

Fast Non-Malleable Commitments

Hai Brenner*

Vipul Goyal†

Silas Richelson‡

Alon Rosen§

Margarita Vald¶

ABSTRACT

The notion of non-malleability in cryptography refers to the setting where the adversary is a man-in-the-middle (MIM) who takes part in two or more protocol executions and tries to use information obtained in one, to violate the security of another. Despite two decades of research, non-malleable commitments (NMCs) have remained too inefficient to be implemented in practice, without some sort of trusted setup.

In this work, we give a fast implementation of NMC in the plain model, based on the DDH assumption being hard over elliptic curve groups. Our main theoretical result is a new NMC scheme which can be thought of as a “high dimensional” generalization of the one in the recent work of [GRRV14]. Central to our efficiency improvements is a method of constraining challenges sent by the receiver. This new approach enables us to obtain dramatically improved parameters over those suggested in [GRRV14]. In particu-

lar, our work opens the door to implementations based on Elliptic Curves.

Our prototype implementation gives evidence of our protocol’s efficiency. Additionally, like the Elgamal commitment it is built on top of, our scheme allows for homomorphic operations on committed values, and is amenable to fast Schnorr proofs of knowledge. Thus, it will work well when used as a building block inside larger cryptographic protocols. As an example of its performance, our protocol allows a committer to commit to a 1.9–KB message using a scheme supporting 2^{20} identities in less than one second.

Categories and Subject Descriptors

F.2.0 [Theory of Computation]: Analysis of Algorithms and Problem Complexity—*General*

General Terms

Cryptography, Theory

Keywords

Non-malleable Commitments; Elliptic Curve Cryptography; Protocols; Practical implementation

1. INTRODUCTION

Secure computation was introduced in the 1980s and still remains an active area of research. Over the years classical feasibility results [Yao86, GMW87, BGW88] have given way to constructions which are both more efficient and which satisfy stronger notions of security. In recent works we have seen a push towards obtaining “practice-oriented” protocols [IKO⁺11, Lin13, FJN⁺13, LR14, AMPR14] (and the references therein). Several concrete secure computation systems building upon these improvements have been implemented (e.g. the JustGarble system [BHKR13]). The goal of this work is to initiate and promote the study of *non-malleable commitments* from the practical efficiency perspective.

1.1 Non-Malleable Commitment

Non-malleable cryptography models the scenario where the adversary is a man-in-the-middle (MIM) who participates in two or more instantiations of a protocol and tries to use information obtained in one execution to harm the security of another. Many tasks in cryptography are susceptible to such an attack, and thus non-malleable security arises

*Efi Arazi School of Computer Science, IDC Herzliya, Israel. Email: haibrenner@gmail.com Research supported by the ERC under the EU’s Seventh Framework Programme (FP/2007-2013) ERC Grant Agreement n. 307952.

†Microsoft Research, Bangalore. Email: vipul@microsoft.com. Part of this work done while visiting IDC Herzliya.

‡UCLA. Email: SiRichel@ucla.edu. Work done while visiting IDC Herzliya. Supported by the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307952

§Efi Arazi School of Computer Science, IDC Herzliya, Israel. Email: alon.rosen@idc.ac.il. Work supported by ISF grant no. 1255/12 and by the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307952.

¶The Blavatnik School of Computer Science, Tel Aviv University, Israel. Email: margarita.vald@cs.tau.ac.il. Work supported in part by ISF grant no. 1255/12.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CCS’15, October 12–16, 2015, Denver, Colorado, USA.

© 2015 ACM. ISBN 978-1-4503-3832-5/15/10 ...\$15.00.

<http://dx.doi.org/10.1145/2810103.2813721>.

naturally in many settings (e.g. commitment, encryption, coin-flipping witness-indistinguishable and zero-knowledge proofs, etc.). Interest in non-malleable security is motivated both by the strong security guarantees it provides, and by the unfortunate reality that many widely used protocols are actually highly malleable.

Non-malleable commitment (NMC), introduced by Dolev, Dwork and Naor [DDN91], is especially well studied, and has proven to be an immensely useful primitive. Very briefly we say that a commitment scheme is *non-malleable* if for every message m , no MIM adversary, intercepting a commitment $\text{Com}(m)$ and modifying it at will, is able to efficiently generate a commitment to a related message \tilde{m} . NMC is extremely versatile and is often used as a building block in more complex protocols. For example, it is known how to use NMC to construct several other non-malleable primitives such as zero-knowledge proofs.

Because of its applications across cryptography, the efficiency of NMC has been studied extensively. Beginning with [DDN91], who gave a protocol for NMC with $\mathcal{O}(\log n)$ rounds, research over the last two decades has improved the round complexity from logarithmic to constant [Bar02, PR05, Goy11, LP11] and recently to just 4 rounds [GRRV14]. However common to most of these works is that they use computationally heavy tools (such as generic zero-knowledge proofs) leading to quite inefficient protocols. The focus of our work is on obtaining *efficient* non-malleable commitments, usable in practice.

1.2 Motivation

Constant Round Multi-Party Computation. Non-malleable commitments are important building blocks in round-efficient multi-party computation (MPC) protocols. Without some form of non-malleability, any MPC protocol must have round complexity that is at least linear in the number of parties, since in protocols with sublinear round complexity, multiple parties must commit to their inputs in parallel. When the fraction of corrupt parties is less than $1/2$, one can use information theoretic verifiable secret sharing (VSS). However when the corruption threshold increases above $1/2$, information theoretic techniques fail and cryptographic subprotocols are needed. It is known how to construct constant round MPC protocols with high corruption tolerance using constant round NMC [Pas04, Wee10, Goy11] (essentially by replacing the VSS with NMC). Traditionally VSS is thought to be much more efficient than NMC, however in this work we give evidence that this isn't the case. Our prototype implementation has computation and communication complexity which is comparable with that of widely used VSS schemes. In this way our work allows one to increase the corruption tolerance of an efficient MPC scheme above $1/2$ without incurring a penalty.

Composable Security. Another application of NMC is to the area of secure protocol composition [DNS98, Fei90, Can01]. So far, all work on efficient secure computation that we are aware of focuses on the standalone setting. Designing protocols which are secure in the more demanding (though realistic) models which allow for protocol composition is vastly more difficult. Over the past decade [CLOS02, BS05, Goy12] has shown that non-malleability lies at the heart of protocol composition; MIM attacks seem to be the

most devastating. Our work takes an important step toward the goal of attaining efficient protocols which are provably secure in composition.

