primitives. We discuss more in detail the relation with our work in the next section, when we describe our improvements over the DDN paper.

To the best of our knowledge, the only other work providing a related definition (only for encryption) is the one of Katz [29]. However, our definition is stronger, as we place more restriction on the attacker to declare that the attacker ‘acts as a wire’. Moreover, the solutions given in [29] use so called timing assumptions, while our constructions are in the standard model.

1.1 Our Results

We now describe our motivations and results in more detail.

DEFINITIONAL FRAMEWORK. Our first goal was to extend the short and elegant definitions of non-interactive encryption/signatures to the interactive setting, without making the definitions long and tedious. Unfortunately, in the interactive setting things *are* more complicated, as issues of concurrency

³ Of course, we obtain this by changing the model to allow for interaction, which is the reason for the quotation marks.

⁴ Due to the massive scope of [19], the DDN paper did not give formal definitions and proofs for the encryption results (saying they are “outside the scope” of their paper; see page 32), and only sketched the definition/proof for the authentication case.

and state, among others, must be dealt with. The way we managed to achieve our goal, was to split our definitions into two parts. The first (somewhat boring) part is *independent* of the particular primitive (e.g., PKE/PKMA), and simply introduces the bare minimum of notions/notation to deal with interaction. For example (see Section 2 for details), we define (a) what it means to have (concurrent) oracle access to an *interactive party* under attack; and (b) what it means to ‘act as a wire’ between two honest parties (for brevity, we call this trivial, but unavoidable, attack a ‘ping-pong’ attack). Once the notation is developed, however, our actual definitions of possibly interactive PKE/PKMA are *as short and simple as in the non-interactive setting* (see Definitions 5 and 6). E.g., in the PKMA setting (Definition 6) the attacker \mathcal{A} has (concurrent) oracle access to the honest signer (as defined in (a)), and simultaneously tries to convince an honest verifier (i.e., “challenger”). \mathcal{A} wins if the challenger accepts, and \mathcal{A} ’s forgery was not a ‘ping-pong’ of one of its conversations with the signer (as defined in (b)).⁵ Overall, the definition consists of the same couple of lines as in the non-interactive setting! And the same holds for the encryption case in Definition 5, which naturally generalizes the notion of CCA-security to the interactive setting.

BETTER EFFICIENCY/ASSUMPTIONS VIA INTERACTION. Turning from definitions to constructions, we show how a *single* extra round of interaction (i.e., a 2-round protocol) can help “solve” (in a sense explained in Footnote 3) two of the arguably toughest open problems in the areas of non-interactive PKE and PKMA, respectively.

In the area of PKE, the question is to build a CCA-secure PKE from a CPA-secure PKE. In principle, such constructions are known using an appropriately-chosen notion of non-interactive zero-knowledge proofs [19,36,37,39,31]. However, all these constructions are generally inefficient and considered somewhat unsatisfactory, since they use the code of the given CPA-secure encryption scheme. In particular, the question of finding so called *black-box* constructions of CCA-encryption from CPA-encryption remains open. In fact, although several partial progress along both positive (e.g., [12,11,32]) and negative (e.g., [25]) fronts was made, the general question remains elusive. In contrast, we show a relatively simple, black-box, 2-round CCA-secure encryption from CPA-secure encryption (in fact, the two primitives are *equivalent*). We notice that the DDN paper [19] itself already made a similar conclusion, by presenting a 3-round black-box protocol from any CPA-secure PKE.⁶ Thus, aside from presenting a formal model for interactive CCA-secure encryption, our result can be viewed as improving the round efficiency of the DDN paper from 3 to 2.

In the area of PKMA, the “theory” question was settled pretty quickly, by showing that the strongest security notion for signature schemes — strong existential unforgeability against chosen message attack — can be realized assuming the mere existence of one-way functions [26,35,3,4,38]. Unfortunately, the generic constructions were primarily of theoretical interest, and did not result in practical enough signature schemes. In fact, practical signature schemes (outside of the random oracle model [7,23,40,27]) are only constructed from a handful of “not-too-standard” number-theoretic assumptions, such as ‘strong RSA’ [14,24] and ‘Bilinear-Diffie-Hellman’ [41], and even these ‘practical’ constructions were generally somewhat slow, requiring generation of primes, long keys or bilinear maps. In particular, one of the main open questions is to build an *efficient* digital signature scheme from a standard assumption, such as factoring, CDH or DDH.

In contrast, we show very efficient, 2-round, strongly unforgeable PKMA schemes from virtually all standard assumptions, including factoring, CDH and DDH. In fact, although we are not aware of any paper explicitly claiming this result, *it follows in a relatively simple manner by combining various prior works*.⁷ Let us explain. With a different motivation in mind, the DDN paper [19] showed a simple 3-

⁵ This generalizes the notion of *strong* unforgeability, as opposed to regular unforgeability, as was done in DDN [19].

⁶ Although they do not give a formal definition/proof, their construction is easily seen to be secure in our model (see Appendix B).

⁷ Except only establishing regular unforgeability, but the actual constructions are easily seen to be strongly unforgeable.

round⁸ PKMA scheme from any CCA-secure encryption. At the time, CCA-secure PKE was considered a very ‘advanced’ primitive, so the construction was not considered ‘efficient’. Over the years, though, many truly efficient CCA-secure schemes were constructed from virtually all popular assumptions, including factoring [28], CDH [42] and DDH [13] (despite the fact that no such efficient signature schemes are known from these assumptions!). Thus, these results, if combined, immediately yield an efficient 3-round PKMA from all these assumptions. Moreover, it is well known (e.g., see [30,1]) that one can reduce the number of rounds from 3 to 2 by also using a message authentication code (MAC), in addition to CCA-secure encryption. Of course, in theory, a MAC is implied by a CCA-secure PKE, albeit in an inefficient manner. Moreover, until recently, even direct efficient MAC constructions from concrete assumptions, such as DDH [33] and factoring [34], required long keys (quadratic in security parameter). Fortunately, Dodis et al. [17] recently observed an elementary efficient (probabilistic) MAC construction from any CCA-secure scheme. This gives an efficient 2-round PKMA scheme from any CCA-secure encryption. We also manage to further optimize the resulting construction, and obtain a really simple (new!) 2-round protocol, depicted in Figure 3. In turn, this gives efficient 2-round PKMA from a variety of standard assumptions, including factoring, CDH and DDH.⁹

DUALITY BETWEEN INTERACTIVE PKE AND PKMA. Interestingly, our 2-round CCA-secure PKE uses a *signing* key as its long-term “decryption secret” (and generates several ephemeral keys for the CPA-secure scheme), while our 2-round strongly unforgeable PKMA scheme uses a *decryption* key for a CCA-secure encryption as its long term “authentication secret”. We show that this duality is not a coincidence. In fact, our 2-round results follow as corollaries of two more general schemes, depicted in Figures 1 and 2: an interactive CCA-secure scheme from any (interactive or not) strongly unforgeable PKMA scheme (plus any CPA-secure PKE¹⁰), and an interactive strongly unforgeable PKMA scheme from any (interactive or not) CCA-secure PKE. The ‘duality’ of our results (authentication using encryption and vice-versa) shows that, perhaps, the practical/theoretical distinction between *interactive* encryption and authentication is not as great as one could have guessed by looking at what is known in the *non-interactive* setting.

OVERCOMING IMPOSSIBILITY RESULTS. We also show that our simple definitional framework (one generic/reusable ‘long-and-boring’ part followed by many application-specific ‘short-and-intuitive’ parts) easily lends itself to various extensions. Examples include various notions of privacy and/or authenticity which are *impossible* in the non-interactive setting, such as forward-secure PKE and the notion of replay-secure PKE/PKMA that we introduce in this work.

1.2 Preliminaries and Notation.

We denote by $\lambda \in \mathbb{N}$ the security parameter. A function $\epsilon(\lambda)$ is said *negligible* if it is a positive function that vanishes faster than the inverse of any polynomial in λ . If X is a set, let $x \stackrel{\$}{\leftarrow} X$ denote the process of selecting x uniformly at random in X . An algorithm \mathcal{A} is called *PPT* if it is a probabilistic Turing machine whose running time is bounded by some polynomial in λ . If \mathcal{A} is a PPT algorithm, then $y \stackrel{\$}{\leftarrow} \mathcal{A}(x)$ denotes the process of running \mathcal{A} on input x and assigning its output to y . Definitions of standard cryptographic primitives are recalled in Appendix A.

⁸ The construction becomes 2-round if the verifier knows the authenticated message in advance.

⁹ Of course, in practice one should not use CCA-secure encryption to build a MAC (instead, one should use practical MACs such as CBC-MAC or HMAC), but here we use it to establish *efficient feasibility results* from concrete number-theoretic assumptions.

¹⁰ For the sake of elegance and modularity, we use a slightly stronger notion of so called “1-bounded CCA-secure PKE”, but the latter can be built from any CPA-secure PKE [12].

1.3 Organization of the paper

The paper is organized as follows. In Section 2 we introduce our definitional framework for message transmission protocols with the security notions of (possibly interactive) iCCA-secure PKE and iCMA-secure PKMA. Next, in Section 3, we focus on realizations of these protocols, notably, iCCA-secure PKE from iCMA-secure PKMA and CPA-secure encryption, and iCMA-secure PKMA from iCCA-secure PKE. Finally, Section 4 discusses advanced security properties for (interactive) PKE and PKMA. We postpone to the appendix the definitions of standard cryptographic primitives (Appendix A), the 3-round interactive encryption of Dolev, Dwork and Naor [19], the construction of a 1-bounded-IND-CCA-secure PKE scheme from an IND-CPA-secure one (Appendix C), the extension of (interactive) PKE to support “labels” and its application to PKMA (Appendices D–E), and a few simple proofs (Appendix F).

2 Defining Message Transmission Protocols

In this section we introduce *message transmission protocols*: we define their syntax as well as suitable notions of confidentiality (called iCCA security) and authenticity (called iCMA security).

We give a generic definition of message transmission protocols involving two parties: a sender S and a receiver R , such that the goal of S is to send a message m to R while preserving certain security properties on m . In particular, in the next two sections we consider arguably the most basic security properties: the confidentiality/authenticity of the messages sent by S to R . Formally, a message transmission protocol Π consists of algorithms (Setup, S, R) defined as follows:

$\text{Setup}(1^\lambda)$: on input the security parameter λ , the setup algorithm generates a pair of keys $(\text{sendk}, \text{recvk})$.

In particular, these keys contain an implicit description of the message space \mathcal{M} .

$S(\text{sendk}, m)$: is a possibly interactive algorithm that is run with the sender key sendk and a message $m \in \mathcal{M}$ as private inputs.

$R(\text{recvk})$: is a possibly interactive algorithm that takes as private input the receiver key recvk , and whose output is a message $m \in \mathcal{M}$ or an error \perp .

When S and R are jointly run, they exchange messages in a specific order, e.g., S starts by sending M_1 , R sends M_2 , S sends M_3 , and so on and so forth until they both terminate. We say that Π is an n -round protocol if the number of messages exchanged between S and R during a run of the protocol is n . If Π is 1-round, then we say that Π is *non-interactive*. Since the sender gets no output, we assume without loss of generality that the *sender always speaks last*.¹¹ Thus, in an n -round protocol, R (resp. S) speaks first if n is even (resp. odd). For compact notation, we denote with $\langle S(\text{sendk}, m), R(\text{recvk}) \rangle = m'$ the process of running S and R on inputs (sendk, m) and recvk respectively, and assigning R 's output to m' . In our notation, we will use $m \in \mathcal{M}$ for messages (aka plaintexts), and capital M for protocol messages.

Definition 1 (Correctness). *A message transmission protocol $\Pi = (\text{Setup}, S, R)$ is correct if for all honestly generated keys $(\text{sendk}, \text{recvk}) \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda)$, and all messages $m \in \mathcal{M}$, we have that $\langle S(\text{sendk}, m), R(\text{recvk}) \rangle = m$ holds with all but negligible probability.*

Defining Security: Man-in-the-Middle Adversaries. In our work, we assume that the sender and the receiver speak in the presence of powerful adversaries that have full control of the communication channel, i.e., the adversary can eavesdrop the content of the communication, and it can stop/delay/alter the messages passing over the channel. Roughly speaking, the goal of an adversary is to violate a given security property (say confidentiality or authenticity) in a run of the protocol that we call the *challenge*

¹¹ This holds wlog because if Π is an n -round protocol in which R sends the last message, then we can construct an $(n - 1)$ -round protocol Π' by simply dropping the last message from R to S (and that part of the code of R generating the last message is run internally).

session. Formally, this session is a protocol execution $\langle S(\text{sendk}, m), \mathcal{A}^{R(\text{recvk})} \rangle$ or $\langle \mathcal{A}^{S(\text{sendk}, \cdot)}, R(\text{recvk}) \rangle$ where the adversary runs with a honest party (S or R). By writing \mathcal{A}^P , we mean that the adversary has oracle access to *multiple* honest copies of party P (where $P = R$ or $P = S$), i.e., \mathcal{A} can start as many copies of P as it wishes, and it can run the message transmission protocol with each of these copies. This essentially formalizes the fact that the adversary can “sit in the middle of two honest parties” relaying messages between them in an active way. Sometimes, we also write \mathcal{A}^{P_k} to denote that the adversary can start at most k copies of party P. In our model we assume that whenever \mathcal{A} sends a message to the oracle P, then \mathcal{A} always obtains P’s output. In particular, in the case of the receiver oracle, when \mathcal{A} sends the last protocol message to R, \mathcal{A} obtains the (private) output of the receiver, i.e., a message m or \perp .

Since all these protocol sessions can be run in a concurrent way, the adversary might entirely replay the challenge session by using its oracle. This is something that we would like to prevent in our definitions. To formalize this idea, we take an approach similar to the one introduced by Bellare and Rogaway [6] in the context of key exchange, which is based on the idea of “matching conversations”. First of all, we introduce a notion of time during the execution of the security experiment. We stress that this is done for ease of analysis of the security model: there is no need to keep track of global timing in the real protocols. Let t be a global counter which is progressively incremented every time a party (including the adversary) sends a message, and assume that every message sent by a party (S, R or \mathcal{A}) gets timestamped with the current time t . Note that this includes all messages of the sessions established by the adversary using its oracle. Using this notion of time, we define the transcript of a protocol session as follows:

Definition 2 (Protocol Transcript). *The transcript of a protocol session between two parties is the timestamped sequence of messages exchanged by the parties during a run of the message transmission protocol Π . If Π is n -round, then a transcript T is of the form $T = \langle (M_1, t_1), \dots, (M_n, t_n) \rangle$, where M_1, \dots, M_n are the exchanged messages, and t_1, \dots, t_n are the respective timestamps.*

In a protocol run $\langle S(\text{sendk}, m), \mathcal{A}^{R(\text{recvk})} \rangle$ (resp. $\langle \mathcal{A}^{S(\text{sendk}, \cdot)}, R(\text{recvk}) \rangle$) we have a transcript T^* of the challenge session between S and \mathcal{A} (resp. \mathcal{A} and R), and Q transcripts T_1, \dots, T_Q , one for each of the Q sessions established by \mathcal{A} with R (resp. S) via the oracle.

While we postpone to the next two sections the definition of specific security properties of message transmission (e.g., confidentiality and authenticity), our goal here is to formalize in a generic fashion which adversaries are effective for “uninteresting”/unavoidable reasons. Namely, when the challenge session is obtained by entirely replaying one of the oracle sessions: what we call a “ping-pong” attack, that we formalize via the following notion of matching transcripts.

