

Distinguisher-Dependent Simulation in Two Rounds and its Applications

Abhishek Jain ^{*}, Yael Tauman Kalai, Dakshita Khurana, and Ron Rothblum ^{**}

¹ Department of Computer Science, Johns Hopkins University, Baltimore, USA.

² Microsoft Research, Cambridge, USA.

³ Department of Computer Science, UCLA, USA.

⁴ Department of Computer Science, MIT, Cambridge, USA.

Abstract. We devise a novel simulation technique that makes black-box use of the adversary as well as the distinguisher. Using this technique we construct several round-optimal protocols, many of which were previously unknown even using non-black-box simulation techniques:

- Two-round witness indistinguishable (WI) arguments for NP from different assumptions than previously known.
- Two-round arguments and three-round arguments of knowledge for NP that achieve strong WI, witness hiding (WH) and distributional weak zero knowledge (WZK) properties in a setting where the instance is only determined by the prover in the last round of the interaction. The soundness of these protocols is guaranteed against adaptive provers.
- Three-round two-party computation satisfying input-indistinguishable security as well as a weaker notion of simulation security against malicious adversaries.
- Three-round extractable commitments with guaranteed correctness of extraction from polynomial hardness assumptions.

Our three-round protocols can be based on DDH or QR or N^{th} residuosity and our two-round protocols require quasi-polynomial hardness of the same assumptions. In particular, prior to this work, two-round WI arguments for NP were only known based on assumptions such as the existence of trapdoor permutations, hardness assumptions on bilinear maps, or the existence of program obfuscation; we give the first construction based on (quasi-polynomial) DDH or QR or N^{th} residuosity. Our simulation technique bypasses known lower bounds on black-box simulation [Goldreich-Krawczyk'96] by using the distinguisher's output in a meaningful way. We believe that this technique is likely to find additional applications in the future.

1 Introduction

The notion of zero-knowledge (ZK) proofs [38] is fundamental to cryptography. Intuitively, zero-knowledge proofs guarantee that the proof of a statement

^{*} Supported in part by a DARPA/ARL Safeware Grant W911NF-15-C-0213.

^{**} Partially supported by the grants: NSF MACS - CNS-1413920, DARPA IBM - W911NF-15-C-0236 and SIMONS Investigator award Agreement Dated 6-5-12.

does not reveal anything beyond the validity of the statement. This seemingly paradoxical requirement is formalized via the *simulation* paradigm, namely, by requiring the existence of an efficient simulator that simulates the view of a malicious verifier, without access to any witness for the statement.

Over the years, ZK proofs (and arguments) have been integral to the design of numerous cryptographic protocols, most notably general-purpose secure computation [36], as well as specific tasks such as coin-tossing, equivocal and/or extractable commitments and non-malleable protocols [25]. Even protocols satisfying weaker notions of ZK such as strong witness indistinguishability and witness hiding (WH)[29], are typically constructed only via a ZK protocol⁵. In particular, the round complexity of ZK determines the round complexity of known constructions for these tasks.

Goldreich and Krawczyk (GK) [35] established that three round ZK arguments for NP with black-box simulation do not exist for languages outside BPP. Furthermore, all known non-black-box simulation techniques [3] require more than three rounds.⁶ This has acted as a barrier towards achieving round-efficient protocols for many of the aforementioned tasks. In this work, we investigate the possibility of overcoming this barrier.

(When) Is ZK Necessary? ZK proofs are typically used to enforce “honest behaviour” for participants of a cryptographic protocol. The zero-knowledge property additionally ensures privacy of the inputs of honest parties. However, many applications of ZK described above do not themselves guarantee simulation-based security but only weaker indistinguishability-based security. As such, it is not immediately clear whether the “full” simulation power of ZK is necessary for such applications.

For example, *strong witness indistinguishability* requires that for two indistinguishable statement distributions $\mathcal{X}_1, \mathcal{X}_2$, a proof (or argument) for statement $x_1 \leftarrow \mathcal{X}_1$ must be indistinguishable from a proof (or argument) for statement $x_2 \leftarrow \mathcal{X}_2$. All known constructions of strong witness indistinguishable protocols rely on ZK arguments with standard simulation – and therefore end up requiring at least as many rounds as ZK arguments. Similar issues arise in constructing input-hiding/input-indistinguishable secure computation, witness hiding arguments and proofs, and extractable (or other sophisticated) commitment schemes. However, it is unclear whether ZK is actually *necessary* in these settings.

This raises the question of whether it is possible to devise “weaker” simulation strategies in three rounds or less that can be used to recover several applications of ZK. In this work, we implement such a black-box simulation strategy in only *two* rounds.

Distinguisher-Dependent Simulation. Our starting observation is that for any cryptographic protocol that only aims to achieve indistinguishability-based

⁵ The work of Bitansky and Paneth [10] constructing 3 round witness-hiding and weak zero-knowledge from variants of auxiliary-input point obfuscation, is an exception.

⁶ Here we only refer to *explicit* simulation, and not non-explicit simulation via knowledge assumptions [41,5].

security, the security reduction has access to an efficient *distinguisher*. In such scenarios, one can hope to argue security via a (weaker) simulation strategy that potentially makes use of the distinguisher in a non-trivial manner.

The idea of distinguisher-dependent simulation is not new and has previously been studied in the context of interactive proofs, where it is referred to as weak zero knowledge (WZK) [28]⁷. Informally, WZK says that any bit of information that can be learned by the verifier by interacting with the prover can be simulated given only the instance. As such, WZK suffices for many applications of ZK, and in particular, implies meaningful weaker notions such as WH and WI [29].

The immediate question is whether distinguisher-dependent simulation can be realized in three rounds or less. At first, the answer seems to be negative since the lower bound of GK also extends to WZK (this was already noted in [10]).

A key insight in our work is that in many applications of ZK proofs, the statement being proven is chosen by the prover from a (public) distribution. Suppose that the proof system is *delayed-input* [48], namely, where the instance and witness are only required for computing the last prover message. In this case, it is to an honest prover’s advantage to reveal the instance to the verifier only in the last round. This does not violate correctness due to the delayed input property, but “weakens” a malicious verifier, and in particular, ensures that a malicious verifier’s messages are independent of the instance. Interestingly, we observe that the lower bound of GK no longer holds in this case⁸.

At a high-level, this is because in this setting, a simulator may be able to learn non-trivial information about the distinguisher’s behavior by observing its output on different samples created using possibly different instances from the same distribution. This observation is, in fact, not limited to delayed-input proofs and extends to a large class of important two-party functionalities including coin-tossing, generating common reference strings and oblivious PRFs.

This observation opens doors to the possibility of constructing proof systems and secure computation in three rounds or less with meaningful simulation-based and indistinguishability-based security guarantees.

A New Black-box Simulation Technique. We devise a new distinguisher-dependent black-box simulation technique that only requires two-rounds of communication. Roughly, we show that a single bit of information (of whether the proof is accepted or rejected by the distinguisher) can be used to learn information about the (possibly) malicious verifier and distinguisher, in a bit-by-bit fashion, and that this information can later be used to efficiently simulate the proof.

We remark that the ability to learn a bit of information based on whether the protocol execution is accepted or rejected has in the past been viewed as a

⁷ Recall that standard ZK requires that for any adversarial verifier, there exists a simulator that can produce a view that is indistinguishable from the real one to every distinguisher. WZK relaxes this notion by reversing the order of quantifiers, and allowing the simulator to depend on the distinguisher.

⁸ Indeed, the GK proof strategy crucially uses a verifier that chooses its protocol message as a function of the instance. See Section 1.2 for further discussion.

source of insecurity in cryptographic protocols. For example, in the delegation of computation schemes of [33,18], an adversarial prover can successfully cheat if it is able to observe the verifier’s output over multiple executions. For similar reasons, special care is taken to prevent “input-dependent aborts” in the design of many secure computation protocols.