Sigma Protocols. Probably the most useful aspect of the Elgamal commitment scheme is its seamless compatibility with Schnorr proofs of knowledge, thus allowing a committer to prove knowledge of his committed value, or that two committed values are equal. In fact, the recent work of [AMPR14] shows that these plus several other techniques “go all the way” to an efficient protocol for general 2PC. Our NMC prototype is instantiated on top of the Elgamal commitment scheme and inherits all of the fast Σ -protocols available to the original. NMC is most commonly used as a subprotocol, so compatibility with usual cryptographic operations is an important feature. Such compatibility is not offered by most schemes which use a random oracle.

Comparison with NMC in the Random Oracle Model. Prior to this work, we are not aware of any attempts to implement non-malleable commitments in the plain model. In part this is because such a simple solution exists in the random oracle model: $\text{Com}(m; r) = H(m, r)$ where H is the random oracle. However our construction has some advantages over this trivial one. First, our implementation is provably secure under well understood and widely accepted assumptions. This is in contrast to the random oracle heuristic which is known to break down as soon as the random oracle is replaced by a specific hash function. Second, as mentioned above, our protocol is compatible with fast proofs of knowledge of committed values and Σ -protocols for proving that these values satisfy certain relations. The above and other random oracle constructions (such as the recent one from [CJS14]) need to use cut-and-choose techniques for such proofs, which introduces extra rounds of interaction and requires more communication than our protocol.

1.3 Results

The recent work of [GRRV14] gives a relatively simple, round-efficient NMC. However (as remarked by the authors), the proof of non-malleability requires choosing parameters which are much too large to be used in practice. In fact, attempting to instantiate their protocol with more reasonable parameter choices opens the door to actual attacks. Exact communication/computation complexities are shown in Table 2, but just to give an idea: one execution of the protocol instantiated over the DDH group \mathbb{Z}_q with 2^{-80} security and supporting 2^{16} identities (for background on the identities, see Section 2.4) requires the committer to compute about 4000 exponentiations in \mathbb{Z}_q where q is a 6000-bit prime. In this work we give a new NMC scheme which is similar to that of [GRRV14] (in particular, it is provably secure in the plain model) but has much better performance. For example, for the security and identity parameters mentioned, our scheme requires C to compute fewer than 600 exponentiations in elliptic curve over $GF(p)$, where p can be few hundred bit prime of our choice.¹

THEOREM 1. *There exists a four-round, statistically binding, non-malleable commitment scheme which, when instan-*

¹R has to compute many more exponentiations than C in the naïve protocol. However, since these are during the verification phase, optimization techniques such as batching are available to reduce R 's workload considerably.

iated over any DDH group, requires C (resp. R) to perform $18k$ (resp. $4k^2$) exponentiations where k is the length of identities supported (in bits). Furthermore, all of R’s exponentiations take place during the decommitment phase (and so may be done offline, and are subject to optimizations).

Our prototype implementation is instantiated over an elliptic curve group G and is secure assuming that DDH is hard over G . This is the first implementation of a NMC scheme whose proof of security does not use random oracles, or some other form of trusted setup. Our implementation capitalizes on several optimizations. For example, we use batching to reduce the computation complexity of the receiver during verification to below what are shown in Table 2. Specific performance numbers for a variety of parameters is shown in Table 1.

We stress that our scheme can be instantiated over any DDH group; it does not place restrictions on the size of the group. See Table 2 for more efficiency-related information. Additionally, we note that our new scheme is general (it can be instantiated from any one-way function), and though we have chosen to focus on a DDH-based instantiation, it could be implemented, giving similar results, starting from any homomorphic commitment scheme.

Next we increase the message space of our new NMC scheme. The basic protocol allows a committer C to commit to a vector $\mathbf{m} \in \mathbb{Z}_q^\ell$ where q is any prime specified by the protocol and ℓ depends linearly on the length of identities the scheme can support (for most choices of q , $\ell = 2k$ where the scheme can support 2^k identities is natural). The overall communication complexity of our scheme, however, is $\Omega(k^2 \log q)$, meaning that our *rate* (i.e., the message length divided by the communication complexity) is $\frac{1}{\Omega(k)}$. This is lower than we would like, especially since there are well known ordinary commitment schemes with constant rate.

It turns out that increasing the rate of NMC can be done trivially: simply commit non-malleably to a short seed s , then use s to encrypt (malleably) a longer message m . This construction, however, destroys the structure of the NMC scheme in the sense that any fast Σ -protocols available to the original NMC needn’t (and likely will not) be compatible with the high rate NMC. For example, if the NMC protocol of Theorem 1 is instantiated on top of the Elgamal commitment scheme it inherits homomorphic operations and Schnorr-like Σ -protocols for proving that the committed value satisfies a linear relation. However, if the amortization technique just mentioned is used on top to increase the message size, these fast proofs are lost. This lack of compatibility with common cryptographic operations severely limits the utility of this NMC when used as a building block for larger protocols. We show that our scheme is amenable to an amortization technique which is as efficient as the general one, and which also preserves the Σ -protocols available to the original.

1.4 Central Idea: From Subsets to Subspaces

Our techniques expand on the algebraic techniques put forth in [GRRV14]. Briefly recall their scheme.

1. C \rightarrow R commitments to the coefficients of n linear polynomials all with the same constant term. We

denote the i -th polynomial $f_i(x) = r_i x + m$ where $m \in \mathbb{Z}_q$ is the value C is committing to non-malleably, and the $r_i \in \mathbb{Z}_q$ are random.

2. R \rightarrow C a random query vector $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_q^n$.
3. C \rightarrow R response vector $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{Z}_q^n$ where $a_i = f_i(\alpha_i)$.
4. C proves to R in ZK that the values sent in step 1 are valid commitments, consistent with the responses in step 3.

Crucial for the proof of non-malleability is the fact that the challenges α_i are actually drawn from proper subsets $V_i \subset \mathbb{Z}_q$, the sizes of which depend on C’s identity. Since C’s identity and M’s identity are distinct, M will be asked queries from different challenge spaces than he will be allowed to use to query C. It is then possible to show that by carefully controlling the V_i , M will not be able to use C’s answer to construct his own, and thus the protocol is non-malleable.

Unfortunately, the requirements needed on the V_i for the proof to work mean that, for example, the size of V_{i+1} is at least a constant multiple of the size of V_i for each i . Since in [GRRV14], each $V_i \subset \mathbb{Z}_q$, this and other requirements on the V_i necessitate $q = 2^{\Omega(k^2)}$. The main advantage of our new protocol is that for us each V_i will be a subset of the vector space \mathbb{Z}_q^ℓ for some ℓ which may be controlled. This allows us to keep q fixed to a small value and still make sure that $|V_i|$ is large.