Definition 3 (Matching Transcripts). *Let $T = \langle (M_1, t_1), \dots, (M_n, t_n) \rangle$ and $T^* = \langle (M_1^*, t_1^*), \dots, (M_n^*, t_n^*) \rangle$ be two protocol transcripts. We say that T matches T^* ($T \equiv T^*$, for short) if $\forall i = 1, \dots, n$, $M_i = M_i^*$ and the two timestamp sequences are “alternating”, i.e., $t_1 < t_1^* < t_2 < t_2^* < t_3 < \dots < t_{n-1} < t_{n-1}^* < t_n < t_n^*$ if R speaks first, or $t_1^* < t_1 < t_2^* < t_2 < t_3^* < \dots < t_{n-1}^* < t_{n-1} < t_n^*$ if S speaks first. We remark that the notion of match is not commutative.*

Given all the definitions above, we can define the notion of ping-pong adversary:

Definition 4 (Ping-pong Adversary). *Consider a run of the protocol Π involving \mathcal{A} and a honest party (it can be either $\langle S(\text{sendk}, m), \mathcal{A}^{R(\text{recvk})} \rangle$ or $\langle \mathcal{A}^{S(\text{sendk}, \cdot)}, R(\text{recvk}) \rangle$), and let T^* be the transcript of the challenge session, and T_1, \dots, T_Q be the transcripts of all the oracle sessions established by \mathcal{A} . Then we say that \mathcal{A} is a ping-pong adversary if there is a transcript $T \in \{T_1, \dots, T_Q\}$ such that T matches T^* , i.e., $T \equiv T^*$.*

2.1 Interactive Chosen-Ciphertext-Secure Encryption.

Here we propose a suitable notion of confidentiality for message transmission protocols that we call *interactive chosen ciphertext security* (iCCA). Our notion is designed as a very natural generalization of the classical notion of IND-CCA security to the interactive setting. In fact, IND-CCA security is a special case of iCCA security for 1-round (i.e., non-interactive) protocols (we leave this check to the reader). Roughly speaking, in the IND-CCA definition the adversary has to distinguish whether a given “challenge” ciphertext encrypts a message m_0 or m_1 while having access to a decryption oracle. To make the definition non-trivial, the adversary is denied to query the decryption oracle on the challenge ciphertext. Our notion of iCCA security is obtained similarly: the adversary \mathcal{A} interacts with a sender which sends either m_0 or m_1 and \mathcal{A} has to tell the two cases apart while having access to the receiver (instead of the decryption oracle); the restriction on the challenge ciphertext is replaced by requiring that \mathcal{A} cannot be ping-pong. The formal definition follows.

Let $\Pi = (\text{Setup}, \text{S}, \text{R})$ be a message transmission protocol and \mathcal{A} be an adversary. To define iCCA security, consider the following experiment:

Experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda)$

$b \xleftarrow{\$} \{0, 1\}$

$(\text{sendk}, \text{recvk}) \xleftarrow{\$} \text{Setup}(1^\lambda)$

$(m_0, m_1) \leftarrow \mathcal{A}^{\text{R}(\text{recvk})}(\text{sendk})$

$b' \leftarrow \langle \text{S}(\text{sendk}, m_b), \mathcal{A}^{\text{R}(\text{recvk})}(\text{sendk}) \rangle$

If $b' = b$ and \mathcal{A} is not “ping-pong”, then output 1. Else output 0.

Definition 5 (iCCA security). For any $\lambda \in \mathbb{N}$, we define the advantage of an adversary \mathcal{A} in breaking iCCA security of a message transmission protocol Π as $\text{Adv}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda) = \Pr[\text{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda) = 1] - \frac{1}{2}$, and we say that Π is iCCA-secure if for any PPT \mathcal{A} , $\text{Adv}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda)$ is negligible. We call a message transmission protocol satisfying this notion an interactive encryption scheme.

As we mentioned in the introduction, it is worth noting that our definition is similar to the one proposed by Katz in [29]. The main difference is in the restrictions applied to the adversary in the security game. In [29] this is realized by means of a notion of “equality of transcripts” which considers only equality of messages (in the same order) and leaves any time constraint to the specific protocols realizations. Our security definition, instead, directly takes into account time constraints in the security experiment via the notion of matching transcripts. To see the difference between the two definitions with an example, consider an adversary who creates an oracle session having the same transcript as the one of the challenge session, but where the timestamps of the messages are not correctly alternating. Such an adversary would not be legal according to the definition of [29], but is legal (i.e., not ping-pong) according to ours.

Later, in Section 4.3, we strengthen our requirements by considering an orthogonal security property for interactive encryption that we call “Replay Security”. This will allow to model MiM attacks more easily and from a different perspective. It will also turn out to be realizable under our basic iCCA notion. However, since a replay attack is *always* possible in any non-interactive solution, replay security will only be realizable by an *interactive* protocol.

q -bounded-iCCA Security. In this work we also consider a weaker notion of iCCA security, called *q -bounded-iCCA*, in which the adversary is restricted to complete at most q sessions with the oracle R , i.e., in $\text{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda)$ we run \mathcal{A}^{R^q} . This is the analogous of (non-interactive) q -bounded IND-CCA security [12].

Labeled (Interactive) Encryption. We define an extension of interactive encryption in which both the sender and the receiver algorithms take a public string—a label—as an additional input (similarly

to the non-interactive setting [8]). Very intuitively, the idea is that the receiver working with label L decrypts correctly only if the sender works with the same label L . In Appendix D we provide a full formalization of *labeled* iCCA encryption and show that it can be generically constructed from ‘plain’ iCCA-secure encryption. Here we briefly recall that for any label L we use S^L to denote that an algorithm S takes as input L .

2.2 Interactive Chosen-Message Secure Public Key Message Authentication.

Here we propose a suitable notion of authenticity for message transmission protocols that we call *interactive unforgeability under chosen message attacks* (iCMA). Our notion is designed as a very natural generalization to the interactive setting of the standard notion of strong unforgeability (suf-cma) for digital signatures. In fact, suf-cma security is a special case of iCMA security for 1-round protocols. Roughly speaking, in strong unforgeability the adversary has to produce a valid signature while having access to the signer. In order for the definition to be non-trivial, however, such signature has to be “new”, i.e., not obtained from the signing oracle. Our notion of iCMA security naturally extends suf-cma as follows: the adversary \mathcal{A} has to convince a receiver while having oracle access to the sender (instead of the signing oracle); the requirement that the signature must be new is replaced by requiring that \mathcal{A} is not ping-pong. The formal definition follows.

Let $\Pi = (\text{Setup}, S, R)$ be a message transmission protocol and \mathcal{A} be an adversary. To define iCMA security, consider the following experiment:

Experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{iCMA}}(\lambda)$

$(\text{sendk}, \text{recvk}) \xleftarrow{S} \text{Setup}(1^\lambda)$

$m^* \leftarrow \langle \mathcal{A}^{S(\text{sendk}, \cdot)}(\text{recvk}), R(\text{recvk}) \rangle$

If $m^* \neq \perp$ and \mathcal{A} is not “ping-pong”, then output 1. Else output 0.

Definition 6 (iCMA security). For any $\lambda \in \mathbb{N}$, the advantage of \mathcal{A} in breaking the iCMA security of a message transmission protocol Π is $\text{Adv}_{\Pi, \mathcal{A}}^{\text{iCMA}}(\lambda) = \Pr[\text{Exp}_{\Pi, \mathcal{A}}^{\text{iCMA}}(\lambda) = 1]$, and we say that Π is iCMA-secure if for any PPT \mathcal{A} , $\text{Adv}_{\Pi, \mathcal{A}}^{\text{iCMA}}(\lambda)$ is negligible. We call a message transmission protocol satisfying this notion a public key message authentication (PKMA) protocol.

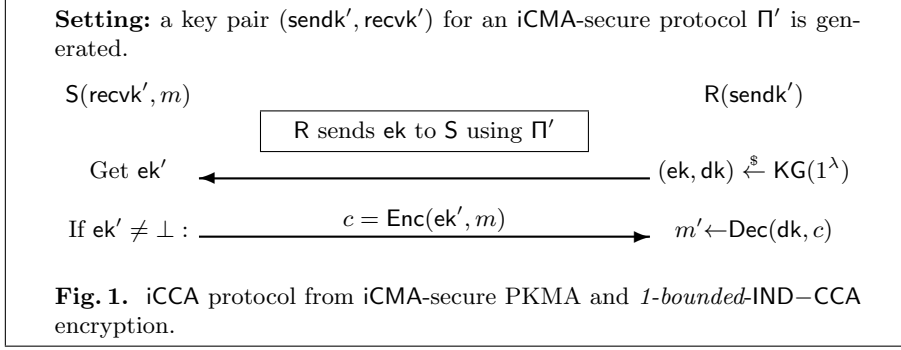
Later, we further motivate the importance of interactive PKMA protocols by considering additional security notions that can be realized *only* in the interactive setting. This is the case for our notion of Replay Security (see Section 4.3), and for the *deniability* property (see Section 4.1). The former notion intuitively captures the idea of obtaining security by denying the adversary to access the sender oracle during the challenge session. Deniability, instead, allows the authenticator to deny his active participation in a protocol, i.e., it denies the verifier to transfer the information authenticated by S .

3 Basic Constructions

In this section we propose realizations of message transmission protocols satisfying our iCCA and iCMA security notions. Interestingly, our constructions show that iCCA security is implied by the iCMA and IND-CPA notions, whereas (somehow vice-versa) iCMA security is directly implied by iCCA. Our results thus show that in the interactive setting—and with a minimum level of interaction (i.e., 2 rounds) indeed!—the notions of confidentiality and authenticity present somewhat surprising and interesting relations unknown in the non-interactive case.

3.1 iCCA Encryption from IND-CPA Encryption and iCMA PKMA

Our result is a simple interactive encryption protocol that is based on a PKMA protocol Π' and a public key encryption scheme $\mathcal{E} = (\text{KG}, \text{Enc}, \text{Dec})$. The idea is simple and is illustrated in Figure 1: the receiver



sends a “fresh” public key ek authenticated using Π' , and the sender encrypts the message using ek . As we show in our theorem below, the PKMA protocol has to be iCMA-secure, while the PKE scheme \mathcal{E} needs only to be 1-bounded-IND-CCA-secure. Concretely, Π' can be a strongly unforgeable signature and \mathcal{E} can be constructed using an IND-CPA-secure encryption (as shown by Cramer et al. [12] and recalled in Appendix C), thus yielding an *optimal* 2-round encryption protocol that is iCCA-secure based only on IND-CPA security. A more precise description follows.

Let $\Pi' = (\text{Setup}', S', R')$ be a PKMA protocol, and $\mathcal{E} = (\text{KG}, \text{Enc}, \text{Dec})$ be a (non-interactive) public key encryption scheme. We build protocol $\Pi = (\text{Setup}, S, R)$ as follows:

Setup (1^λ) : run $(\text{sendk}', \text{recvk}') \xleftarrow{\$} \text{Setup}'(1^\lambda)$ and output $\text{sendk} = \text{recvk}'$ and $\text{recvk} = \text{sendk}'$.

S (sendk, m) : first run the PKMA protocol Π' with R by playing the role of the receiver, i.e., S runs $R'(\text{recvk}')$. If Π' terminates correctly with output ek , then send $c = \text{Enc}(\text{ek}, m)$ to R . Otherwise, stop running.

R (recvk) : generate $(\text{ek}, \text{dk}) \xleftarrow{\$} \text{KG}(1^\lambda)$, run $S'(\text{sendk}', \text{ek})$ to send ek with authenticity to S , and keep dk as private information. After Π' terminates, on input a message c from S , the receiver algorithm computes $m \leftarrow \text{Dec}(\text{dk}, c)$ and returns the message m as its private output.

Theorem 1. *If \mathcal{E} is 1-bounded-IND-CCA-secure and Π' is iCMA-secure, then Π is iCCA-secure.*

Proof. Consider the experiment $\mathbf{Exp}_{\text{PKE}, \mathcal{A}}^{\text{iCCA}}$ in which T_1, \dots, T_Q and T^* are the transcripts of the oracle sessions and the challenge session. Each transcript T can be written as $T = (T', T_c)$ such that T' is the transcript of the PKMA protocol Π' and T_c is the portion containing the last message c . Let **Forge** be the event that in $\mathbf{Exp}_{\text{PKE}, \mathcal{A}}^{\text{iCCA}}$ the adversary \mathcal{A} is not ping-pong in the subprotocol Π' (i.e., $T'_i \neq T'^*$ $\forall i = 1, \dots, Q$), and that in the challenge session the sender accepts (i.e., R' terminates correctly returning $\text{ek}^* \neq \perp$). It is not hard to see that

$$\mathbf{Adv}_{\mathcal{A}, \Pi}^{\text{iCCA}}(\lambda) \leq \Pr[\mathbf{Exp}_{\text{PKE}, \mathcal{A}}^{\text{iCCA}} = 1 \mid \overline{\text{Forge}}] - \frac{1}{2} + \Pr[\text{Forge}].$$

The security of our protocol then follows from showing that: (1) **Forge** occurs with negligible probability under the assumption that Π' is iCMA-secure, and (2) if **Forge** does not occur, then any adversary winning in $\mathbf{Exp}_{\text{PKE}, \mathcal{A}}^{\text{iCCA}}$ can be used to break the 1-bounded-IND-CCA security of the encryption scheme. We formally prove these facts in the following claims.

Claim 1 *If Π' is iCMA-secure, then $\Pr[\text{Forge}]$ is negligible.*

Claim 2 *If \mathcal{E} is 1-bounded IND-CCA-secure, then $\Pr[\mathbf{Exp}_{\text{PKE}, \mathcal{A}}^{\text{iCCA}} = 1 \mid \overline{\text{Forge}}] - \frac{1}{2}$ is negligible.*

The proof of Claim 1 is straightforward and appears in Appendix F.1. The proof of Claim 2 is given below.

Proof (Claim 2). Assume by contradiction there exists an efficient adversary \mathcal{A} such that

$$\Pr[\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda) = 1 \mid \overline{\text{Forge}}] - \frac{1}{2} \geq \epsilon$$

is non-negligible. Then we build a PPT adversary \mathcal{B} that has non-negligible advantage in breaking the 1-bounded-IND-CCA security of the encryption scheme $\mathcal{E} = (\text{KG}, \text{Enc}, \text{Dec})$.