In this work, we turn this apparent weakness into a positive by using it to devise a new black-box simulation strategy. Using this strategy, we obtain several new results on proof systems and secure computation. Most of our results were previously unknown even using non-black-box simulation techniques.

Our Setting. In order to prove privacy, we must sometimes restrict ourselves to a setting where the prover has the *flexibility* to sample instances and witnesses in the last round of the argument. More specifically, our simulator will require knowledge of any witnesses that are fixed (implicitly or explicitly) before the last message is sent; however, it will not require knowledge of witnesses fixed in the last round.

1.1 Our Results

We now proceed to describe our results. We start with our results on interactive proof systems and then describe their applications to secure two-party computation and extractable commitment schemes. All of these results rely on our new black-box simulation strategy.

I. Delayed-Input Interactive Proofs. We study two and three round *delayed-input* interactive proof systems where the instance to be proven can be chosen by the prover in the last round, and soundness holds even against adaptive cheating provers who choose the instance depending upon the verifier’s message. First studied by [48], delayed-input protocols have found numerous applications over the years in the design of round-efficient cryptographic protocols for a variety of tasks such as secure computation [47,31,44], resettable security [24,60], non-malleable commitments [58,20], improved Σ -protocols [21,22,50], and so on.

In the context of establishing various privacy notions, we consider both *adaptive* verifiers, who receive the instance at the beginning of the protocol, and hence may choose their message based on this instance, and *non-adaptive* verifiers, who receive the instance only in the last round of the protocol, and hence their message is independent of the instance. As we discuss later, guaranteeing privacy against non-adaptive verifiers suffices for many natural applications of delayed-input proof systems.

(I). TWO ROUND ARGUMENT SYSTEMS. Our first contribution is a two-round delayed-input argument system that achieves witness-indistinguishability (WI) against *adaptive* verifiers, and strong WI, witness hiding (WH) and distributional weak zero-knowledge (WZK) against *non-adaptive* verifiers.

Theorem 1 (Informal). *Assuming the existence of two-round oblivious transfer that is secure against malicious PPT receivers and quasi-polynomial time*

semi-honest senders, there exists a two-round delayed-input interactive argument system for NP with adaptive soundness and the following privacy guarantees:

- *WI against adaptive verifiers.*
- *Strong WI, WH and distributional WZK against non-adaptive verifiers.*

Oblivious transfer (OT) protocols as required in the above theorem can be constructed based on quasi-polynomial hardness of Decisional Diffie-Hellman (DDH) [51] or N 'th Residuosity or Quadratic Residuosity [45,43].

Comparison with Prior Work. If we know an a priori super-polynomial bound on the hardness of the language, then two-round WH can be obtained from two-round ZK with super-polynomial time simulators (SPS) [53]. However, no constructions of two-round WH or distributional WZK for NP against non-uniform verifiers were previously known. (We refer the reader to Section 1.3 for a more thorough discussion.)

WI proofs in two rounds (or less) were previously only known based on either trapdoor permutations⁹ [27], or the decision linear assumption on bilinear groups [40], or indistinguishability obfuscation [11]. Our result in Theorem 1 substantially adds to the set of standard assumptions that suffice for two-round WI. We remark that unlike previous protocols, our WI protocol is not publicly verifiable.

Privacy Amplification via Round Compression. We obtain Theorem 1 by “compressing” any Σ -protocol¹⁰ [23] into a two-round private-coin argument using OT. Our compiler follows the approach of [1,46], except that we use a maliciously secure OT as opposed to a computational PIR [17].

Interestingly, our approach of compressing a Σ -protocol into a two-round argument results in amplifying its privacy guarantees. Indeed, standard Σ -protocols are not known to be WZK. Furthermore, [42,54] proved that such protocols cannot be proven WH using black-box reductions.

Avoiding NP Reductions. An added benefit of our approach is that given a Σ -protocol for a language L , we obtain a two-round private-coin argument system with the security guarantees stated in Theorem 1 for the same language L , *without* using expensive NP reductions. To the best of our knowledge, no such two-round argument system was previously known.

(II). **THREE ROUND ARGUMENTS OF KNOWLEDGE.** Our second contribution is a three-round delayed-input interactive *argument of knowledge* system that achieves WH and distributional WZK against non-adaptive verifiers. This protocol uses only polynomial assumptions, but requires an extra round.

Theorem 2 (Informal). *Assuming the existence of two-round oblivious transfer (OT) that is secure against malicious PPT receivers and semi-honest PPT*

⁹ Presently, the only known candidates for trapdoor permutations are based on factoring or indistinguishability obfuscation [12,32].

¹⁰ Very roughly, a Σ -protocol is a three round protocol that is honest verifier zero-knowledge, and has a strong soundness guarantee. We refer the reader to Definition 1.

senders, as well as dense cryptosystems, there exists a three-round interactive argument of knowledge for NP that achieves soundness against adaptive (unbounded) provers and Strong WI, WH and distributional WZK against non-adaptive PPT verifiers.

Comparison with Prior Work. Three-round ZK arguments are known either based on non-standard “knowledge assumptions” [41,5], or against adversaries with *bounded* non-uniformity [9,7]. In this work, we consider security against adversaries with non-uniform advice of arbitrarily large polynomial length, based on standard cryptographic assumptions. Prior to our work, three-round WH and WZK arguments for NP were known from non-black-box techniques that rely on auxiliary input point obfuscation assumptions [10]. These protocols, unlike ours, guarantee privacy also against adaptive verifiers. However, some of their underlying assumptions have recently been shown to be implausible [14,6]. (See Section 1.3 for a more detailed discussion.)

II. Secure Two-Party Computation. We next study two-party computation against malicious adversaries in the plain model without trusted setup assumptions. In this setting, the state of the art result is due to Katz and Ostrovsky [47] who constructed a four-round protocol for general functions in the setting where only one party receives the output. We refer to the output recipient as the *receiver* and the other party as the *sender*.

As an application of our new simulation technique, we obtain two new results on two-party computation in *three* rounds. Our first result achieves input-indistinguishable security [49] against malicious receivers, while our second result achieves distinguisher-dependent simulation security against malicious receivers. In both of these results, we achieve standard simulation security against malicious senders. We elaborate on these results below.

(I). **THREE ROUND INPUT-INDISTINGUISHABLE COMPUTATION.** The notion of input-indistinguishable computation (IIC) was introduced by Micali, Pass and Rosen [49] as a weakening of standard simulation-based security notion for secure computation while still providing meaningful security. (See also [30,52].) Roughly, input-indistinguishable security against malicious receivers guarantees¹¹ that for any function f and a pair of inputs (x_1, x_2) for the sender, a malicious receiver cannot distinguish whether the sender’s input is x_1 or x_2 as long as the receiver’s “implicit input” y in the execution is such that $f(x_1, y) = f(x_2, y)$.¹²

We construct the first three-round IIC protocol for general functions based on polynomial hardness assumptions. In fact, our protocol achieves standard simulation-based security against malicious senders and input-indistinguishable security against malicious receivers.

¹¹ Security against malicious senders can be defined analogously.

¹² The formal security definition of IIC is much more delicate, and we refer the reader to the technical sections for details.

