Alternatives to Non-Malleability: Definitions, Constructions and Applications
(Extended Abstract)

Philip MacKenzie[1], Michael K. Reiter[2], and Ke Yang[2]

[1] Bell Labs, Lucent Technologies, Murray Hill, NJ, USA;
philmac@research.bell-labs.com
[2] Carnegie Mellon University, Pittsburgh, PA, USA;
reiter@cmu.edu,yangke@cs.cmu.edu

Abstract. We explore whether non-malleability is necessary for the applications typically used to motivate it, and propose two alternatives. The first we call weak non-malleability (wnm) and show that it suffices to achieve secure contract bidding (the application for which non-malleability was initially introduced), despite being strictly weaker than non-malleability. The second we call tag-based non-malleability (tnm), and show that it suffices to construct an efficient universally-composable secure message transmission (SMT) protocol, for which the only previous solution was based on a public key encryption functionality whose security is equivalent to non-malleability. We also demonstrate constructions for wnm and tnm encryption schemes that are simpler than known constructions of non-malleable encryption schemes.

1 Introduction

Non-malleability [11] is a security condition for encryption schemes that requires, informally, that an attacker given a challenge ciphertext be unable to produce another, different ciphertext so that the plaintexts underlying the two ciphertexts are "meaningfully related" to each other. Non-malleability is the strongest commonly considered notion of security for encryption, being strictly stronger than indistinguishability [14] under chosen-plaintext or indifferent chosen-ciphertext ("lunchtime") attacks, and being equivalent to indistinguishability under adaptive chosen-ciphertext attacks [1].

In this paper we revisit the definition of non-malleability with an eye toward whether it is *necessary* for applications commonly used to motivate it. Our contributions in this study are twofold. First, we identify alternatives to non-malleability that suffice for applications where previously non-malleability seemed warranted. Second, we identify encryption schemes that implement these variants and that are conceptually simpler than known non-malleable schemes.

The alternative definitions that we propose deviate from non-malleability in different ways. The first notion, which we call *weak non-malleability* (wnm), identifies a point in the space of definitions strictly between non-malleability

and indistinguishability (in those cases where there is room between them, i.e., under chosen-plaintext and lunchtime attacks). Informally, wnm allows mauling of a ciphertext c, but such that this mauling does not benefit the adversary. In particular, a mauling that produces a valid ciphertext c' would imply that the adversary has successfully guessed the plaintext corresponding to c, and thus for many natural applications, this mauling would not be useful. In other words, in such applications, wnm should suffice in place of non-malleability. As an example, we show that a wnm encryption scheme suffices to implement a secure contract bidding auction in the spirit of that originally used to (informally) motivate non-malleability [11]. Still, wnm does allow an adversary to produce a ciphertext c' that has a (very restricted) dependence of a given ciphertext c, and we can in fact show that wnm is a *strictly* weaker property than non-malleability. In addition, we show that this weaker property may be satisfied by very simple encryption schemes similar to those used in Bellare and Rogaway [2] to achieve the (even less stringent) property of indistinguishability under chosen-plaintext attacks [2].[1] These schemes assume p is a prime, H is a hash function (modeled by a random oracle in our security analyses) with range a group X with group operation "$\cdot$", and f denotes a trapdoor permutation that constitutes the public key (with the trapdoor being the private key):

Mult-Range scheme The encryption of m is $E(m) = <f(r), H(r) \cdot m>$ where r is chosen randomly (per encryption) from the domain of f, the plaintext space is an integer range $[a, b]$ satisfying $0 < a < b < p$, $a > (b-a)^2$ and $p > 2b^2$, and $X = \mathbb{Z}_p^*$ with $\cdot$ being multiplication in $\mathbb{Z}_p^*$.

Mult-Adjacent scheme The encryption of m is $E(m) = <f(r), H(r) \cdot (m, m+1)>$ where r is chosen randomly (per encryption) from the domain of f, the plaintext space is $\mathbb{Z}_p^* \setminus \{p-1\}$, and $X = \mathbb{Z}_p^* \times \mathbb{Z}_p^*$ with group operation $\cdot$ being component-wise multiplication in $\mathbb{Z}_p^*$, i.e., $(x_0, x_1) \cdot (y_0, y_1) = (x_0 y_0, x_1 y_1)$.

Add-Square scheme The encryption of m is $E(m) = <f(r), H(r) \cdot (m, m^2)>$, where the plaintext space is $\mathbb{Z}_p^*$, and $X = \mathbb{Z}_p \times \mathbb{Z}_p$ with group operation $\cdot$ being component-wise addition in $\mathbb{Z}_p$, i.e., $(x_0, x_1) \cdot (y_0, y_1) = (x_0 + y_0, x_1 + y_1)$.

For some intuition behind weak non-malleability, consider the Mult-Range scheme above. Without the range restriction on the plaintext space, this scheme would be completely malleable (similar to the first scheme introduced in [2]). However, simply by restricting the range of plaintexts (as opposed to, e.g., adding an additional hash for verification/redundancy, as is done in [2] to achieve non-malleability) we are able to achieve wnm. Informally, this is because any modification of a ciphertext (v, w) to (v, w') implies a multiplying factor w'/w for which there is only a single plaintext in the range that would be transformed into another plaintext in the range.

[1] While there exist efficient encryption systems that implement indistinguishability under adaptive chosen-ciphertext attacks (and thus non-malleability under these attacks, e.g., [2, 8]), we are unaware of prior constructions that, like those listed here, so *simply* implement a property strictly stronger than indistinguishability (in this case, weak non-malleability) under chosen-plaintext and lunchtime attacks.

The second alternative to non-malleability that we propose is called *tag-based non-malleability* (tnm). Here, we structurally modify encryption and decryption to take an additional public string argument called a *tag*. Informally, tnm dictates that an adversary be unable to create a (ciphertext,tag) pair with plaintext related to that of the challenge ciphertext and with the tag being different from the challenge tag, even though it is able to obtain decryptions of (ciphertext,tag) pairs with any tag different from the challenge tag. We demonstrate the utility of tnm by using it to implement the "secure message transmission functionality" in the universal composability framework of [5], replacing the use of non-malleable encryption there, and arguably providing a more natural implementation. tnm also admits exceedingly simple implementations, e.g.:

Tag-based scheme The encryption of m with tag t is $E(m,t) = <f(r), H(r,t) \cdot m>$ where r is chosen randomly (per encryption) from the domain of f. The plaintext space is $\mathbb{Z}_p^*$, and $X = \mathbb{Z}_p^*$ with $\cdot$ being multiplication in $\mathbb{Z}_p^*$.

We also present a tnm construction that is a (simpler) variation of the Cramer-Shoup encryption scheme [8, 9]. The change in structure for encryption and decryption (specifically due to the tag) does not permit us to argue that tnm is definitionally weaker than non-malleability. However, given a non-malleable encryption scheme, it is trivial to implement a tnm scheme using it with no additional assumptions or loss in security. We also show how to implement a non-malleable scheme using a tnm scheme and a one-time signature scheme.

2 Preliminaries

Trapdoor Permutations [2, 15] A *permutation generator* G_* is a probabilistic polynomial time algorithm that takes as input 1^k and outputs three polynomial-time algorithms (f, f^{-1}, d), the first two being deterministic, and the last being probabilistic. The range of $d(1^k)$ is required to be a subset of $\{0,1\}^k$, and f, f^{-1} are permutations over the range of $d(1^k)$, and are inverses of each other. G_* is a *trapdoor permutation generator* if it is a permutation generator such that for all non-uniform polynomial-time algorithms $\mathcal{A}$, $\Pr[(f, f^{-1}, d) \leftarrow G_*(1^k); x \leftarrow d(1^k); y \leftarrow f(x) : \mathcal{A}(f, d, y) = x]$ is negligible. It is commonly assumed that, for example, RSA is a trapdoor permutation.

Encryption schemes An *encryption scheme* Π is a triple (G, E, D) of algorithms, the first two being probabilistic, and all running in polynomial time. G takes as input 1^k and outputs a public key pair (pk, sk), i.e., $(pk, sk) \leftarrow G(1^k)$. E takes a public key pk and a message m as input and outputs an encryption c for m; we denote this $c \leftarrow E_{pk}(m)$. D takes a private key sk and a ciphertext c as input and returns either a message m such that c is a valid encryption of m, if such an m exists, and otherwise returns $\perp$; we denote this $m \leftarrow D_{sk}(c)$.