The proof of non-malleability in [GRRV14] is delicate enough to where even this small change threatens to destroy the proof and much care is needed. The key is setting the $V_i \subset \mathbb{Z}_q^\ell$ to be vector subspaces rather than just proper subsets. This allows most of the algebraic ideas in the proof to carry over unchanged. Instead of the $f_i(\cdot)$, we use maps $\mathbf{v}_i \mapsto \langle \mathbf{z}_i, \mathbf{v}_i \rangle$ where the vector $\mathbf{z}_i = (\mathbf{m}, r_i) \in \mathbb{Z}_q^\ell$. Our proof uses similar ideas such as collinearity testing, though our setting is more complicated as we work in vector spaces of high dimension.

2. PRELIMINARIES

For positive $n \in \mathbb{N}$, let $[n] = \{1, \dots, n\}$. A function $\varepsilon : \mathbb{N} \rightarrow \mathbb{R}^+$ is *negligible* if it tends to 0 faster than any inverse polynomial i.e., for all constants c there exists $n_c \in \mathbb{N}$ such that for every $n > n_c$ it holds that $\varepsilon(n) < n^{-c}$. We use $\text{negl}(\cdot)$ to specify a generic negligible function. We abbreviate “probabilistic polynomial time” with PPT. We assume familiarity with computational indistinguishability and zero-knowledge proofs (and related protocols).

2.1 Commitment schemes

Commitment schemes are protocols which enable a party, known as the committer C, to commit himself to a value while keeping it secret from the (potentially cheating) receiver, R. This property is known as hiding. Additionally, upon receiving the commitment from C, R is ensured that even if C cheated, there is at most one value that C can decommit to during a later, decommitment phase (binding). In this work, we consider commitment schemes that are statistically-binding which means that the hiding property only holds against computationally bounded adversaries.

Table 1: Implementation Time Comparisons in DDH Implementation. The complexity is measured in Giga clock-cycles. For example, considering identities domain of size 2^{16} with 2^{-96} security our proposed implementation runs in about 0.07 seconds on the committer’s side, and about 0.41 seconds on the receiver’s side. This is highly efficient in comparison to [GRRV14] whose duration is about 50 seconds for the committer and about 46 seconds for the receiver (given matching parameters of identity space and slightly worse security). Measurements were taken on a Core i7 Intel Ivy Bridge architecture (with 2.9 GHz frequency) using only a single core. We note that for id ’s of size larger than 2^{30} the [GRRV14] scheme is already infeasible.

Scheme	IDs	Prime Size ($\log q$)	Msg size	Clock Cycles (C + R)	Security
[GRRV14]	2^{16}	6400	6400	146G + 134G	2^{-80}
this paper	2^{16}	192	12288	0.2G + 1.15G	2^{-96}
this paper	2^{32}	192	24576	0.4G + 4.1G	2^{-96}
this paper	2^{32}	224	28672	0.82G + 8.7G	2^{-112}

Table 2: Complexity Comparisons in DDH Implementations. Communication is computed in terms of field elements.

Scheme	IDs	Prime Size	Msg Space	Comm.	Exponentiations (C+R)
[GRRV14]	2^k	$\log q = \Omega(k^2)$	\mathbb{Z}_q	$96k$	$32k + 96k$
this paper	2^k	any	\mathbb{Z}_q^{4k}	$2k^2$	$18k + 4k^2$
this paper	2^k	any	$\mathbb{Z}_q^K, K = \Omega(k^2)$	$14K + 2k^2$	$6K + 2Kk$

DEFINITION 1 (Statistically Binding). Let $\langle C, R \rangle$ be an interactive protocol between C and R . We say that $\langle C, R \rangle$ is a statistically binding commitment scheme if the following properties hold:

Correctness: If C and R do not deviate from the protocol, then R should accept (with probability 1) during the decommit phase.

Binding: For every C^* , there exists a negligible function $\text{negl}(\cdot)$ such that C^* succeeds in the following game with probability at most $\text{negl}(\lambda)$: On security parameter 1^λ : C^* first interacts with R in the commit phase to produce commitment c . Then C^* outputs two decommitments (c, m_0, d_0) and (c, m_1, d_1) , and succeeds if $m_0 \neq m_1$ and R accepts both decommitments.

Hiding: For every PPT receiver R^* and every two messages m_0, m_1 , the view of R^* after participating in the commitment phase, where C committed to m_0 is indistinguishable from its view after participating in a commitment to m_1 .

[Nao91] gives a 2-round, statistically binding bit commitment scheme that can be built from any OWF [HILL99].

2.2 Non-malleable commitments

We wish for our commitment scheme to be impervious to a MIM adversary, M , who takes part in two protocol executions (in the left interaction M acts as the receiver while in the right, M plays the role of the committer), and tries to use the left interaction to affect the right. The security property we desire can be summarized:

For any MIM adversary M , there exists a standalone machine who plays only one execution as the committer, yet whose commitment is indistinguishable from M ’s commitment on the right.

At first glance, non-malleability seems impossible as surely nothing can be done to protect against a MIM who simply copies messages from one protocol execution to another. For this reason, non-malleable security offers protection only

against any MIM who tries to change messages in a meaningful way.

On the Existence of Identities. In this work, just as in [DDN91], we assume that the committer has an identity $id \in \{0, 1\}^k$. In order to perform a successful mauling attack, a MIM has to maul a commitment corresponding to C ’s identity into a commitment of his own, distinct identity. Though this sounds like a strong assumption on the network, essentially requiring that “you know who you are talking to”, for our purposes, it is actually equivalent to the requirement discussed above, that the MIM does something other than simply copy messages. This is because our protocol is interactive, and the first committer message contains a statistically binding commitment to m . This means that if we set the committer’s identity to be the first committer message, C ’s and M ’s identities will be distinct unless M copied C ’s first message.

Moving forward, we assume that the committer’s id is externally given and we require that non-malleability holds only in the case when C and M ’s identities are different. We also assume for simplicity that player identities are known before the protocol begins, though strictly speaking this is not necessary, as the identities do not appear in the protocol until after the first committer message. We point out that M can choose his identity adversarially, as long as it is not equal to C ’s.

Definition of Non-Malleable Commitments. In this work, we consider the notion of non-malleability with respect to commitment and we will frequently refer to the “message committed to by a MIM adversary M during the commitment phase”. We note that this is uniquely defined, as all commitment schemes in this work are statistically binding, and so for all but a negligible fraction of the possible transcripts T of the interaction between M and an honest receiver R , there exists at most one message m that is consistent with T (i.e., for which there exist random coin tosses which give T). We recall the definition of non-malleable commitments of Lin et al *et al.* [LPV08].

The man-in-the-middle execution. In the man-in-the-middle execution, the MIM adversary M is simultaneously participating in two interactions called the left and the right interaction. In the left interaction M is the receiver and interacts with a honest committer whereas in the right interaction M is the committer and interacts with a honest receiver. We define a random variable $\mathbf{MIM}_{(C,R)}(m, z)$ describing (\tilde{m}, v) : the value M commits to in the right interaction, and M 's view in the full experiment. Specifically, M has auxiliary information z and interacts on the left with an honest committer C with input message m and identity id and on the right with honest receiver R . M attempts to commit to a value \tilde{m} that is related to m using an identity \tilde{id} of its choice. If the right commitment (as determined by the transcript) is invalid or undefined, or $id = \tilde{id}$ its value is set to \perp .