Theorem 3 (Informal). *Assuming the existence of two-round oblivious transfer that is secure against malicious PPT receivers and semi-honest PPT senders, along with dense cryptosystems, there exists a three-round secure two-party computation protocol for general functions between a sender and a receiver, where only the receiver obtains the output, with standard simulation security against malicious senders and input-indistinguishable security against malicious receivers.*

(II). **THREE ROUND TWO-PARTY COMPUTATION WITH DISTINGUISHER DEPENDENT SIMULATION.** We also consider a weak simulation-based security notion for two-party computation that is defined analogously to distributional WZK by allowing the simulator to depend (non-uniformly) upon the distinguisher and the distribution over the public input to the adversary. We refer to this as distributional distinguisher-dependent simulation secure two-party computation. While this generalizes the notion of distributional WZK, it also implies distinguisher-dependent simulation security for all functionalities where the honest party's input can be efficiently sampled (without the need for non-uniform advice) even if the input of the malicious party and any common input is already fixed.

We show that the same protocol as in Theorem 3 also satisfies distributional distinguisher-dependent security for all functionalities. In particular, we obtain three round distinguisher-dependent simulation secure two party computation for inherently distributional functionalities such as coin-tossing, generating common *reference* strings and oblivious PRFs.

Theorem 4 (Informal). *Assuming the existence of two-round oblivious transfer that is secure against malicious PPT receivers and semi-honest PPT senders, as well as dense cryptosystems, there exists a three-round protocol for secure two-party computation for any function between a sender and receiver, where only the receiver obtains the output, with standard simulation security against a malicious sender and distributional distinguisher-dependent simulation security against a malicious receiver. This implies distinguisher-dependent simulation secure two-party computation for any function where the sender's input can be efficiently sampled even if the receiver's input (and any common input) is already fixed.*

A Two-round Protocol. We also remark that our three-round two-party computation protocol can be downgraded to a two-round protocol that achieves distributional distinguisher-dependent simulation security or input-indistinguishable security against malicious receivers and quasi-polynomial time simulation security against malicious senders (or polynomial-time simulation security against semi-honest senders).

Outputs for Both Parties. Theorem 3 and Theorem 4 consider the case where only one party, namely the receiver, learns the output. As observed in [47], such a protocol can be easily transformed into one where both parties receive the output by computing a modified functionality that outputs signed values. Now the output recipient can simply forward the output to the other party who accepts it only if the signature verifies.

This adds a round of communication, making the protocol four rounds in total. Because we consider distinguisher-dependent simulation security (or input-indistinguishable security), this bypasses the lower bound of [47] who proved that coin-tossing cannot be realized with standard simulation-based security in less than five rounds when both parties receive output.

III. Extractable Commitments. We finally discuss application of our techniques to *extractable* commitments. A commitment scheme is said to be extractable if there exists a PPT extractor that can extract the committed value with *guaranteed correctness of extraction*. In particular, if the commitment is not “well-formed” (i.e., not computed honestly), then the extractor must output \perp , while if the commitment is well-formed, then the extractor must output the correct committed value. Extractable commitments are very useful in the design of advanced cryptographic protocols, in particular, to facilitate the extraction of the adversary’s input in tasks such as secure computation, non-malleable commitments, etc.

A standard way to construct extractable commitment schemes is to “compile” a standard commitment scheme with a ZKAoK, namely, by having a committer commit to its value using a standard commitment and additionally give a ZKAoK to prove knowledge of the decommitment value. The soundness property of ZKAoK guarantees the well-formedness of commitment, which in turn guarantees correctness of extraction of the committed value using the AoK extractor for ZKAoK, while the ZK property preserves the hiding of the underlying commitment. This approach yields a four round extractable commitment scheme starting from any four round ZKAoK. However, in the absence of three-round ZKAoK, constructing three-round extractable commitments from *polynomial* hardness assumptions have so far proven to be elusive.¹³

The main challenge here is to enforce honest behavior on a malicious committer, while at the same time guaranteeing privacy for honest committers. Indeed, natural variations of the above approach (e.g., using weaker notions such as WIPOK that are known in three rounds) seem to only satisfy one of these two requirements, but not both.

As an application of Theorem 2, we construct the first three-round extractable commitment scheme based on standard polynomial-time hardness assumptions.

Theorem 5 (Informal). *Assuming the existence of two-round oblivious transfer that is secure against malicious PPT receivers and semi-honest PPT senders, as well as dense cryptosystems, there exists a three-round extractable commitment scheme.*

¹³ All known constructions of three-round extractable commitments from polynomial-hardness assumptions (such as [55,56]) only satisfy a weak extraction property where either the extractor outputs (with non-negligible probability) a non \perp value when the commitment is not well-formed, or it fails to output the correct value when the commitment is well-formed. It is, however, possible to construct extractable commitments using quasi-polynomial hardness [39] or using three round zero-knowledge with super-polynomial simulation [53].

Roughly, our construction of extractable commitments follows the same approach as described above. Our main observation is that the hiding property of the extractable commitment can be argued if the AoK system satisfies a *strong* WI property (instead of requiring full-fledged ZK).

1.2 Discussion

Non-adaptive Verifiers. Our results on distributional WZK, WH and strong WI are w.r.t. non-adaptive verifiers who learn the statement in the last round of the protocol. To the best of our knowledge, privacy against non-adaptive verifiers has not been studied before, and therefore, it is natural to ask whether it is a meaningful notion of privacy.

We argue that privacy against non-adaptive verifiers is very useful. Our main observation is that in many applications of delayed-input proof systems, the verifier is already non-adaptive, or can be made non-adaptive by design. Two concrete examples follow:

- We construct a three-round extractable commitment scheme by combining a standard commitment with a three-round delayed-input strong WIAoK of correctness of the committed value, that achieves security against non-adaptive verifiers. By sending the commitment in the last round, we automatically make the verifier non-adaptive.
- In secure computation using garbled circuits (GCs) [59], a malicious sender must prove correctness of its GC. In this case, the instance (i.e., the GC) can simply be sent together with the last prover message, which automatically makes the verifier non-adaptive. This does not affect the security of the receiver if the proof system achieves adaptive soundness (which is true for our constructions). Indeed, our construction uses exactly this approach.

We anticipate that the notion of privacy against non-adaptive verifiers will find more applications in the future.

Bypassing GK and GO Lower Bounds. We now elaborate on the reasons why we are able to bypass the lower bounds of [35] and [37]. The black-box impossibility result of [35] for three-round ZK crucially uses an adaptive verifier. More specifically, they consider a verifier that has a random seed to a pseudo-random function hard-wired into it, and for any instance and first message sent by the prover, it uses its PRF seed, to answer honestly with fresh-looking randomness. It is then argued that a black-box simulator can be used to break soundness. Very roughly, this is because a cheating prover can simply run the black-box simulator; if the simulator rewinds the verifier, then the cheating prover answers it with a random message on behalf of the verifier. This proof also extends to WZK because any query made by the simulator to the distinguisher can simply be answered with “reject.”

Note, however, that in the non-adaptive setting, the verifier is not allowed to generate different messages for different instances, and hence the simulator has more power than a cheating prover, since it can fix the first message of the prover

and then test whether the distinguisher accepts or not with various instances and various third round messages. Indeed, we exploit exactly this fact to design a distinguisher-dependent simulator for our protocols.

We next explain why we are able to overcome the lower bound of [37] for two-round ZK. A key argument in the proof of [37] is that no (possibly non-black-box) simulator can simulate the prover’s message for a false statement (even when the protocol is privately verifiable). For ZK, this is argued by setting the verifier’s auxiliary input to be an honestly generated first message and providing the corresponding private randomness to the distinguisher, who is chosen *after* the simulator. Now, if the simulator succeeds, then we can break soundness of the protocol. However, in WZK, since the distinguisher is fixed in advance, the above approach does not work. In particular, if the distinguisher is given the private randomness then the simulator is given it as well (and hence can simulate), and otherwise, the simulator can succeed by simulating a rejecting transcript.