As discussed in Section 1, indistinguishability [14] is the most commonly studied goal for encryption. Here we adopt definitions ind-cpa, ind-cca1, and ind-cca2 from [1]. Below we give the definition of non-malleability from Dolev, Dwork

and Naor [11], as written explicitly as the simulator-based non-malleable (snm) definition in Bellare and Sahai [4].[2] In this definition and throughout, we use atk to denote one of {cpa, cca1, cca2} and define oracles $\mathcal{O}_1$ and $\mathcal{O}_2$ as follows:

$$\begin{aligned}
\mathsf{atk} = \mathsf{cpa} &\Rightarrow \mathcal{O}_1(\cdot) = \epsilon, \mathcal{O}_2(\cdot) = \epsilon \\
\mathsf{atk} = \mathsf{cca1} &\Rightarrow \mathcal{O}_1(\cdot) = D_{sk}(\cdot), \mathcal{O}_2(\cdot) = \epsilon \\
\mathsf{atk} = \mathsf{cca2} &\Rightarrow \mathcal{O}_1(\cdot) = D_{sk}(\cdot), \mathcal{O}_2(\cdot) = D_{sk}(\cdot)
\end{aligned}$$

Definition 1 (snm-cpa, snm-cca1, snm-cca2). *Let $\Pi = (G, E, D)$ be an encryption scheme, let R be a relation, let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary, and let $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ be an algorithm (the "simulator"). For $k \in \mathbb{N}$ define* $\mathsf{Adv}^{\mathsf{snm\text{-}atk}}_{\mathcal{A},\mathcal{S},\Pi}(R,k) \stackrel{\text{def}}{=} \Pr[\mathsf{Expt}^{\mathsf{snm\text{-}atk}}_{\mathcal{A},\Pi}(R,k) = 1] - \Pr[\mathsf{Expt}^{\mathsf{snm\text{-}atk}}_{\mathcal{S},\Pi}(R,k) = 1]$, *where*

$\mathsf{Expt}^{\mathsf{snm\text{-}atk}}_{\mathcal{A},\Pi}(R,k)$:	$\mathsf{Expt}^{\mathsf{snm\text{-}atk}}_{\mathcal{S},\Pi}(R,k)$:
$(pk, sk) \leftarrow G(1^k)$	$(pk, sk) \leftarrow G(1^k)$
$(M, s_1, s_2) \leftarrow \mathcal{A}_1^{\mathcal{O}_1}(pk)$	$(M, s_1, s_2) \leftarrow \mathcal{S}_1(pk)$
$x \leftarrow M$	$x \leftarrow M$
$y \leftarrow E_{pk}(x)$	
$\mathbf{y} \leftarrow \mathcal{A}_2^{\mathcal{O}_2}(s_2, y)$	$\mathbf{y} \leftarrow \mathcal{S}_2(s_2)$
$\mathbf{x} \leftarrow D_{sk}(\mathbf{y})$	$\mathbf{x} \leftarrow D_{sk}(\mathbf{y})$
Return 1 iff $y \notin \mathbf{y} \wedge R(x, \mathbf{x}, M, s_1)$	Return 1 iff $R(x, \mathbf{x}, M, s_1)$

We say that Π is secure in the sense of snm-atk *for if for every polynomial $q(k)$, every R computable in time $q(k)$, every $\mathcal{A}$ that runs in time $q(k)$ and outputs a valid message space M samplable in time $q(k)$, there exists a polynomial-time algorithm $\mathcal{S}$ such that $\mathsf{Adv}^{\mathsf{snm\text{-}atk}}_{\mathcal{A},\mathcal{S},\Pi}(R,k)$ is negligible.*

Technically, for our definitions to hold with respect to random oracles we would need to explicitly include a random oracle in our experiments. However, this can be done in a standard way, and for readability it is not included.

3 Weak non-malleability

3.1 Definition

Here we propose a definition for weak non-malleable (wnm) encryption schemes. As in Definition 1, a wnm-secure encryption scheme requires the existence of a simulator $\mathcal{S}$ (not given a challenge ciphertext y) that has roughly the same probability as an adversary $\mathcal{A}$ (given y) of generating a vector $\mathbf{y}$ of ciphertexts for which the plaintext vector $\mathbf{x}$ bears some relationship R with the plaintext x of y. In the wnm definition, the adversary experiment will take exactly the same form as that in Definition 1. The difference lies in the simulator experiment and the form of $\mathcal{S}$. Specifically, $\mathcal{S}$ is permitted to make each element y_i of $\mathbf{y}$

[2] Actually we slightly modify the definition of [4] so as to not require that every element of $\mathbf{y}$ decrypt to a valid plaintext. This is needed for the equivalences stated in [4] to hold.

contingent upon a "guess" z_i as to the value of x. That is, relation R tests x against a vector $\mathbf{x}$ where each element x_i is the plaintext of the corresponding y_i in $\mathbf{y}$ if either $\mathcal{S}$ guessed x or offered no guess (i.e., guessed $\perp$), and where x_i is $\perp$ otherwise.

It is easy to see that any snm-secure encryption scheme is also wnm-secure, since the wnm-simulator is simply given more power. It is perhaps not as easy to see that this power is sufficient to allow a wnm-secure scheme that is not snm-secure, but we will show that in fact this is the case. For example, the wnm-schemes presented in the introduction are not snm-secure in the random oracle model.[3]

The precise definition of wnm security is as follows.

Definition 2 (wnm-cpa, wnm-cca1, wnm-cca2). *Let $\Pi = (G, E, D)$ be an encryption scheme, let R be a relation, let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary, and let $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ be an algorithm ("simulator"). For $k \in \mathbb{N}$ define* $\mathsf{Adv}^{\mathsf{wnm\text{-}atk}}_{\mathcal{A},\mathcal{S},\Pi}(R,k) \stackrel{\text{def}}{=} \Pr[\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{A},\Pi}(R,k) = 1] - \Pr[\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{S},\Pi}(R,k) = 1]$, *where*

$\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{A},\Pi}(R,k)$:	$\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{S},\Pi}(R,k)$:
$(pk, sk) \leftarrow G(1^k)$	$(pk, sk) \leftarrow G(1^k)$
$(M, s_1, s_2) \leftarrow \mathcal{A}_1^{\mathcal{O}_1}(pk)$	$(M, s_1, s_2) \leftarrow \mathcal{S}_1(pk)$
$x \leftarrow M$	$x \leftarrow M$
$y \leftarrow E_{pk}(x)$	
$\mathbf{y} \leftarrow \mathcal{A}_2^{\mathcal{O}_2}(s_2, y)$	$(\mathbf{y}, \mathbf{z}) \leftarrow \mathcal{S}_2(s_2)$
$\mathbf{x} \leftarrow D_{sk}(\mathbf{y})$	$\mathbf{x} \leftarrow D'_{sk}(\mathbf{y}, \mathbf{z}, x)$
Return 1 iff $(y \notin \mathbf{y}) \wedge R(x, \mathbf{x}, M, s_1)$	Return 1 iff $R(x, \mathbf{x}, M, s_1)$

and $D'_{sk}(\mathbf{y}, \mathbf{z}, x)$ returns the decryption of each $y_i \in \mathbf{y}$ for which $z_i = x$ or $z_i = \perp$, and returns $\perp$ for each other index. We say that Π is wnm-atk*-secure if for every polynomial $q(k)$, and every $\mathcal{A}$ that runs in time $q(k)$ and outputs a valid message space M samplable in time $q(k)$, there exists a polynomial-time algorithm $\mathcal{S}$ such that for every R computable in time $q(k)$, $\mathsf{Adv}^{\mathsf{wnm\text{-}atk}}_{\mathcal{A},\mathcal{S},\Pi}(R,k)$ is negligible.*

The proofs of the following lemmas will appear in the full version of the paper.

Lemma 1. *For any* atk $\in \{$cpa, cca1, cca2$\}$, snm-atk $\Rightarrow$ wnm-atk $\Rightarrow$ ind-atk.

Lemma 2 (ind-cca1 $\not\Rightarrow$ wnm-cpa). *If there exists an* ind-cca1*-secure encryption scheme, then there exists an* ind-cca1*-secure encryption scheme that is not* wnm-cpa*-secure.*

Lemma 3 (wnm-cca1 $\not\Rightarrow$ snm-cpa). *If there exists an* snm-cca1*-secure encryption scheme, then there exists a* wnm-cca1*-secure encryption system that is not* snm-cpa*-secure.*

[3] Actually, it is much easier to see that they are not comparison-based non-malleable (cnm) [4], and then use the result in [4] that simulation-based non-malleability implies comparison-based non-malleability. Also, note that our separation result in Lemma 3 holds not just in the random oracle model, but in the standard model.

3.2 Constructions

In Section 1, we introduced several constructions for wnm-secure encryption, denoted "Mult-Range", "Mult-Adjacent", and "Add-Square". Our goal in this section will be to prove Lemma 4.

Lemma 4. *The Mult-Range, Mult-Adjacent, and Add-Square schemes are all* wnm-atk *secure, for* $\mathsf{atk} \in \{\mathsf{cpa}, \mathsf{cca1}\}$.

In fact, we prove a more general result. We show a general construction of weakly non-malleable encryption schemes, of which the three constructions above are special cases. We first introduce a notion called "uniquely identifiable subset," which we will use in our general construction.