The simulated execution. In the simulated execution a simulator \mathcal{S} interacts with an honest receiver R . \mathcal{S} receives security parameter 1^λ and auxiliary information z and interacts with the honest receiver R . Let $\mathbf{SIM}_{(C,R)}^{\mathcal{S}}(1^\lambda, z)$ denote the random variable describing (\tilde{m}, v) : the value \mathcal{S} commits to in the right interaction, and \mathcal{S} 's view during the entire experiment. If the commitment produced by \mathcal{S} is invalid or undefined, its value is set to \perp .

DEFINITION 2. Let $\langle C, R \rangle$ be a commitment scheme. $\langle C, R \rangle$ is non-malleable with respect to commitment if for every PPT MIM adversary M , there exists a PPT simulator \mathcal{S} such that the following ensembles are indistinguishable for all $m \in \{0, 1\}^\lambda$:

$$\{\mathbf{MIM}_{(C,R)}(m, z)\}_{z \in \{0,1\}^*}; \{\mathbf{SIM}_{(C,R)}^{\mathcal{S}}(1^\lambda, z)\}_{z \in \{0,1\}^*, id \in \{0,1\}^k}$$

2.3 Σ -protocols

The notion of a Σ -protocol generalizes the identification protocols by Schnorr [Sch91], and Guillou and Quisquater [GQ88], as well as many others. A Σ -protocol is required to be zero-knowledge against an honest verifier only. Despite this limitation, Σ -protocols are building blocks in many protocols due to their remarkable efficiency.

Let $R = \{(x, w)\} \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be a binary relation, where x denotes the common input to the prover and the verifier, and w denotes a witness, which is the private input to the prover. Let $L_R = \{x : \exists w \text{ such that } (x, w) \in R\}$ denote the language corresponding to relation R .

DEFINITION 3. A Σ -protocol for relation R is a protocol between a prover P and a verifier V of the following 3-move form:

Satisfying the following three properties:

3-move form. Given a common input x to P, V and a w such that $(x, w) \in R$ is private input to P :

1. P sends a message a .
2. V sends a random t -bit string e .
3. P sends a reply z , and V decides to accept or reject based on x, a, e, z .

Completeness. if P, V follow the protocol on input x and private input w to P where $(x, w) \in R$, the verifier always accepts.

Special soundness. From any input x and any pair of accepting conversations on input x , (a, e, z) and (a, e', z') with $e \neq e'$, one can efficiently compute witness w such that $(x, w) \in R$.

Special honest-verifier zero-knowledge . There exists a polynomial-time simulator S , which on input x and a random e outputs an accepting conversation of the form (a, e, z) , with the same probability distribution as conversations between the honest P, V on input x

2.4 Challenge Space Generation

In this section, we describe how to derive the tags from C 's identity, highlighting the properties we will use moving forward. Let $id \in \{0, 1\}^k$ be C 's identity and $\{t_i\}_{i=1, \dots, n}$ the tags resulting from it in as described below. Note that we have already touched on the role that the tags play in our protocol: the size of the challenge space in coordinate i is determined by t_i . In particular, V_i will be a t_i -dimensional vector space over \mathbb{Z}_q , and so its order is q^{t_i} . Of particular importance is the quantity $|V_{i+1}|/|V_i|$. We will always have $t_i < t_{i+1}$ so this quantity will always be at least $q = 2^{160}$ (our general protocol can be tuned to handle other sizes of q using techniques presented in [GRRV14]; however since $q \approx 2^{160}$ in our implementation, we do not discuss this point further). In particular, $|V_i|/|V_{i+1}|$ is negligible for all i . The other property we need from the tags is the following: if $i \neq j$ then $t_i \neq t_j$; moreover $t_i < \tilde{t}_i$ holds for at least one $i \in [n]$ (as does $t_i > \tilde{t}_i$).

The tags are generated according to the ‘‘DDN trick’’, namely $t_i = 2i + id_i$ where id_i is the i -th bit of C 's identity. This ensures that if $id \neq \tilde{id}$ that $t_i \neq \tilde{t}_i$ for some i . To arrange that $t_i > \tilde{t}_i$ for some i , we define a final tag $t_{k+1} = 2(k+1) - |id|$ where $|id|$ is the Hamming weight of id . Our final parameters therefore, are $n = k + 1$ and $\ell = 2(k + 1)$.

3. THE BASIC PROTOCOL

In this section, we describe our protocol given a large prime $q \approx 2^{160}$ and vector spaces $V_1, \dots, V_n \subset \mathbb{Z}_q$ as described in Section 2.4. We use Naor's two round, statistically binding bit commitment scheme [Nao91] as a building block.² We use boldface to denote vectors, and $\langle \cdot, \cdot \rangle$ to denote inner product of vectors. In particular we denote our challenge vector set as $\{\mathbf{v}_i\}_{i=1, \dots, n}$ and our response set $\{w_i\}$. We write **Com** for the entire first commitment message. Our non-malleable commitment scheme $\langle C, R \rangle_{\text{BASIC}}$ shown in Figure 1. The decommitment phase is done by having the committer C send m and the randomness it used during the protocol.

PROPOSITION 1. The commitment scheme $\langle C, R \rangle_{\text{BASIC}}$ is computationally hiding and statistically binding.

PROOF. Computational hiding follows from the hiding of **Com**, the ZK property and having having $n + \ell - 1$ variables and only n equations, where each equation has an independent, random component. Statistical binding follows from

²Briefly recall Naor's scheme: 1) R sends random initialization message σ , and 2) C responds with $\text{Com}_\sigma(m; s)$, a commitment to $m \in \{0, 1\}$ using randomness s (we will feel free to just write $\text{Com}(m)$, suppressing σ and s for simplicity)

Public Parameters: A large prime q , an integer ℓ and vector spaces $V_1, \dots, V_n \subset \mathbb{Z}_q^\ell$ obtained from identities as described in Section 2.4.

Committer's Private Input: $\mathbf{m} \in \mathbb{Z}_q^{\ell-1}$ to be committed to.

