

Black-Box, Round-Efficient Secure Computation via Non-Malleability Amplification

Hoeteck Wee
Queens College, CUNY
hoeteck@cs.qc.cuny.edu

Abstract

We present round-efficient protocols for secure multi-party computation with a dishonest majority that rely on black-box access to the underlying primitives. Our main contributions are:

- a $O(\log^* n)$ -round protocol that relies on black-box access to dense cryptosystems, homomorphic encryption schemes, or lossy encryption schemes. This improves upon the recent $O(1)^{\log^* n}$ -round protocol of Lin, Pass and Venkatasubramanian (STOC 2009) that relies on non-black-box access to a smaller class of primitives.
- a $O(1)$ -round protocol requiring in addition, black-box access to a one-way function with sub-exponential hardness, improving upon the recent work of Pass and Wee (Eurocrypt 2010).

These are the first black-box constructions for secure computation with sublinear round complexity. Our constructions build on and improve upon the work of Lin and Pass (STOC 2009) on non-malleability amplification, as well as that of Ishai et al. (STOC 2006) on black-box secure computation.

In addition to the results on secure computation, we also obtain a simple construction of a $O(\log^* n)$ -round non-malleable commitment scheme based on one-way functions, improving upon the recent $O(1)^{\log^* n}$ -round protocol of Lin and Pass (STOC 2009). Our construction uses a novel transformation for handling arbitrary man-in-the-middle scheduling strategies which improves upon a previous construction of Barak (FOCS 2002).

Keywords- secure multi-party computation, round complexity, black-box constructions, non-malleable commitments.

1. Introduction

Secure multi-party computation (MPC) allows several mutually distrustful parties to perform a joint computation without compromising, to the greatest extent possible, the privacy of their inputs or the correctness of the outputs. The early work of Goldreich, Micali and Wigderson [13] showed that we may realize secure multi-party computation with a dishonest majority under general cryptographic assumptions. Over the last decade, substantial progress was made towards improving the round complexity and computational efficiency of these protocols in two separate lines of works, culminating in (1) *constant-round* protocols for secure computation [18, 23, 27, 31] as well as (2) *black-box* constructions that avoid the use of (typically expensive)

general NP reductions [8, 14, 16, 17, 19]. However, simultaneously achieving both of these efficiency guarantees has so far remained quite elusive; the state-of-the-art for black-box constructions is a $O(n)$ -round protocol where n is the number of parties.¹ This raises the following natural question:

Does there exist a black-box, $o(n)$ -round protocol for secure multi-party computation, or is there an inherent trade-off between round complexity and computational efficiency?

Before stating our results, we provide some additional context and motivation.

Round-efficient secure computation. In the GMW protocol for secure computation, each player takes turns to sequentially commit to its input (along with a “proof of knowledge”); any non-trivial improvement in round complexity will require interweaving these input commitments, which could potentially allow an adversary to violate input independence via a man-in-the-middle attack. For this reason, improvements in round complexity for secure computation has often paralleled results on *non-malleability* [9]. Constant-round MPC protocols were first obtained by Katz, Ostrosky and Smith [18] (relying on [2]) and by Pass [27] based on the existence of enhanced trapdoor permutations and in addition, collision-resistant hash functions. More recently, Lin, Pass and Venkatasubramanian [20, 23] showed that the latter assumption can be eliminated while still maintaining almost constant – specifically, $O(1)^{\log^* n}$ – round complexity. In follow-up work, Pass and Wee [31] gave a constant-round protocol, assuming in addition one-way functions with sub-exponential hardness. An advantage of these latter two works is that they avoid the use of non-black-box simulation techniques [1] along with the sophisticated machinery (e.g. the PCP theorem) associated with them.

1. Throughout the introduction, we use n to denote the number of parties; in particular, the round complexity of all the protocols we discuss here depends only on the number of parties, and is independent of the security parameter.

Black-box secure computation. The general question of whether we can securely realize cryptographic tasks via black-box access to a general primitive is of great theoretical and practical interest. In particular, black-box constructions (namely, those that refer only to the input/output behavior of the underlying primitive) are typically more efficient in terms of both computational and communication complexity, and also more suited for implementation as compared to non-black-box constructions. As such, non-black box constructions traditionally only serve as “feasibility” results, and indeed, a series of recent works on secure computation [8, 16, 17, 25] views black-box constructions as an important step towards making MPC more “practical”. Previous works on black-box constructions have also introduced new techniques, such as the use of randomized encodings in [5, 19], and provided new conceptual and technical insights into the original non-black-box constructions.

For MPC without an honest majority, there is a fairly large gap between the round complexity of black-box constructions and that of non-black-box constructions. Specifically, the round complexity of the existing black-box constructions grows linearly with the number of parties.

1.1. Our Results

In this work, we present the first black-box constructions of general MPC protocols in the standard model with a sub-linear number of rounds:

THEOREM (INFORMAL). There exists a $O(\log^* n)$ -round protocol for securely computing any n -party functionality against a malicious adversary corrupting any number of parties that relies on black-box access to enhanced trapdoor permutations.

This construction (and the next) may be extended to a larger class of assumptions, such as dense cryptosystems and lossy encryption schemes. This improves upon the recent $O(1)^{\log^* n}$ -round protocol Lin, Pass and Venkatasubramanian [20, 23], which relies on non-black-box access to a smaller class of assumptions. Next, we show that we can also obtain constant-round protocols by relying on an additional assumption, namely one-way functions secure against sub-exponential size circuits, improving upon the non-black-box construction of Pass and Wee [31].