We say a sequence of sets $X = \{X_k\}_{k>0}$, $X_k \subseteq \{0,1\}^*$, is *efficient* if there exists a polynomial $p(\cdot)$ such that membership in X_k can be tested in time $p(k)$. For simplicity, we often abuse notation by referring to the sequence $\{X_k\}$ as "the efficient set X" and omitting the subscript k, although it should be understood that X is a sequence of sets. We extend this notation to groups, too, i.e., when we say "X is an efficient finite group," it should be understood that $X = \{X_k\}$ is in fact a sequence of finite groups, whose membership can be efficiently determined. Furthermore, for efficient sets X and S, we use the phrase "S is a subset of X" as shorthand for "for every k, S_k is a subset of X_k."

Definition 3 (Unique Identifiability). *Let X be an efficient finite group with identity element e, and let S be an efficient subset of X. We say S is a* uniquely identifiable subset of X*, if for every $\lambda \in X \backslash \{e\}$, there exists at most one $x_\lambda \in S$, such that $\lambda \cdot x_\lambda \in S$ and for any other $x \in S, x \neq x_\lambda$, $\lambda \cdot x \notin S$. Here "$\cdot$" is the group operation. We call x_λ the* soft spot *for λ. If no such x_λ exists, we write this as $x_\lambda = \perp$. We denote the soft spot of λ by* $\mathsf{ss}(\lambda)$.

Furthermore, we say S is an efficient uniquely identifiable subset of X*, if there exists a polynomial-time algorithm A that outputs x_λ on input λ.*

Putting the definition in our context, X is the space of all messages and S is the set of all "valid" messages. The group operation "$\cdot$" is a "mauling" function the converts an encryption of x to an encryption of $\lambda \cdot x$, and we call λ the "mauling factor." The unique identifiability indicates, therefore, for every mauling factor λ, there is at most one valid message x_λ that can be mauled into another valid one (all other valid messages are mapped to invalid ones). For an efficient uniquely identifiable subset, one can in fact find x_λ efficiently.

Next, we give several examples of efficient uniquely identifiable subsets, which are closely related to the Mult-Range, Mult-Adjacent, and the Add-Square schemes.

Example 1 (Mult-Adjacent). Let $X = \mathbb{Z}_p^* \times \mathbb{Z}_p^*$ with the group operation being component-wise multiplication in $\mathbb{Z}_p^*$, i.e., $(x_0, x_1) \cdot (y_0, y_1) = (x_0 \cdot y_0,\ x_1 \cdot y_1)$. Let $S = \{(x, x+1) \mid x \in \mathbb{Z}_p^*\}$.

Example 2 (Add-Square). Let $X = \mathbb{Z}_p \times \mathbb{Z}_p$, with the group operation being component-wise addition in $\mathbb{Z}_p$, i.e., $(x_0, x_1) \cdot (y_0, y_1) = (x_0 + y_0,\ x_1 + y_1)$. Let $S = \{(x, x^2) \mid x \in \mathbb{Z}_p\}$.

Example 3 (Mult-Range). Let $X = \mathbb{Z}_p^*$ with multiplication as the group operation. Let $S = \{a, ..., b\}$, where $a > (b-a)^2$ and $p > 2b^2$.

Lemma 5. *All three examples above are efficient uniquely identifiable systems.*

The proof of Lemma 5 is straightforward for Mult-Adjacent and Add-Square; Mult-Range is not straightforward, however. The proof will be provided in the full version of the paper.

Now we present our general construction of wnm encryption schemes.

Construction 1 *Let X be an efficient finite group. Let S be an efficient uniquely identifiable subset of X, and $H : \{0,1\}^* \to X$ be a hash function. Let G_* be a trapdoor permutation generator. We construct an encryption scheme as follows. G runs G_* to get (f, f^{-1}, d), and sets $pk = <f, d>$, and $sk = f^{-1}$. The plaintext space of E_{pk} is S.[4] To encrypt a message m, $E_{pk}(m)$ generates $r \leftarrow d(1^k)$ and returns $<f(r), H(r) \cdot m>$, where "$\cdot$" is the group operation in X. To decrypt a ciphertext $c = (\alpha, \beta)$, $D_{sk}(c)$ computes $m = \beta \cdot (H(f^{-1}(\alpha))^{-1})$, returns m if $m \in S$, and $\perp$ otherwise.*

Lemma 6. *Following the notation in Construction 1, if $H(\cdot)$ is a random oracle, then Construction 1 is* wnm-atk *secure, for* atk $\in \{$cpa, cca1$\}$.

The proof of this result is in Appendix A.1.

3.3 Applications

In this section we show that weak non-malleability suffices to implement a secure contract bidding system between two bidders. Intuitively, in an ideal contract bidding system, each of two bidders would submit its bid to a trusted authority through a secure channel (so that the messages are both secret and authenticated). In a real contract bidding system, however, it may be the case that a dishonest bidder may be able to see the encrypted bid from an honest bidder before it submits its own bid. In either case, we assume there is a public "award" function over these input bids. Depending on the application, the award function varies. For example, the simplest award function can be $\mathsf{Award}((x_0, x_1)) = (y_0, y_1)$, where $y_i = x_i$ if $x_i = \min\{x_0, x_1\}$ and $y_i = 0$ otherwise. This indicates the rule that the lowest bidder wins, with the award being his bid, and the other bidder loses and thus has zero award. (We assume a unique minimum between the bids, otherwise nobody wins.) Other forms of the award function exist.

We specify our contract bidding system as follows:

Setup: A bidding system consisting of two bidders B_0, B_1 and an award function Award. There is also a bidding upper bound $U > 0$, such that the only valid bids are integers between 0 and U. Both bidders are given U and the award function Award.

[4] More precisely, we assume a one-to-one correspondence between plaintexts and elements of S, and efficient encoding and decoding functions to map plaintexts to and from elements of S.

Award function: The function $\mathsf{Award} : \{\perp, 0, 1..., U\}^2 \rightarrow \{0, 1, ..., U\}^2$ takes the bids from the bidders and computes their awards, respectively.[5] We say an award function is *fair*, if for any $\mathbf{x} = (x_0, x_1)$ and any $i \in \{0, 1\}$, $\mathsf{Award}(\mathbf{x}|_{\perp \rightarrow [i]})[i] \leq \mathsf{Award}(\mathbf{x})[i]$, and $\mathsf{Award}(\mathbf{x}|_{x_{1-i} \rightarrow [i]})[i] \leq \mathsf{Award}(\mathbf{x})[i]$. Here we use $\mathbf{x}|_{y \rightarrow [i]}$ to indicate the vector obtained by replacing the ith entry of $\mathbf{x}$ by y and we use $\mathbf{x}[i]$ to indicate the ith entry of $\mathbf{x}$. Intuitively, the fairness indicates that B_i would not gain any advantage in profit by changing his bid to $\perp$ or to B_{1-i}'s bid. We note that fairness is a reasonable requirement for bidding systems to be "useful."

Real Adversary: To model security, we consider an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that corrupts bidder B_1. $\mathcal{A}_1$ receives the public key pk and U, and outputs a polynomial-time samplable distribution M of bids, from which a bid bid_0 is chosen for B_0. $\mathcal{A}_2$ is then given the ciphertext of bid_0 and outputs encrypted bid ebid_1 for B_1. The profit of the adversary is the award of B_1.

Definition 4 (Secure Contract Bidding). *Let* CBS *be a contract bidding system with bidding upper bound U and encryption scheme $\Pi = (G, E, D)$.* CBS *is secure if for every fair award function* Award*, every polynomial $q(k)$, every adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that runs in time $q(k)$, there exists a polynomial-time simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ such that $\mathsf{Adv}^{\mathsf{profit}}_{\mathcal{A},\mathcal{S},\mathsf{CBS}}(k)$ is negligible,[6] where $\mathsf{Adv}^{\mathsf{profit}}_{\mathcal{A},\mathcal{S},\mathsf{CBS}}(k) \stackrel{\text{def}}{=} E[\mathsf{Expt}^{\mathsf{real}}_{\mathcal{A},\mathsf{CBS}}(k) - \mathsf{Expt}^{\mathsf{ideal}}_{\mathcal{S},\mathsf{CBS}}(k)]$, and*

$\mathsf{Expt}^{\mathsf{real}}_{\mathcal{A},\mathsf{CBS}}(k)$:	$\mathsf{Expt}^{\mathsf{ideal}}_{\mathcal{S},\mathsf{CBS}}(k)$:
$(pk, sk) \leftarrow G(1^k)$	
$(M, s) \leftarrow \mathcal{A}_1(pk, U)$	$(M, s) \leftarrow \mathcal{S}_1(U)$
$\mathsf{bid}_0 \leftarrow M$	$\mathsf{bid}_0 \leftarrow M$
$\mathsf{ebid}_0 \leftarrow E_{pk}(\mathsf{bid}_0)$	
$\mathsf{ebid}_1 \leftarrow \mathcal{A}_2(\mathsf{ebid}_0, s)$	$\mathsf{bid}_1 \leftarrow \mathcal{S}_2(s)$
$\mathsf{bid}_1 \leftarrow D_{sk}(\mathsf{ebid}_1)$	
return $\mathsf{Award}((\mathsf{bid}_0, \mathsf{bid}_1))[1]$	return $\mathsf{Award}((\mathsf{bid}_0, \mathsf{bid}_1))[1]$

It is clear that if the encryption scheme Π is malleable, then the system might not be secure. For example, consider the scheme where message m is encrypted as $<f(r), H(r) + m \bmod p>$, where $f(\cdot)$ is a trapdoor permutation and H a random oracle. It is an ind-cpa-secure scheme, but the real bidding system is not secure, since an adversary seeing the bid $<\alpha, \beta>$ from bidder B_0 can submit bid $<\alpha, \beta - 1>$, and underbid B_0 by 1. It is also obvious that if Π is snm-cpa-secure, then the bidding system is secure. The next theorem shows that in fact wnm-cpa-security suffices.