Commit Phase:

1. $R \rightarrow C$: Send the first message σ of the Naor commitment scheme.
2. $C \rightarrow R$: Choose random values $r_1, \dots, r_n \in \mathbb{Z}_q$. This defines vectors $\mathbf{z}_1, \dots, \mathbf{z}_n \in \mathbb{Z}_q^\ell$ where $\mathbf{z}_i = (r_i, \mathbf{m})$. C sends commitments to every coordinate of the \mathbf{z}_i .
3. $R \rightarrow C$: Send random challenge vectors $\{\mathbf{v}_i\}_{i=1, \dots, n}$ where each $\mathbf{v}_i \in V_i \subset \mathbb{Z}_q^\ell$.
4. $C \rightarrow R$: Send evaluations $\{w_i\}$, where each $w_i = \langle \mathbf{z}_i, \mathbf{v}_i \rangle \in \mathbb{Z}_q$.
5. $C \longleftrightarrow R$ **Consistency proof:** Parties engage in a zero-knowledge argument protocol where C proves to R that

$$\exists ((m_1, s_1), \dots, (m_{\ell-1}, s_{\ell-1}), (r_1, s'_1), \dots, (r_n, s'_n))$$
 such that:
 - $\mathbf{Com} = (\text{Com}_\sigma(m_1; s_1), \dots, \text{Com}_\sigma(m_{\ell-1}; s_{\ell-1}), \text{Com}_\sigma(r_1; s'_1), \dots, \text{Com}_\sigma(r_n; s'_n))$; and
 - $w_i = \langle \mathbf{z}_i, \mathbf{v}_i \rangle \forall i = 1, \dots, n$.

Figure 1: Non-malleable commitment scheme $\langle C, R \rangle_{\text{BASIC}}$.

the statistical binding property of the underlying commitment scheme Com . \square

THEOREM 2. *The commitment scheme $\langle C, R \rangle_{\text{BASIC}}$ is non-malleable.*

3.1 Discussion of the Proof of Non-Malleability

We now mention some of the key points from the proof of non-malleability. See the full version for more details. We first define the notion of ε -dependence.

DEFINITION 4. *Fix a commitment message \mathbf{Com} . We say that $\mathbf{v}_{i'}$ is ε -dependent on $\tilde{\mathbf{v}}_i$ if*

$$\Pr_{\{\tilde{\mathbf{v}}_i\}}(\mathbf{v}_{i'} \in \text{Span}\{\mathbf{v}_{i'}\} | \tilde{\mathbf{v}}_i \in \text{Span}\{\tilde{\mathbf{v}}_i\}) \geq \varepsilon.$$

Definition 4 implicitly assumes we are given a completed transcript with first message \mathbf{Com} and right (resp. left) challenges $\{\tilde{\mathbf{v}}_i\}$ (resp. $\{\mathbf{v}_i\}$). Then M is rewound to the beginning of the second message and given new challenges $\{\tilde{\mathbf{v}}_i\}$ with $\tilde{\mathbf{v}}_i \in \text{Span}\{\tilde{\mathbf{v}}_i\}$, and we are interested in the probability that the resulting left challenges $\{\mathbf{v}_{i'}\}$ are such that $\mathbf{v}_{i'} \in \text{Span}\{\mathbf{v}_{i'}\}$. Intuitively, $\mathbf{v}_{i'}$ being dependent on $\tilde{\mathbf{v}}_i$ is a result of M performing a mauling attack. Suppose that M mauls $\text{Com}(\mathbf{z}_{i'})$ in order to obtain $\text{Com}(\tilde{\mathbf{z}}_i)$. Then M does not know $\tilde{\mathbf{z}}_i$ and so cannot hope to answer $\tilde{\mathbf{v}}_i$ except by mauling C 's answer to $\mathbf{v}_{i'}$. Therefore, if M is rewound to the beginning of the right session's query phase and asked a different query $\{\tilde{\mathbf{v}}_i\}$ such that $\tilde{\mathbf{v}}_i \in \text{Span}\{\tilde{\mathbf{v}}_i\}$, M will have

to ask $\{\mathbf{v}_{i'}\}$ such that $\mathbf{v}_{i'} \in \text{Span}\{\mathbf{v}_{i'}\}$ if he wants to answer successfully. The proof of non-malleability proceeds in cases, depending on how the right queries can depend on the left queries. This is exactly the same high-level structure of the proof of non-malleability in [GRRV14], though some of the specifics differ.

Consider the bipartite graph $G = (V, E)$ with vertex set $V = [n] \times [n]$ and an edge $(i', i) \in E$ exactly when $\mathbf{v}_{i'}$ is ε -dependent on $\tilde{\mathbf{v}}_i$. We consider four cases:

- **IND:** There exists i such that $(i', i) \notin E$ for all i' .
- **L < R:** There exist $i' < i$ such that $(i', i) \in E$.
- **L > R:** There exist $i' > i$ such that $(i', i) \in E$.
- **L = R:** $E = \{(i, i) : i = 1, \dots, n\}$.

We think of **IND**, **L < R**, **L > R**, **L = R** as events whose probability of occurring depends on M . Notice however that regardless of M 's behavior,

$$\Pr(\mathbf{IND} \text{ or } \mathbf{L} < \mathbf{R} \text{ or } \mathbf{L} > \mathbf{R} \text{ or } \mathbf{L} = \mathbf{R}) = 1.$$

Therefore, it suffices to prove that if there exists a PPT M such that

$$\Pr(M \text{ mauls } \langle C, R \rangle_{\text{BASIC}} \ \& \ \mathbf{E}) = \frac{1}{\text{poly}}$$

for any $\mathbf{E} \in \{\mathbf{IND}, \mathbf{L} < \mathbf{R}, \mathbf{L} > \mathbf{R}, \mathbf{L} = \mathbf{R}\}$, then M breaks the hiding of Com . This is the same approach taken in [GRRV14] and many supporting lemmas carry over directly to our setting. In particular, the case of $\mathbf{E} = \mathbf{IND}$ is handled in Claim 9 of [GRRV14]. Additionally, Claims 5 through 7 of [GRRV14] show that the cases for the other values of \mathbf{E} follow from the same lemma which we state below and prove in Appendix 3.1. Our main technical lemma is much the same as Claim 8 from [GRRV14]. We prove it formally, however, as it is the point in which the differences between our high-dimensional setting and their linear setting are most apparent.

Note that once \mathbf{Com} is fixed, M defines a map from right challenges $\{\tilde{\mathbf{v}}_i\}$ to left challenges $\{\mathbf{v}_i\}$. Fix (as a function of λ), $\omega = \omega(1)$. Given a transcript with left challenges $\{\mathbf{v}_i\}$, say that the event **SUPER-POLY** occurs if the preimage of $\{\mathbf{v}_i\}$ has superpolynomial size: $|\mathbf{M}^{-1}(\{\mathbf{v}_i\})| \geq \lambda^\omega$.

LEMMA 1. *Fix non-negligible σ . If*

$$\Pr(M \text{ mauls } \langle C, R \rangle_{\text{BASIC}} \ \& \ \mathbf{SUPER-POLY}) \geq \sigma$$

then there exists a PPT algorithm \mathcal{A} who breaks the hiding of $\langle C, R \rangle_{\text{BASIC}}$.