THEOREM (INFORMAL). There exists a $O(1)$ -round protocol for securely computing any n -party functionality against a malicious adversary corrupting any number of parties that relies on black-box access to enhanced trapdoor permutations and a one-way function secure against sub-exponential size circuits.

1.2. Our Constructions and Techniques

Both of our constructions follow the same high-level framework, which we will describe in the context of our first result.

Basic Commitment Scheme. The starting point is a constant-round commitment scheme for $O(1)$ parties, that guarantees (many-many) non-malleability against a synchronizing adversary. The trivial construction wherein each party *sequentially* commits to its input using an extractable commitment achieves such a guarantee [18]. For our $O(1)$ -round MPC protocol, we begin with a black-box variant of the constant-round non-malleable commitments in [31] for $\log \log \log n + O(1)$ parties, based on one-way functions with sub-exponential hardness.

Non-Malleability Amplification. Next, we provide a black-box transformation of a commitment scheme that is (many-many) non-malleable for t parties into one for 2^{t-1} parties, while incurring only a constant *additive* blow-up in the number of rounds; applying this transformation $O(\log^* n)$ times to our basic scheme yields a $O(\log^* n)$ -round non-malleable commitment scheme for n parties. Our transformation simplifies and improves upon the earlier construction of Lin and Pass [20] which is non-black-box and incurs a constant multiplicative blow-up in the number of rounds.²

OT Compiler. In the next step, we use our n -party non-malleable commitment scheme to realize an n -party oblivious transfer (OT) functionality. Our construction starts from a two-party OT protocol Π that is secure against a malicious sender and semi-honest receiver; many semi-honest OT protocols such as those where the sender encrypts both its inputs (c.f. [6, 10, 11, 32]) already have this property. In order to “boost” the security of Π to tolerate malicious receivers, we rely on a recent construction of Ishai et al. [7, 14, 16], which may in turn be viewed as a cut-and-choose variant of the “GMW compiler”. However, this cut-and-choose compiler relies on a commitment scheme that is extractable and equivocal, and moreover, must remain non-malleable while simulating an equivocal commitment in the left interaction. Our main technical contribution for this step lies in eliminating the equivocality requirement.

MPC from OT. In the last step, we combine our round-efficient n -party OT protocol with the constant-round MPC protocol of Ishai, Prabhakaran and Sahai in the OT-hybrid model [17, Theorem 3]. Here, we rely on the composition theorem for the stand-alone model in [4].

2. Subsequent to our work, Lin and Pass observed that they may also achieve a constant additive overhead with a minor modification of their transformation (c.f. full version of [20]).

Next, we provide an overview of the two novel building blocks in our construction, namely the non-malleability amplification protocol and the OT compiler.

Improved, simpler non-malleability amplification. Given a many-many non-malleable commitment scheme tagCom for identities of length $\log t + 1$, we construct a many-many non-malleable commitment scheme for identities of length t with a constant *additive* blow-up in the number of rounds. Our construction (shown in Fig 2), roughly speaking, proceeds as follows: to commit to a string v with identity $\text{ID} = (\text{ID}_1, \dots, \text{ID}_t) \in \{0, 1\}^t$:

- Commit to v using tagCom with identities $(1, \text{ID}_1), \dots, (t, \text{ID}_t)$ a total of t times in *parallel*.
- Prove using a zero-knowledge argument of knowledge that all t committed values are equal.

We argue, informally, that the new scheme is many-many non-malleable. Consider for simplicity the stand-alone setting, where the adversary receives a single commitment to v on the left with identity ID and tries to commit to a related value \tilde{v} with identity $\tilde{\text{ID}} \neq \text{ID}$. There must exist some i for which $\tilde{\text{ID}}_i \neq \text{ID}_i$ and thus $(i, \tilde{\text{ID}}_i)$ is different from all of $(1, \text{ID}_1), \dots, (t, \text{ID}_t)$. By many-many non-malleability of tagCom , the committed value for $(i, \tilde{\text{ID}}_i)$ is independent of all the left commitments. Furthermore, by soundness of the argument of knowledge, this value determines \tilde{v} , and thus \tilde{v} must be independent of v . This argument extends naturally to the setting where there are multiple commitments on the right, which in turn implies non-malleability with multiple commitments on both the left and on the right [22, 28].

We point out here that the overall approach of using multiple commitments to the same value and then providing a zero-knowledge proof of consistency is reminiscent of the constructions of CCA2-secure and non-malleable encryption schemes [5, 9, 29]. Our analysis considers explicitly an “alternative opening phase”, which is inspired by the notion of an “alternative decryption oracle” in the literature on encryption.

To obtain a fully black-box construction that uses black-box access to a statistically binding commitment scheme Com (and thus any one-way function [15, 26]), we combine the previous construction with the message encoding technique from [5, 30] (see Fig 3). Here, we rely crucially on the fact that in our construction, the zero-knowledge argument is used to enforce *equality* amongst committed values. Indeed, we do not know how to directly obtain a black-box variant of the Lin-Pass non-malleability amplification protocol [20] because the zero-knowledge arguments therein are used to enforce that committed values satisfy a more complex relation.