Theorem 1. *Let $\Pi = (G, E, D)$ be a* wnm-cpa*-secure encryption scheme, with a domain that includes the integer range $[0, U]$ where U is polynomially bounded by k. Then a contract bidding system* CBS *with bidding upper bound U and encryption scheme Π is secure.*

[5] We insist that the award function be a positive function. However, this is entirely arbitrary, since one can always "shift" the award function by a constant without changing its nature.

[6] It may be negative, in which case we also consider it to be negligible.

The proof of this result is in Appendix A.2. We mention that our result only applies to the case of a single auction, and specifically does not claim that repeated auctions will be secure if they use the same encryption scheme. Obviously, for repeated auctions to be secure, we would need some kind of cca2 security for our encryption scheme.

We also mention that the result does not apply to contract bidding schemes with multiple bidders that may collude. Intuitively, this is because they may each make guesses which cover the possible choices of the honest bidder, and a wrong guess for one party does not reduce the award of the party that guesses correctly. To solve the problem with multiple bidders using a wnm-secure cryptosystem, one could either allow *randomization* in the bids (e.g., each bid would be of the form (bid, r), where $r \leftarrow \{0,1\}^k$, which would ensure that the adversary has a negligible chance of guessing the full plaintext), or one could change the model to levy *penalties* for invalid bids.

4 Tag-based non-malleability

In this section, we introduce tag-based non-malleability as an alternative to standard non-malleability. Informally, in a tag-based encryption system, the encryption and decryption operations take an additional "tag." A tag is simply a binary string of appropriate length (i.e., its length has to be polynomially bounded by the security parameter), and need not have any particular internal structure. We define security for tag-based encryption in manners analogous to security for standard encryption systems. In particular, we define *tag-based non-malleability* (Definition 5) and *tag-based indistinguishability* (Definition 6) with respect to cpa, cca1, and cca2 attacks. The only changes we make to the definitions for standard encryption are: (i) in a cca2 attack, instead of requiring that the adversary $\mathcal{A}$ not query the decryption oracle with the ciphertext y that $\mathcal{A}$ receives as a challenge, we require that $\mathcal{A}$ not query the decryption oracle with a (ciphertext,tag) pair using the same tag with which y was encrypted; (ii) for tag-based non-malleability, instead of requiring that $\mathcal{A}_2$ not output the ciphertext y it receives, we require that $\mathcal{A}_2$ not output any (ciphertext,tag) pair for decryption using the tag with which y was encrypted. Informally, one simply changes the "equality of two ciphertexts" in the standard definitions to "equality of *the tags of* two ciphertexts," and we have a tag-based definition.

4.1 Definition

Tag-based encryption schemes A *tag-based encryption scheme* Π is a triple (G, E, D) of algorithms, the first two being probabilistic, and all running in expected polynomial time. G takes as input 1^k and outputs a public key pair (pk, sk), i.e., $(pk, sk) \leftarrow G(1^k)$. E takes a public key pk, a message m, and a tag t as input and outputs an encryption c for m associated with t; we denote this $c \leftarrow E_{pk}(m, t)$. D takes a private key sk, a ciphertext c, and a tag t as input and returns either a message m such that c is a valid encryption of m associated with t, if such an m exists, and otherwise returns $\perp$; we denote this $m \leftarrow D_{sk}(c, t)$.

Definition 5 (tnm-cpa, tnm-cca1, tnm-cca2). *Let $\Pi = (G, E, D)$ be an encryption scheme, let R be a relation, let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary, and let $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ be an algorithm (the "simulator"). For $k \in \mathbb{N}$ define* $\mathsf{Adv}^{\mathsf{tnm\text{-}atk}}_{\mathcal{A},\mathcal{S},\Pi}(R,k) \stackrel{\text{def}}{=} \Pr[\mathsf{Expt}^{\mathsf{tnm\text{-}atk}}_{\mathcal{A},\Pi}(R,k) = 1] - \Pr[\mathsf{Expt}^{\mathsf{tnm\text{-}atk}}_{\mathcal{S},\Pi}(R,k) = 1]$, *where*

$\mathsf{Expt}^{\mathsf{tnm\text{-}atk}}_{\mathcal{A},\Pi}(R,k)$:	$\mathsf{Expt}^{\mathsf{tnm\text{-}atk}}_{\mathcal{S},\Pi}(R,k)$:
$(pk, sk) \leftarrow G(1^k)$	$(pk, sk) \leftarrow G(1^k)$
$(M, t, s_1, s_2) \leftarrow \mathcal{A}_1^{\mathcal{O}_1}(pk)$	$(M, t, s_1, s_2) \leftarrow \mathcal{S}_1(pk)$
$x \leftarrow M$	$x \leftarrow M$
$y \leftarrow E_{pk}(x, t)$	
$(\mathbf{y}, \mathbf{t}) \leftarrow \mathcal{A}_2^{\mathcal{O}_2}(s_2, y, t)$	$(\mathbf{y}, \mathbf{t}) \leftarrow \mathcal{S}_2(s_2, t)$
$\mathbf{x} \leftarrow D_{sk}(\mathbf{y}, \mathbf{t})$	$\mathbf{x} \leftarrow D_{sk}(\mathbf{y}, \mathbf{t})$
Return 1 iff $(t \notin \mathbf{t}) \wedge R(x, \mathbf{x}, M, s_1)$	Return 1 iff $R(x, \mathbf{x}, M, s_1)$

We require that $\mathcal{O}_2$ not be queried with the t given to $\mathcal{A}_2$. We say that Π is secure in the sense of tnm-atk *if for every polynomial $q(k)$, every R computable in time $q(k)$, and every $\mathcal{A}$ that runs in time $q(k)$ and outputs a valid message space M samplable in time $q(k)$, there exists a polynomial-time algorithm $\mathcal{S}$ such that $\mathsf{Adv}^{\mathsf{tnm\text{-}atk}}_{\mathcal{A},\mathcal{S},\Pi}(R,k)$ is negligible.*

Definition 6 (tind-cpa,tind-cca1,tind-cca2). *Let $\Pi = (G, E, D)$ be a tag-based encryption scheme, and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary. For $k \in \mathbb{N}$ define* $\mathsf{Adv}^{\mathsf{tind\text{-}atk}}_{\mathcal{A},\Pi}(k) \stackrel{\text{def}}{=} 2 \cdot \Pr[\mathsf{Expt}^{\mathsf{ind\text{-}atk}}_{\mathcal{A},\Pi}(k) = 1] - 1$ *where*

$\mathsf{Expt}^{\mathsf{tind\text{-}atk}}_{\mathcal{A},\Pi}(k)$:
$(pk, sk) \leftarrow G(1^k)$
$(x_0, x_1, t, s) \leftarrow \mathcal{A}_1^{\mathcal{O}_1}(pk)$
$b \stackrel{R}{\leftarrow} \{0,1\}$
$y \leftarrow E_{pk}(x_b, t)$
$b' \leftarrow \mathcal{A}_2^{\mathcal{O}_2}(x_0, x_1, t, s, y)$
Return 1 iff $b = b'$

We require that $|x_0| = |x_1|$, and that $\mathcal{O}_2$ is not queried with tag t. We say that Π is secure in the sense of tind-atk *if for every polynomial $q(k)$ and every adversary $\mathcal{A}$ that runs in time $q(k)$, $\mathsf{Adv}^{\mathsf{tind\text{-}atk}}_{\mathcal{A},\Pi}(k)$ is negligible.*

Theorem 2 (tnm-atk $\Rightarrow$ tind-atk). *If an encryption scheme is* tnm-atk*-secure, then it is* tind-atk*-secure, for* atk $\in \{$cpa, cca1, cca2$\}$.

4.2 Constructions

We give two constructions of tag-based encryption schemes, both achieving tnm-cca2-security. The first one is based one-way trapdoor permutations in the random oracle model. It is similar to the semantically secure (ind-cpa) encryption scheme from Bellare and Rogaway [2], but enjoys a higher level of security. The second is a modification of the Cramer-Shoup scheme [8, 9], but simpler.