1.3 Related Work

Concurrent Work. Concurrent to our work, Badrinarayanan et al. [2] construct protocols that are similar to our two-round protocols. However their focus is on super-polynomial simulation, whereas we focus on polynomial time distinguisher-dependent simulation. They also give other instantiations of two-round OT, which can be combined with our results to obtain two-round delayed-input distributional weak zero-knowledge from additional assumptions.

Proof Systems. We mention two related works on two-round ZK proofs that overcome the lower bound of [37] in different ways. A recent work of [19] constructs a two-round (T, t, ϵ) -ZK proof system for languages in statistical zero-knowledge, where roughly, (T, t, ϵ) ZK requires the existence of a simulator that simulates the view of the verifier for any distinguisher running in time t and distinguishing probability ϵ . The running time T of the simulator depends upon t and ϵ . In another recent work, [9] construct a two-round ZK argument system against verifiers with auxiliary inputs of a priori bounded size.

Three-round ZK proofs are known either based on non-standard “knowledge assumptions” [41,5], or against adversaries that receive auxiliary inputs of a priori bounded size [9,7]. In contrast, in this work, we consider security against adversaries with non-uniform advice of arbitrarily polynomial size, based on standard cryptographic assumptions.

Finally, we discuss WI, WH and WZK in three rounds. While three round WI is known from injective one-way functions [29], WH and WZK are non-trivial to realize even in three rounds. In particular, [42] proved a lower bound for three-round public-coin WH w.r.t. a natural class of black-box reductions. More recently, [54] extended their result to rule out all black-box reductions. Presently, the only known constructions of three-round WH and WZK for NP require either “knowledge assumptions” [41,5], or rely on the assumption of auxiliary-input point obfuscation (AIPO) and auxiliary-input multi-bit point obfuscation (AIMPO), respectively, with an additional “recognizability” property [10]. For

general auxiliary inputs, however, AIMPO was recently proven to be impossible w.r.t. general auxiliary inputs [14], assuming the existence of indistinguishability obfuscation [4]. Further, one of the assumptions used by [10] to build recognizable AIPO, namely, strong DDH assumption [15], was recently shown to be impossible w.r.t. general auxiliary inputs [6], assuming the existence of virtual grey-box obfuscation [8].

Secure Computation. Katz and Ostrovsky [47] constructed a four-round two-party computation protocol for general functions where only one party receives the output. A recent work of Garg et al. [31] extends their result to the simultaneous-message model to obtain a four-round protocol where both parties receive the outputs.

The notion of input-indistinguishable computation (IIC) was introduced by Micali, Pass and Rosen [49] as a weakening of standard simulation-based security notion for secure computation while still providing meaningful security. (See also [30,52].) We provide the first three-round protocol that provides input-indistinguishable security.

A recent work of Döttling et al. [26] constructs a two-round two-party computation protocol for oblivious computation of cryptographic functionalities. They consider semi-honest senders and malicious receivers, and prove game-based security against the latter. We remark that our three-round two-party computation protocol can be easily downgraded to a two-round protocol that achieves weak simulation security against malicious receivers and super-polynomial time simulation security against malicious senders (or polynomial-time simulation against semi-honest senders). We note that our result is incomparable to [26], because we consider a restricted class of distributions (such as product distributions), albeit any functionality, whereas [26] considers the class of cryptographic functionalities.

1.4 Organization

The rest of this paper is organized as follows. We begin with an overview of our techniques in Section 2. In Section 3, we describe important relevant preliminaries including Σ -protocols and oblivious transfer. In Section 4, we recall definitions of adaptive soundness, witness indistinguishability, distributional weak-ZK and witness hiding against non-adaptive verifiers. In Section 5, we describe our two-round protocol, which uses any Σ -protocol with a special structure, together with 2-message OT. In the same section, we describe how to modify our protocol so as to rely on *any* Σ -protocol, and also show how to base security on polynomial hardness assumptions at the cost of adding an extra round. Due to lack of space, we defer additional details of our three round protocols and their applications to the full version of the paper.

2 Technical Overview

We now give an overview of our main ideas and techniques.

2.1 Argument Systems

We construct a two-round argument system, which we prove is both witness indistinguishable (against all malicious verifiers), and is distributional ϵ -weak zero-knowledge (against non-adaptive malicious verifiers). Our protocol makes use of two components:

- Any Σ -protocol consisting of three messages (a, e, z) that is secure against unbounded provers,
- Any two-message oblivious transfer protocol, denoted by $(\text{OT}_1, \text{OT}_2)$, which is secure against malicious PPT receivers, and malicious senders running in time at most $2^{|z|}$. For receiver input b and sender input messages (m_0, m_1) , we denote the two messages of the OT protocol as $\text{OT}_1(b)$ and $\text{OT}_2(m_0, m_1)$. We note that $\text{OT}_2(m_0, m_1)$ also depends on the message $\text{OT}_1(b)$ sent by the receiver. For the sake of simplicity, we omit this dependence from the notation.

For simplicity, throughout most of the paper, we assume that the Σ -protocol is a parallel repetition of Σ -protocols with a single-bit challenge and constant soundness¹⁴. Namely, we assume that the Σ -protocol contains three messages, denoted by (a, e, z) and that these messages can be parsed as $a = (a_1, \dots, a_\kappa)$, $e = (e_1, \dots, e_\kappa)$, and $z = (z_1, \dots, z_\kappa)$, where for each $i \in [\kappa]$, the triplet (a_i, e_i, z_i) are messages corresponding to an underlying Σ -protocol with a single-bit challenge (i.e., where $e_i \in \{0, 1\}$). We denote by f_1 and f_2 the functions that satisfy $a_i = f_1(x, w; r_i)$ and $z_i = f_2(x, w, r_i, e_i)$, for answers provided by the honest prover, and where r_i is uniformly chosen randomness.

We show how to convert any such Σ -protocol into a two-round protocol (P, V) using OT. Our transformation is essentially the same as the one suggested by Aeillo et. al. [1], and used by Kalai and Raz [46], to reduce rounds in interactive protocols, except that we use an OT scheme rather than a computational PIR scheme (since as opposed to [1,46] we are not concerned with compressing the length of the messages). Specifically, given any such Σ -protocol and OT protocol, our two-round protocol (P, V) , proceeds as follows.

- For $i \in [\kappa]$, V picks $e_i \xleftarrow{\$} \{0, 1\}$, and sends $\text{OT}_{1,i}(e_i)$ in parallel. Each e_i is encrypted with a fresh OT instance.
- For $i \in [\kappa]$, P computes $a_i = f_1(x, w; r_i)$, $z_i^{(0)} = f_2(x, w, r_i, 0)$, $z_i^{(1)} = f_2(x, w, r_i, 1)$. The prover P then sends $a_i, \text{OT}_{2,i}(z_i^{(0)}, z_i^{(1)})$ in parallel for all $i \in [\kappa]$.
- The verifier V recovers $z_i^{(e_i)}$ from the OT, and accepts if and only if for every $i \in [\kappa]$, the transcript $(a_i, e_i, z_i^{(e_i)})$ is an accepting transcript of the underlying Σ -protocol.

Soundness. It was proven in [46] that such a transformation from any public-coin interactive proof to a two-round argument preserves soundness against PPT

¹⁴ We later describe how garbled circuits can be used in order to modify our construction to work with any Σ -protocol.

provers. We extend their proof to show that the resulting two-round protocol also satisfies *adaptive* soundness, i.e., is sound against cheating provers that may adaptively choose some instance x as a function of the verifier message.