On input a public key $\bar{\text{ek}}$, \mathcal{B} works as follows:

- Generate a key pair $(\text{sendk}', \text{recvk}') \xleftarrow{\$} \text{Setup}'(1^\lambda)$, set $\text{sendk} = \text{recvk}'$ and run $\mathcal{A}(\text{sendk})$.
- Initialize a counter $j \leftarrow 0$ for the number of new sessions opened by \mathcal{A} with R during the experiment.
- Choose a random index $\mu \xleftarrow{\$} \{1, \dots, Q\}$ where $Q = \text{poly}(\lambda)$ is an upper bound on the number of sessions opened by \mathcal{A} with the oracle receiver R . Since Forge does not occur, \mathcal{A} is ping-pong in the subprotocol Π' and thus in the challenge session \mathcal{A} will re-use one of the public keys $\text{ek}_1, \dots, \text{ek}_Q$ obtained by the oracle R . Therefore, μ represents a guess for the index of such public key ek_μ .
- For every oracle query asking to interact with a new copy of R : first, increment j by 1. Now, let j be the index of the current query. If $j \neq \mu$, then \mathcal{B} generates a new encryption key pair $(\text{ek}_j, \text{dk}_j) \xleftarrow{\$} \text{KG}(1^\lambda)$, and runs $S'(\text{sendk}', \text{ek}_j)$ to simulate all the oracle answers of this session corresponding to the run of Π' . Otherwise, if $j = \mu$, it sets $\text{ek}_\mu = \bar{\text{ek}}$, and runs $S'(\text{sendk}', \bar{\text{ek}})$.
- When \mathcal{A} queries R with the last protocol message c_i on the i -th opened session: let ek_i be the public key previously generated in the above step. If $i \neq \mu$, then \mathcal{B} knows the corresponding decryption key dk_i , (\mathcal{B} generated it by itself), and it answers by computing $m_i \leftarrow \text{Dec}(\text{dk}_i, c_i)$. If $i = \mu$, \mathcal{B} asks c_i to the decryption oracle, obtains \tilde{m} , and answers \tilde{m} . Notice that for $i = \mu$ such a query can occur only once, as c_i is the last message of the protocol (the session is then closed).
- Let (m_0, m_1) be the message pair returned by the adversary \mathcal{A} . Then the challenge session starts and \mathcal{A} is expected to “speak first”, by sending ek^* using Π' . Since Forge does not occur, we have that either $\text{ek}^* = \perp$, or $\text{ek}^* \neq \perp$ and \mathcal{A} is ping-pong, i.e., $\text{ek}^* \in \{\text{ek}_1, \dots, \text{ek}_Q\}$. In the first case \mathcal{B} returns an error (this is a correct simulation by protocol’s construction). In the second case: if $\text{ek}^* \neq \text{ek}_\mu$, then \mathcal{B} aborts the simulation and outputs a random bit. Otherwise, it continues as described below. \mathcal{B} forwards (m_0, m_1) to its challenger, gets back a ciphertext c^* , and sends c^* to \mathcal{A} in the challenge session.
- \mathcal{B} answers oracle queries as before.
- Finally, let b' be \mathcal{A} ’s output, then \mathcal{B} outputs the same bit.

Let Abort be the event that \mathcal{B} aborts during the experiment. If Abort occurs, then $\Pr[\mathbf{Exp}_{\mathcal{E}, \mathcal{B}}^{1\text{-IND-CCA}}(\lambda) = 1 \mid \text{Abort}] = 1/2$. Moreover, as long as Abort does not occur the distribution of the public keys simulated by \mathcal{B} is identical to the one in the real experiment, and thus the index μ is perfectly hidden. Hence, we have that $\Pr[\overline{\text{Abort}}] = 1/Q$ and

$$\text{Adv}_{\mathcal{E}, \mathcal{B}}^{1\text{-IND-CCA}}(\lambda) \geq \frac{1}{Q} \cdot \left(\Pr[\mathbf{Exp}_{\mathcal{E}, \mathcal{B}}^{1\text{-IND-CCA}}(\lambda) = 1 \mid \overline{\text{Abort}}] - \frac{1}{2} \right).$$

To complete the proof we show that in the case Abort does not occur \mathcal{B} ’s simulation of the iCCA game to \mathcal{A} is perfect. In particular, we show that as long as \mathcal{A} is not ping-pong (as it must be by definition of iCCA security) \mathcal{B} can answer correctly to all queries made by \mathcal{A} . Precisely, the tricky case that needs to be checked is that \mathcal{B} can answer with the correct decryption when the adversary sends the last message on sessions that were already opened. Let $T = (T', T_c)$ be the transcript of the queried session where ek is the corresponding public key and c is the ciphertext sent by \mathcal{A} , and let $T^* = (T^{*'}, T_c^*)$ be the transcript of the challenge session. If $T' \not\equiv T^{*'}$, it essentially means that \mathcal{B} generated ek and thus it can decrypt. If $T' \equiv T^{*'}$, since \mathcal{A} is not ping-pong, then it must be that either (I) $c \neq c^*$ (in this case \mathcal{B} forwards c to the decryption oracle), or (II) $c = c^*$ and the corresponding timestamps are not alternating, i.e., $t_n^* > t_n$, that is c was sent before \mathcal{B} asked (and sent) the challenge ciphertext. Thus \mathcal{B}

could have asked c to its decryption oracle recall that in the 1-bounded-IND-CCA game such decryption query is legal.

In conclusion, we obtain: $\mathbf{Adv}_{\mathcal{E}, \mathcal{B}}^{1\text{-IND-CCA}}(\lambda) \geq \frac{\epsilon}{Q}$. \square

As an interesting consequence of Theorem 1 we obtain an equivalence between the notions of IND-CPA and iCCA security:

Corollary 1. *2-round-iCCA encryption exists if and only if (non-interactive) IND-CPA PKE exists.*

Proof. The first direction (IND-CPA \Rightarrow 2-round-iCCA) follows by observing that: (i) 1-bounded-IND-CCA is implied by IND-CPA security (see [12] and Appendix C); (ii) iCMA security can be realized using digital signatures, and thus from one-way functions (see our Corollary 3). The second direction follows from the following simple Lemma:

Lemma 1. *2-round-iCCA \Rightarrow IND-CPA*

Proof. Let $\Pi = (\text{Setup}, \text{S}, \text{R})$ be a 2-round iCCA protocol (recall that wlog we assume that R speaks first). We construct a non-interactive encryption scheme $\mathcal{E} = (\text{KG}, \text{Enc}, \text{Dec})$ from Π as follows:

$\text{KG}(1^\lambda)$: run $(\text{sendk}, \text{recvk}) \xleftarrow{\$} \text{Setup}(1^\lambda)$, and $(R; \rho) \xleftarrow{\$} \text{R}(\text{recvk})$ where we denote with R the message sent by R, and by ρ its private coins. Output $\text{ek} = (\text{sendk}, R)$ and $\text{dk} = (\text{recvk}, \rho)$.

$\text{Enc}(\text{ek}, m)$: run the sender algorithm $\text{S}(\text{sendk}, m)$ with input message R . Let C be the message generated by S. Output C as the ciphertext.

$\text{Dec}(\text{dk}, C)$: run $m \leftarrow \text{R}(\text{DK})$ with input message C and private random coins ρ , and output m .

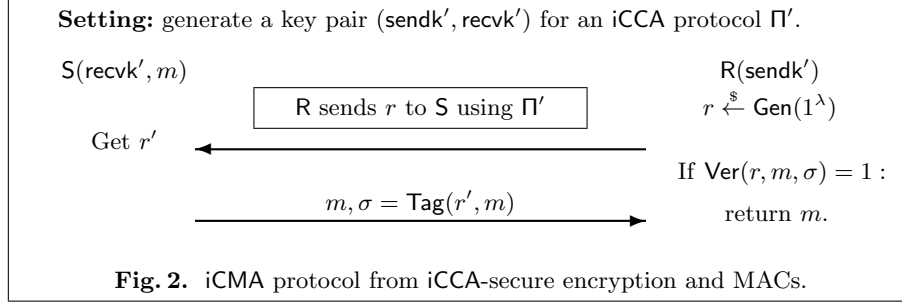
PROOF OF SECURITY. Assume there exists an efficient \mathcal{A} that breaks the IND-CPA security of \mathcal{E} with non-negligible advantage ϵ , i.e., $\mathbf{Adv}_{\mathcal{A}, \mathcal{E}}^{\text{IND-CPA}}(\lambda) \geq \epsilon$. We build an adversary \mathcal{B} that uses \mathcal{A} to obtain non-negligible advantage in breaking the iCCA security of protocol Π . \mathcal{B} is run on input the public sender key sendk and has oracle access to R. First, \mathcal{B} queries $\text{R}(\text{recvk})$ to start a new session, and it obtains R . \mathcal{B} sets $\text{ek} = (\text{sendk}, R)$ and runs $(m_0, m_1) \xleftarrow{\$} \mathcal{A}(\text{ek})$. \mathcal{B} outputs the same pair of messages, and then runs the challenge session with $\text{S}(\text{sendk}, m_b)$, to which it sends R and receives C^* back. \mathcal{B} finally runs $b' \xleftarrow{\$} \mathcal{A}(C^*)$ and outputs the same bit b' . It is easy to see that the simulation of the IND-CPA experiment is perfect, and thus $\mathbf{Adv}_{\mathcal{B}, \Pi}^{\text{iCCA}}(\lambda) = \mathbf{Adv}_{\mathcal{A}, \mathcal{E}}^{\text{IND-CPA}}(\lambda) \geq \epsilon$. \square

For the case of non-interactive public key encryption, it is an open problem to understand whether the notion of IND-CPA security implies the one of IND-CCA security (in a fully black-box sense), and there have been provided some evidences that it may not be the case [25]. In contrast, our result shows that by adding a *single* round of communication the basic notion of IND-CPA-security is sufficient to realize 2-round iCCA-secure encryption.

It is worth noting that a 3-round encryption protocol based on IND-CPA encryption was earlier proposed by Dolev, Dwork and Naor [19]. While the protocol in [19] can be proven in our formalization of interactive encryption (see Appendix B), our contribution is a protocol which is optimal (2 rounds) and is based on the same weaker assumption (IND-CPA security).

3.2 iCMA-secure PKMA from iCCA security

We show how to realize iCMA-secure message transmission from any message transmission protocol that is iCCA-secure, and any strongly unforgeable MAC. The protocol is based on an old idea of realizing public key authentication via CCA-secure encryption and message authentication codes. The basic idea is that the authenticator shows its ability to decrypt: the verifier encrypts a MAC key r for the authenticator, who decrypts and sends back the MAC of message m using the key r . This protocol (briefly sketched in Figure 2) implicitly appeared first in [30] and was proven secure in [1]. Here we generalize



this construction in the framework of our definitions, i.e., by using (possibly interactive) iCCA-secure encryption in place of (non-interactive) IND-CCA-secure encryption. Furthermore, in Appendix E we generalize a 3-round protocol earlier proposed by Dolev, Dwork and Naor [19] that is based only on iCCA security.

Let $\text{MAC} = (\text{Gen}, \text{Tag}, \text{Ver})$ be a strongly unforgeable MAC with message space \mathcal{M} and key space \mathcal{K} , and let $\Pi' = (\text{Setup}', S', R')$ be an encryption protocol with message space \mathcal{K} . We build a PKMA protocol $\Pi = (\text{Setup}, S, R)$ as follows.

Setup (1^λ) : run $(\text{sendk}', \text{recvk}') \xleftarrow{\$} \text{Setup}'(1^\lambda)$ and output $\text{recvk} = \text{sendk}'$ and $\text{sendk} = \text{recvk}'$. The message space of the protocol Π is the message space \mathcal{M} of the MAC.

S (sendk, m) : run the encryption protocol Π' with R with reversed roles, i.e., S runs the receiver algorithm R' of Π' while R will run the sender algorithm S' . Let r be the output of R' at the end of the run of protocol Π' . Then S sends $(m, \sigma = \text{Tag}(r, m))$ to R as the last message.

R (recvk) : generate a fresh MAC key $r \xleftarrow{\$} \text{Gen}(1^\lambda)$ and send r to S using the encryption protocol Π' , i.e., R runs $S'(\text{sendk}', r)$. Once the encryption protocol is over, R waits for a message (m, σ') from S. If $\text{Ver}(r, m, \sigma') = 1$ then R outputs m . Otherwise, it outputs \perp .

We can now state the following theorem.

Theorem 2. *If Π' is iCCA-secure and MAC is strongly-unforgeable, then Π is an iCMA-secure message transmission protocol.*

Proof. To prove the theorem we define the following hybrid games and we denote with G_i the event that the outcome of Game i , run with \mathcal{A} , is 1.

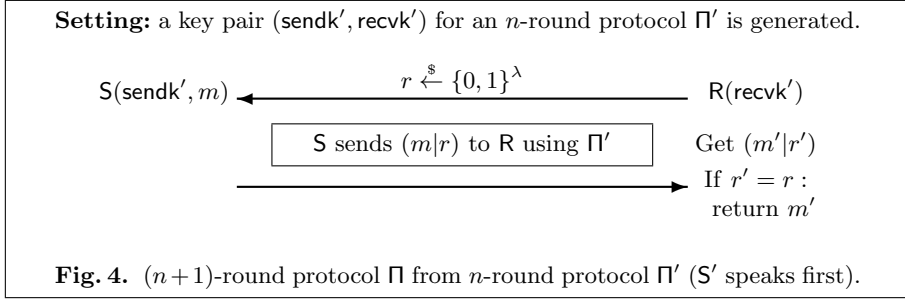
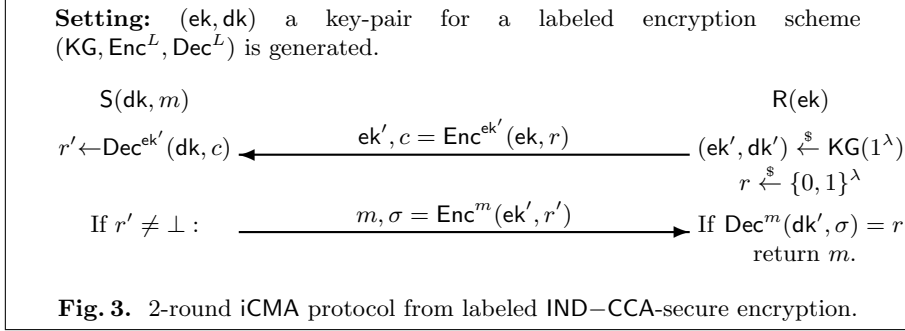
Game 0: this is the real iCMA game.

Game 1: this is the same as Game 0 except that in the challenge session the receiver generates a random key $r^* \xleftarrow{\$} \text{Gen}(1^\lambda)$ but runs the encryption protocol Π' with the message 0, i.e., R runs $S'(\text{sendk}', 0)$. We point out that the receiver keeps using r^* as the random string associated with the run of the encryption protocol in the challenge session. This means that:

1. the receiver uses r^* to check whether $\text{Ver}(r^*, m^*, \sigma^*) = 1$, where (m^*, σ^*) is the forgery returned by the adversary (i.e., the last message of \mathcal{A} in the challenge session);
2. the receiver uses r^* to compute $\sigma = \text{Tag}(r^*, m)$ in all those queries to $S(\text{sendk}, m)$ where the plaintext is m and all previous messages of the session are exactly the same as those of the challenge session (basically, the adversary replayed all messages of the encryption protocol but asked for a different plaintext $m \neq m^*$).

Via a straightforward reduction to the iCCA-security of the encryption protocol, it is possible to show that there exists an adversary \mathcal{B} such that:

$$\Pr[G_0] - \Pr[G_1] \leq 2 \cdot \text{Adv}_{\Pi', \mathcal{B}}^{\text{iCCA}}(\lambda).$$



Finally, if we consider Game 1 it is not hard to see that by the strong unforgeability of the MAC we have that $\Pr[G_1] \leq \mathbf{Adv}_{\mathcal{A}, \text{MAC}}^{\text{suf-cmva}}(\lambda)$. \square

Remark 1. While the protocol uses a fresh MAC key for every session, we stress that a one-time MAC (e.g., a pairwise independent hash function) is not sufficient to prove iCMA security. Intuitively, the reason is that the adversary may fully-replay the first portion of the protocol (i.e., the one related to Π') from the challenge session to many copies of the sender, each initialized with a different plaintext $m \neq m^*$, thus obtaining several MACs under the same key.

If we instantiate the above construction with a 1-round (aka non-interactive) iCCA-secure encryption scheme, and one of the constructions of MACs from IND-CCA encryption proposed in [17] (that have the advantage of having a ‘compact’ secret key), we then obtain an elegant and efficient 2-round PKMA protocol based only on IND-CCA security. Moreover, by directly observing the MAC of [17] and the resulting protocol, we managed to further optimize this protocol: we notice that the ephemeral secret key dk' (which is part of the MAC key with r) is only used for verification, and there is no need to encrypt it inside c ; instead, we use labels to bind ek' with c . The resulting optimized protocol is presented in Figure 3.