OT compiler. We use the OT compiler in [7, 14, 16] to transform a two-party OT protocol Π secure against a semi-

honest receiver into one that is secure against a malicious receiver. The idea is to run multiple copies of Π and rely on cut-and-choose to guarantee that in most of these executions, the malicious OT receiver is behaving consistently with Π (see [7, Section 2] for an overview). The random n -bit challenge for the cut-and-choose phase is determined via a coin-tossing protocol as follows: (1) the sender first commits to a random n -bit string q_S (using our extractable non-malleable commitment); (2) the receiver then responds with a random n -bit string q_R ; (3) the sender opens its commitment and the challenge is given by $q_S \oplus q_R$.

The OT compiler guarantees that there is at most one random challenge q^* that allows the malicious receiver to cheat in the cut-and-choose phase. If the sender’s commitment is equivocal, then the probability of cheating is negligible since the probability that a random equivocation equals $q_R \oplus q^*$ is 2^{-n} . To eliminate the equivocality requirement while bounding the probability of cheating, we rely on the simulator from [7] which has the property that q^* is efficiently computable (this in turn relies on extractability of the receiver’s commitment in an earlier stage of the protocol). Now suppose the simulator cheats with non-negligible probability; then, with roughly the same probability, the sender’s committed value must equal $q_R \oplus q^*$, which contradicts the hiding guarantee of the commitment scheme (amidst extraction).

MPC from non-malleable commitments. We note that our approach for deriving round-efficient MPC protocols from non-malleable commitments is quite different from that used in previous protocols with sub-logarithmic round complexity [18, 23, 27]. One limitation of the previous approaches is that the ensuing constructions rely on enhanced trapdoor permutations, or similar primitives with an “oblivious sampling” requirement; in particular, we do not know how to extend these constructions to work with lossy trapdoor functions [33]. This is in fact an inherent limitation in the techniques underlying previous constructions. Roughly speaking, these protocols all entail the use of a coin-tossing protocol to “obliviously sample” a random challenge, whereas in the simulation, this challenge is generated in a non-oblivious manner along with some trapdoor.

1.3. Additional Results

In this work, we also present new results on “full-fledged” non-malleable commitments against general, non-synchronizing adversaries, whereas the results in the previous section only address synchronizing adversaries.

$O(\log^* n)$ -round non-malleable commitments. The first is a simple construction of non-malleable commitments from one-way functions with better round complexity:

Protocol	Rounds	Assumptions	Black-Box?
GMW	$O(n)$	none	no
IKLP	$O(n)$	none	yes
KOS	$O(\log n)$	none	no
LP/LPV	$O(1)^{\log^* n}$	none	no
KOS/Pass	$O(1)$	CRHF	no
PW/LPV	$O(1)$	sub-exp OWF	no
this work	$O(\log^* n)$	none	yes
this work	$O(1)$	sub-exp OWF	yes

Figure 1: Summary of MPC protocols [3, 13, 16, 18, 20, 23, 24, 27, 30, 31]. The third column lists any additional assumptions apart from TDPs and the fourth column indicates whether the construction is black-box.

THEOREM (INFORMAL). Suppose there exists one-way functions. Then, there exists a $O(\log^* n)$ -round non-malleable commitment scheme with a black-box proof of security.

This improves upon the previous $O(1)^{\log^* n}$ -round protocol of Lin and Pass [20]. As noted in the previous section, applying our non-malleability amplification procedure a total of $O(\log^* n)$ to the trivial commitment scheme already yields a $O(\log^* n)$ non-malleable commitment scheme for n -bit identities and a synchronizing adversary.

We then provide a simple and general transformation of non-malleable commitment schemes that are secure against synchronizing adversaries into one that are secure against arbitrary scheduling strategies, with an additive increase in round complexity. This construction improves upon a previous transformation of Barak [2, Theorem 6.1], which in turn requires constant-round perfectly hiding commitments. As with [2], our transformation proceeds by creating multiple rewinding opportunities; the difference is that we add rewinding slots to the sender (as with the transformation of one-one non-malleable commitments into many-many non-malleable commitments in [20]) as opposed to the receiver.

Comparison with [20]. We highlight several technical differences between our non-malleability amplification protocol with that in [20]:

- The construction in [20] comprises of two steps: the first (implicit in [9]) transforms a many-many non-malleable commitment scheme for t parties into a one-one non-malleable commitment scheme for 2^{t-1} parties; the second transforms a one-one non-malleable commitment scheme into a many-many non-malleable commitment

scheme (with a multiplicative overhead in round complexity). Both steps address a non-synchronizing adversary.

- Our construction is “one-shot”, directly transforming a many-many non-malleable commitment scheme for t parties into a many-many non-malleable commitment scheme for 2^{t-1} parties (with an additive overhead). The basic construction only handles a synchronizing adversary.
- Unlike the construction in [20] as well as the MPC protocol in [23], our non-malleability amplification procedure and MPC protocol do not require that the underlying non-malleable commitment scheme be robust, that is, non-malleable with respect to arbitrary constant-round protocols in the left interaction. This is because we only handle synchronizing adversaries and clarifies the role of robustness in [20, 21, 23].

We note that our non-malleability amplification procedure may be modified to handle a non-synchronizing adversary by having the sender first commit to v using a statistically binding commitment, and in the zero-knowledge argument, prove that all $t + 1$ committed values are equal. We require here that the underlying non-malleable commitment be robust.