PROOF. We give a reduction from a MIM who mauls $\langle C, R \rangle_{\text{BASIC}}$ given that the event **SUPER-POLY** occurs, to a PPT \mathcal{A} who breaks the hiding of Com . Our \mathcal{A} proceeds as follows.

- \mathcal{A} chooses random $\mathbf{m}_0, \mathbf{m}_1 \in \mathbb{Z}_q^{\ell-1}$ and begins the hiding game, sending $(\mathbf{m}_0, \mathbf{m}_1)$ to the challenger \mathcal{C} . Then \mathcal{A} instantiates M and runs two sessions of $\langle C, R \rangle_{\text{BASIC}}$ forwarding the messages it receives as C to \mathcal{C} . In the left interaction, \mathcal{C} commits to \mathbf{m}_u for unknown $u \in \{0, 1\}$. Let the resulting transcript be \mathbb{T} , and let $\{\mathbf{v}_i\}$ and $\{w_i\}$ be the query and response messages in the left interaction. Similarly $\{\tilde{\mathbf{v}}_i\}$ and $\{\tilde{w}_i\}$ are the query and response messages in the right.

- \mathcal{A} chooses random $u' \in \{0, 1\}$ and sets vectors $\mathbf{z}_i \in \mathbb{Z}_q^\ell$ to be the unique such vectors which correspond to $\mathbf{m}_{u'}$ and so that $\langle \mathbf{z}_i, \mathbf{v}_i \rangle = w_i$ for all i .
- \mathcal{A} chooses random $\{\tilde{\mathbf{v}}'_i\}, \{\tilde{\mathbf{v}}''_i\}$ such that $\{\tilde{\mathbf{v}}_i, \tilde{\mathbf{v}}'_i, \tilde{\mathbf{v}}''_i\}$ are collinear for all i . He rewinds M twice back to the beginning of the right execution's second message asking $\{\tilde{\mathbf{v}}'_i\}$ and $\{\tilde{\mathbf{v}}''_i\}$, receiving left queries $\{\mathbf{v}'_i\}$ and $\{\mathbf{v}''_i\}$. \mathcal{A} responds to M 's queries with $\{w'_i\}$ and $\{w''_i\}$ where $w'_i = \langle \mathbf{z}_i, \mathbf{v}'_i \rangle$ for all i (and similarly for w''_i), receiving right responses $\{\tilde{w}'_i\}$ and $\{\tilde{w}''_i\}$.
- \mathcal{A} checks whether $\{(\tilde{\mathbf{v}}_i, \tilde{w}_i), (\tilde{\mathbf{v}}'_i, \tilde{w}'_i), (\tilde{\mathbf{v}}''_i, \tilde{w}''_i)\}$ are collinear for all i . \mathcal{A} outputs u' if so, $1 - u'$ if not.

The following is Fact 1 from [GRRV14].

FACT 1. Let \mathbf{E} be an event such that

- $\Pr(\mathbf{E}) \geq \xi$;
- $\Pr(\{(\tilde{\mathbf{v}}_i, \tilde{w}_i), (\tilde{\mathbf{v}}'_i, \tilde{w}'_i), (\tilde{\mathbf{v}}''_i, \tilde{w}''_i)\} \text{ col. } \forall i | u' = u \ \& \ \mathbf{E}) \geq \xi'$;
- $\Pr(\{(\tilde{\mathbf{v}}_i, \tilde{w}_i), (\tilde{\mathbf{v}}'_i, \tilde{w}'_i), (\tilde{\mathbf{v}}''_i, \tilde{w}''_i)\} \text{ col. } \forall i | u' \neq u \ \& \ \mathbf{E}) \leq \xi''$,

for non-negligible values ξ, ξ', ξ'' satisfying $\xi'' = \mathcal{O}(\xi\xi')$. Then \mathcal{A} breaks the hiding of $\langle C, R \rangle_{\text{BASIC}}$.

So it suffices to construct such an \mathbf{E} . We have already defined the event SUPER-POLY. We additionally define the events HONEST and EXT as follows:

- HONEST occurs if $\{\tilde{\mathbf{v}}_i\}$ is such that M gives correct answers $\{\tilde{w}_i\}$ provided his queries $\{\mathbf{v}_i\}$ are answered correctly;
- EXT occurs if $\{\tilde{\mathbf{v}}_i\}$ is such that M gives correct answers with non-negligible probability even he receives random answers to his queries $\{\mathbf{v}_i\}$.

Note that HONEST, EXT and SUPER-POLY are all determined by \mathbf{Com} and $\{\tilde{\mathbf{v}}_i\}$. Since \mathcal{A} fixes \mathbf{Com} and rewinds M asking three challenges, we will think of these events as depending only on $\{\tilde{\mathbf{v}}_i\}$. We may assume $\Pr(\text{HONEST}) = \frac{1}{\text{poly}}$ because otherwise M is not completing the protocol successfully (and so cannot be mauling) except with negligible probability. It can also be shown (Claim 2 in [GRRV14]) that $\Pr(M \text{ mauls } \& \text{ EXT}) = \text{negl}$. This is because if M answers $\{\tilde{\mathbf{v}}_i\}$ correctly with non-negligible probability given random answers to his queries $\{\mathbf{v}_i\}$ then he knows his message $\tilde{\mathbf{m}}$ (it can be extracted in polynomial time). It therefore cannot depend on C 's commitment \mathbf{m} , which is computationally hidden from M .

We use the event

$$\mathbf{E} = \text{HONEST} \ \& \ (\neg \text{EXT}) \ \& \ \text{SUPER-POLY}.$$

It remains to prove

$$\Pr(\{(\tilde{\mathbf{v}}_i, \tilde{w}_i), (\tilde{\mathbf{v}}'_i, \tilde{w}'_i), (\tilde{\mathbf{v}}''_i, \tilde{w}''_i)\} \text{ col. } \forall i | u' = u \ \& \ \mathbf{E}) = \frac{1}{\text{poly}}$$

$$\Pr(\{(\tilde{\mathbf{v}}_i, \tilde{w}_i), (\tilde{\mathbf{v}}'_i, \tilde{w}'_i), (\tilde{\mathbf{v}}''_i, \tilde{w}''_i)\} \text{ col. } \forall i | u' \neq u \ \& \ \mathbf{E}) = \text{negl}.$$

To prove the first, note that if $u' = u$ then M receives correct answers to his queries $\{\tilde{\mathbf{v}}_i\}$, $\{\tilde{\mathbf{v}}'_i\}$, and $\{\tilde{\mathbf{v}}''_i\}$, so by definition

of HONEST he will provide correct answers $\{\tilde{w}_i\}$, $\{\tilde{w}'_i\}$, and $\{\tilde{w}''_i\}$ with non-negligible probability. However, if M gives correct answers then certainly $\{(\tilde{\mathbf{v}}_i, \tilde{w}_i), (\tilde{\mathbf{v}}'_i, \tilde{w}'_i), (\tilde{\mathbf{v}}''_i, \tilde{w}''_i)\}_i$ are collinear for all i .