By instantiating the result of Theorem 2 (and our optimization above) with known constructions of CCA-secure encryption from Factoring [28], DDH [13], or CDH [42], we obtain the following corollary.

Corollary 2. *If the Factoring (resp. DDH, CDH) assumption holds, there exists an efficient 2-round PKMA.*

3.3 Secure Round Extension of Message Transmission Protocols

The basic idea for constructing an $(n+1)$ message transmission protocol Π from an n -round Π' is sketched in Figure 4, and consists in letting the party who speaks first send a random nonce r , and then running the n -round protocol with plaintext $m|r$. Finally, the receiver will terminate correctly only if the n -round protocol returns a plaintext $m|r$, where r is the *same* nonce sent in the first message. A formal description of our round-extension construction follows.

If $\Pi' = (\text{Setup}', S', R')$ is an n -round message transmission protocol, then we construct an $(n+1)$ -round protocol $\Pi = (\text{Setup}, S, R)$ as follows. The key generation algorithm is the same, i.e., $\text{Setup} =$

Setup' and $\text{sendk} = \text{sendk}'$, $\text{recvk} = \text{recvk}'$. To build algorithms S and R we distinguish two cases according to which party speaks first in Π' : S' or R' .

1. **R speaks first.** In the new protocol Π' , it will be S' who will speak first.
 $S'(\text{sendk}', m)$: first, choose $r \xleftarrow{\$} \{0, 1\}^\lambda$ and send r to R' . To react to the later messages, run $S(\text{sendk}, m|r)$ (i.e., on the message obtained by concatenating m and r).
 $R'(\text{recvk}')$: wait for the message r from S' . Next, store r and run $R(\text{recvk})$. At the end of a session, let $m = (m'|r')$ be the the message returned by R. If $r' = r$, then return m' , otherwise, output \perp .
2. **S speaks first.** In the new protocol Π' , it will be R' who will speak first.
 $S'(\text{sendk}', m)$: wait for the message r from R' and then run $S(\text{sendk}, m|r)$.
 $R'(\text{recvk}')$: first, choose a random $r \xleftarrow{\$} \{0, 1\}^\lambda$ and send r to S' . Next, run $R(\text{recvk})$, and let $m = (m'|r')$ be the the message returned by R at the end of its execution. If $r' = r$, then return m' , otherwise output \perp .

In the following theorem we prove that the above construction preserves iCCA and iCMA security.

Theorem 3. *For any $n \geq 1$, if Π' is an iCCA (resp. iCMA) secure n -round protocol, there exists an $(n + 1)$ -round protocol Π that is iCCA (resp. iCMA) secure.*

The proof of the Theorem appears in Appendix F.2 and F.3. In what follows we show the following interesting corollary, where (a) is obtained by applying our observation that 1-round iCMA-secure PKMA is equivalent to strongly unforgeable signatures, and (b) follows from our Theorem 1.

Corollary 3. *(a) For any $n \geq 1$, one-way functions are sufficient to build an n -round PKMA protocol that is iCMA-secure. (b) For any $n \geq 2$, IND-CPA-secure PKE is sufficient to build an n -round encryption protocol that is iCCA-secure.*

4 Advanced Security Properties from the Power of Interaction

We discuss three advanced security properties of message transmission protocols, each *requiring interaction*: deniability, forward security, and replay security.

While the first two properties have been already considered in previous work in the context of encryption, key exchange, and message authentication, the last one, replay security, is new and aims at obtaining more intuitive security definitions of message transmission. Interestingly, we show that replay security can be achieved only by interactive protocols, and that with enough interaction our notions of iCCA and iCMA security already provide replay-secure protocols.

Below we start with an informal discussion of these advanced security notions. Then, we proceed by defining and discussing these properties in more detail: deniability (Section 4.1), forward-security (Section 4.2), and replay-security (Section 4.3).

DENIABILITY. We already mentioned that interactive PKE/PKMA might achieve advanced security properties which are impossible in the non-interactive setting. One such (well studied [19,21,20,15]) notion is that of deniable authentication, which was actually the original motivation of the DDN paper. Since this is not the main topic of this work, we only define the weakest notion of *passive deniability* (and its extension called ‘passive forward deniability’ [15]), and observe that our optimized 2-round variant from non-interactive CCA-secure PKE is passively forward deniable.¹²

FORWARD SECURITY. Another example is the notion of *forward security*, which (intuitively) states that ‘old’ message transmissions should remain private even if the party’s long term secret key is later compromised. Prior to our work, forward security has been extensively studied in the KE literature

¹² The construction can be made ‘actively deniable’, with more rounds, using the techniques developed by [21].

(in fact, in many cases being a mandatory part of a ‘secure’ KE). On the other hand, forward security is (obviously) impossible for non-interactive PKE without “changing the model”; e.g. by introducing global time periods, and periodically refreshing the secret key [9]. In contrast, no such impossibility exists for interactive PKE, and, indeed, our interactive PKE schemes are forward-secure, since all of them use ephemeral keys to actually encrypt the message.

REPLAY-SECURITY. Yet another limitation of non-interactive PKE/PKMA schemes is that they necessarily suffer from what we informally (for now) term “replay” attacks. In the case of encryption, for example, an attacker can always record an ‘old’ ciphertext, and then manage to decrypt it much later. Similarly, a verifier can always pass an ‘old’ signature to another verifier in the future. Motivated by this impossibility, we formalize (to the best of our knowledge, for the first time) the notion of *replay-secure* (necessarily) interactive PKMA/PKE schemes. For the former, a honest verifier is assured that the “current” message is actually being authenticated by the secret key owner “now”, as opposed to some time in the past. For the latter, a honest encryptor is similarly assured that the “current” message can only be decrypted by the secret key owner “now”, as opposed to some time in the future. We then show that *any* interactive PKE/PKMA scheme which has at least 2 rounds is already replay-secure.¹³ For example, we automatically get replay-secure PKE/PKMA schemes, by using the 2-round solutions in this paper.¹⁴ We also notice that a very special case of our replay-secure PKMA, when the message space has cardinality 1, essentially corresponds to the strongest security notion for identification schemes, called impersonation security under concurrent attacks [5]. Here the attacker has concurrent oracle access to the prover (i.e., ‘signer of a fixed message’), then loses this oracle access, and, finally, has to convince an honest verifier. In fact, our 2-round PKMA protocols, when specialized to this trivial case, essentially “collapse” to well-known challenge-response identification protocols from CCA-encryption and signature schemes, respectively. Of course, by having an extra “non-trivial” message, we think that replay-secure PKMA schemes should have more applications than concurrently secure identification schemes.

4.1 Deniability

Informally speaking, a PKMA protocol is deniable if the authenticator S can authenticate a message m to the receiver R in such a way that R *cannot* use the transcript of their conversation as evidence to later convince third parties about the fact that S took part in the protocol and authenticated m .

The area of deniable authentication has attracted a lot of attention [21,20,29,15] and has several variants depending on the exact attack scenario. While a detailed exploration of this area is beyond the scope of our work, to illustrate the potential of interactivity we show that our PKMA protocols based on iCCA encryption already satisfies the weakest form of (perfect) passive deniability, which we define below.

Definition 7 (Passive Deniability). *A PKMA protocol $\Pi = (\text{Setup}, S, R)$ is passive deniable if there exists a PPT simulator Sim such that for all honestly generated keys $(\text{recvk}, \text{sendk}) \xleftarrow{\$} \text{Setup}(1^\lambda)$, and for any message $m \in \mathcal{M}$, $\text{Sim}(\text{recvk}, m)$ generates a transcript π_S that is (computationally, statistically, or perfectly) indistinguishable from a transcript π of a real execution $\langle S(\text{sendk}, m), R(\text{recvk}) \rangle$ of the protocol.*

FORWARD DENIABILITY. As noted by Di Raimondo and Gennaro [15], the standard definition of deniability provides guarantees only to the sender. However, if the sender later changes its mind, it might be able to exhibit some witness (e.g., its private key) that proves its participation in the protocol. To rule out even this possibility, Di Raimondo and Gennaro introduced the notion of *forward deniability*, and showed that every protocol which is deniable in a statistical or perfect sense is also forward deniable.

¹³ For encryption, we also define a stronger notion of replay-security, which requires at least 3 rounds, and is realized by any 3-round CCA secure scheme.

¹⁴ This includes already mentioned 2-round PKMA from CCA-encryption and 2-round PKE from signatures, as well as simple 2-round PKE/PKMA obtained by “extending” 1-round PKE/PKMA schemes.

Remark 2 (Strong Deniability). It is possible to also consider a stronger form of deniability [21], in which the receiver might be dishonest. The resulting definition is essentially very similar to the one for dishonest verifier zero-knowledge. We refer to [21] for more details.

Building Forward Deniable PKMA. Let Π_{mac} and Π_{lab} be our PKMA protocols in Section 3.2 and Appendix E based on iCCA encryption and MACs and on labeled iCCA encryption, respectively, when instantiated with non-interactive, i.e., 1-round, (labeled) iCCA protocols. Then we can state the following theorem:

Theorem 4. *The protocols Π_{mac} and Π_{lab} are passive forward deniable.*

For protocol Π_{lab} , the proof simply follows by observing that the following PPT simulator Sim satisfies Definition 7. $\text{Sim}(\text{recvk}, m)$ chooses a random $r \xleftarrow{\$} \tilde{\mathcal{M}}$ (where $\tilde{\mathcal{M}}$ is the message space of the labeled iCCA encryption protocol), and outputs $\pi_S = \langle m, \mathcal{S}_m(\text{recvk}, r), r \rangle$. It is easy to observe that π_S is exactly distributed as the transcript of a real execution of the protocol. For the protocol Π_{mac} the proof is essentially the same except that the simulated transcript is $\pi_S = \langle \mathcal{S}(\text{recvk}, r), \text{Tag}_r(m) \rangle$. Finally, note that since the transcripts are perfectly indistinguishable the protocols are also forward deniable.

As already noticed in previous work [21,16], both protocols Π_{mac} and Π_{lab} can be modified by adding a challenge-response subprotocol in order to make them strong deniable for sequential executions.

4.2 Forward Security

Intuitively, forward security guarantees that any leak of secret information at some time t should not affect the security of protocol runs that occurred in the past, i.e., at any time $t' < t$. Forward security is a desirable property that is not known to be achieved by standard non-interactive public key encryption. Here we formalize suitable definitions of forward security for interactive encryption protocols, and we show that our constructions satisfy this property.

We consider both a weak and a strong version of forward security. To define *weak forward security*, we introduce an oracle Corrupt , which outputs the secret key recvk of the receiver R. Then we define the experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{wFS}}(\lambda)$ to be the same as $\text{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda)$ except that \mathcal{A} is additionally given access to the oracle Corrupt that can be queried only *after* the challenge session is completed.

Definition 8 (Weak Forward Security). *We define the advantage of an adversary \mathcal{A} in breaking weak forward security (wFS) of protocol Π as $\text{Adv}_{\Pi, \mathcal{A}}^{\text{wFS}}(\lambda) = \Pr[\text{Exp}_{\Pi, \mathcal{A}}^{\text{wFS}}(\lambda) = 1] - \frac{1}{2}$, and we say that Π is weak forward secure (wFS) if for any PPT \mathcal{A} , $\text{Adv}_{\text{iPKE}, \mathcal{A}}^{\text{wFS}}(\lambda)$ is negligible.*

To define *strong forward security* we introduce another oracle, StateCorrupt , which outputs the receiver's secret key recvk as well as R's private state, consisting of random coins and private information generated during the run of the protocol. In order to make the game non-trivial for the adversary, we restrict the revealed state only to sessions that are at the time of the query not completed, and whose transcript T does not *partially match* with the transcript T^* of the challenge session. We say that a transcript T *partially matches* with T^* if T matches T^* in the old sense up to the first n' messages, where n' is the length of T , i.e., if T^* has length n and T has length $n' \leq n$, we consider the first portion \tilde{T}^* of length n' of the transcript T^* , and we apply the previous definition of match, i.e., $T \equiv \tilde{T}^*$. This notion of partial match is introduced to formalize the fact that the adversary should obtain not even a small portion of the private state concerning the challenge session. Although this may appear a restriction, we make the following two observations. First, it follows our intuition for the security of interactive encryption in which sensitive transmissions should be particularly secured so as not to leak private state information. Second, it is reasonable to think that once the decryptor loses its control and reveals *all* its private state to the adversary (even the state of uncompleted sessions), then there

should be no specific security guarantees for every session which is significantly related with the revealed private state.

Formally, we define the experiment $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{sFS}}(\lambda)$ to be the same as $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}(\lambda)$ except that \mathcal{A} is additionally given access to the oracle `StateCorrupt`, that can be queried only *after* the challenge session is completed.

Definition 9 (Strong Forward Security). *We define the advantage of \mathcal{A} in breaking the strong forward security (sFS) of Π as $\mathbf{Adv}_{\Pi, \mathcal{A}}^{\text{sFS}}(\lambda) = \Pr[\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{sFS}}(\lambda) = 1] - \frac{1}{2}$, and we say that Π is strong forward secure (sFS) if for any PPT \mathcal{A} , $\mathbf{Adv}_{\text{iPKE}, \mathcal{A}}^{\text{sFS}}(\lambda)$ is negligible.*

Building Forward Secure (Interactive) PKE. Below we show that our construction of interactive encryption from IND–CPA-secure encryption proposed in Section 3.1 achieves strong forward security.

Intuitively, this follows from observing that the long-term private key of R is the signing key sk whereas its private state consists of one-time decryption keys dk . In more detail, we show that our proof of iCCA security (Theorem 1) can be adapted to the case of strong forward security as follows. First, we can define the same event `Forge`, and observe that the proof of Claim 1 to bound $\Pr[\text{Forge}]$ remains the same. Indeed, in this proof `StateCorrupt` queries do not have to be simulated as the entire simulation will stop *before* the challenge session is completed. Second, the remaining part of the proof of Theorem 1 can be easily changed as follows to enable the simulator \mathcal{B} answer `StateCorrupt` queries. \mathcal{B} knows the secret key $\text{recvk} = \text{sendk}'$ (it is the secret key generated by \mathcal{B} itself), which can thus be returned in output. Then, by definition of strong forward security the only sessions whose state is to be revealed are those sessions that do not partially match with the challenge session: essentially those sessions for which \mathcal{B} already knows the one-time decryption key dk .

Remark 3. While we showed that our construction of Section 3.1 satisfies strong forward security, we observe that a construction obtained by applying our round-extension Theorem 3 instead cannot be proven strong forward secure because of the following attack (here we consider the case when S speaks first).

- The adversary starts the challenge session with to get the first message r^* from the challenger.
- Next, it chooses some $r \neq r^*$ and queries R by sending r to start a new session.
- It receives back ek from R (at the end of Π'), and forwards ek to the challenger as the second message.
- Let c^* be the message received by the challenger. Now, \mathcal{A} queries `StateCorrupt` which will return the secret key dk corresponding to ek . Notice that such a query is legal as this session does not partially match with the challenge one as $r \neq r^*$.
- Finally, \mathcal{A} uses dk to decrypt the ciphertext c^* .

The above issue stems from the fact that our round extension transformation does not preserve forward secrecy. However, nothing is lost as the issue can be fixed, thus achieving strong forward security even for iCCA protocols obtained via round-extension. To do this, we can modify our generic transformation as follows: when R receives the random r from S , R signs *every* protocol message together with r , and includes such signatures along with the messages. Clearly, the above attack no longer applies as \mathcal{A} will not be able to create valid protocol messages for an $r \neq r^*$. A formal proof easily follows the intuition above, and is omitted.