Our general transformation for handling non-synchronizing adversaries also requires robustness. This reinforces an observation used in [21] that robust commitments may be used in place of constant-round statistically-hiding commitments in many constructions of non-malleable protocols; the advantage is that the former may be based on one-way functions whereas the latter requires collision-resistant hash functions.

Black-box non-malleable commitments. The preceding construction can also be made black-box, which partially addresses in the affirmative an open problem posed by Pass (namely, whether the $O(1)^{\log^* n}$ protocol in [20] can be made black-box).

THEOREM (INFORMAL). There exists a (fully) black-box construction of a $O(\log^* n)$ -round commitment scheme that is extractable and non-malleable, starting from any one-way function.

However, the black-box construction only realizes a weaker notion of non-malleability w.r.t extraction. Roughly speaking, this means that in the man-in-the-middle setting, the values output by the extractor for the right interactions (which is the same as the committed value whenever the commitment has a valid opening and may be arbitrary when the commitment opens to \perp) are independent of the committed values in the left interaction. We stress that this weaker notion is sufficient for secure MPC and also implies the notion of non-malleability in [2, 9].

Organization. In Section 3, we present our non-malleability amplification protocol for synchronizing adversaries. We defer the remaining constructions to the full version of this paper.

2. Preliminaries and Definitions

We use Com to denote a non-interactive statistically binding commitment scheme. Our constructions may be easily extended to handle the 2-message statistically binding commitment scheme based on one-way functions from [15, 26], where the first message can be fixed “once and for all”. We also use WIPOK to denote 3-round witness-distinguishable proofs of knowledge for NP with special soundness (assuming Com) [12].

Non-malleable commitments. We recall the definition of many-many non-malleability from [22], which builds upon those in [9, 28]. Let $\text{TagCom} = (\mathcal{C}, \mathcal{R})$ be a commitment scheme with identities, and 1^n be the security parameter. In the man-in-the-middle execution, the adversary \mathcal{A} is participating m left interactions and m right interactions. In the left interactions, \mathcal{A} interacts with \mathcal{C} receiving a commitment to m values v_1, \dots, v_m , using identities $\text{ID}^1, \dots, \text{ID}^m$ of its choice. In the right interactions, \mathcal{A} interacts with \mathcal{R} attempting to commit to a sequence of m related values $\tilde{v}_1, \dots, \tilde{v}_m$, again using identities $\tilde{\text{ID}}^1, \dots, \tilde{\text{ID}}^m$ of its choice. \mathcal{A} also receives an auxiliary z . In general, we allow \mathcal{A} complete control over the scheduling of the messages, although we will also refer to an *synchronizing adversary* that always sends the i 'th messages in each of the right sessions immediately after it receives the i 'th messages in all of the left sessions and vice versa (i.e. it sends the j 'th messages in each of the left sessions immediately after it receives the j 'th messages in all of the right sessions.) If any of the right commitments as determined by the transcript are invalid or undefined, its value is set to \perp . For any i such that $\tilde{\text{ID}}^i \in \{\text{ID}^1, \dots, \text{ID}^m\}$, the value \tilde{v}_i is also set to \perp (that is, any commitment where adversary uses the same identity as that in one of the left interactions is considered invalid). We write $\text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v_1, \dots, v_m), \mathcal{R})$ to denote a random variable comprising the view of \mathcal{A} along with the m -tuple of values $(\tilde{v}_1, \dots, \tilde{v}_m)$.

Definition 2.1. A commitment scheme $(\mathcal{C}, \mathcal{R})$ is many-many non-malleable (w.r.t. opening) if for every PPT \mathcal{A} and every polynomial $m = m(n)$ and every pair of m values $(v_1^0, \dots, v_m^0), (v_1^1, \dots, v_m^1)$ along with any $z \in \{0, 1\}^*$, the distributions

$$\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v_1^0, \dots, v_m^0), \mathcal{R}) \right\} \text{ and } \left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v_1^1, \dots, v_m^1), \mathcal{R}) \right\}$$

are computationally indistinguishable.

We will also consider a restricted notion of many-many non-malleability where in the left and right interactions, the adversary \mathcal{A} may only use identities of length at most d . In addition, we will refer to relaxed notions of many-many non-malleability: one-many and one-one non-malleability. In the former, the adversary participates in one interaction on the left and m interactions on the right, and in the latter, the adversary participates in one interaction on the left and one interaction on the right. As shown in [22], any commitment scheme that is one-many non-malleable is also many-many non-malleable.

Proposition 2.2 ([22]). *Let $(\mathcal{C}, \mathcal{R})$ be a one-many non-malleable commitment (resp. w.r.t. synchronizing adversaries). Then, $(\mathcal{C}, \mathcal{R})$ is also a many-many non-malleable commitment (resp. w.r.t. synchronizing adversaries).*

3. Improved Non-Malleability Amplification

We present our construction for non-malleability amplification in Fig 2.