Construction 2 *Let G_* be a trapdoor permutation generator. Let X be a finite group and $H : \{0,1\}^* \to X$ a hash function. We construct an encryption scheme as follows. G runs G_* to get (f, f^{-1}, d), and sets $pk = <f, d>$, and $sk = f^{-1}$. All messages are restricted to be elements in X. To encrypt a message m with tag t, $E_{pk}(m)$ generates $r \leftarrow d(1^k)$ and returns $<f(r), H(r,t) \cdot m>$, where "$\cdot$" is the group operation in X. To decrypt a ciphertext $c = (\alpha, \beta)$, $D_{sk}(c)$ returns $m = \beta \cdot (H(f^{-1}(\alpha), t)^{-1})$.*

Lemma 7. *Let H be a random oracle. If f is a trapdoor permutation, then the scheme in Construction 2 is* tnm-cca2*-secure.*

The intuition behind the proof of Lemma 7 is that the simulator can simulate the decryption oracles using the knowledge of the random oracle queries, and the fact that the adversary cannot make an $\mathcal{O}_2(\cdot)$ query with the same tag as in the challenge encryption. Details will be given in the full version.

Construction 3 *Let G_q be a finite group in which the DDH assumption holds.*[7] *We define an encryption scheme as follows.*

$G_{CS}(G_q)$**:** *Let g be the generator of G_q (included in the description of G_q). Generate $g_2 \stackrel{R}{\leftarrow} G_q$ and $a, b, c, d, e \stackrel{R}{\leftarrow} \mathbb{Z}_q$, and set $U \leftarrow g^a (g_2)^b$, $V \leftarrow g^c (g_2)^d$, and $W \leftarrow g^e$. Let the public key be $<g, g_2, U, V, W>$ and the secret key be $<a, b, c, d, e>$.*

$E_{<g,g_2,U,V,W>}(m,t)$**:** *Generate $r \stackrel{R}{\leftarrow} \mathbb{Z}_q$ and $x \leftarrow g^r$, $y \leftarrow (g_2)^r$, $w \leftarrow W^r m$, and $v \leftarrow U^r V^{rt}$. Return $<x, y, w, v>$ as the ciphertext.*

$D_{<a,b,c,d,e>}(<x,y,w,v>,t)$**:** *If $v \neq x^{a+ct} y^{b+dt}$, return $\perp$, else return w/x^e.*

Informally, our construction removes the collision-resistant hash function from the original Cramer-Shoup construction, and replaces the hash value $\alpha = H(x, y, w)$ by the tag t.[8]

Lemma 8. *The encryption scheme in Construction 3 is* tnm-cca2*-secure.*

The proof of this lemma almost directly follows the proof of security for the original Cramer-Shoup construction; we omit it here.

4.3 Applications

Intuitively, tag-based encryption schemes (and in particular, tnm schemes) are useful in systems that already have authentication, i.e., in systems where Bob cannot impersonate Alice and send messages using Alice's identity. We stress that even with authentication, we still need non-malleability. For example, in the contract-bidding scenario in both [11] and the previous section, we still need

[7] Note that one possible group G_q may be found by generating a large prime p such that q divides $p-1$, and letting G_q be the subgroup of order q in $\mathbb{Z}_p^*$.

[8] We assume that $t \in \mathbb{Z}_q$. Otherwise, we would need a collision-resistant hash function to hash the tag.

to make sure that Bob cannot underbid Alice by mauling her message. With a tnm system, we can use the sender's identity as the tag to achieve this goal. Suppose Alice sends a encrypted message $c = E_{pk}(m, \mathsf{Alice})$ to Charlie. A malicious Bob may be able to maul c into another ciphertext with the same tag, i.e., Alice — this is allowed in the definition — but this would not be useful for him since he cannot fake Alice's identity. Bob needs to produce some message with the tag Bob, but tnm stipulates that Bob will not have any advantage in doing so. To demonstrate this, we show how to use a tnm-cca2 scheme (in fact, a tind-cca2 scheme) to construct a protocol that realizes the *secure message transmission* functionality in the $\mathcal{F}_{\text{AUTH}}$-hybrid model, in the universal composability framework. Previously, this was done using an ind-cca2 encryption scheme [5].

Universal-composability framework The universal composability framework was proposed by Canetti [5] for defining the security and composition of protocols. To define security one first specifies an *ideal functionality* using a trusted party that describes the desired behavior of the protocol. Then one proves that a particular protocol operating in a real-life model securely realizes this ideal functionality. Here we briefly summarize the framework.

A (real-life) protocol π is defined as a set of n interactive Turing Machines $P_1, \ldots, P_n$, designating the n parties in the protocol. It operates in the presence of an environment $\mathcal{Z}$ and an adversary $\mathcal{A}$, both of which are also modeled as interactive Turing Machines. The environment $\mathcal{Z}$ provides inputs and receives outputs from honest parties, and may communicate with $\mathcal{A}$. $\mathcal{A}$ controls (and may view) all communication between the parties. (Note that this models asynchronous communication on open point-to-point channels.) We will assume that messages are authenticated, and thus $\mathcal{A}$ may not insert or modify messages between honest parties. (This feature could be added to an unauthenticated model using a message authentication functionality as described in [5].) $\mathcal{A}$ also may corrupt parties, in which case it obtains the internal state of the party.

The ideal process with respect to a functionality $\mathcal{F}$, is defined for n parties $P_1, \ldots, P_n$, an environment $\mathcal{Z}$, and an (ideal-process) adversary $\mathcal{S}$. However, $P_1, \ldots, P_n$ are now dummy parties that simply forward (over secure channels) inputs received from $\mathcal{Z}$ to $\mathcal{F}$, and forward (again over secure channels) outputs received from $\mathcal{F}$ to $\mathcal{Z}$. Thus the ideal process is a trivially secure protocol with the input-output behavior of $\mathcal{F}$.

UC secure message transmission The functionality $\mathcal{F}_{\text{M-SMT}}$ is given in Figure 1. Intuitively, this functionality allows multiple parties to send messages securely to a single receiver. Both the secrecy and the integrity of the messages are guaranteed. See [5] for more discussions.

Canetti [5] constructed a protocol that securely realizes this functionality in the $(\mathcal{F}_{\text{AUTH}}, \mathcal{F}_{\text{PKE}})$-hybrid model. He also showed that any ind-cca2 encryption scheme can securely realize the $\mathcal{F}_{\text{PKE}}$ functionality. Therefore, one can construct a protocol using an ind-cca2 encryption scheme to securely realize $\mathcal{F}_{\text{M-SMT}}$ in

$\mathcal{F}_{\text{M-SMT}}$ proceeds as follows, running with parties $P_1, \ldots, P_n$, and an adversary $\mathcal{A}$:

- In the first activation, expect to receive a value $(\textbf{receiver}, id)$ from some party P_i. Then send $(\textbf{receiver}, id, P_i)$ to the all parties and the adversary. ¿From now on, ignore all $(\textbf{receiver}, id)$ values.
- Upon receiving a value $(\textbf{send}, id, m)$ from some party P_j, send (id, P_j, m) to P_i and $(id, P_j, |m|)$ to the adversary.

Fig. 1. Functionality $\mathcal{F}_{\text{M-SMT}}$

the $\mathcal{F}_{\text{AUTH}}$-hybrid model. Here, we show that one can instead use a tag-based tind-cca2 encryption scheme.

Given a tind-cca2 encryption scheme $\Pi = (G, E, D)$, the protocol σ runs as follows. In this description, we include the identity of the receiver in the session identifier. (i) When a party P_i receives an input $(\textbf{receiver}, id|P_i)$, it runs $(pk, sk) \leftarrow G(1^k)$, and sends $(\mathsf{key}, id|P_i, pk)$ to all other parties using $\mathcal{F}_{\text{AUTH}}$. Any messages of this type with an identifier not in the correct format are ignored. (ii) On receiving the first message $(P_{i'}, P_j, (\mathsf{key}, id|P_{i'}, pk'))$ from $\mathcal{F}_{\text{AUTH}}$, P_j records $(P_{i'}, id, pk')$ and outputs $(\textbf{receiver}, id|P_{i'}, P_{i'})$. Any messages of this type with an identifier not in the correct format are ignored. Subsequent messages of this type with identifier $id|P_{i'}$ are ignored. (iii) After this, when P_j receives an input $(\textbf{send}, id|P_{i'}, m)$, P_j runs $c \leftarrow E_{pk'}(m, P_j)$, and invokes $\mathcal{F}_{\text{AUTH}}$ to send $(\mathsf{msg}, id|P_{i'}, c)$ to $P_{i'}$. (iv) On receiving a message $(P_j, P_i, (\mathsf{msg}, id|P_i, c))$ from $\mathcal{F}_{\text{AUTH}}$, P_i runs $m \leftarrow D_{sk}(c, P_j)$ and if $m \neq \perp$, outputs $(id|P_i, P_j, m)$. Intuitively, the protocol uses the identity of the senders as the tag for the encryption.

Theorem 3. *The protocol σ securely realizes the SMT functionality in the $\mathcal{F}_{\text{AUTH}}$ hybrid model, assuming static corruptions.*

The proof of Theorem 3 is in Appendix A.3.