4.3 Replay Security

In what follows we formalize the notion of replay security, and then we show that traditional non-interactive protocols (e.g., CCA encryption and signatures) *cannot* be replay-secure. Intuitively, the reason is that a ciphertext or a digital signature can always be “replayed” after its transmission (this also explains our choice for the name of this notion).

TIME INTERVALS AND CONCURRENT SESSIONS. To formalize our definitions and argue more easily about concurrent sessions, we refine our notion of time and we introduce an intuitive terminology. If t and t' are time instants such that $t < t'$, then we denote with $[t, t']$ the *time interval* between t and t' , i.e., the sequence $\langle t, t + 1, t + 2, \dots, t' \rangle$. For every protocol transcript $T = \langle (M_1, t_1), \dots, (M_n, t_n) \rangle$ (i.e., for every session) of an n -round protocol there exists a corresponding time interval $[t_1, t_n]$ in which the session starts and ends. Let $[t_1^*, t_n^*]$ and $[t_1, t_n]$ be the time intervals of two protocol sessions. We say that $[t_1^*, t_n^*]$ and $[t_1, t_n]$ *overlap* if $[t_1, t_n] \cap [t_1^*, t_n^*] \neq \emptyset$. Moreover, we say that \mathcal{A} is an *overlapping* adversary if it generates an oracle session $[t_1, t_n]$ that overlaps with the challenge session $[t_1^*, t_n^*]$.

In the following lemma we show that in any protocol with at least two rounds, any ping-pong adversary is also overlapping. We use this general statement to prove that any 2-round secure message-transmission protocol is also replay-secure.

Lemma 2. *Let $n \geq 2$ and Π be an n -round message transmission protocol. If \mathcal{A} is a ping-pong adversary against Π , then \mathcal{A} is overlapping.*

Proof. Assume by contradiction that \mathcal{A} is not overlapping, then we show that \mathcal{A} is not ping-pong. Let T^* and $[t_1^*, t_n^*]$ be the transcript and time interval of the challenge session. Since \mathcal{A} is not overlapping, no oracle session T overlaps with T^* , i.e., for every oracle session with time interval $[t_1, t_n]$ it holds $[t_1^*, t_n^*] \cap [t_1, t_n] = \emptyset$. However, since $n \geq 2$, it is easy to see that if $[t_1^*, t_n^*] \cap [t_1, t_n] = \emptyset$ then the timestamps of these two sessions are not alternating, and thus $T \neq T^*$. \square

Replay-Secure Public-Key Message Authentication. Informally speaking, a PKMA protocol is replay-secure if all oracle sessions initialized with the challenge plaintext m^* do not overlap with the challenge session. This essentially means that the adversary loses access to the legitimate signer (on message m^*) *before* starting to forge m^* .

For a more formal definition, consider the experiment $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCMA}}$ (defined in Section 2.1) and let $[t_1^*, t_n^*]$ be the time interval of the challenge session, and $\{[t_1^i, t_n^i]\}_{i=1, \dots, Q}$ be the time intervals of all the oracle sessions established by \mathcal{A} . Moreover, let m^i be the plaintext used to initialize the sender oracle $S(\text{sendk}, m^i)$ in the i -th oracle session. Then we call \mathcal{A} a *replay adversary* if there exists $i \in \{1, \dots, Q\}$ such that $m^i = m^*$ and $[t_1^i, t_n^i] \cap [t_1^*, t_n^*] \neq \emptyset$. Replay security for PKMA is defined as follows:

Definition 10 (Replay Secure PKMA). *Let $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{RSMA}}$ be the same experiment as $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCMA}}$, except that \mathcal{A} is required not to be a replay adversary (instead of not being ping-pong). Then we say that an interactive protocol Π is a replay secure PKMA (RSMA for short) if for any PPT \mathcal{A} , its advantage $\mathbf{Adv}_{\Pi, \mathcal{A}}^{\text{RSMA}}(\lambda) = \Pr[\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{RSMA}}(\lambda) = 1]$ is negligible.*

First, in the following theorem we show that 1-round PKMA protocols cannot be replay-secure. Its proof is rather simple and follows from the fact that in 1-round protocols there is clearly no overlap between different sessions. Hence, the dummy adversary who replays a signature received from the legitimate signer is a valid adversary in the RSMA experiment.

Theorem 5. *Any 1-round iCMA-secure PKMA protocol is not RSMA-secure.*

It is worth noting that one can obtain 1-round replay-secure solutions in different models, e.g., by introducing global time periods and requiring the signer to always sign the message with the current timestamp. However, such a solution falls outside the pure non-interactive model considered by our work.

Therefore, while 1-round replay-secure PKMA cannot be achieved, in the following theorem we show that with at least *two* rounds of interaction any iCMA-secure protocol is a replay-secure PKMA.

Theorem 6. *For $n \geq 2$, any n -round iCMA-secure protocol Π is RSMA-secure.*

Proof. The proof follows by observing that if \mathcal{A} is not a replay adversary in $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{RSMA}}$ (for Π with at least 2 rounds), then \mathcal{A} is also not ping-pong in $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCMA}}$. Namely, we show that for a non-replay \mathcal{A} we have that for all $i = 1, \dots, Q$, $T_i \not\equiv T^*$. For every i , there are only two possible cases: $m^i = m^*$ or $m^i \neq m^*$. If $m^i = m^*$, then $T_i \not\equiv T^*$ follows by Lemma 2. In the case of sessions i where $m^i \neq m^*$, assume by contradiction that $T_i \equiv T^*$. Then by definition of match the sessions must share the same protocol messages in the same order (i.e., $M_j^i = M_j^*$) which, by correctness, implies that $m^i = m^*$, a contradiction. \square

From Theorem 6 we obtain two interesting results that we summarize in the following corollary:

Corollary 4. (1) *There exists a simple 2-round replay-secure PKMA protocol based on any strongly unforgeable signature scheme (and thus on one-way functions); (2) there exists a simple 2-round replay-secure PKMA protocol based on any IND-CCA-secure PKE.*

The construction (1) follows by combining Theorem 6 with our round-extension result (Theorem 3) applied to any strongly unforgeable signature (aka 1-round iCMA-secure PKMA). The construction (2) from IND-CCA-secure PKE is instead obtained by applying Theorem 6 to our 2-round construction of Theorem 2 (instantiated with a non-interactive IND-CCA-secure PKE scheme).

Remark 4 (Relation to Concurrent-Secure Identification Schemes). Notice that in the special case when the message space has cardinality 1, our notion of replay-secure PKMA essentially corresponds to the strongest security notion for identification schemes, called impersonation security under concurrent attacks [5]. In this case (i.e., message space of cardinality 1), a PKMA can be indeed seen as an identification scheme. Moreover, by considering authentication with an empty message space, the 2-round PKMA protocols mentioned in Corollary 4 recover well known 2-round, signature-based and encryption-based, identification schemes (see, e.g., [2], and notice that outputting the secret key is indeed a secure MAC for an empty message space).

Replay-Secure Public-Key Encryption. Informally speaking, a PKE protocol is *replay-secure* if there is no overlap between the challenge session and all oracle sessions in which the plaintext revealed by the receiver is one of the two challenge plaintexts. In essence, this means that *during* the challenge session the adversary loses access to the legitimate decryptor, but only on the challenge plaintexts m_0 or m_1 . More formally, consider the experiment $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}$ (defined in Section 2.1) and let $[t_1^*, t_n^*]$ be the time interval of the challenge session, and $\{[t_1^i, t_n^i]\}_{i=1, \dots, Q}$ be the time intervals of all the oracle sessions established by \mathcal{A} . Moreover, let m^i be the plaintext revealed by the receiver in the i -th oracle session (wlog we only consider completed sessions). Then, we call \mathcal{A} a *replay adversary* if there exists $i \in \{1, \dots, Q\}$ such that $m^i = m_0$ or $m^i = m_1$, and $[t_1^i, t_n^i]$ overlaps with $[t_1^*, t_n^*]$. Replay security for PKE is defined as follows:

Definition 11 (Replay-Secure Encryption). *Let $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{RSE}}$ be the same as experiment $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}$, except that \mathcal{A} is required not to be a replay adversary (instead of denying it to be ping-pong). Then we say that an interactive protocol Π is a replay-secure encryption (RSE) if for any PPT \mathcal{A} , its advantage $\mathbf{Adv}_{\Pi, \mathcal{A}}^{\text{RSE}}(\lambda) = \Pr[\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{RSE}}(\lambda) = 1] - \frac{1}{2}$ is negligible.*

It is worth mentioning that the notion of replay-secure PKE is similar to the notion of Replayable CCA-secure encryption (RCCA) introduced by Canetti, Krawczyk and Nielsen [10]. In RCCA security the adversary is allowed to submit any ciphertext c to the decryption oracle, except that if c decrypts to m_0 or m_1 the adversary gets a special string `test` as response.

Similarly to replay-secure PKMA, in the following theorems we show that replay-secure PKE requires at least 2-rounds of interaction to be achieved.

Theorem 7. *Any 1-round iCCA-secure PKE protocol is not RSE-secure.*

The proof is obtained by considering the adversary who replays the challenge ciphertext to the receiver.

Theorem 8. *For $n \geq 2$, any n -round iCCA-secure protocol Π is RSE-secure.*

Proof. The proof follows by observing that if \mathcal{A} is not a replay adversary in $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{RSE}}$ where Π has at least 2 rounds, then \mathcal{A} is also not ping-pong in $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}$. Namely, we show that for a non-replay adversary \mathcal{A} we have that for all $i = 1, \dots, Q$, it holds $T_i \not\equiv T^*$. For every i , there are only two possible cases: (1) $m^i = m_0$ or $m^i = m_1$, and (2) $m^i \neq m_0, m_1$. In the first case, note that $T_i \not\equiv T^*$ follows by Lemma 2, i.e., if the \mathcal{A} is not replay, then \mathcal{A} is not ping-pong w.r.t. these sessions. In the case of sessions i where $m^i \neq m_0, m_1$, assume by contradiction that $T_i \equiv T^*$. Then by definition of match, these sessions must share the same protocol messages in the very same order (i.e., $M_j^i = M_j^*$). However, by correctness, this implies $m^i = m_b$ where b is the secret bit chosen in the experiment, which is a contradiction as we assumed $m^i \neq m_0, m_1$. \square

From Theorem 8 we obtain a collection of nice results summarized in the following Corollary:

Corollary 5. *(1) IND-CPA-secure PKE is sufficient to build RSE-secure PKE; (2) there exists an efficient 2-round RSE-secure PKE protocol based on 1-bounded-IND-CCA-secure PKE and signature schemes (see Theorem 1); (3) there exists an efficient 2-round RSE-secure PKE protocol based on any IND-CCA-secure PKE (via our round-extension transformation of Theorem 3).*

Strong Replay-Secure Encryption. We notice that our definition of replay security (for both PKMA and PKE) does not allow the adversary to suspend “critical” sessions (e.g., sessions authenticating m^* , or sessions which decrypt to m_0 or m_1) during the challenge session, and to later resume these sessions once the challenge session is over. While allowing such suspended sessions would not make sense for PKMA (indeed observe that the outcome of the security experiment is determined upon the end of the challenge session), it might be a reasonable strengthening for replay-secure PKE. Here we define strong-replay-secure PKE, and we show that two rounds of interaction are insufficient to achieve strong replay-security with suspended sessions (see Theorem 9 below), but *three* rounds are enough and indeed any (at least) 3-round iCCA-secure PKE is strong replay-secure (cf. Theorem 10).

Towards defining strong replay security more formally, let us say that $[t_1, t_n]$ is *off* during $[t_1^*, t_n^*]$ if for all $j = 1, \dots, n$ we have that $t_j \notin [t_1^*, t_n^*]$. We call \mathcal{A} a *strong replay adversary* if there exists $i \in \{1, \dots, Q\}$ such that $m^i = m_0$ or $m^i = m_1$, and $[t_1^i, t_n^i]$ is *not off* during $[t_1^*, t_n^*]$.

Definition 12 (Strong Replay-Secure Encryption). *Let $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{sRSE}}$ be the same experiment as $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{iCCA}}$, except that \mathcal{A} is required not to be a strong replay adversary (instead of denying it to be ping-pong). Then we say that an interactive protocol Π is a strong replay secure encryption (sRSE) if for any PPT \mathcal{A} , its advantage $\mathbf{Adv}_{\Pi, \mathcal{A}}^{\text{sRSE}}(\lambda) = \Pr[\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{sRSE}}(\lambda) = 1] - \frac{1}{2}$ is negligible.*

Theorem 9. *Any 2-round iCCA-secure PKE protocol is not sRSE-secure.*

Proof. To prove the theorem we show that there exists an adversary \mathcal{A} who is not a strong replay adversary in $\mathbf{Exp}_{\Pi, \mathcal{A}}^{\text{sRSE}}$ but is a ping-pong adversary. Then, being ping-pong, \mathcal{A} can trivially obtain a decryption of the challenge plaintext.

First, observe that in any 2-round protocol we have the first message from R to S. So, consider the following adversary \mathcal{A} that plays the role of the receiver in the challenge session:

- \mathcal{A} queries R on a new session and obtains M_1 with timestamp $t_1 = t$.
- \mathcal{A} sends M_1 to the honest sender S as the first protocol message in the challenge session (here M_1 gets timestamp $t_1^* = t + 1$). \mathcal{A} receives back M_2 from S, where M_2 gets timestamp $t_2^* = t + 2$.
- \mathcal{A} forwards M_2 to R in the session previously opened in step 1 (this message gets timestamp $t_2 = t + 3$).

While \mathcal{A} is clearly ping-pong according to Definition 4, \mathcal{A} is still not a strong replay adversary in $\text{Exp}_{\Pi, \mathcal{A}}^{\text{sRSE}}$ since $t_1, t_2 \notin [t+1, t+2]$. This is exactly a case in which the adversary suspended a session. \square

Theorem 10. *For $n \geq 3$, any n -round iCCA-secure protocol Π is sRSE-secure.*

The proof is basically the same as that of Theorem 8, except that here we use the following lemma to show that for 3-round protocols any ping-pong adversary generates an oracle session (decrypting to one of the challenge plaintexts) which is not off during the challenge session.

Lemma 3. *Let $n \geq 3$, and let T , $[t_1, t_n]$ and T^* , $[t_1^*, t_n^*]$ be the transcripts and time intervals of two sessions of an n -round protocol. If $[t_1, t_n]$ is off during $[t_1^*, t_n^*]$, then T does not match with T^* .*

Proof. Recall that $[t_1, t_n]$ is off during $[t_1^*, t_n^*]$ if for all $j = 1, \dots, n$ we have that $t_j \notin [t_1^*, t_n^*]$. This means that either one of the following cases occurs: (1) $[t_1^*, t_n^*] \cap [t_1, t_n] = \emptyset$, (2) $[t_1^*, t_n^*] \cap [t_1, t_n] \neq \emptyset$. Case (1) is identical to that of Lemma 2. In case (2), we have that the two sessions overlap, and we also know that $t_j \notin [t_1^*, t_n^*], \forall j = 1, \dots, n$, that is $t_1 < t_1^*$ and $t_n > t_n^*$. Since there are at least 3 rounds, there exists at least a distinct timestamp t_2 such that $t_1 < t_2 < t_n$ and such that t_2 satisfies either one of the following conditions: $t_1 < t_2 < t_1^*$, or $t_n > t_2 > t_n^*$. However, one can again check that in neither one of these cases the timestamps are correctly alternating. Therefore, $T \neq T^*$. \square

Acknowledgements. The authors would like to thank Adam O’Neill, Victor Shoup and Stefano Tessaro for valuable discussions on this work. The research of Yevgeniy Dodis is partially supported by gifts from VMware Labs and Google, and NSF grants 1319051, 1314568, 1065288, 1017471. The research of Dario Fiore is partially supported by the European Commission Seventh Framework Programme Marie Curie Cofund Action AMAROUT-II (grant no. 291803), and the Madrid Regional Government under project PROMETIDOS-CM (ref. S2009/TIC1465).