Proposition 3.1 (Non-malleability amplification with synchronization). *For every $t = t(n) \geq 4$, if tagCom is one-many non-malleable for identities of length $\log t + 1$ w.r.t. synchronizing adversaries, then $(\mathcal{C}, \mathcal{R})$ as shown in Fig 2 is one-many non-malleable for identities of length t w.r.t. synchronizing adversaries.*

3.1. Proof overview

We want to show that $(\mathcal{C}, \mathcal{R})$ is one-many non-malleable w.r.t. any synchronizing adversary \mathcal{A} . Let ID denote the identity on the left, and fix some identity $\tilde{\text{ID}} \neq \text{ID}$ on the right. Following the informal argument given in Section 1.2, the key in the analysis is to examine the committed value in the tagCom commitment with identity $(i, \tilde{\text{ID}}_i)$ for which $\tilde{\text{ID}}_i \neq \text{ID}_i$. Towards formalizing this argument, it is helpful to consider an “alternative open phase” for the man-in-the-middle execution corresponding to a “receiver” \mathcal{R}_0^* and \mathcal{R}^* .

More formally, we first write $\text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v), \mathcal{R}_0^*)$ to denote a random variable that is the same as $\text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v), \mathcal{R})$, except the m -tuple of values $(\tilde{v}_1, \dots, \tilde{v}_m)$ is defined as follows: for $\tilde{\text{ID}} \in \{\tilde{\text{ID}}^1, \dots, \tilde{\text{ID}}^m\}$, we set the corresponding committed value \tilde{v} as follows:

- if $\tilde{\text{ID}} = \text{ID}$, then we set \tilde{v} to \perp ,
- if $\tilde{\text{ID}} \neq \text{ID}$, let $i \in [t]$ be the first index such that $\tilde{\text{ID}}_i \neq \text{ID}_i$ and set \tilde{v} to be the committed value in Stage 1 corresponding to the tag $(i, \tilde{\text{ID}}_i)$.

Next, we write $\text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v), \mathcal{R}^*)$ to denote a random variable that is the same as $\text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v), \mathcal{R}_0^*)$, except each committed value \tilde{v} is set to \perp whenever the corresponding WIPOK in Stage 2 is rejecting. We will argue that the

Common input : security parameter 1^n and an identity $ID = (ID_1, \dots, ID_t) \in \{0, 1\}^t$.

Sender's input : a value $v \in \{0, 1\}^{\text{poly}(n)}$.

.....

COMMIT PHASE.

Stage 0: \mathcal{R} sends a random $s = f(r)$.^a \mathcal{C} responds with a dummy message.^b

Stage 1: \mathcal{C} commits to v using tagCom with tags $(1, ID_1), \dots, (t, ID_t)$. That is, \mathcal{C} executes $\text{tagCom}(id_i, v)$ in parallel for $i = 1, 2, \dots, t$, where $id_i = (i, ID_i)$.

Stage 2: \mathcal{C} proves a WIPOK of the statement:

all t commitments in Stage 1 are commitments to the same value or $s \in f(\{0, 1\}^n)$

using as witness v along with the randomnesses used for the commitments in Stage 1.

.....

OPEN PHASE.

- \mathcal{C} opens the first commitment to v in Stage 1 (the one using $(1, ID_1)$).^c

a. Following [20], \mathcal{R} should send a witness hiding proof that $s \in f(\{0, 1\}^n)$ after Stage 1.

b. The dummy message is essential for technical reasons, to ensure that \mathcal{R} 's first message in tagCom is always sent after it sends the random challenge $f(r)$.

c. Note that if the WIPOK in Stage 2 is not accepting, then the committed value corresponds to \perp .

Figure 2: Commitment scheme $\text{TagCom} = (\mathcal{C}, \mathcal{R})$.

committed values are essentially the same whether we refer to \mathcal{R} or \mathcal{R}^* . Looking ahead, we highlight two properties of the intermediate \mathcal{R}_0^* that will come in handy later:

- Property A: We can efficiently compute the output of $\text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v), \mathcal{R}^*)$ given that of $\text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v), \mathcal{R}_0^*)$; this is because we can check whether the WIPOK in Stage 2 is accepting given the transcript of the commit phase.
- Property B: The committed value according to \mathcal{R}_0^* is completely determined upon the completion of Stage 1 on the right. (In contrast, the committed values according to \mathcal{R} and \mathcal{R}^* depend also on the outcome of Stage 2.)

3.2. The hybrid argument

We begin with an overview of the hybrid argument used to establish one-many non-malleability of $(\mathcal{C}, \mathcal{R})$:

STEP 1: SWITCHING TO \mathcal{R}^* . By the soundness of the WIPOK and the one-way'ness of f , we may deduce that

$$\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^0), \mathcal{R}) \right\} \cong \left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^0), \mathcal{R}^*) \right\}$$

STEP 2: SWITCHING TO $\mathcal{C}^*(v^0)$. We change the WIPOK on the left to use the trapdoor witness r , i.e. we replace $\mathcal{C}(v^0)$ in the left execution with $\mathcal{C}^*(v^0)$ where \mathcal{C}^* is the following (computationally unbounded) sender that on input v , behaves exactly like $\mathcal{C}(v)$ in Stages 0 and 1, and proceeds as follows in Stage 2:

- (Stage 2) Compute $r \in f^{-1}(s)$ via brute force (where s is the challenge sent by \mathcal{A} in Stage 0 on the left) and complete the WIPOK using r as the witness.