To prove the second, note that if $u' \neq u$ then M receives random answers $\{w'_i\}$ to his queries $\{\mathbf{v}'_i\}$, so by definition of EXT, his responses $\{\tilde{w}'_i\}$ to $\{\tilde{\mathbf{v}}'_i\}$ on the right will be incorrect with high probability. This means that the lines spanned by the points $\{(\tilde{\mathbf{v}}'_i, \tilde{w}'_i), (\tilde{\mathbf{v}}''_i, \tilde{w}''_i)\}_i$ are not contained in the ℓ -planes $\{(\tilde{\mathbf{v}}, \langle \tilde{\mathbf{v}}, \mathbf{z}_i \rangle)\}$ to which the points $\{(\tilde{\mathbf{v}}_i, \tilde{w}_i)\}_i$ belong. It follows that the intersections of these lines and these ℓ -planes is at most a single point. By definition of SUPER-POLY, the probability that this $\{\tilde{\mathbf{v}}_i\}$ in the intersection is chosen out of all the superpolynomially many preimages of $\{\mathbf{v}_i\}$ is negligible.

□

4. AMORTIZING NMC

We now show that the scheme from the previous section $\langle C, R \rangle_{\text{BASIC}}$ can be amortized as efficiently as the trivial construction mentioned in the introduction, while still preserving any fast proofs available to the original scheme. Our protocol $\langle C, R \rangle_{\text{AMOR}}$ is shown in Figure 2.

The proof that $\langle C, R \rangle_{\text{AMOR}}$ is non-malleable is almost exactly the same as the proof of non-malleability for $\langle C, R \rangle_{\text{BASIC}}$. Recall that most of the proof for $\langle C, R \rangle_{\text{BASIC}}$ involved viewing M as a map from right challenges $\{\tilde{\mathbf{v}}_i\}$ to left challenges $\{\mathbf{v}_i\}$ and proving that in all cases, M cannot be mauling. The challenge space for $\langle C, R \rangle_{\text{AMOR}}$ is the same as the challenge space for $\langle C, R \rangle_{\text{BASIC}}$ and so the same analysis for M as a function from right to left challenge spaces carries over. The only thing that must be changed is that the extractor must be modified in order to extract all of the $\tilde{\mathbf{m}}_i$. This, however, is a minor change which we omit.

Note that when k is large, the communication complexity of $\langle C, R \rangle_{\text{AMOR}}$ is dominated by the second message when C sends commitments to the $\tilde{\mathbf{m}}_i$. If \mathbf{Com} has constant rate, then so does $\langle C, R \rangle_{\text{AMOR}}$. Finally, if \mathbf{Com} is homomorphic then so is $\langle C, R \rangle_{\text{AMOR}}$, and similarly, any fast Σ -protocols available to \mathbf{Com} for proving knowledge of committed values or that these values satisfy a linear relation will likewise be available to $\langle C, R \rangle_{\text{AMOR}}$.

5. IMPLEMENTATION

We now discuss some of the specifics of our prototype implementation. The two cryptographic building blocks needed for our protocol are statistically binding homomorphic commitment and zero-knowledge proof of knowledge. Any statistically binding commitment can be used, but for the sake of efficiency, we instantiate our protocol using ElGamal commitments [EG85]. Zero-knowledge is instantiated with Schnorr protocol [Sch91] as follows: a receiver proves knowledge of a trapdoor followed by a schnorr based Σ -OR protocol [CDS94], where the committer proves knowledge of either the committed message or receiver's trapdoor.

Optimized performance is achieved by aggregating many of the operations of the atomic protocols together. Our main most valuable tool is the modular exponentiation with

Public Parameters: A large prime $q = \lambda^{\omega(1)}$, integers ℓ, k and vector spaces $V_1, \dots, V_n \subset \mathbb{Z}_q^\ell$ obtained from C 's identity.

Committer's Private Input: Messages $\mathbf{m}_1, \dots, \mathbf{m}_k \in \mathbb{Z}_q^{\ell-1}$ to be committed to.

Commit Phase:

1. $R \rightarrow C$: Send the first message σ of the Naor commitment scheme.
2. $C \rightarrow R$: Choose random values $\mathbf{r}_1, \dots, \mathbf{r}_n \in \mathbb{Z}_q^k$. This defines vectors $\mathbf{z}_{i,j} \in \mathbb{Z}_q^\ell$ for $i = 1, \dots, n$ and $j = 1, \dots, k$ where $\mathbf{z}_{i,j} = (r_{i,j}, \mathbf{m}_j)$. C sends commitments to every coordinate of the $\mathbf{z}_{i,j}$ using Com .
3. $R \rightarrow C$: Send random challenge vectors $\{\mathbf{v}_i\}_{i=1, \dots, n}$ where each $\mathbf{v}_i \in V_i \subset \mathbb{Z}_q^\ell$.
4. $C \rightarrow R$: Send evaluation vectors $\{\mathbf{w}_i\}$, where each $\mathbf{w}_i \in \mathbb{Z}_q^k$ is such that $w_{i,j} = \langle \mathbf{z}_{i,j}, \mathbf{v}_i \rangle \in \mathbb{Z}_q$.
5. $C \longleftrightarrow R$ **Consistency proof:** Parties engage in a zero-knowledge argument protocol where C proves to R that

$$\exists (\{(m_{j,l}, s_{j,l})\}_{1 \leq j \leq k, 1 \leq l \leq \ell-1}, \{(r_{i,j}, s'_{i,j})\}_{1 \leq i \leq n, 1 \leq j \leq k})$$
 such that:
 - $\text{Com} = (\{\text{Com}_\sigma(m_{j,l}, s_{j,l})\}_{1 \leq j \leq k, 1 \leq l \leq \ell-1}, \{\text{Com}_\sigma(r_{i,j}, s'_{i,j})\}_{1 \leq i \leq n, 1 \leq j \leq k})$; and
 - $w_{i,j} = \langle \mathbf{z}_{i,j}, \mathbf{v}_i \rangle \quad \forall 1 \leq i \leq n, 1 \leq j \leq k$.

Figure 2: Constant Rate Non-malleable commitment scheme $\langle C, R \rangle_{\text{AMOR}}$.

precomputation [BGMW92]. This provides a tremendous improvement for both parties. We also use batch verification [LL94, BGR98] for each of the parties' verification of repeated equations. In particular, we use the bucket test with small exponents. For further optimization we use the Comba multiplication [Com90].