References

1. M. Bellare, R. Canetti, and H. Krawczyk. A modular approach to the design and analysis of authentication and key exchange protocols (extended abstract). In *30th Annual ACM Symposium on Theory of Computing*, pages 419–428. ACM Press, May 1998.
2. M. Bellare, M. Fischlin, S. Goldwasser, and S. Micali. Identification protocols secure against reset attacks. In B. Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 495–511. Springer, May 2001.
3. M. Bellare and S. Micali. How to sign given any trapdoor function (extended abstract). In *20th Annual ACM Symposium on Theory of Computing*, pages 32–42. ACM Press, May 1988.
4. M. Bellare and S. Micali. How to sign given any trapdoor function. *Journal of the ACM*, 39(1):214–233, 1992.
5. M. Bellare and A. Palacio. GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In M. Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 162–177. Springer, Aug. 2002.
6. M. Bellare and P. Rogaway. Entity authentication and key distribution. In D. R. Stinson, editor, *Advances in Cryptology – CRYPTO’93*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249. Springer, Aug. 1993.
7. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73. ACM Press, Nov. 1993.
8. J. Camenisch and V. Shoup. Practical verifiable encryption and decryption of discrete logarithms. In D. Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 126–144. Springer, Aug. 2003.
9. R. Canetti, S. Halevi, and J. Katz. A forward-secure public-key encryption scheme. In E. Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 255–271. Springer, May 2003.
10. R. Canetti, H. Krawczyk, and J. B. Nielsen. Relaxing chosen-ciphertext security. In D. Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 565–582. Springer, Aug. 2003.
11. S. G. Choi, D. Dachman-Soled, T. Malkin, and H. Wee. Black-box construction of a non-malleable encryption scheme from any semantically secure one. In R. Canetti, editor, *TCC 2008: 5th Theory of Cryptography Conference*, volume 4948 of *Lecture Notes in Computer Science*, pages 427–444. Springer, Mar. 2008.

12. R. Cramer, G. Hanaoka, D. Hofheinz, H. Imai, E. Kiltz, R. Pass, A. Shelat, and V. Vaikuntanathan. Bounded CCA2-secure encryption. In K. Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007*, volume 4833 of *Lecture Notes in Computer Science*, pages 502–518. Springer, Dec. 2007.
13. R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In H. Krawczyk, editor, *Advances in Cryptology – CRYPTO’98*, volume 1462 of *Lecture Notes in Computer Science*, pages 13–25. Springer, Aug. 1998.
14. R. Cramer and V. Shoup. Signature schemes based on the strong RSA assumption. In *ACM CCS 99: 6th Conference on Computer and Communications Security*, pages 46–51. ACM Press, Nov. 1999.
15. M. Di Raimondo and R. Gennaro. New approaches for deniable authentication. In V. Atluri, C. Meadows, and A. Juels, editors, *ACM CCS 05: 12th Conference on Computer and Communications Security*, pages 112–121. ACM Press, Nov. 2005.
16. M. Di Raimondo, R. Gennaro, and H. Krawczyk. Deniable authentication and key exchange. In A. Juels, R. N. Wright, and S. Vimercati, editors, *ACM CCS 06: 13th Conference on Computer and Communications Security*, pages 400–409. ACM Press, Oct. / Nov. 2006.
17. Y. Dodis, E. Kiltz, K. Pietrzak, and D. Wichs. Message authentication, revisited. In D. Pointcheval and T. Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 355–374. Springer, Apr. 2012.
18. D. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography (extended abstract). In *23rd Annual ACM Symposium on Theory of Computing*, pages 542–552. ACM Press, May 1991.
19. D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.
20. C. Dwork and M. Naor. Zaps and their applications. In *41st Annual Symposium on Foundations of Computer Science*, pages 283–293. IEEE Computer Society Press, Nov. 2000.
21. C. Dwork, M. Naor, and A. Sahai. Concurrent zero-knowledge. In *30th Annual ACM Symposium on Theory of Computing*, pages 409–418. ACM Press, May 1998.
22. C. Dwork and A. Sahai. Concurrent zero-knowledge: Reducing the need for timing constraints. In H. Krawczyk, editor, *Advances in Cryptology – CRYPTO’98*, volume 1462 of *Lecture Notes in Computer Science*, pages 442–457. Springer, Aug. 1998.
23. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *Advances in Cryptology – CRYPTO’86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, Aug. 1986.
24. R. Gennaro, S. Halevi, and T. Rabin. Secure hash-and-sign signatures without the random oracle. In J. Stern, editor, *Advances in Cryptology – EUROCRYPT’99*, volume 1592 of *Lecture Notes in Computer Science*, pages 123–139. Springer, May 1999.
25. Y. Gertner, T. Malkin, and S. Myers. Towards a separation of semantic and CCA security for public key encryption. In S. P. Vadhan, editor, *TCC 2007: 4th Theory of Cryptography Conference*, volume 4392 of *Lecture Notes in Computer Science*, pages 434–455. Springer, Feb. 2007.
26. S. Goldwasser, S. Micali, and R. L. Rivest. A “paradoxical” solution to the signature problem (abstract) (impromptu talk). In G. R. Blakley and D. Chaum, editors, *Advances in Cryptology – CRYPTO’84*, volume 196 of *Lecture Notes in Computer Science*, page 467. Springer, Aug. 1984.
27. L. C. Guillou and J.-J. Quisquater. A “paradoxical” indentity-based signature scheme resulting from zero-knowledge. In S. Goldwasser, editor, *Advances in Cryptology – CRYPTO’88*, volume 403 of *Lecture Notes in Computer Science*, pages 216–231. Springer, Aug. 1988.
28. D. Hofheinz and E. Kiltz. Practical chosen ciphertext secure encryption from factoring. In A. Joux, editor, *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 313–332. Springer, Apr. 2009.
29. J. Katz. Efficient and non-malleable proofs of plaintext knowledge and applications. In E. Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 211–228. Springer, May 2003.
30. H. Krawczyk. Skeme: a versatile secure key exchange mechanism for internet. In *Network and Distributed System Security, 1996., Proceedings of the Symposium on*, pages 114–127, feb 1996.
31. Y. Lindell. A simpler construction of cca2-secure public-key encryption under general assumptions. In E. Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 241–254. Springer, May 2003.
32. S. Myers and A. Shelat. Bit encryption is complete. In *50th Annual Symposium on Foundations of Computer Science*, pages 607–616. IEEE Computer Society Press, Oct. 2009.
33. M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *38th Annual Symposium on Foundations of Computer Science*, pages 458–467. IEEE Computer Society Press, Oct. 1997.
34. M. Naor, O. Reingold, and A. Rosen. Pseudo-random functions and factoring (extended abstract). In *32nd Annual ACM Symposium on Theory of Computing*, pages 11–20. ACM Press, May 2000.
35. M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *21st Annual ACM Symposium on Theory of Computing*, pages 33–43. ACM Press, May 1989.

36. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd Annual ACM Symposium on Theory of Computing*, pages 427–437. ACM Press, May 1990.
37. C. Rackoff and D. R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In J. Feigenbaum, editor, *Advances in Cryptology – CRYPTO’91*, volume 576 of *Lecture Notes in Computer Science*, pages 433–444. Springer, Aug. 1991.
38. J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd Annual ACM Symposium on Theory of Computing*, pages 387–394. ACM Press, May 1990.
39. A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th Annual Symposium on Foundations of Computer Science*, pages 543–553. IEEE Computer Society Press, Oct. 1999.
40. C.-P. Schnorr. Efficient identification and signatures for smart cards. In G. Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 239–252. Springer, Aug. 1989.
41. B. R. Waters. Efficient identity-based encryption without random oracles. In R. Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 114–127. Springer, May 2005.
42. H. Wee. Efficient chosen-ciphertext security via extractable hash proofs. In T. Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 314–332. Springer, Aug. 2010.

A Standard Cryptographic Primitives

A.1 (Non-Interactive) CCA-secure Public-Key Encryption

A public key encryption scheme \mathcal{E} is a tuple of algorithms $(\text{KG}, \text{Enc}, \text{Dec})$ defined as follows:

$\text{KG}(1^\lambda)$ on input the security parameter, the key generation returns a public key ek and a secret key dk .

$\text{Enc}(\text{ek}, m)$ on input the public key ek and a message m , it outputs a ciphertext c .

$\text{Dec}(\text{dk}, c)$ given the secret key dk and a ciphertext c , it outputs a message m or an error symbol \perp .

Consider the following experiment involving the scheme \mathcal{E} and an adversary \mathcal{A} :

Experiment $\mathbf{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{IND-CCA}}(\lambda)$

$b \xleftarrow{\$} \{0, 1\}$

$(\text{ek}, \text{dk}) \xleftarrow{\$} \text{KG}(1^\lambda)$

$(m_0, m_1) \leftarrow \mathcal{A}^{\text{Dec}(\text{dk}, \cdot)}(\text{ek})$

$c^* \xleftarrow{\$} \text{Enc}(\text{ek}, m_b)$

$b' \leftarrow \mathcal{A}^{\text{Dec}(\text{dk}, \cdot)}(c^*)$

If $b' = b$ and \mathcal{A} is “legal” output 1

Else output 0.

In the above experiment, \mathcal{A} is called “legal” if it does not query the decryption oracle $\text{Dec}(\text{dk}, \cdot)$ on the challenge ciphertext c^* (after \mathcal{A} receives c^*).

The advantage of an adversary \mathcal{A} in breaking the IND–CCA security of an encryption scheme \mathcal{E} is

$$\mathbf{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{IND-CCA}}(\lambda) = \left| \Pr[\mathbf{Exp}_{\mathcal{E}, \mathcal{A}}^{\text{IND-CCA}}(\lambda) = 1] - \frac{1}{2} \right|$$

Definition 13 (IND–CCA security). *An encryption scheme \mathcal{E} is IND–CCA-secure if for any PPT \mathcal{A} , $\mathbf{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{IND-CCA}}(\lambda)$ is negligible.*

A weaker notion of IND–CCA security that we consider in our work is *q-bounded* IND–CCA security [12]. This notion is defined as IND–CCA security except that the adversary is restricted to query the decryption oracle at most q times (where q is a pre-fixed bound).

A further weaker notion of security for public key encryption is semantic security, or indistinguishability against chosen-plaintext attacks (IND–CPA). Its definition is the same as IND–CCA security except that the adversary does not get access to any decryption oracle.

A.2 Digital Signatures

A digital signature scheme consists of a triple of algorithms $\Sigma = (\Sigma.\text{kg}, \text{Sign}, \text{Ver})$ working as follows:

$\Sigma.\text{kg}(1^\lambda)$ the key generation takes as input a security parameter λ and returns a pair of keys (sk, vk) .
 $\text{Sign}(\text{sk}, m)$ on input a signing key sk and a message m , the signing algorithm produces a signature σ .
 $\text{Ver}(\text{vk}, m, \sigma)$ given a triple vk, m, σ the verification algorithm tests if σ is a valid signature on m with respect to verification key vk .

For security we define the following experiment:

Experiment $\mathbf{Exp}_{\mathcal{A}, \Sigma}^{\text{uf-cma}}(\lambda)$
 $(\text{sk}, \text{vk}) \xleftarrow{\$} \Sigma.\text{kg}(1^\lambda)$
 $(m^*, \sigma^*) \xleftarrow{\$} \mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}(\text{vk})$
 If $\text{Ver}(\text{vk}, m^*, \sigma^*) = 1$ and (m^*, σ^*) is “new” then output 1
 Else Output 0

We say that the forgery (m^*, σ^*) is “new” if it is different from all the pairs (m_i, σ_i) obtained from the signing oracle $\text{Sign}(\text{sk}, \cdot)$. We define the advantage of an adversary \mathcal{A} in breaking the strong unforgeability against chosen-message attacks (suf-cma) of Σ as $\mathbf{Adv}_{\mathcal{A}, \Sigma}^{\text{suf-cma}}(\lambda) = \Pr[\mathbf{Exp}_{\mathcal{A}, \Sigma}^{\text{suf-cma}}(\lambda) = 1]$.

Definition 14 (suf-cma security). *A digital signature scheme Σ is suf-cma-secure if for any PPT \mathcal{A} , $\mathbf{Adv}_{\mathcal{A}, \Sigma}^{\text{suf-cma}}(\lambda)$ is negligible.*

A weaker notion of security is (simple) unforgeability against chosen-message attacks (uf-cma), which is defined as the strong version above, except that (m^*, σ^*) is considered “new” if only the message m^* (instead of the pair) is different from all messages m_i queried to the signing oracle.

A.3 Message Authentication Codes

A message authentication code consists of a triple of algorithms $\text{MAC} = (\text{Gen}, \text{Tag}, \text{Ver})$ working as follows:

$\text{Gen}(1^\lambda)$: the key generation algorithm takes as input the security parameter λ and returns a key $k \in \mathcal{K}$.
 $\text{Tag}(k, m)$: on input a secret key $k \in \mathcal{K}$ and a message $m \in \mathcal{M}$, the authentication algorithm produces an authentication tag σ .
 $\text{Ver}(k, m, \sigma)$: given the secret key k , a message m and an authentication tag σ , the verification algorithm tests if σ correctly authenticates m .

A scheme MAC is correct if for all $\lambda \in \mathbb{N}$ and $m \in \mathcal{M}$, the probability $\Pr[\text{Ver}(k, m, \sigma) = 1 : k \xleftarrow{\$} \text{Gen}(1^\lambda), \sigma \xleftarrow{\$} \text{Tag}(k, m)]$ is overwhelming. The security is defined via the following experiment:

Experiment $\mathbf{Exp}_{\mathcal{A}, \text{MAC}}^{\text{suf-cmva}}(\lambda)$
 $k \xleftarrow{\$} \text{Gen}(1^\lambda)$
 Run $\mathcal{A}^{\text{Tag}(k, \cdot), \text{Ver}(k, \cdot)}(\lambda)$
 If \mathcal{A} makes a verification query (m^*, σ^*) such that $\text{Ver}(k, m^*, \sigma^*) = 1$ and (m^*, σ^*) is “new”, then output 1.
 Else output 0

where for (m^*, σ^*) being “new” we mean that it must be different from all the pairs (m_i, σ_i) obtained from the tag oracle $\text{Tag}(k, \cdot)$. The advantage of an adversary \mathcal{A} in breaking the strong unforgeability against chosen-message and chosen verification queries attacks (suf-cmva) of MAC is $\mathbf{Adv}_{\mathcal{A}, \text{MAC}}^{\text{suf-cmva}}(\lambda) = \Pr[\mathbf{Exp}_{\mathcal{A}, \text{MAC}}^{\text{suf-cmva}}(\lambda) = 1]$.

Definition 15 (suf-cmva security). *A message authentication code MAC is suf-cmva-secure if for any PPT \mathcal{A} , $\mathbf{Adv}_{\mathcal{A}, \text{MAC}}^{\text{suf-cmva}}(\lambda)$ is negligible.*

B DDN 3-round iCCA-Secure Encryption from IND–CPA Security

In this section we recall the 3-round iCCA-secure PKE proposed by Dolev, Dwork and Naor [19], which is based on IND–CPA-secure PKE and signature schemes.