By witness-indistinguishability, switching to \mathcal{C}^* does not change the distribution of the transcripts. Moreover, since the adversary is synchronizing, changing the distribution of Stage 2 on the left does not affect the distribution of the committed values in Stage 1 on the right. This means that

$$\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^0), \mathcal{R}^*) \right\} \cong \left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^0), \mathcal{R}^*) \right\}$$

STEP 3: SWITCHING TO $\mathcal{C}^*(v^1)$. We switch the left commitment in Stage 1 to v^1 (i.e. we replace $\mathcal{C}^*(v^0)$ on the left with $\mathcal{C}^*(v^1)$) and exploit many-many non-malleability of tagCom to argue that

$$\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^0), \mathcal{R}^*) \right\} \cong \left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^1), \mathcal{R}^*) \right\}$$

STEP 4: SWITCHING TO $\mathcal{C}(v^1)$. This is analogous to Step 2.

$$\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^1), \mathcal{R}^*) \right\} \cong \left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^1), \mathcal{R}^*) \right\}$$

STEP 5: SWITCHING BACK TO \mathcal{R} . This is analogous to Step 1.

$$\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^1), \mathcal{R}^*) \right\} \cong \left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^1), \mathcal{R}) \right\}$$

Next, we sketch the proofs for steps 1, 2 and 3. Looking ahead, we note that for steps 2 and 3, it suffices to establish indistinguishability of the distributions where we replace every instance of \mathcal{R}^* with \mathcal{R}_0^* . This is because we can

efficiently compute the committed values according to \mathcal{R}^* from that of \mathcal{R}_0^* together with the transcript (c.f. Property A).

Switching to \mathcal{R}^* (step 1). Here, we just need to argue that for each of the right sessions, if the Stage 2 WIPOK is accepting, then the t committed values in Stage 1 are equal (and whenever this holds, the committed values are the same whether we consider \mathcal{R} or \mathcal{R}^*). Suppose otherwise, that is, there exists a MIM adversary \mathcal{A} that with non-negligible probability, produces an accepting right execution in which the t committed values in Stage 1 are not all equal. Now, we may incorporate the left execution into \mathcal{A} (by honestly committing to v^0) to obtain a stand-alone cheating prover \mathcal{P}^* for the WIPOK in that particular right execution. Then, rewinding and extracting from \mathcal{P}^* must yield a witness for $s \in f(\{0,1\}^n)$, which contradicts one-wayness of f . We note that this is the only step of the hybrid argument (apart from the analogous Step 5) that requires rewinding or extraction.

Switching to \mathcal{C}^* (step 2). As noted above, it suffices to establish the following claim:

Lemma 3.2 (exploiting WIPOK).

$$\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^0), \mathcal{R}_0^*) \right\} \cong \left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^0), \mathcal{R}_0^*) \right\}$$

We begin with the observation that the only difference between these two distributions is the witness used in the WIPOK used in Stage 2 on the left.

Proof: Let $\Phi_1(\mathcal{A}, z)$ denote the distribution of all joint views τ of \mathcal{A} and the receivers on the right up to the point before Stage 2 on the left begins (i.e., just after the completion of Stage 1 on the right). In addition, we add to $\Phi_1(\mathcal{A}, z)$ the following values: (1) v_0 and the randomness σ used for all of the Stage 1 commitments on the left in τ ; (2) $r \in f^{-1}(s)$ where s is the Stage 0 challenge on the left in τ ; and (3) the m committed values $(\tilde{v}_1, \dots, \tilde{v}_m)$ in m executions on the right as determined by \mathcal{R}_0^* (here, we use the fact that to determine the committed values according to \mathcal{R}_0^* , we only need to look at the transcript up to the completion of Stage 1, c.f. Property B). We do not require that these latter values $((v_0, \sigma), s, (\tilde{v}_1, \dots, \tilde{v}_m))$ be efficiently computable.

Now, consider a WIPOK prover \mathcal{P} for the statement either all t commitments in Stage 1 are commitments to the same value or $s \in f(\{0,1\}^n)$ (the commitments and s refer to those for the left interaction embedded in the view τ).

against a cheating verifier \mathcal{V}^* that receives as auxiliary input $\Phi_1(\mathcal{A}, z)$. It is straight-forward to construct \mathcal{V}^* such that

- if \mathcal{P} uses the witness (v_0, σ) , then the output of \mathcal{V}^* has the same distribution as $\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}(v^0), \mathcal{R}_0^*) \right\}$; and

- if \mathcal{P} uses the witness $r \in f^{-1}(s)$, then the output of \mathcal{V}^* has the same distribution as $\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^0), \mathcal{R}_0^*) \right\}$.

Roughly speaking, \mathcal{V}^* upon receiving the auxiliary input from $\Phi_1(\mathcal{A}, z)$ (i.e. the view τ together with the values $(v_0, \sigma), s, (\tilde{v}_1, \dots, \tilde{v}_m)$), proceeds by simulating $\mathcal{A}(z)$ internally, using the messages from \mathcal{P} for the messages from \mathcal{C} or \mathcal{C}^* in the left interaction and internally simulating the receiver in Stage 2 for the m right interactions; the committed values $(\tilde{v}_1, \dots, \tilde{v}_m)$ for the m right interactions are provided as part of \mathcal{V}^* 's auxiliary input. The claim then follows from witness indistinguishability. \square

Exploiting non-malleability of tagCom (step 3). Again, it suffices to show that $\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^0), \mathcal{R}_0^*) \right\}$ and $\left\{ \text{mim}_{\mathcal{A}(z)}(\mathcal{C}^*(v^1), \mathcal{R}_0^*) \right\}$ are indistinguishable. We begin with the observation that the only difference between these two distributions lies in Stage 1 on the left; in the former, they comprise t commitments to v^0 using tagCom and in the latter, they comprise t commitments to v^1 using tagCom. To carry out the reduction to the non-malleability of tagCom, we consider a ‘‘cut-off’’ point as in [22].