We implement our protocol over elliptic curves. Note that [GRRV14] required the group size to be exponential in the tag values, and therefore required approximately 6000 bits group size to achieve decent security. Our protocol, on the other hand, allows using any group size (as long as it is secure) regardless of the tags values. This allows us to use more standard elliptic curves groups. We choose to use NIST GF(P) elliptic curves, and construct our implementation in MIRACL SDK – an open source C software library for elliptic curve cryptography. Our implementation is generic for any EC parameters. Concrete measurements of performance for our protocol are presented in table 3. For example, to commit to a message of size 24,576 bits with El-Gamal modulus of size 192 bits, tags of size 32 bits, it takes approximately 4.5 giga clock cycles. That is, on average for every committed bit we run approximately 0.19 mega clock cycles.

IDs	Prime Size (log q)	Msg size	Clock Cycles (C+R)	Sec.
2^{16}	192	12288	0.2 G + 1.15 G	2^{-96}
2^{32}	192	24576	0.4 G + 4.1 G	2^{-96}
2^{64}	192	49152	0.79 G + 15.3 G	2^{-96}
2^{16}	224	14336	0.4 G + 2.46 G	2^{-112}
2^{32}	224	28672	0.82 G + 8.7 G	2^{-112}

Table 3: Implementation time measurements in DDH Implementation. The complexity is measured in Giga clock-cycles. Measurements were taken on a Core i7 Intel Ivy Bridge architecture (with 2.9 GHz frequency) using only a single core.

6. REFERENCES

- [AMPR14] Arash Afshar, Payman Mohassel, Benny Pinkas, and Ben Riva. Non-interactive secure computation based on cut-and-choose. In *Eurocrypt*, 2014.
- [Bar02] Boaz Barak. Constant-Round Coin-Tossing with a Man in the Middle or Realizing the Shared Random String Model. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, FOCS '02*, pages 345–355, 2002.
- [BGMW92] Ernest F. Brickell, Daniel M. Gordon, Kevin S. McCurley, and David Bruce Wilson. Fast exponentiation with precomputation (extended abstract). In Rainer A. Rueppel, editor, *Advances in Cryptology - EUROCRYPT '92, Workshop on the Theory and Application of Cryptographic Techniques, Balatonfired, Hungary, May 24-28, 1992, Proceedings*, volume 658 of *Lecture Notes in Computer Science*, pages 200–207. Springer, 1992.
- [BGR98] Mihir Bellare, Juan A. Garay, and Tal Rabin. Fast batch verification for modular exponentiation and digital signatures. pages 236–250. Springer-Verlag, 1998.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation (Extended Abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, STOC '88*, pages 1–10, 1988.
- [BHKR13] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. Efficient garbling from a fixed-key blockcipher. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 478–492. IEEE Computer Society, 2013.
- [BS05] Boaz Barak and Amit Sahai. How To Play Almost Any Mental Game Over The Net - Concurrent Composition via Super-Polynomial Simulation. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 543–552, 2005.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. pages 136–147, 2001. Preliminary full version

- available as Cryptology ePrint Archive Report 2000/067.
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Desmedt [Des94], pages 174–187.
- [CJS14] Ran Canetti, Abhishek Jain, and Alessandra Scafuro. Practical UC security with a global random oracle. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 597–608, 2014.
- [CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing, STOC '02*, pages 494–503, 2002.
- [Com90] Paul G. Comba. Exponentiation cryptosystems on the ibm pc. *IBM systems journal*, 29(4):526–538, 1990.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-Malleable Cryptography (Extended Abstract). In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, STOC '91*, pages 542–552, 1991.
- [Des94] Yvo Desmedt, editor. *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings*, volume 839 of *Lecture Notes in Computer Science*. Springer, 1994.
- [DNS98] Cynthia Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. In *STOC*, pages 409–418, 1998.
- [EG85] Taher El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Proceedings of CRYPTO 84 on Advances in Cryptology*, pages 10–18, New York, NY, USA, 1985. Springer-Verlag New York, Inc.
- [Fei90] Uriel Feige. Ph.d. thesis, alternative models for zero knowledge interactive proofs. Weizmann Institute of Science, 1990.
- [FJN⁺13] Tore Kasper Frederiksen, Thomas Pelle Jakobsen, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. Minilego: Efficient secure two-party computation from general assumptions. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 537–556. Springer, 2013.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 218–229, 1987.
- [Goy11] Vipul Goyal. Constant Round Non-malleable Protocols Using One-way Functions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, STOC '11*, pages 695–704. ACM, 2011.
- [Goy12] Vipul Goyal. Positive results for concurrently secure computation in the plain model. In *FOCS*, 2012.
- [GQ88] Louis C. Guillou and Jean-Jacques Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory. In Christoph G. Günther, editor, *Advances in Cryptology - EUROCRYPT '88, Workshop on the Theory and Application of of Cryptographic Techniques, Davos, Switzerland, May 25-27, 1988, Proceedings*, volume 330 of *Lecture Notes in Computer Science*, pages 123–128. Springer, 1988.
- [GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *FOCS*, 2014.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A Pseudorandom Generator from any One-way Function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [IKO⁺11] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, and Amit Sahai. Efficient non-interactive secure computation. In *Eurocrypt*, 2011.
- [Lin13] Yehuda Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2013.
- [LL94] Chae Hoon Lim and Pil Joong Lee. More flexible exponentiation with precomputation. In Desmedt [Des94], pages 95–107.
- [LP11] Huijia Lin and Rafael Pass. Constant-round Non-malleable Commitments from Any One-way Function. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, STOC '11*, pages 705–714, 2011.
- [LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent Non-malleable Commitments from Any One-Way Function. In *Theory of Cryptography, 5th Theory of Cryptography Conference, TCC 2008*, pages 571–588, 2008.
- [LR14] Yehuda Lindell and Ben Riva. Cut-and-choose yao-based secure computation in the online/offline and batch settings. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA*,

- USA, August 17-21, 2014, *Proceedings, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 476–494. Springer, 2014.
- [Nao91] Moni Naor. Bit Commitment Using Pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [Pas04] Rafael Pass. Bounded-Concurrent Secure Multi-Party Computation with a Dishonest Majority. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, STOC '04, pages 232–241, 2004.
- [PR05] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, STOC '05, pages 533–542, 2005.
- [Sch91] Claus-Peter Schnorr. Efficient signature generation by smart cards. *J. Cryptology*, 4(3):161–174, 1991.
- [Wee10] Hoeteck Wee. Black-Box, Round-Efficient Secure Computation via Non-malleability Amplification. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science*, pages 531–540, 2010.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *FOCS*, 1986.