Let $\Sigma = (\Sigma.\text{kg}, \text{Sign}, \text{Ver})$ be a (regular) signature scheme, $\Sigma' = (\Sigma.\text{kg}', \text{Sign}', \text{Ver}')$ be a one-time signature scheme, and $\mathcal{E} = (\text{KG}, \text{Enc}, \text{Dec})$ be a (non-interactive) public key encryption scheme. Using our syntax, the DDN protocol $\Pi_{DDN} = (\text{Setup}, \text{S}, \text{R})$ is as follows:

Setup(1^λ): run $(\text{sk}, \text{vk}) \xleftarrow{\$} \Sigma.\text{kg}(1^\lambda)$ and output $\text{sendk} = \text{vk}$ and $\text{recvk} = \text{sk}$.

S(sendk, m): first generate a fresh one-time signing key pair $(\text{sk}_S, \text{vk}_S) \xleftarrow{\$} \Sigma.\text{kg}'(1^\lambda)$ and send vk_S to R.

Upon receiving the second message from R, the sender checks that $\text{Ver}(\text{sendk}, \text{vk}_S | \text{ek}, \sigma_R) = 1$, and if so computes $c \xleftarrow{\$} \text{Enc}(\text{ek}, m)$ and $\sigma_S \xleftarrow{\$} \text{Sign}'(\text{sk}_S, c)$, and sends (c, σ_S) .

R(recvk): upon receipt of the first message vk_S from S, generate a fresh encryption key pair $(\text{ek}, \text{dk}) \xleftarrow{\$} \text{KG}(1^\lambda)$, compute $\sigma_R \xleftarrow{\$} \text{Sign}(\text{recvk}, \text{vk}_S | \text{ek})$ and send (ek, σ_R) to S.

Once the message (c, σ_S) is obtained, the receiver checks that $\text{Ver}'(\text{vk}_S, c, \sigma_S) = 1$, and if so returns $m \leftarrow \text{Dec}(\text{dk}, c)$ as its private output.

Below we sketch the proof of security of Π_{DDN} under our iCCA notion.

Theorem 11. *If \mathcal{E} is IND–CPA-secure, Σ is uf-cma-secure, and Σ' is one-time uf-cma-secure, then Π_{DDN} is iCCA-secure.*

Proof (Sketch). Consider the experiment $\mathbf{Exp}_{\Pi_{DDN}, \mathcal{A}}^{\text{iCCA}}$, and let vk_S^* and (c^*, σ_S^*) be the first and the third protocol messages, respectively, generated by the honest sender in the challenge session. Let Forge_1 be the event that in the challenge session the adversary \mathcal{A} outputs a second protocol message including a valid signature σ_R that was *not* obtained by the R oracle. It is not hard to see that under the assumption that the signature scheme Σ is secure we have that $\Pr[\text{Forge}_1]$ is negligible. Also, let Forge_2 be the event that the adversary queries the receiver oracle with a third protocol message (c, σ_S) such that: the first protocol message of the given session is vk_S^* (i.e., the same verification key as in the challenge session), but c is “new”, in particular $c \neq c^*$. Again, it is possible to show that under the assumption that Σ' is a one-time signature the probability $\Pr[\text{Forge}_2]$ is negligible.

Finally, it is possible to show that $\Pr[\mathbf{Exp}_{\Pi_{DDN}, \mathcal{A}}^{\text{iCCA}} \mid \overline{\text{Forge}_1} \wedge \overline{\text{Forge}_2}]$ is negligible under the assumption that \mathcal{E} is IND–CPA-secure. The main observation is that one can simulate answers of the receiver oracle to third messages (i.e., decryptions). This follows from the fact that when Forge_1 and Forge_2 do not occur, a non-ping-pong adversary cannot create a valid third message that uses the same one-time signature/encryption keys of the challenge session. \square

C A 1-bounded IND–CCA-Secure Encryption Scheme

Here we recall a black-box construction of a 1-bounded IND–CCA-secure encryption scheme from an IND–CPA-secure one. The scheme is a (simplified) special case of the q -bounded one of Cramer et al. [12], but it is also very similar to the scheme of Dolev, Dwork and Naor [19] without NIZK proofs.

Let $\mathcal{E}' = (\text{KG}', \text{Enc}', \text{Dec}')$ be a semantic secure public key encryption scheme. The IND–CCA scheme $\mathcal{E} = (\text{KG}, \text{Enc}, \text{Dec})$ is defined as follows:

KG(1^λ): Generate $2n$ keys $(\text{ek}_i^b, \text{dk}_i^b) \xleftarrow{\$} \text{KG}'(1^\lambda) \forall i = 1, \dots, n, b = 0, 1$, and a random universal one-way hash function h . Output the public key $\text{ek} = (h, \text{ek}_1^0, \text{ek}_1^1, \dots, \text{ek}_n^0, \text{ek}_n^1)$ and the secret key $\text{dk} = (\text{dk}_1^0, \text{dk}_1^1, \dots, \text{dk}_n^0, \text{dk}_n^1)$.

Enc(ek, m): First, generate a pair of keys (sk, vk) for a one-time signature, and compute $v = h(\text{vk})$ (write $v = v_1 \cdots v_n$ using its bitwise representation). Next, generate random messages m_1, \dots, m_n such that $m_1 \oplus m_2 \oplus \dots \oplus m_n = m$. Then, for $i = 1$ to n , compute $c_i \xleftarrow{\$} \text{Enc}(\text{ek}_i^{v_i}, m_i)$, create a signature σ on (c_1, \dots, c_n) using sk , and output $C = (\text{vk}, \sigma, c_1, \dots, c_n)$.

$\text{Dec}(\text{dk}, C)$: Verify that σ is a valid signature on (c_1, \dots, c_n) using verification key vk . Then compute $v = h(\text{vk})$ and decrypt each ciphertext computing $m_i \leftarrow \text{Dec}(\text{dk}_i^{v_i})$. If the signature is valid output $m = m_1 \oplus m_2 \oplus \dots \oplus m_n$. Otherwise output \perp (reject).

D iCCA-Secure Interactive Encryption with Labels

In an encryption scheme with “labels” the encryption and decryption algorithms are assumed to take a public string called label as an additional input. While this notion has been already introduced in the non-interactive setting [8], here we propose its natural extension to the interactive scenario. We will show an elegant application of labeled iCCA encryption to our notion of public key message authentication (see Section 2.2).

An interactive encryption protocol with labels is basically the same as the one defined in Section 2.1, except that both sender and receiver work with a public label L (i.e., an arbitrary binary string) as additional input. For any label L we denote the algorithms with $S_L(\text{sendk}, m)$ and $R_L(\text{recvk})$. For correctness, we require that for all honestly generated keys $(\text{sendk}, \text{recvk})$, all messages $m \in \mathcal{M}$ and all labels L , $\langle S_L(\text{sendk}, m), R_L(\text{recvk}) \rangle = m$ holds with all but negligible probability.

The iCCA security notion is extended to the labeled case as follows: when the adversary queries the receiver oracle, it also specifies a label L , i.e., \mathcal{A} interacts with $R_L(\text{recvk})$. The adversary \mathcal{A} is also required to output a label L^* along with the message pair (m_0, m_1) , and the challenge session is $\langle S_{L^*}(\text{sendk}, m_b), \mathcal{A}^{R_{(\cdot)}(\text{recvk})} \rangle$. The restriction on \mathcal{A} not to be ping-pong is extended in such a way that for every oracle session with transcript T and label L , it must hold $L \neq L^*$ or $T \neq T^*$.

Building iCCA Encryption with Labels from iCCA Encryption. We show that iCCA encryption with labels can be realized from iCCA-secure encryption. The idea of the construction is the same as for the non-interactive case: if L is the label, then the encryption protocol is run with plaintext $m_L = L|m$, i.e., the concatenation of the label and the plaintext m ; when the decryptor obtains a plaintext $m'_L = L'|m'$ it checks whether it obtained the label L , i.e., if $L' = L$.

A formal description of this construction follows. Let $\Pi = (\text{Setup}, S, R)$ be an interactive encryption protocol. We build an interactive encryption protocol with labels $\Pi' = (\text{Setup}', S', R')$ as follows. The key generation is the same $\text{Setup}' = \text{Setup}$, thus $\text{sendk}' = \text{sendk}, \text{recvk}' = \text{recvk}$. The sender algorithm $S'_L(\text{sendk}, m)$ simply runs $S(\text{sendk}, L|m)$ (i.e., on the concatenation of the message m and the label L). The receiver algorithm $R'_L(\text{recvk})$ runs $R(\text{recvk})$. At the end of a session, if $m'_L = (L'|m')$ is the the message returned by R , then R' returns m' only if $L' = L$. Otherwise, it returns \perp .

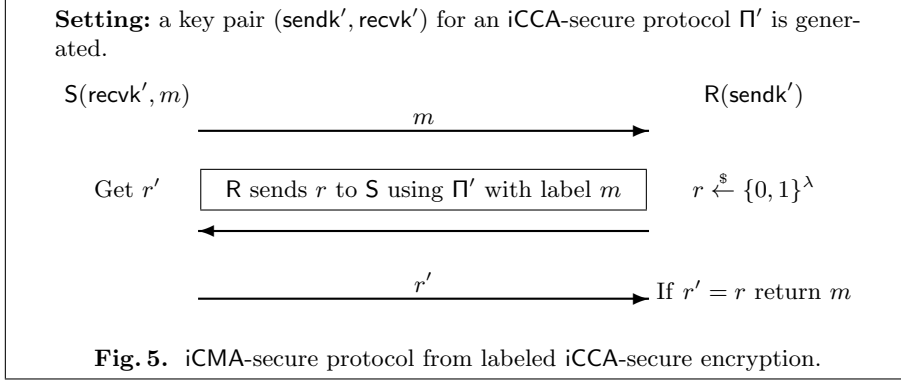
We show the security of this construction via the following theorem.

Theorem 12. *If Π is a iCCA-secure encryption protocol, then Π' is a iCCA-secure encryption protocol with labels.*

Proof. Assume by contradiction that there exists a PPT adversary \mathcal{A} that breaks the iCCA security of the labeled scheme Π' , then we show how to build a PPT algorithm \mathcal{B} that uses \mathcal{A} as a subroutine to break the iCCA security of Π .

\mathcal{B} is given sendk , it runs $\mathcal{A}(\text{sendk})$ and answers oracle queries as follows. When \mathcal{A} will query R' on a new session, \mathcal{A} has to provide a label L . So, \mathcal{B} stores L , queries its oracle, R , and returns its answer to \mathcal{A} . All the oracle queries to R' for sessions that already started are simply forwarded by \mathcal{B} to its R oracle. When \mathcal{A} queries R' on the last message, \mathcal{B} first forwards this query to R , obtains $m = (L'|m')$, and returns to \mathcal{A} m' , if $L = L'$ (where L is the label stored for this session), and \perp otherwise.

When \mathcal{A} outputs a message pair (m_0, m_1) and a label L^* , \mathcal{B} starts its challenge session by returning $(L^*|m_0, L^*|m_1)$ as its message pair. Next, for all subsequent messages \mathcal{B} simply forwards them to its challenge session, and forwards back to \mathcal{A} the respective responses. Let $T^* = \langle (M_1^*, t_1^*), \dots, (M_n^*, t_n^*) \rangle$ be the transcript of the challenge session between \mathcal{B} (who is simulating $S'_{L^*}(\text{sendk}, m_b)$) and \mathcal{A} .



In this second phase, \mathcal{B} answers all oracle queries as before except for the following change. Assume that \mathcal{A} sends the last message M_n in some session whose label is L and whose transcript is $T = \langle (M_1, t_1), \dots, (M_n, t_n) \rangle$. If $T \equiv T^*$ but $L \neq L^*$, then \mathcal{B} returns \perp as the R' output. Note that by correctness of Π , R would decrypt T to some message of the form $L^*|m_b$. Hence, as $L \neq L^*$ R' would reject this message, that is \mathcal{B} 's answer is correctly distributed. If b' is the output of \mathcal{A} at the end of the simulation, \mathcal{B} outputs the same value. To complete the proof, we observe that if \mathcal{A} is not ping-pong, then \mathcal{B} is not ping-pong. \square

E iCMA-secure PKMA from labeled iCCA-secure encryption.

Here we consider another protocol that achieves public key message authentication based on iCCA-secure encryption. The basic protocol was first proposed by Dolev, Dwork, and Naor in [19], and later discussed by Dwork, Naor, and Sahai [21,22]. Here we revisit this construction by using the abstraction of encryption with labels, and by using our notion of (possibly) interactive iCCA-secure encryption (see Section D). A sketch of the protocol is summarized in Figure 5: (1) S sends the message m to R ; (2) R picks a random string r and confidentially transmits r with label m to S by using the encryption protocol Π' ; (3) S obtains r' and hands r' to the receiver. If R obtains the same r , then it accepts the message m .

To improve on round complexity, one can use even a 1-round (aka non-interactive) iCCA-secure labeled encryption scheme. In this case, our construction yields a 3-round iCMA-secure protocol. However, we note that the first round, where S sends m to R , can be dropped in all those applications where the message is implicitly known to the receiver (e.g., it is some public information). This way, one obtains an optimal 2-round protocol.

More formally, if $\Pi' = (\text{Setup}', S', R')$ is an encryption protocol with labels with message space $\tilde{\mathcal{M}}$ such that $|\tilde{\mathcal{M}}| \approx 2^\lambda$, then we build a PKMA protocol $\Pi = (\text{Setup}, S, \text{Ver})$ as follows.

Setup(1^λ): run $(\text{sendk}', \text{recvk}') \xleftarrow{\$} \text{Setup}'(1^\lambda)$ and output $\text{sendk} = \text{recvk}'$ and $\text{recvk} = \text{sendk}'$.

S(sendk, m): as first protocol message, send m . Next, run the encryption protocol Π' by playing the role of the receiver. More precisely, use m as the label, and run $R'_m(\text{recvk}')$. Let r' be the output of R' at the end of the encryption protocol. Finally, send r' to R as the last message.

R(recvk): first, wait for message m from S . Next, choose a random string $r \xleftarrow{\$} \tilde{\mathcal{M}}$ and run the encryption protocol Π' by playing the role of the sender. More precisely, use m as the label and run $S'_m(\text{sendk}', r)$. If the encryption protocol terminates correctly, then wait for the last message r' from S . If $r' = r$ then output m . Otherwise, output \perp .

We can now prove the following theorem.

Theorem 13. *If Π' is iCCA-secure encryption with labels, with message space of size at least 2^λ , then the protocol Π described above is iCMA-secure.*

Proof. To prove the theorem we define the following hybrid games and we denote with G_i the event that the outcome of Game i , run with \mathcal{A} , is 1.

Game 0: this is the real iCMA game.

Game 1: this is the same as Game 0 except that in the challenge session the challenger generates two random strings $r^* \xleftarrow{\$} \tilde{\mathcal{M}}$, and it runs the encryption protocol Π' with a fixed string r_1^* (e.g., $r_1^* = 0^\lambda$), i.e., \mathcal{B} runs $S'_{m^*}(\text{sendk}', r_1^*)$. However, the challenger keeps using r_0^* as the random string associated with the run of the encryption protocol in the challenge session, i.e., the remaining part of the game (outside the encryption protocol) remains the same as Game 0. Precisely, this means that the challenger uses r_0^* to check the validity of the last message r' provided by the adversary in the challenge session, i.e., it checks whether $r' = r_0^*$.