4.4 Relation to standard definitions

We study the relation between the tag-based definitions and the standard ones. First, we note that they are not directly comparable, due to the structural difference in encryption and decryption. However, given a standard encryption scheme $\Pi = (G, E, D)$, it is straightforward to construct a tag-based scheme $\Pi' = (G', E', D')$ with the same security as follows. G' is the same as G; $E'_{pk}(m, t)$ calls $E_{pk}(m \circ t)$, where $x \circ y$ denotes a canonical encoding of the concatenation of two binary strings that can be uniquely parsed; $D'_{sk}(c, t)$ calls $(m, t') \leftarrow D_{sk}(c)$ to and returns m if $t = t'$ and $\perp$ otherwise. It is easy to check that Π' enjoys the same level of security (in the sense of Definition 5) as Π (in the sense of Definition 1).

Interestingly, the other direction also holds: given a tag-based scheme, one can construct a standard scheme, using a strong one-time signature scheme [20].

Construction 4 *Let* $\Pi = (G, E, D)$ *be a tag-based encryption scheme. Let* $\mathsf{SIG} = (\mathsf{sig_gen}, \mathsf{sig_sign}, \mathsf{sig_verify})$ *be a strong one-time signature scheme. We construct a standard scheme* $\Pi' = (G', E', D')$ *as follows.* $G' = G$*. To encrypt massage* m *using* pk*, generate a signing/verification key pair* $(\mathsf{sig_vk}, \mathsf{sig_sk}) \leftarrow \mathsf{sig_gen}(1^k)$*; encrypt* m *using* $\mathsf{sig_vk}$ *as the tag, i.e.,* $c \leftarrow E_{pk}(m, \mathsf{sig_vk})$*; sign* c *using* $\mathsf{sig_sk}$*, i.e.,* $s \leftarrow \mathsf{sig_sign}(sk, c)$*; and output* $(\mathsf{sig_vk}, c, s)$ *as the encryption. To decrypt a ciphertext* $(\mathsf{sig_vk}, c, s)$*, verify that* s *is a valid signature of* c *with respect to* $\mathsf{sig_vk}$*; if not, output* $\perp$*; if so, return* $D_{sk}(c, \mathsf{sig_vk})$*.*

Theorem 4. *For* $\mathsf{atk} \in \{\mathsf{cpa}, \mathsf{cca1}, \mathsf{cca2}\}$*: if* Π *is* $\mathsf{tnm\text{-}atk}$ *secure, then* Π' *is* $\mathsf{snm\text{-}atk}$ *secure; and if* Π *is* $\mathsf{tind\text{-}atk}$ *secure, then* Π' *is* $\mathsf{ind\text{-}atk}$ *secure.*

The proof of this result will be given in the full version.

Construction 4 is essentially the construction first shown in [11] and later used in [20, 10, 17, 6] to obtain non-malleable encryption schemes, except that we explicitly separate the underlying tag-based scheme from the "wrapper" that uses the one-time signature scheme. Thus, in each of these papers, there is an implicit tag-based non-malleable encryption scheme.[9] We illustrate this with the scheme of Lindell [17], which we denote as Π_L. In Π_L, an encryption of message m is a tuple $<c_0, c_1, pk_0, pk_1, r, \mathsf{sig_vk}, \sigma, s>$. Here c_0 and c_1 are two encryptions of m using two $\mathsf{ind\text{-}cpa}$ systems with public keys pk_0 and pk_1, respectively; $\mathsf{sig_vk}$ is a "fresh" verification key of a strong one-time signature scheme; r is a random string; σ is an NIZK proof that either c_0 and c_1 are the encryption of the same message, or r is the commitment of $\mathsf{sig_vk}$; s is a signature of the tuple $<c_0, c_1, pk_0, pk_1, \mathsf{sig_vk}, r, \sigma>$. Then in the underlying tag-based encryption scheme Π, an encryption of message m with tag t is the tuple $<c_0, c_1, pk_0, pk_1, r, t, \sigma>$, where c_0, c_1, pk_0, pk_1, and r are all the same as before, and σ becomes an NIZK proof that either c_0 and c_1 are the encryptions of the same message, or r is the commitment of t. It is easy to verify that Π is $\mathsf{tnm\text{-}cca2}$-secure. In fact, one can prove the security for Π almost exactly the same way as for the security proof of Π_L, observing that the use of the strong one-time signature in Π_L is solely for enforcing that an adversary will not make a query to the decryption oracle with a ciphertext having the same verification key. Since in the tag-based system Π, the verification key is replaced by the tag, by definition, the adversary cannot query the decryption oracle with a ciphertext having the same tag. So in fact the proof for the security of Π is even simpler than the proof for for Π_L. Furthermore, Π_L is exactly the transformed version of protocol Π under Construction 4. Therefore, one could obtain an alternative proof of security for Π_L by pluggin Π into Theorem 4.

In the full version of this paper, we will show that many known relations between standard security definitions translate to analogous relations between tag-based definitions.

[9] In the (independent and concurrent) result of [6], there is actually an explicit *identity-based encryption (IBE) scheme* which corresponds to our tag-based non-malleable encryption scheme. They essentially prove the cca2 case of Theorem 4. (Note: their cpa-secure IBE scheme corresponds to our cca2-secure tnm scheme.)

References

1. M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations among notions of security for public-key encryption schemes. In *Advances in Cryptology—CRYPTO '98* (Lecture Notes in Computer Science 1462), pp. 26–45, 1998.
2. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In 1st *ACM Conf. on Comp. and Comm. Security*, pp. 62–73, 1993.
3. M. Bellare and P. Rogaway. Optimal asymmetric encryption. In *Advances in Cryptology—EUROCRYPT '94* (Lecture Notes in Computer Science 950), pp. 92–111, 1995.
4. M. Bellare and A. Sahai. Non-Malleable Encryption: Equivalence between Two Notions, and an Indistinguishability-Based Characterization. In *Advances in Cryptology—CRYPTO '99* (Lecture Notes in Computer Science 1666), pp. 519–536, 1999.
5. R. Canetti. Universally Composable Security: A new paradigm for cryptographic protocols. `http://eprint.iacr.org/2000/067`, Extended abstract in 42nd FOCS, 2001.
6. R. Canetti, S. Halevi and J. Katz. Chosen-Ciphertext Security from Identity-Based Encryption. In *ePrint archive*, Report 2003/182.
7. R. Canetti, Y. Lindell, R. Ostrovsky and A. Sahai. Universally composable two-party computation. In *STOC 02*, pp. 494–503, 2002. The full version in *ePrint archive*, Report 2002/140.
8. R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *Advances in Cryptology—CRYPTO '98* (Lecture Notes in Computer Science 1462), pp. 13–25, 1998.
9. R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *Advances in Cryptology—EUROCRYPT 2002* (Lecture Notes in Computer Science 2332), pp. 45–64, 2002.
10. A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano and A. Sahai. Robust non-interactive zero knowledge. In *Advances in Cryptology – CRYPTO 2001* (LNCS 2139), 566–598, 2001.
11. D. Dolev, C. Dwork and M. Naor. Non-malleable cryptography. *SIAM J. on Comput.*, 30(2):391–437, 2000. An earlier version appeared in *23rd ACM Symp. on Theory of Computing*, pp. 542–552, 1991.
12. T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. on Info. Theory* 31:469–472, 1985.
13. S. Even, O. Goldreich, and S. Micali. On-line/Off-line digital signatures. *J. Cryptology* 9(1):35-67 (1996).
14. S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences* 28:270–299, 1984.
15. S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. on Computing* 17(2):281–308, Apr. 1988.
16. C. Rackoff and D. Simon. Noninteractive zero-knowledge proof of knowledge and chosen ciphertext attack. In *Advances in Cryptology—CRYPTO '91*, (Lecture Notes in Computer Science 576), pp. 433–444, 1991.
17. Y. Lindell. A simpler construction of CCA2-secure public-key encryption under general assumption. In *Advances in Cryptology—EUROCRYPT 2003*, (LNCS 2656), pp. 241–254, 2003.

18. R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystem. *Comm. of the ACM* 21(2):120–126, Feb. 1978.
19. A. M. Rockett and P. Szüsz. Continued Fractions. World Scientific Publishing Co. Pte. Ltd. ISBN 981-02-1047-7, 1992.
20. A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th IEEE Symp. on Foundations of Computer Sci.*, pp. 543–553, 1999.
21. V. Shoup and R. Gennaro. Securing threshold cryptosystems against chosen ciphertext attack. In *Advances in Cryptology—EUROCRYPT '98* (Lecture Notes in Computer Science 1403), pp. 1–16, 1998.

A Proofs

A.1 Proof of Lemma 6

We prove the lemma for $\mathsf{atk} = \mathsf{cca1}$, which will imply the case $\mathsf{atk} = \mathsf{cpa}$. We use the notation of Definition 2 and Construction 1.