To prove soundness, we rely on the following special-soundness property of Σ -protocols: There exists a polynomial-time algorithm A that given any instance x of some NP language L with witness relation R_L , and a pair of accepting transcripts $(a, e, z), (a, e', z')$ for x with the same first prover message, where $e \neq e'$, outputs w such that $w \in R_L(x)$. In particular, this means that for any $x \notin L$, for any fixed message a , there exists at most *one* unique value of receiver challenge e , for which there exists z such that (a, e, z) is an accepting transcript (as otherwise the algorithm A would output a witness $w \in R_L(x)$, which is impossible).

Going back to our protocol – suppose a cheating prover, on input the verifier message $\text{OT}_1(e^*)$, outputs $x^* \notin L$, together with messages $a^*, \text{OT}_2(z^*)$, such that the verifier accepts with non-negligible probability. Since, for any $x^* \notin L$ and any a^* , there exists at most one unique value of receiver challenge e , for which there exists a z that causes the verifier to accept – intuitively, this means that a^* encodes the receiver challenge e^* .

Thus, for fixed a^* , a reduction can enumerate over all possible values of z (corresponding to all possible e), and check which single e results in an accepting transcript. Then, this would allow a reduction to break receiver security of the oblivious transfer. Since such a reduction would require time at least $2^{|z|}$, we need the underlying oblivious transfer to be $2^{|z|}$ -secure (or, sub-exponentially secure). If z can be scaled down to be of size poly-logarithmic in the security parameter, we can rely on an oblivious transfer protocol which is quasi-polynomially secure against malicious receivers.

A New Extraction Technique for Proving Weaker Notions of Zero-Knowledge. We now proceed to describe our main ideas for proving the privacy guarantees of our protocol. For simplicity, consider a single repetition of the protocol outlined above. That is, consider a protocol where the verifier picks a random **bit** $e \leftarrow \{0, 1\}$ and sends $r = \text{OT}_1(e)$ to the prover. The prover then sends $a, \text{OT}_2(z^{(0)}, z^{(1)})$ to the verifier, where $(a, z^{(0)}, z^{(1)})$ are computed similarly as before.

By the security of the underlying OT scheme against malicious receivers (see Definition 2 and discussion therein), the following holds: For any malicious verifier (i.e. malicious receiver of the OT scheme) there exists a (possibly inefficient) simulator that interacts with an ideal OT functionality and is able to simulate the view of the verifier. This means that for any PPT distinguisher \mathcal{D}_V (that obtains as input the view of the verifier and additional auxiliary information), its output distribution when the prover sends $(a, \text{OT}_2(z^{(0)}, z^{(1)}))$ is indistinguishable from one of the following:

- Its output distribution when the prover sends $(a, \text{OT}_2(z^{(0)}, z^{(0)}))$ (implicitly corresponding to receiver choice bit 0).
- Its distribution output when the prover sends $(a, \text{OT}_2(z^{(1)}, z^{(1)}))$ (implicitly corresponding to receiver choice bit 1).

Suppose the message of the verifier, $\text{OT}_1(e)$ is generated independently of the instance x , and suppose that the instance x is generated according to some distribution \mathcal{D} . Then an extractor \mathcal{E} , given the message $\text{OT}_1(e)$, can guess e (if the distinguisher “knows” e), up to ϵ -error in time $\text{poly}(1/\epsilon)$, as follows: The extractor will generate $\text{poly}(1/\epsilon)$ many instance-witness pairs $(x, w) \in R_L$, where each x is distributed independently from \mathcal{D} (\mathcal{E} will have these instance-witness pairs hard-wired if they are hard to sample). Then for each such instance-witness pair the extractor will generate $(a, z^{(0)}, z^{(1)})$, and will observe the distinguisher’s output corresponding to the prover’s message $(a, \text{OT}_2(z^{(0)}, z^{(0)}))$, $(a, \text{OT}_2(z^{(1)}, z^{(1)}))$, and $(a, \text{OT}_2(z^{(0)}, z^{(1)}))$. If the distinguisher cannot distinguish between these three distributions then the extractor outputs \perp (indicating that the distinguisher does not know e). If the extractor outputs \perp , the distinguisher is (distributionally) insensitive to the prover’s response, so we can behave as if it was approximated to 0.

However, if the distinguisher can distinguish between $(a, \text{OT}_2(z^{(0)}, z^{(1)}))$ and $(a, \text{OT}_2(z^{(b)}, z^{(b)}))$, then the distinguisher will guess $e = 1 - b$. In this way, the extractor can approximate (up to ϵ -error) whether the implicit receiver choice bit is 0 or 1, while running in time $\text{poly}(1/\epsilon)$. This idea forms the basis of our new extraction technique.

Witness Indistinguishability. Since witness indistinguishability is known to compose under parallel repetition, it suffices to prove WI for a single repetition of the protocol outlined above. In fact, we will try to prove something even stronger.

As explained above, there exists a distinguisher-dependent simulator $\text{Sim}_{\mathcal{D}_V}$, that, given a fixed receiver message r , can try to approximate the verifier’s implicit challenge bit e , by observing the distinguisher’s output corresponding to various sender messages, up to error ϵ . Once $\text{Sim}_{\mathcal{D}_V}$ has successfully extracted the verifier’s challenge, it can use the honest-verifier zero-knowledge simulator of the underlying Σ -protocol.

Of course, to even begin the extraction process, $\text{Sim}_{\mathcal{D}_V}$ needs to observe the output of the distinguisher on $(a, \text{OT}_2(z^{(0)}, z^{(1)}))$. However, even computing $(a, \text{OT}_2(z^{(0)}, z^{(1)}))$ correctly, requires access to a witness! This is because a correctly compute tuple $(a, z^{(0)}, z^{(1)})$ actually *encodes a witness*.

In the case of witness indistinguishability, this is not a problem – since an “intermediate” simulator for witness indistinguishability has access to both witnesses in question, and therefore *can* generate valid messages $(a, \text{OT}_2(z^0, z^1))$ using both witnesses. It can use these transcripts to learn the verifier’s challenge bit, and then use the bit it learned, to generate a simulated transcript for the same receiver message r (where the simulated transcript uses neither of the two witnesses). We mainly rely on OT security to show that the distinguisher \mathcal{D}_V cannot distinguish between the view generated by such a simulator $\text{Sim}_{\mathcal{D}_V}$ and the real view of the verifier, when he interacts with an honest prover that uses only one of the witnesses.

There are additional subtleties in the proof, for instance, in ensuring that the extracted values when the simulator uses one particular witness for learning, do

not contradict the values extracted when it uses the other witness. We refer the reader to Section 5.3 for a detailed proof.

Distributional Weak Zero-Knowledge. We prove that the same protocol satisfies distributional weak zero-knowledge against non-adaptive verifiers (which can also be easily seen to imply witness-hiding against non-adaptive verifiers). Distributional weak zero-knowledge is a “distributional” relaxation of the standard notion of zero-knowledge where the simulator is additionally allowed to depend on the distribution of instances, and on the distinguisher. This notion roughly requires that for every distribution \mathcal{X} over instances, every verifier V and distinguisher \mathcal{D}_V that obtains the view of V , every $\epsilon = \frac{1}{\text{poly}(\kappa)}$ for some polynomial $\text{poly}(\cdot)$, there exists a simulator $\text{Sim}_{\mathcal{D}_V}$ that runs in time $\text{poly}(1/\epsilon)$ and outputs a view, such that the distinguisher \mathcal{D}_V has at most ϵ -advantage in distinguishing the real view of V from the simulated view.

Fix the first message of the verifier (since the verifier is non-adaptive, this is fixed independently of the instance). The simulator $\text{Sim}_{\mathcal{D}_V}$ obtains as (non-uniform) advice, $\text{poly}(1/\epsilon)$ *randomly chosen* instance-witness pairs from the distribution in question.¹⁵ It then uses these pairs together with the extraction strategy \mathcal{E} described above, to “learn” an approximation to the verifier’s implicit challenge string in the fixed verifier message. However, distributional weak zero-knowledge is not known to be closed under parallel composition. Therefore, we modify the simple extraction strategy described previously for a single repetition, so as to extract *all* bits of the verifier’s challenge, while still remaining efficient in $\text{poly}(1/\epsilon)$.