Let $\Phi_2(\mathcal{A}, z)$ denote the distribution of all joint views τ of \mathcal{A} and the receivers on the right up to the point immediately after \mathcal{A} sends the dummy messages in Stage 0 in the right interactions. In addition, we add to $\Phi_2(\mathcal{A}, z)$ the value $r \in f^{-1}(s)$ where s is the Stage 0 challenge on the left in τ .

Lemma 3.3 (reduction to tagCom). *For all ppt \mathcal{A} , there exists a ppt \mathcal{B} and D such that for all z, v :*

$$\left\{ \text{mim}_{\mathcal{A}(z)}^{\text{TagCom}}(\mathcal{C}^*(v), \mathcal{R}_0^*) \right\} \cong \left\{ D(\text{mim}_{\mathcal{B}(z^*)}^{\text{tagCom}}(\mathcal{C}(\overbrace{v, \dots, v}^{t \text{ times}}), \mathcal{R})) : z^* \leftarrow \Phi_2(\mathcal{A}, z) \right\}$$

are statistically indistinguishable. Note that the first distribution refers to TagCom and the second refers to tagCom with t left interactions, all committing to v .

Once we establish this lemma, our claim follows readily from the many-many non-malleability of tagCom, which guarantees that

$$\text{mim}_{\mathcal{B}(z^*)}^{\text{tagCom}}(\mathcal{C}(v^0, \dots, v^0), \mathcal{R}) \cong \text{mim}_{\mathcal{B}(z^*)}^{\text{tagCom}}(\mathcal{C}(v^1, \dots, v^1), \mathcal{R})$$

Proof: The high-level idea is to construct a machine \mathcal{B} that on input z^* runs internally a copy of \mathcal{A} and simulates the view of \mathcal{A} in the experiment $\text{mim}_{\mathcal{A}(z)}^{\text{tagCom}}(\mathcal{C}^*(v), \mathcal{R}_0^*)$ while participating in $\text{mim}_{\mathcal{B}(z^*)}^{\text{tagCom}}(\mathcal{C}(v, \dots, v), \mathcal{R})$. The machine D will essentially ‘‘post process’’ the committed values in the second distribution according to \mathcal{R}_0^* . Note that \mathcal{B} will run t interactions of tagCom on the left, and tm interactions of tagCom on the right.

We first describe how to simulate the messages from $\mathcal{C}^*(v)$ in the left interaction in the view of \mathcal{A} :

Common input : security parameter 1^n and an identity $\text{ID} = (\text{ID}_1, \dots, \text{ID}_t) \in \{0, 1\}^t$.

Sender's input : a value $v \in \{0, 1\}^{\text{poly}(n)}$.

.....

COMMIT PHASE.

Stage 0: \mathcal{R} commits to a random subset $\Gamma \subset [10n]$ of size n . \mathcal{C} responds with a dummy message.

Stage 1: \mathcal{C} computes shares \mathbf{s} of v using a n -out-of- $10n$ secret-sharing scheme and commits to the shares using tagCom with tags $(1, \text{ID}_1), \dots, (t, \text{ID}_t)$.

- \mathcal{C} picks a random degree n polynomial p over $\text{GF}(2^{|v|})$ whose constant term is v , and computes $\mathbf{s} = (s_1, \dots, s_{10n}) = (p(1), \dots, p(10n))$.
- \mathcal{C} executes $\text{tagCom}(\text{id}_i, s_1), \dots, \text{tagCom}(\text{id}_i, s_{10n})$ in parallel, for $i = 1, 2, \dots, t$ and $\text{id}_i = (i, \text{ID}_i)$.

Stage 3: \mathcal{C} proves consistency of the commitments by opening to the shares indexed by Γ .

- \mathcal{R} opens the commitment to Γ .
- \mathcal{C} opens all t commitments to s_j in Stage 1 for all $j \in \Gamma$.
- \mathcal{R} checks that all t commitments to s_j are consistent for all $j \in \Gamma$.

.....

OPEN PHASE.

- \mathcal{C} sends v and opens the commitment to \mathbf{s} corresponding to the tag $(1, \text{ID}_1)$ in Stage 1.
 - \mathcal{R} computes codeword \mathbf{w} that is 0.9-close to \mathbf{s} .
 - \mathcal{R} checks that \mathbf{w} is a codeword corresponding to v and that \mathbf{w} and \mathbf{s} agree on all positions in Γ .
-

Figure 3: Black-box non-malleability amplification.

- (Stage 0) Stage 0 is embedded in τ .
- (Stage 1) \mathcal{B} chooses identities $(1, \text{ID}_1), \dots, (t, \text{ID}_t)$ for the t left interactions (scheduled in parallel), and forwards the messages from the external $\mathcal{C}(v, \dots, v)$ to \mathcal{A} as if coming from $\mathcal{C}^*(v)$.
- (Stage 2) \mathcal{B} computes the prover's messages in the WIPOK by using the witness r which is part of its auxiliary input z^* .