For an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, we construct simulators $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ as follows. The simulator $\mathcal{S}_1$ runs $\mathcal{A}_1$ and simulates both the random oracle $H(\cdot)$ and the decryption oracle D_{sk} in a quite standard way. More specifically, $\mathcal{S}_1$ maintains a "query list" L consisting of pairs (α, t), such that $H(f^{-1}(\alpha)) = t$. L is initially $\emptyset$. When $\mathcal{A}_1$ makes a query r to H, $\mathcal{S}_1$ checks if $(f(r), t) \in L$ for some t, and replies with t if so; otherwise, $\mathcal{S}_1$ picks a random $t \leftarrow X$, adds $(f(r), t)$ to L, and replies with t. When $\mathcal{A}_1$ makes a query $y = (\alpha, \beta)$ to D_{sk}, $\mathcal{S}_1$ checks if $(\alpha, t) \in L$ for some t, and replies with $\psi(\beta \cdot t^{-1})$ if so; otherwise, $\mathcal{S}_1$ picks a random $t \leftarrow X$, adds (α, t) to L, and replies with $\psi(\beta \cdot t^{-1})$. Finally, when $\mathcal{A}_1$ outputs (M, s_1, s_2), $\mathcal{S}_1$ outputs $(M, s_1, (s_2, L))$.

Upon invocation, the simulator $\mathcal{S}_2$, generates $r \leftarrow d(1^k)$, $\alpha \leftarrow f(r)$, $\beta \leftarrow X$, $y \leftarrow (\alpha, \beta)$. Then $\mathcal{S}_2$ invokes $\mathcal{A}_2$ with parameters (s_2, y), and simulates the random oracle for $\mathcal{A}_2$ in the same way as $\mathcal{S}_1$ does for $\mathcal{A}_1$, using the list L passed from $\mathcal{S}_1$. When $\mathcal{A}_2$ outputs $\mathbf{y} \leftarrow \mathcal{A}_2(s_2, y)$, $\mathcal{S}_2$ aborts if $y \in \mathbf{y}$. Otherwise, we assume that $\mathbf{y} = (y_1, y_2, .., y_\ell)$, where $y_i = (\alpha_i, \beta_i)$ for $i = 1, 2, ..., \ell$. $\mathcal{S}_2$ generates $\mathbf{z}$ as follows. For each i, if $\alpha_i \neq \alpha$, then set $z_i \leftarrow \perp$; otherwise compute $\lambda_i \leftarrow \beta_i \cdot (\beta)^{-1}$, compute its soft spot $x_i \leftarrow \mathsf{ss}(\lambda_i)$, and then set $z_i \leftarrow x_i$. Finally $\mathcal{S}_2$ sets $\mathbf{z} = (z_1, z_1, ..., z_\ell)$, and outputs $(\mathbf{y}, \mathbf{z})$.

Next, we prove that $\mathcal{S}$ is a valid simulator, i.e., that $\Pr[\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{A},\Pi}(R,k) = 1] - \Pr[\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{S},\Pi}(R,k) = 1]$ is negligible in k. In order to do so, we introduce a new experiment called Mix. Informally, $\mathsf{Mix}^{\mathsf{wnm\text{-}atk}}_{\mathcal{A},\mathcal{S},\Pi}(R,k)$ is the same as $\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{A},\Pi}(R,k)$, except using the simulator $\mathcal{S}$ to simulator the random oracle and the decryption. Now let $p_{\mathcal{A}}$, p_{Mix} and $p_{\mathcal{S}}$ be the probabilities of success in the real, Mix, and ideal experiments, respectively. We shall prove that both $p_{\mathcal{A}} - p_{\mathsf{Mix}}$ and $p_{\mathsf{Mix}} - p_{\mathcal{S}}$ are negligible k, and the lemma will follow directly.

To see that $p_{\mathcal{A}} - p_{\mathsf{Mix}}$ is negligible, note that the simulation of the random oracle and decryption in Mix will be valid except in the case where $\mathcal{A}_1$ has queried H with r or D_{sk} with (α, β') for some β'. Since r and α are chosen randomly after $\mathcal{A}_1$ is executed, the probability of this is obviously negligible.

To see that $p_{\mathsf{Mix}} - p_{\mathcal{S}}$ is negligible, note that the two experiments only differ when $\mathcal{A}_2$ queries $H(r)$, and the probability of this is negligible by the security of f. (Using unique identifiability and by viewing random oracle queries, $\mathcal{S}_2$ is able to simulate the decryption exactly.)

Details will be provided in the full version of the paper.

A.2 Proof of Theorem 1

First the intuition. If CBS were not secure, then there would be an adversary $\mathcal{A}$ that breaks it, meaning that for some fair Award, no simulator could achieve an expected award negligibly less than $\mathcal{A}$. But we will construct an adversary $\mathcal{B}$ for Π out of $\mathcal{A}$, and by the wnm-security of Π, there is a simulator $\mathcal{S}'$ that approximates $\mathcal{B}$ for any relation. Then we will use $\mathcal{S}'$ to build a simulator $\mathcal{S}$ for CBS that does achieve an expected award negligibly less than $\mathcal{A}$, which is a contradiction.

Now the details. Assume CBS is not secure. Then for some $q(k)$ there exists an adversary $\mathcal{A}$ that runs in time $q(k)$ and such that for every simulator $\mathcal{S}$, there exists a non-negligible $r(k)$ and an infinite number of k's in which $\mathsf{Adv}^{\mathsf{profit}}_{\mathcal{A},\mathcal{S},\mathsf{CBS}}(k) \geq r(k)$. Let $\mathsf{Edge}_{\mathcal{A},\mathcal{S},\mathsf{CBS}}(k,c) \stackrel{\text{def}}{=} \Pr[\mathsf{Expt}^{\mathsf{real}}_{\mathcal{A},\mathsf{CBS}}(k) \geq c] - \Pr[\mathsf{Expt}^{\mathsf{ideal}}_{\mathcal{S},\mathsf{CBS}}(k) \geq c]$. Then using the definition of expectation, there is a c such that for an infinite number of k's, $\mathsf{Edge}_{\mathcal{A},\mathcal{S},\mathsf{CBS}}(k,c) \geq r(k)/U$. Without loss of generality, we may assume that $\mathcal{A}_2$ never outputs ebid_0, since by the fairness of the Award function, this cannot increase its advantage.

Define relation $R_c(x, \mathbf{x}, M, s_1)$ to return 1 iff $|\mathbf{x}| = 1$ and $\mathsf{Award}(x, \mathbf{x}[0])[1] \geq c$. Consider the following adversary $\mathcal{B}$ for the wnm-cpa-security of Π.

$\mathcal{B}_1(pk)$:	$\mathcal{B}_2(s, y)$:
$(M, s) \leftarrow \mathcal{A}_1(pk, U)$	$\mathsf{ebid}_1 \leftarrow \mathcal{A}_2(y, s)$
return $(M, \perp, s)$	return ebid_1

Since Π is wnm-cpa-secure, there exists a simulator $\mathcal{S}' = (\mathcal{S}'_1, \mathcal{S}'_2)$ such that $\mathsf{Adv}^{\mathsf{wnm\text{-}cpa}}_{\mathcal{B},\mathcal{S}',\Pi}(R_c, k)$ is negligible for all c. Because $R_c(x, \mathbf{x}, M, s_1)$ returns 1 only if $|\mathbf{x}| = 1$, we assume without loss of generality that $\mathcal{S}'_2$ returns one-element vectors $\mathbf{y}$ and $\mathbf{z}$, i.e., values y and z. Now let a simulator $\mathcal{S}'' = (\mathcal{S}''_1, \mathcal{S}''_2)$ for the contract bidding system be defined as follows.

$\mathcal{S}''_1(U)$:	$\mathcal{S}''_2(s)$:
$(pk, sk) \leftarrow G(1^k)$	
$(M, s_1, s_2) \leftarrow \mathcal{S}'_1(pk)$	$(y, z) \leftarrow \mathcal{S}'_2(s)$
return (M, s_2)	return $D_{sk}(y)$

Note that by the fairness of Award, the award can never decrease when bid_1 is changed from $\perp$ to a valid bid, and so it is easy to see that $\Pr[\mathsf{Expt}^{\mathsf{ideal}}_{\mathcal{S}'',\mathsf{CBS}}(k) \geq c] \geq \Pr[\mathsf{Expt}^{\mathsf{wnm\text{-}atk}}_{\mathcal{S}',\Pi}(R_c, k) = 1]$. Using this fact, one can see that for all c, $\mathsf{Edge}_{\mathcal{A},\mathcal{S}'',\mathsf{CBS}}(k,c) \leq \mathsf{Adv}^{\mathsf{wnm\text{-}cpa}}_{\mathcal{B},\mathcal{S}',\Pi}(R_c, k)$, and thus by the discussion above, for all c, $\mathsf{Edge}_{\mathcal{A},\mathcal{S}'',\mathsf{CBS}}(k,c)$ is negligible. This is a contradiction, so CBS must be secure.