This is done inductively: at any time-step $i \in [\kappa]$, the simulator $\text{Sim}_{\mathcal{D}_V}$ has extracted an approximation for the first $(i-1)$ bits of the verifier’s challenge, and is now supposed to extract the i^{th} bit. At a high level, the extraction strategy of $\text{Sim}_{\mathcal{D}_V}$ is as follows:

- It generates a “fake” output for the first $(i-1)$ parallel repetitions as follows: for $j \in [i-1]$, if the j^{th} bit of the verifier’s challenge was approximated to 0, respond with $a_j, (z_j^0, z_j^0)$ in the j^{th} repetition (and similarly, if it was approximated to 1, respond with $a_j, (z_j^1, z_j^1)$).
- For all $j \in [i+1, \kappa]$ it responds honestly with $a_j, (z_j^0, z_j^1)$ in the j^{th} repetition.
- With outputs for all $j < i$ set to “fake” according to approximated challenge, and for all $j > i$ set to honest, at $j = i$, $\text{Sim}_{\mathcal{D}_V}$ uses the extraction strategy \mathcal{E} described above. That is, for $j = i$, it sets the output to $a_i, (z_i^0, z_i^1)$, $a_i, (z_i^0, z_i^0)$, and $a_i, (z_i^1, z_i^1)$, and checks whether the output of the distinguisher when given inputs corresponding to $a_i, \text{OT}_{2,i}(z_i^0, z_i^1)$ is close to its output when given inputs corresponding to $a_i, \text{OT}_{2,i}(z_i^0, z_i^0)$ or to $a_i, \text{OT}_{2,i}(z_i^1, z_i^1)$. It uses this to approximate the i^{th} bit of the verifier’s challenge.

¹⁵ In most cryptographic applications, and in all our applications, it is possible for the simulator to efficiently sample random instance-witness pairs from the distribution on its own, without the need for any non-uniform advice.

Via an inductive hybrid argument, we prove that with high probability, the approximation computed by $\text{Sim}_{\mathcal{D}_V}$ has at most $\Theta(\epsilon)$ -error when $\text{Sim}_{\mathcal{D}_V}$ runs in time $\text{poly}(1/\epsilon)$. Once $\text{Sim}_{\mathcal{D}_V}$ has successfully extracted the verifier’s challenge, it can use the honest-verifier zero-knowledge simulator of the underlying Σ -protocol as before.

Note that in order to perform extraction, the simulator is required to generate various $a_i, \text{OT}_{2,i}(z_i^0, z_i^1)$ tuples, which it does using the instance-witness pairs it sampled or obtained as advice. $\text{Sim}_{\mathcal{D}_V}$ then uses the challenge it extracted to generate fake proofs for various other $x \leftarrow \mathcal{X}$. Non-adaptivity of the verifier ensures that the simulator can, for a fixed verifier messages, generate proofs for several other statements in the distribution while observing the output of the distinguisher. We refer the reader to Section 5.4 for a complete proof.

Three Round Protocols from Polynomial Hardness Assumptions. We also describe how quasi-polynomial assumptions can be avoided at the cost of an extra round. The need for quasi-polynomial assumptions in our two-round protocols is to guarantee soundness: roughly, we require that a cheating prover should be unable to “maul” the receiver’s challenge while providing his message. In the two-round setting, this is achieved by ensuring (via complexity leveraging) that the security of receiver OT message is stronger than the security of the prover’s response. Three rounds, however, give an opportunity to *rewind* and extract the value inside the prover’s message, while relying on (polynomial) hiding of the receiver OT message.

We assume here that the first round of the Σ -protocol consists of commitments to certain values, and the third round consists of decommitments to a subset of these commitments, together with additional auxiliary information (for instance, the Blum protocol for Graph Hamiltonicity satisfies this requirement). We modify the protocol to have the prover send extractable commitments (instead of standard commitments) to commit to the values needed for the first round of the Σ -protocol.

Consider a PPT cheating prover that generates a proof for $x \notin L$. A reduction can obtain the receiver OT message externally as an encryption of some n -bit challenge, and then extract the values committed by the prover. Because the underlying Σ -protocol is special-sound against unbounded provers, any accepting proof for $x \notin L$, will allow recovering the receiver challenge directly by observing the values committed by the prover. We must, however, ensure that adding such extractable commitments does not harm privacy – since our simulator is required to generate several actual proofs before it is able to output a simulated proof. To accomplish this, we design a special kind of (weakly) extracting commitments, details of which can be found in Section 5.7. We note here that over-extraction suffices, in particular, we only care about extracting the values committed by provers that generate accepting transcripts.

2.2 Applications

We now describe some applications of our proof systems. As a first step, we describe a transformation from our three-round distributional WZK argument

system to an *argument of knowledge*¹⁶ (that retains the distributional weak ZK/strong WI property against non-adaptive verifiers).

Weak ZK/Strong WI Argument of Knowledge. We begin with the following simple idea for a distributional weak ZKAoK, for instances $x \leftarrow \mathcal{X}$: Let us use a delayed-input witness indistinguishable adaptive proof of knowledge (WIPoK), for instance the Lapidot-Shamir proof [48], to prove the following statement:

Either $x \in L$, OR, \exists randomness r such that $c = \text{com}(1^\kappa; r)$.

Here, the commitment string c is also chosen and sent in the last round, together with instance x . Furthermore, to ensure that a (cheating) prover indeed uses the witness for $x \in L$, the prover must also give a weak ZK proof, for the same string c that $\exists r$ such that $c = \text{com}(0^\kappa; r)$. The argument of knowledge property of this protocol now follows from the proof of knowledge property of WIPoK, the soundness of the weak ZK argument, and the statistical binding property of the commitment scheme. Specifically, by adaptive soundness of the weak ZK proof, c must indeed be constructed as a commitment to 0^κ ; moreover, by the statistical binding property of the commitment scheme, the same string c cannot be a commitment to 1^κ . Therefore, the only possible witness that can be extracted from the WIPoK is indeed a witness for the instance x .

To prove weak ZK/strong WI property for the same protocol, we would ideally like to have the following sequence of hybrid arguments: First, we start simulating the weak ZK proof, by observing the output of the distinguisher on several different instances from the distribution \mathcal{X} , while using correct witnesses for these instances. We then use the information learned to simulate the weak ZK proof for c obtained externally in the main transcript. Since the string c is not used in the main thread at all, we change it so that $\text{com}(0^\kappa; r)$ for uniformly random r . Next, we must begin using (c, r) as witnesses in the WIPoK, instead of using the witness for x .

It is in this step that there arises a subtle issue, because of the way our simulator works. In each experiment, before it can generate a simulated proof, it must first generate several real proofs for other random instances. We require the WIPoK to maintain witness indistinguishability, *even when the simulator provides multiple proofs for different instances using the same first two messages*. This is in general, not true for proof systems such as Lapidot-Shamir [48]. This is also not as strong a requirement as resettable-WI [16] since the verifier’s message is fixed and remains the same for all proofs.

We refer to this property as *reusable* WI and construct an adaptively sound argument of knowledge satisfying this property. The argument of knowledge works by the prover sending two three-round extractable commitments (with “over” extraction) [55,56] to random strings, encrypting the witness with each of these strings using standard private key encryption, and sending a three-round delayed-input reusable WI argument (this does not need to be an argument

¹⁶ Despite using a variant of extractable commitments, the three-round argument described in the previous section not a standard AoK in the delayed-input setting.

of knowledge, and could be instantiated with a ZAP, or with our three round arguments) to establish that one of the two commitments is a valid extractable commitment, and the corresponding ciphertext correctly encrypts the witness. The use of private key encryption gives us the additional desired property of reusability.