Via a straightforward reduction to the iCCA-security of the encryption protocol, it is possible to show that there exists an adversary \mathcal{B} such that

$$\Pr[G_0] - \Pr[G_1] \leq 2 \cdot \text{Adv}_{\Pi', \mathcal{B}}^{\text{iCCA}}(\lambda)$$

Finally, if we analyze Game 1, it is not hard to see that in order to let Game 1 output 1, the adversary has to correctly guess the value of r_0^* . However, in Game 1, the string r_0^* is randomly chosen and is not used in the encryption protocol. Hence, its value is information-theoretically hidden to any adversary \mathcal{A} . Hence, we have that $\Pr[G_1] \leq \frac{1}{2^\lambda}$, which is negligible, as desired. \square

F Postponed Proofs

F.1 Proof of Claim 1

Proof (Claim 1). Assume by contradiction that there exists an efficient adversary \mathcal{A} such that $\Pr[\text{Forge}] \geq \epsilon$ for a non-negligible value ϵ . Then we show how to build a PPT adversary \mathcal{B} that has non-negligible advantage against the iCMA security of the message transmission protocol Π' .

\mathcal{B} is given in input the public receiver key recvk' and it has oracle access to the sender $S'(\text{sendk}', \cdot)$. \mathcal{B} proceeds as follows:

- Run $\mathcal{A}(\text{sendk})$ with $\text{sendk} = \text{recvk}'$.
- For every oracle query asking to interact with a new copy of R: generate a new encryption key pair $(\text{ek}, \text{dk}) \xleftarrow{\$} \text{KG}(1^\lambda)$, and query $S'(\text{sendk}', \text{ek})$ to simulate the first part of the protocol.
- When the adversary makes the last query to R on a session that is already opened: let ek be the first message previously generated in the above step, and let c be the message sent by \mathcal{A} for this query. \mathcal{B} answers by computing $m \leftarrow \text{Dec}(\text{dk}, C)$ (where dk is the decryption key generated by \mathcal{B} together with ek).
- When \mathcal{A} outputs the message pair (m_0, m_1) , \mathcal{B} also starts its challenge session, and forwards all \mathcal{A} 's messages of the subprotocol Π' to its challenger. If Π' in the challenge session terminates with ek^* , then \mathcal{B} returns ek^* .

If **Forge** occurs then, by the definition of the event, \mathcal{A} is not ping-pong and thus \mathcal{B} is not ping-pong either. Hence, $\text{Adv}_{\mathcal{B}, \Sigma}^{\text{uf-cma}}(\lambda) = \Pr[\text{Forge}] \geq \epsilon$, which concludes the proof of the claim. \square

F.2 Proof of Theorem 3 (iCCA security)

Assume by contradiction there exists a PPT adversary \mathcal{A} that breaks the iCCA security of Π' , then we show how to build a PPT algorithm \mathcal{B} that uses \mathcal{A} as a subroutine to break the iCCA security of Π .

Here too, we distinguish between the above two cases according to which party speaks first.

1. **R speaks first.** \mathcal{B} is run on input sendk . It runs $\mathcal{A}(\text{sendk})$ and answers oracle queries as follows. Since in Π' it is S' who speaks first, when \mathcal{A} will query R' to start a new session, \mathcal{A} will also send a nonce $r \in \{0, 1\}^\lambda$. \mathcal{B} stores r , queries R and returns its answer to \mathcal{A} . R' oracle queries for sessions that already started before are forwarded by \mathcal{B} to R , except for the following change for the last message. \mathcal{B} forwards the message received from \mathcal{A} to R . Let $(m'|r')$ be its response, and let r be the string stored for this session by \mathcal{B} . If $r = r'$ then \mathcal{B} returns m' . Otherwise, it returns \perp .

When \mathcal{A} outputs a message pair (m_0, m_1) \mathcal{B} chooses a random string $r^* \xleftarrow{\$} \{0, 1\}^\lambda \setminus \{r_1, \dots, r_Q\}$ (where r_1, \dots, r_Q are all the strings sent by \mathcal{A} to the oracle R' in the previous phase), and sends r^* to \mathcal{A} . Note that as long as the r string is sufficiently large, the distribution of this choice of r^* is statistically close to the one in the real experiment. So, \mathcal{B} outputs $(r^*|m_0, r^*|m_1)$ as its message pair, and continues the simulation by forwarding the remaining messages of the challenge session to its challenger.

In this second phase, all oracle queries are answered as before, except for the following change. Assume that the challenge session is completed, let $T^* = \langle (r^*, t_1^*), (M_1^*, t_2^*), \dots, (M_n^*, t_{n+1}^*) \rangle$ be its transcript, and assume that \mathcal{A} queries R' on an already opened session by sending the last message. \mathcal{B} should then answer with the output of R' . Let $T = \langle (r, t_1), (M_1, t_2), \dots, (M_n, t_{n+1}) \rangle$ be the transcript of the queried session. Since \mathcal{A} is not ping-pong, $T \neq T^*$, and we would like to argue that \mathcal{B} must not be ping-pong as well while still being capable to answer queries correctly.

Let us write $T^* = (r^*, t_1^*), T^{*'} and $T = (r, t_1), T'$ (i.e., we explicitly separate the first message from the transcript of the n -round protocol). If $T \neq T^*$, either one of the following cases holds:$

- $T' \neq T^{*'}$: \mathcal{B} is not ping-pong in its game and it can answer the query by forwarding \mathcal{A} 's last message to its oracle R .
- $T' \equiv T^{*'}$, $t_1^* > t_1$ and $r = r^*$: $t_1^* > t_1$ means that r^* was generated after r has been sent from \mathcal{A} . However, by construction of the simulation it must be $r^* \neq r$. So, this case cannot occur.
- $T' \equiv T^{*'}$, $t_1^* < t_1$ and $r^* \neq r$: This means that all timestamps are correctly alternating, and the two sessions differ only in the first message, i.e., $r \neq r^*$. This is the case where \mathcal{B} changes its way of answering: it does not forward the last message to its oracle R , and returns \perp to \mathcal{A} . We argue that \mathcal{B} 's answer is correct. Indeed, since $T' \equiv T^{*'}$, by correctness of Π , R 's private output will be a message of the form $r^*|m_b$. However, since $r^* \neq r$ R' would reject this message

At the end, \mathcal{B} returns the same output of \mathcal{A} .

2. **S speaks first.** \mathcal{B} is run on input sendk . It runs $\mathcal{A}(\text{sendk})$ and answers oracle queries as follows. Since in Π' it is R' who speaks first, when \mathcal{A} will query R' on a new session, \mathcal{A} does not send any message, and \mathcal{B} simulates the answer by returning a randomly chosen string $r \in \{0, 1\}^\lambda$. \mathcal{B} stores r , and then forward all subsequent queries on this session to its oracle R , forwarding the corresponding answers to \mathcal{A} . The only change is in simulating answers to the last message. \mathcal{B} forwards the message received from \mathcal{A} to R . Let $(m'|r')$ be its response, and let r be the string stored for this session by \mathcal{B} . If $r = r'$ then \mathcal{B} returns m' . Otherwise, it returns \perp .

When \mathcal{A} outputs a message pair (m_0, m_1) — recall that \mathcal{A} (who is playing the role of the receiver) is supposed to speak first in the challenge session — \mathcal{B} waits for the message $r^* \in \{0, 1\}^\lambda$ from \mathcal{A} . \mathcal{B} stores r^* and starts forwarding all messages to its own challenge session which is initialized by choosing $(r^*|m_0, r^*|m_1)$ as the challenge message pair.

In this second phase all oracle queries are answered as before, except for the following changes. First, in every query from \mathcal{A} to R' to open a new session, \mathcal{B} chooses the nonce $r \xleftarrow{\$} \{0, 1\}^\lambda \setminus \{r^*\}$. It is easy to see that the distribution of this choice of r is statistically close to the one in the real experiment. Second, assume that the challenge session is completed, let $T^* = \langle (r^*, t_1^*), (M_1^*, t_2^*), \dots, (M_n^*, t_{n+1}^*) \rangle$ be its transcript, and assume that \mathcal{A} queries R' on an already opened session by sending the last message. \mathcal{B} should then answer with the (simulated) output of R' . Let $T = \langle (r, t_1), (M_1, t_2), \dots, (M_n, t_{n+1}) \rangle$ be the transcript of the queried session. Since \mathcal{A} is not ping-pong, $T \neq T^*$, and we would like to argue that \mathcal{B} must not be ping-pong as well while still being capable to answer queries correctly.

Let us write $T^* = (r^*, t_1^*), T^{*'} and $T = (r, t_1), T'$ (i.e., we explicitly separate the first message from the transcript of the n -round protocol). If $T \not\equiv T^*$, either one of the following cases holds:$

- $T' \not\equiv T^{*'}$: \mathcal{B} is clearly not ping-pong and it can answer forwarding this query to its oracle R .
- $T' \equiv T^{*'}$, $t_1 > t_1^*$ and $r = r^*$: $t_1 > t_1^*$ means that r was generated by \mathcal{B} after r^* has been sent by \mathcal{A} . However, by construction of \mathcal{B} (see above) we always have $r^* \neq r$. So this case cannot occur.
- $T' \equiv T^{*'}$, $t_1 < t_1^*$ and $r^* \neq r$: This means that all timestamps are correctly alternating, and the two sessions differ only in the first message, i.e., $r \neq r^*$. In this case \mathcal{B} outputs \perp . By the same reasons in case (1) of the proof, this answer is correct.

At the end, \mathcal{B} returns the same output of \mathcal{A} . □

F.3 Proof of Theorem 3 (iCMA security)

Assume by contradiction there exists a PPT adversary \mathcal{A} that breaks the iCMA security of Π' , then we show how to build a PPT algorithm \mathcal{B} that uses \mathcal{A} as a subroutine in order to break the iCMA security of Π . We distinguish between the following two cases according to which party speaks first.

1. **R speaks first.** \mathcal{B} is run on input recvk . It runs $\mathcal{A}(\text{recvk})$ and answers oracle queries as follows. Since in Π' it is S' who speaks first, when \mathcal{A} queries S' on a new session (providing a plaintext m), \mathcal{B} simulates the first message from S' by returning a random string $r \in \{0, 1\}^\lambda$, which is stored by \mathcal{B} . All subsequent queries on this session are forwarded from \mathcal{B} to its oracle S (initialized with the plaintext $m|r$).

When \mathcal{A} starts the challenge session, \mathcal{A} must start speaking by sending a string r^* . \mathcal{B} stores r^* , starts its challenge session and forwards all messages to and from \mathcal{A} . After the challenge session has started, \mathcal{B} chooses the nonces $r \xleftarrow{\$} \{0, 1\}^\lambda \setminus \{r^*\}$. It is easy to see that the distribution of this choice of r is statistically close to the one in the real experiment.

Observe that by construction of R' , if \mathcal{A} makes $\mathsf{R}'(\text{recvk})$ accept and return a message m^* in the simulated challenge session with \mathcal{B} , then \mathcal{B} (who forwards the very same queries) will make $\mathsf{R}(\text{recvk})$ accept with message $r^*|m^*$ in its challenge session. Therefore, to complete the proof, it is left to show that if \mathcal{A} is not ping-pong, the same holds for \mathcal{B} . Namely, let $T^* = \langle (r^*, t_1^*), (M_1^*, t_2^*), \dots, (M_n^*, t_{n+1}^*) \rangle$ be the transcript of the challenge session between \mathcal{B} (simulating $\mathsf{R}'(\text{recvk})$) and \mathcal{A} , and let $T = \langle (r, t_1), (M_1, t_2), \dots, (M_n, t_{n+1}) \rangle$ be the transcript of the any session between \mathcal{A} and its oracle S' (simulated by \mathcal{B}). Since \mathcal{A} is not ping-pong, $T \not\equiv T^*$. Let us write $T^* = (r^*, t_1^*), T^{*'}$ and $T = (r, t_1), T'$ (i.e., we explicitly separate the first message from the transcript of the n -round protocol). If $T \not\equiv T^*$, either one of the following cases holds:

- $T' \not\equiv T^{*'}$: \mathcal{B} is also not ping-pong in its game.
- $T' \equiv T^{*'}$, $t_1^* > t_1$ and $r = r^*$: $t_1^* > t_1$ means that r^* was generated after r has been sent from \mathcal{A} . However, by construction of the simulation we only have $r \neq r^*$, and thus this case cannot occur.
- $T' \equiv T^{*'}$, $t_1^* < t_1$ and $r^* \neq r$: This means that all timestamps are correctly alternating, and the two sessions differ only in the first message, i.e., $r \neq r^*$. Since $T' \equiv T^{*'}$, by correctness of Π , R would accept in the challenge session with a message of the form $m^*|r^*$, but this implies that R' rejects this message as $r^* \neq r$.

2. **S speaks first.** \mathcal{B} is run on input recvk . It runs $\mathcal{A}(\text{recvk})$ and answer oracle queries as follows. Since in Π' it is R' who speaks first, when \mathcal{A} will query S' on a new session, \mathcal{A} has to provide a string r (in addition to the chosen plaintext m). \mathcal{B} then queries its oracle S with plaintext $m|r$ and starts forwarding all subsequent queries of this session to \mathcal{A} .

When \mathcal{A} starts the challenge session, \mathcal{B} chooses a random string $r^* \xleftarrow{\$} \{0, 1\}^\lambda \setminus \{r_1, \dots, r_Q\}$ (where r_1, \dots, r_Q are all the strings sent by \mathcal{A} to the oracle R' in the previous phase), and sends r^* to \mathcal{A} . Note that as long as the r string is sufficiently large, the distribution of this choice of r^* is statistically

close to the one in the real experiment. Next, in the challenge session \mathcal{B} simply relay all messages between \mathcal{A} and its challenger.

Observe that by construction of \mathcal{R}' , if \mathcal{A} would make $\mathcal{R}'(\text{recvk})$ accept and return a message m^* in the simulated challenge session with \mathcal{B} , then \mathcal{B} (who forwards the very same messages) can make $\mathcal{R}(\text{recvk})$ accept with message $m^*|r^*$ in its challenge session. Therefore, to complete the proof, it is left to show that if \mathcal{A} is not ping-pong, the same holds for \mathcal{B} . Namely, let $T^* = \langle (r^*, t_1^*), (M_1^*, t_2^*), \dots, (M_n^*, t_{n+1}^*) \rangle$ be the transcript of the challenge session between \mathcal{B} (simulating $\mathcal{R}'(\text{recvk})$) and \mathcal{A} , and let $T = \langle (r, t_1), (M_1, t_2), \dots, (M_n, t_{n+1}) \rangle$ be the transcript of any session between \mathcal{A} and its oracle \mathcal{S}' (simulated by \mathcal{B}). Since \mathcal{A} is not ping-pong, $T \not\equiv T^*$. Let us write $T^* = (r^*, t_1^*), T^{*'} and $T = (r, t_1), T'$ (i.e., we explicitly separate the first message from the transcript of the n -round protocol). If $T \not\equiv T^*$, either one of the following cases holds:$

- $T' \not\equiv T^{*'}$: \mathcal{B} is clearly not ping-pong in its game.
- $T' \equiv T^{*'}$, $t_1 > t_1^*$ and $r = r^*$: $t_1 > t_1^*$ means that r was generated after r^* has been sent from \mathcal{A} . However, by construction of the simulation we only have $r \neq r^*$, and thus this case cannot occur.
- $T' \equiv T^{*'}$, $t_1 < t_1^*$ and $r^* \neq r$: This means that all timestamps are correctly alternating, and the two sessions differ only in the first message, i.e., $r \neq r^*$. Since $T' \equiv T^{*'}$, by correctness of Π , \mathcal{R} would accept in the challenge session with a message of the form $m^*|r^*$, but this implies that \mathcal{R}' rejects this message as $r^* \neq r$. □