A.3 Proof of Theorem 3

Let $\mathcal{A}$ be an adversary that interacts with parties running σ in the $\mathcal{F}_{\text{AUTH}}$-hybrid model. We will construct an adversary $\mathcal{S}$ in the ideal process for $\mathcal{F}_{\text{M-SMT}}$ such that no environment $\mathcal{Z}$ can distinguish whether it is interacting with $\mathcal{A}$ and σ in the $\mathcal{F}_{\text{AUTH}}$-hybrid model, or with $\mathcal{S}$ and $\mathcal{F}_{\text{M-SMT}}$ in the ideal process. For simplicity, we assume there exists only one instance of $\mathcal{F}_{\text{M-SMT}}$ with identifier $id|P_i$ for some P_i. It is straightforward to extend the behavior of $\mathcal{S}$ to the case of multiple instances. $\mathcal{S}$ runs a simulated copy of $\mathcal{A}$ and maintains a tuple $(pk^*, sk^*, \mathsf{owner})$, where pk^* is the "session public key", sk^* is the corresponding secret key, and owner is the index of the party who "owns" it. [10] The session key pair (pk^*, sk^*) is initialized to $\perp$. Then $\mathcal{S}$ forwards messages from $\mathcal{Z}$ to $\mathcal{A}$, as well as messages from $\mathcal{A}$ to $\mathcal{Z}$. Furthermore, $\mathcal{S}$ also sees the *public part* (also known as "header" [7]) of all the messages from uncorrupted parties to $\mathcal{F}_{\text{M-SMT}}$ and may decide when and if to forward these messages. We refer the readers to [7] for more detailed discussions. In the case of $\mathcal{F}_{\text{M-SMT}}$, all messages to $\mathcal{F}_{\text{M-SMT}}$ are public, except the "payload message" m in $(\mathbf{send}, id, m)$. $\mathcal{S}$ also simulates the ideal functionality $\mathcal{F}_{\text{AUTH}}$.

Next, we describe the behavior of $\mathcal{S}$ in more detail. Note that $\mathcal{S}$ simulates $\mathcal{F}_{\text{AUTH}}$ as normal except as detailed below.

Simulating Communication with $\mathcal{Z}$: $\mathcal{S}$ directly forwards any messages between $\mathcal{Z}$ and $\mathcal{A}$.

Key Generation: If P_i is uncorrupted and $\mathcal{S}$ sees a message $(\mathbf{receiver}, id|P_i)$ from P_i to $\mathcal{F}_{\text{M-SMT}}$, $\mathcal{S}$ forwards this message to $\mathcal{F}_{\text{M-SMT}}$. If $pk^* \neq \perp$ it does nothing else. Otherwise $\mathcal{S}$ generates $(pk, sk) \leftarrow G(1^k)$, sets $(pk^*, sk^*) \leftarrow (pk, sk)$, and $\mathsf{owner} \leftarrow i$, and simulates $\mathcal{F}_{\text{AUTH}}$ to send $(\mathsf{key}, id|P_i, pk)$ to all other parties.

If P_i is corrupted and $\mathcal{S}$ sees P_i send a message $(\mathsf{key}, id|P_i, pk)$ to $\mathcal{F}_{\text{AUTH}}$, $\mathcal{S}$ simulates $\mathcal{F}_{\text{AUTH}}$. Furthermore, if $pk^* = \perp$ and $pk \neq \perp$, then $\mathcal{S}$ sends message $(\mathbf{receiver}, id)$ to $\mathcal{F}_{\text{M-SMT}}$ on behalf of P_i and sets $\mathsf{owner} \leftarrow i$ and $(pk^*, sk^*) \leftarrow (pk, ?)$. Here "$sk^* =?$" indicates that $\mathcal{S}$ does not know the corresponding secret key.

Delivery of the public key: When $\mathcal{A}$ delivers a message $(P_i, P_j, (\mathsf{key}, id|P_i, pk))$ from $\mathcal{F}_{\text{AUTH}}$ to an uncorrupted party P_j that has not received such a message previously, $\mathcal{S}$ records the tuple $(P_j, (P_i, pk))$ and delivers $(\mathbf{receiver}, id|P_i, P_i)$ from $\mathcal{F}_{\text{M-SMT}}$ to P_j.

Message transfer from an uncorrupted party: If $\mathcal{S}$ sees an uncorrupted party P_j send a message $(\mathbf{send}, id|P_i, -)$ to $\mathcal{F}_{\text{M-SMT}}$, where "$-$" indicates the "private" part of the message that $\mathcal{S}$ does not see, and if $\mathcal{S}$ has stored a tuple $(P_i, (P_j, pk'))$, $\mathcal{S}$ does the following. First $\mathcal{S}$ forwards the **send** message to $\mathcal{F}_{\text{M-SMT}}$, and receives the length ℓ. Next, if P_i is corrupted, then $\mathcal{S}$ receives the message (id, m, P_i) from $\mathcal{F}_{\text{M-SMT}}$ to the ideal P_i, sets $c \leftarrow E_{pk'}(m, P_i)$. If

[10] Since we assume there is only one instance of $\mathcal{F}_{\text{M-SMT}}$ ideal functionality, there is only one instance of protocol σ, and thus there is only one key. Also, in the case of identifier $id|P_i$, $\mathsf{owner} = i$.

P_i is uncorrupted, then $\mathcal{S}$ sets $c \leftarrow E_{pk^*}(0^l, P_i)$. Finally, $\mathcal{S}$ simulates $\mathcal{F}_{\text{AUTH}}$ to send $(id|P_i, c)$ to P_i.

Message transfer from a corrupted party If $\mathcal{S}$ sees a corrupted party P_j (controlled by $\mathcal{A}$) send message $(id|P_i, c)$ to P_i through $\mathcal{F}_{\text{AUTH}}$, we may assume that P_i is uncorrupted, since otherwise $\mathcal{S}$ does not need to do anything. In this case, $\mathcal{S}$ sets $m \leftarrow D_{sk^*}(c, P_j)$ and if $m \neq \perp$, sends message $(\textbf{send}, id, m)$ to $\mathcal{F}_{\text{M-SMT}}$, forwarding the message (id, P_j, m) to the ideal P_i when $\mathcal{A}$ forwards the corresponding message to P_i from $\mathcal{F}_{\text{AUTH}}$.

Now we show that if any $\mathcal{Z}$ can distinguish whether it is interacting with $\mathcal{A}$ and σ in the $\mathcal{F}_{\text{AUTH}}$-hybrid model, or with $\mathcal{S}$ and $\mathcal{F}_{\text{M-SMT}}$ in the ideal process, then this can be used to construct an adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ that breaks the tind-cca2-security of Π.

Intuitively, this is true because the only possible difference between the ideal process and the real world is in the case when an uncorrupted party P_j sends a message m to another uncorrupted party P_i. In the real world, an encryption of m is sent through $\mathcal{F}_{\text{AUTH}}$; in the ideal process, $\mathcal{S}$ simulates this message using a encryption of 0^ℓ, since $\mathcal{S}$ does not know m. Notice that the tag for this encryption is always P_j, the identity of an uncorrupted party. $\mathcal{S}$ also performs decryptions, but only for messages from corrupted parties. Therefore, $\mathcal{S}$ only decrypts messages with corrupted parties' identities as tags, and in particular, no ciphertexts with tag P_j are decrypted by $\mathcal{S}$. Then, by the tind-cca2-security of Π, the simulation of $\mathcal{S}$ is indistinguishable from the real world.

We now describe the proof more formally. $\mathcal{B}$ takes a public key pk and decryption oracle, plays the role of $\mathcal{F}_{\text{M-SMT}}$ and runs $\mathcal{S}$ with the following changes. Assume that l messages are sent using $\mathcal{F}_{\text{M-SMT}}$. $\mathcal{B}_1$ choose $h \xleftarrow{R} \{1, \ldots, l\}$. If an uncorrupted party P_i needs to generate a key pair, pk is used as the public key. Let $id|P_i$ be the associated identifier. Then for the first $h-1$ messages to P_i with id from uncorrupted parties, $\mathcal{B}$ has $\mathcal{S}$ encrypt the actual messages, instead of the all zeros message. On the hth message to P_i with $id|P_i$, say from an uncorrupted P_j, $\mathcal{B}_1$ outputs the all zeros message, the real message, the tag P_j, and its internal state. Then $\mathcal{B}_2$ uses the challenge ciphertext in the message to P_i, and continues to run $\mathcal{S}$ as normal, encrypting all zeros messages again. $\mathcal{B}_1$ and $\mathcal{B}_2$ both call the decryption oracle on messages to P_i from a corrupted P_j. Note that the tag in this case is always different from the tag returned by $\mathcal{B}_1$. Finally, $\mathcal{B}_2$ outputs whatever $\mathcal{Z}$ outputs. Note that if $h = 0$ and the bit chosen in the tind-cca2 experiment is 0, $\mathcal{B}$ runs like $\mathcal{S}$, and if $h = \ell$ and the bit chosen in the tind-cca2 experiment is 1, $\mathcal{B}$ runs like $\mathcal{A}$ in the real protocol. Then by a standard hybrid argument, if $\mathcal{Z}$ distinguishes whether it is interacting with $\mathcal{A}$ and σ in the $\mathcal{F}_{\text{AUTH}}$-hybrid model, or with $\mathcal{S}$ and $\mathcal{F}_{\text{M-SMT}}$ in the ideal process, $\mathcal{B}$ breaks the tind-cca2-security of Π.