Extractable Commitments. Given the weak ZK argument of knowledge, our construction of three-round extractable commitments simply consists of sending a non-interactive statistically binding commitment to the message in the last round, together with a (distributional) weak ZK argument of knowledge to establish knowledge of the committed message and randomness. The weak ZK property helps prove hiding of this scheme, while the proof of knowledge property guarantees correct polynomial-time extraction, with overwhelming probability. We refer the reader to the full version for details.

Three Round, Two Party, Input-Indistinguishable Secure Computation. We begin by considering the following two-round protocol for two-party computation: The receiver generates OT messages corresponding to his inputs, together with the first message of a two-round weak ZK argument. Then, the sender generates garbled circuits corresponding to his own input labels, together with the second message of the two-round weak ZK argument.

This protocol already satisfies input-indistinguishable security against malicious receivers, as well as distinguisher-dependent security against malicious receivers, when an honest sender's input is sampled from some public distribution. Even though our weak ZK proof guarantees hiding against malicious receivers, security is not immediate. Indeed, we must first *extract* an adversarial receiver's input from his OT messages, and weak ZK does not help with that. Thus, apart from simulating the weak ZK, we must use our extraction strategy in this context, in order to (distributionally) learn the receiver's input.

In the full version, we describe applications of our techniques to obtaining input-indistinguishable secure computation, as well as distributional distinguisher-dependent secure computation in three rounds from polynomial assumptions. In particular, we also note that a large class of functionalities such as coin tossing, generating common *reference* strings, oblivious PRFs, etc. (that we call *independent-input functions*) are distributional by definition, and can be realized with distinguisher-dependent polynomial simulation security in three rounds.

3 Preliminaries

Throughout this paper, we will use κ to denote the security parameter, and $\text{negl}(\kappa)$ to denote any function that is asymptotically smaller than $\frac{1}{\text{poly}(\kappa)}$ for any polynomial $\text{poly}(\cdot)$.

Definition 1 (Σ -protocols). Let $L \in \text{NP}$ with corresponding witness relation R_L , and let x denote an instance with corresponding witness $w(x)$. A protocol $\Pi = (P, V)$ is a Σ -protocol for relation R_L if it is a three-round public-coin protocol, and the following requirements hold:

- **Completeness:** $\Pr[\langle P(x, w(x)), V(x) \rangle = 1] = 1 - \text{negl}(\kappa)$, assuming P and V follow the protocol honestly.
- **Special Soundness:** There exists a polynomial-time algorithm A that given any x and a pair of accepting transcripts $(a, e, z), (a, e', z')$ for x with the same first prover message, where $e \neq e'$, outputs w such that $w \in R_L(x)$.
- **Honest verifier zero-knowledge:** There exists a probabilistic polynomial time simulator S_Σ such that

$$\{S_\Sigma(x, e)\}_{x \in L, e \in \{0,1\}^\kappa} \approx_c \{\langle P(x, w(x)), V(x, e) \rangle\}_{x \in L, e \in \{0,1\}^\kappa}$$

where $S_\Sigma(x, e)$ denotes the output of simulator S upon input x and e , and $\langle P(x, w(x)), V(x, e) \rangle$ denotes the output transcript of an execution between P and V , where P has input (x, w) , V has input x and V 's random tape (determining its query) is e .

Definition 2 (Oblivious Transfer). Oblivious transfer is a protocol between two parties, a sender S with messages (m_0, m_1) and receiver R with input a choice bit b , such that R obtains output m_b at the end of the protocol. We let $\langle S(m_0, m_1), R(b) \rangle$ denote an execution of the OT protocol with sender input (m_0, m_1) and receiver input bit b . It additionally satisfies the following properties.

Receiver Security. For any sender S^* , all auxiliary inputs $z \in \{0,1\}^*$, and all $(b, b') \in \{0,1\}$, $\text{View}_{S^*}(\langle S^*(z), R(b) \rangle) \approx_c \text{View}_{S^*}(\langle S^*(z), R(b') \rangle)$.

Sender Security. This is defined using the real-ideal paradigm, and requires that for all auxiliary inputs $z \in \{0,1\}^*$, every distribution on the inputs (m_0, m_1) and any adversarial receiver R^* , there exists a (possibly unbounded) simulator Sim_{R^*} that interacts with an ideal functionality \mathcal{F}_{ot} on behalf of R^* . Here \mathcal{F}_{ot} is an oracle that obtains the inputs (m_0, m_1) from the sender and b from the Sim_{R^*} (simulating the malicious receiver), and outputs m_b to Sim_{R^*} . Then $\text{Sim}_{R^*}^{\mathcal{F}_{\text{ot}}}$ outputs a receiver view V_{Sim} that is computationally indistinguishable from the real view of the malicious receiver $\text{View}_{R^*}(\langle S(m_0, m_1, z), R^* \rangle)$.

We will make use of **two-message** oblivious-transfer protocols with security against malicious receivers and semi-honest senders. Such protocols have been constructed based on the DDH assumption [51], and a stronger variant of smooth-projective hashing, which can be realized from DDH as well as the N^{th} -residuosity and Quadratic Residuosity assumptions [45,43]. Such protocols can also be based on indistinguishability obfuscation (*iO*) together with one-way functions [57].

We will use the following sender security property in our protocols (which is implied by the definition of sender security in Definition 2 above). For any fixed first message generated by a malicious receiver R^* , we require that either of the following statements is true:

- For all m_0, m_1 , $\text{View}_{R^*}(\langle S(m_0, m_1, z), R^* \rangle) \approx_c \text{View}_{R^*}(\langle S(m_0, m_0, z), R^* \rangle)$
- Or, for all m_0, m_1 , $\text{View}_{R^*}(\langle S(m_0, m_1, z), R^* \rangle) \approx_c \text{View}_{R^*}(\langle S(m_1, m_1, z), R^* \rangle)$

This follows from the (unbounded) simulation property, i.e., there exists a simulator that extracts some receiver input b from the first message of R^* , sends it

to the ideal functionality, obtains m_b and generates an indistinguishable receiver view. Then, by the definition of sender security, the simulated view must be close to both $\text{View}_{R^*}(\langle S(m_0, m_1, z), R^* \rangle)$, and $\text{View}_{R^*}(\langle S(m_b, m_b, z), R^* \rangle)$.

We also note that all the aforementioned instantiations of two-message oblivious-transfer are additionally secure against *unbounded* malicious receivers.

4 Definitions

4.1 Proof Systems

Delayed-Input Interactive Protocols. An n -round delayed-input interactive protocol (P, V) for deciding a language L with associated relation R_L proceeds in the following manner:

- At the beginning of the protocol, P and V receive the size of the instance and execute the first $n - 1$ rounds.
- At the start of the last round, P receives an input $(x, w) \in R_L$ and V receives x . Upon receiving the last round message from P , V outputs 1 or 0.

An execution of (P, V) with instance x and witness w is denoted as $\langle P, V \rangle(x, w)$. Whenever clear from context, we also use the same notation to denote the output of V . **Delayed-Input Interactive Arguments.** An n -round delayed-input

interactive argument for a language L must satisfy the standard notion of completeness as well as *adaptive soundness*, where the soundness requirement holds even against malicious PPT provers who choose the statement adaptively, depending upon the first $n - 1$ rounds of the protocol.