Next, we describe how \mathcal{B} simulates the messages from \mathcal{R} in the m executions of TagCom on the right and also how D computes the committed values. Again, let $\tilde{\text{ID}}^1, \dots, \tilde{\text{ID}}^m$ denote the m identities on the right. For each $j = 1, \dots, m$,

- (Stage 0) Stage 0 is embedded in τ .
- (Stage 1) \mathcal{B} uses identities $(1, \tilde{\text{ID}}_1^j), \dots, (t, \tilde{\text{ID}}_t^j)$ for the t executions of tagCom on the right. It forwards the messages from the t external copies of \mathcal{R} to \mathcal{A} .
- (Stage 2) \mathcal{B} simulates the verifier's messages in the WIPOK internally.
- (Committed value) If $\tilde{\text{ID}}^j = \text{ID}$, then D simply outputs \perp . Otherwise, D first computes the first index i for which $\text{ID}_i \neq \tilde{\text{ID}}_i^j$ and outputs as \tilde{v}_j the committed value corresponding to the tag $(i, \tilde{\text{ID}}_i^j)$. (D receives this value as part of the output of $\text{mim}_{\mathcal{B}(z^*)}^{\text{tagCom}}(\mathcal{C}(v, \dots, v), \mathcal{R})$).

This completes the reduction. \square

Acknowledgments. I would like to thank Rafael Pass for numerous insightful discussions and his constant encouragement as well as the staff at Ninth Street Espresso for their hospitality.

References

- [1] B. Barak. How to go beyond the black-box simulation barrier. In *FOCS*, pages 106–115, 2001.
- [2] B. Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In *FOCS*, pages 345–355, 2002.
- [3] D. Beaver, S. Micali, and P. Rogaway. The round complexity of secure protocols. In *STOC*, pages 503–513, 1990.
- [4] R. Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
- [5] 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 *TCC*, pages 427–444, 2008.
- [6] S. G. Choi, D. Dachman-Soled, T. Malkin, and H. Wee. Improved non-committing encryption with applications to adaptively secure protocols. In *ASIACRYPT*, pages 287–302, 2009.
- [7] S. G. Choi, D. Dachman-Soled, T. Malkin, and H. Wee. Simple, black-box constructions of adaptively secure protocols. In *TCC*, pages 387–402, 2009.
- [8] I. Damgård and Y. Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In *CRYPTO*, pages 378–394, 2005.
- [9] D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM J. Comput.*, 30(2):391–437, 2000.
- [10] S. Even, O. Goldreich, and A. Lempel. A randomized protocol for signing contracts. In *CRYPTO*, pages 205–210, 1982.
- [11] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The relationship between public key encryption and oblivious transfer. In *FOCS*, pages 325–335, 2000.
- [12] O. Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, 2001.
- [13] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [14] I. Haitner. Semi-honest to malicious oblivious transfer - the black-box way. In *TCC*, pages 412–426, 2008.
- [15] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [16] Y. Ishai, E. Kushilevitz, Y. Lindell, and E. Petrank. Black-box constructions for secure computation. In *STOC*, pages 99–108, 2006.
- [17] Y. Ishai, M. Prabhakaran, and A. Sahai. Founding cryptography on oblivious transfer - efficiently. In *CRYPTO*, pages 572–591, 2008.
- [18] J. Katz, R. Ostrovsky, and A. Smith. Round efficiency of multi-party computation with a dishonest majority. In *EUROCRYPT*, pages 578–595, 2003.
- [19] J. Kilian. Founding cryptography on oblivious transfer. In *STOC*, pages 20–31, 1988.
- [20] H. Lin and R. Pass. Non-malleability amplification. In *STOC*, pages 189–198, 2009.
- [21] H. Lin, R. Pass, W. D. Tseng, and M. Venkatasubramanian. Concurrent non-malleable zero knowledge proofs. In *CRYPTO*, 2010. To appear.
- [22] H. Lin, R. Pass, and M. Venkatasubramanian. Concurrent non-malleable commitments from any one-way function. In *TCC*, pages 571–588, 2008.
- [23] H. Lin, R. Pass, and M. Venkatasubramanian. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In *STOC*, pages 179–188, 2009.
- [24] Y. Lindell. Parallel coin-tossing and constant-round secure two-party computation. *J. Cryptology*, 16(3):143–184, 2003.
- [25] Y. Lindell and B. Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In *EUROCRYPT*, pages 52–78, 2007.
- [26] M. Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [27] R. Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In *STOC*, pages 232–241, 2004.
- [28] R. Pass and A. Rosen. Concurrent nonmalleable commitments. *SIAM J. Comput.*, 37(6):1891–1925, 2008. Preliminary version in *FOCS '05*.
- [29] R. Pass, A. Shelat, and V. Vaikuntanathan. Construction of a non-malleable encryption scheme from any semantically secure one. In *CRYPTO*, pages 271–289, 2006.
- [30] R. Pass and H. Wee. Black-box constructions of two-party protocols from one-way functions. In *TCC*, pages 403–418, 2009.
- [31] R. Pass and H. Wee. Constant-round non-malleable commitments from sub-exponential one-way functions. In *EUROCRYPT*, pages 638–655, 2010.
- [32] C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571, 2008.
- [33] C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196, 2008.