Definition 3 (Delayed-Input Interactive Arguments). *An n -round delayed-input interactive protocol (P, V) for deciding a language L is an interactive argument for L if it satisfies the following properties:*

- **Completeness:** *For every $(x, w) \in R_L$,*

$$\Pr[\langle P, V \rangle(x, w) = 1] \geq 1 - \text{negl}(\kappa),$$

where the probability is over the random coins of P and V .

- **Adaptive Soundness:** *For every $z \in \{0, 1\}^*$, every PPT prover P^* that chooses $x \in \{0, 1\}^\kappa \setminus L$ adaptively, depending upon the first $n - 1$ rounds,*

$$\Pr[\langle P^*(z), V \rangle(x) = 1] \leq \text{negl}(\kappa),$$

where the probability is over the random coins of V .

Witness Indistinguishability. A proof system is witness indistinguishable if for any statement with at least two witnesses, proofs computed using different witnesses are indistinguishable.

Definition 4 (Witness Indistinguishability). A delayed-input interactive argument (P, V) for a language L is said to be witness-indistinguishable if for every non-uniform PPT verifier V^* , every $z \in \{0, 1\}^*$, and every sequence (x, w_1, w_2) such that $w_1, w_2 \in R_L(x)$, the following two ensembles are computationally indistinguishable:

$$\{\langle P, V^*(z) \rangle(x, w_1)\} \text{ and } \{\langle P, V^*(z) \rangle(x, w_2)\}$$

Non-adaptive Distributional Weak Zero Knowledge. Zero knowledge (ZK) requires that for any adversarial verifier, there exists a simulator that can produce a view that is indistinguishable from the real one to *every* distinguisher. Weak zero knowledge (WZK) relaxes the standard notion of ZK by reversing the order of quantifiers, and allowing the simulator to depend on the distinguisher.

We consider a variant of WZK, namely, distributional WZK [34, 28], where the instances are chosen from some hard distribution over the language. Furthermore, we allow the simulator's running time to depend upon the distinguishing probability of the distinguisher. We refer to this as distributional ϵ -WZK, which says that for every distinguisher D with distinguishing probability ϵ (where ϵ is an inverse polynomial) there exists a simulator with running time polynomial in ϵ . This notion was previously considered in [28, 19].

We define distributional ϵ -WZK property against *non-adaptive* malicious verifiers that receive the instance only in the last round of the protocol.

Definition 5 (Non-adaptive Distributional ϵ -Weak Zero Knowledge). A delayed-input interactive argument (P, V) for a language L is said to be distributional ϵ -weak zero knowledge against non-adaptive verifiers if for every efficiently samplable distribution $(\mathcal{X}_\kappa, \mathcal{W}_\kappa)$ on R_L , i.e., $\text{Supp}(\mathcal{X}_\kappa, \mathcal{W}_\kappa) = \{(x, w) : x \in L \cap \{0, 1\}^\kappa, w \in R_L(x)\}$, every non-adaptive PPT verifier V^* , every $z \in \{0, 1\}^*$, every PPT distinguisher \mathcal{D} , and every $\epsilon = 1/\text{poly}(\kappa)$, there exists a simulator \mathcal{S} that runs in time $\text{poly}(\kappa, \epsilon)$ such that:

$$\left| \Pr_{(x, w) \leftarrow (\mathcal{X}_\kappa, \mathcal{W}_\kappa)} [\mathcal{D}(x, z, \text{View}_{V^*}[\langle P, V^*(z) \rangle(x, w)]) = 1] - \Pr_{(x, w) \leftarrow (\mathcal{X}_\kappa, \mathcal{W}_\kappa)} [\mathcal{D}(x, z, \mathcal{S}^{V^*, D}(x, z)) = 1] \right| \leq \epsilon(\kappa),$$

where the probability is over the random choices of (x, w) as well as the random coins of the parties.

Non-adaptive Witness Hiding. Let L be an NP language and let $(\mathcal{X}, \mathcal{W})$ be a distribution over the associated relation R_L . A proof system is witness hiding w.r.t. $(\mathcal{X}, \mathcal{W})$ if for any $(x, w) \leftarrow (\mathcal{X}, \mathcal{W})$, a proof for x is “one-way” in the sense that no verifier can extract a witness for x from its interaction with the prover. Note that in order for WH to be non-trivial, it is necessary that $(\mathcal{X}, \mathcal{W})$ be a “hard” distribution.

Below, we define witness hiding property against *non-adaptive* malicious verifiers that receive the instance only in the last round of the protocol.

Definition 6 (Hard Distributions). Let $(\mathcal{X}, \mathcal{W}) = (\mathcal{X}_\kappa, \mathcal{W}_\kappa)_{\kappa \in \mathbb{N}}$ be an efficiently samplable distribution on R_L , i.e., $\text{Supp}(\mathcal{X}_\kappa, \mathcal{W}_\kappa) = \{(x, w) : x \in L \cap \{0, 1\}^\kappa, w \in R_L(x)\}$. We say that $(\mathcal{X}, \mathcal{W})$ is hard if for any poly-size circuit family $\{C_\kappa\}$, it holds that:

$$\Pr_{(x,w) \leftarrow (\mathcal{X}_\kappa, \mathcal{W}_\kappa)} [C_\kappa(x) \in R_L(x)] \leq \text{negl}(\kappa).$$

Definition 7 (Non-adaptive Witness Hiding). A delayed-input interactive argument (P, V) for a language L is said to be witness hiding against non-adaptive verifiers w.r.t. a hard distribution $(\mathcal{X}_\kappa, \mathcal{W}_\kappa)$ if for every non-adaptive PPT verifier V^* , every $z \in \{0, 1\}^*$, it holds that:

$$\Pr_{(x,w) \leftarrow (\mathcal{X}_\kappa, \mathcal{W}_\kappa)} [\langle P, V^*(z) \rangle(x) \in R_L(x)] \leq \text{negl}(\kappa).$$

Non-adaptive Strong Witness Indistinguishability

Definition 8 (Non-adaptive Strong Witness Indistinguishability). A delayed-input interactive argument (P, V) for a language L is said to be strong witness indistinguishable against non-adaptive verifiers w.r.t. a pair of indistinguishable distributions $(\mathcal{X}_{1,\kappa}, \mathcal{W}_{1,\kappa}), (\mathcal{X}_{2,\kappa}, \mathcal{W}_{2,\kappa})$ if for every non-adaptive PPT verifier V^* , every $z \in \{0, 1\}^*$, it holds that:

$$\left| \Pr_{(x,w) \leftarrow (\mathcal{X}_{1,\kappa}, \mathcal{W}_{1,\kappa})} [\mathcal{D}(x, z, \text{View}_{V^*}[\langle P, V^*(z) \rangle(x, w)]) = 1] - \Pr_{(x,w) \leftarrow (\mathcal{X}_{2,\kappa}, \mathcal{W}_{2,\kappa})} [\mathcal{D}(x, z, \text{View}_{V^*}[\langle P, V^*(z) \rangle(x, w)]) = 1] \right| \leq \text{negl}(\kappa).$$

Remark 1. A non-adaptive distributional weak ZK argument of knowledge is an argument of knowledge that satisfies the distributional weak ZK property against non-adaptive verifiers. Similarly, a non-adaptive strong WI argument of knowledge is an argument of knowledge that satisfies the strong WI property against non-adaptive verifiers. Finally, a non-adaptive witness hiding argument of knowledge can be defined similarly as an argument of knowledge that satisfies the witness hiding property against non-adaptive verifiers.

5 Two Round Argument Systems

5.1 Construction

We show how to use two-message malicious-secure oblivious transfer (OT) to convert any three-message Σ -protocol according to Definition 1, into a two-message argument system. We then prove soundness of the resulting argument

system, assuming sub-exponential security of oblivious transfer. We also prove that this protocol is witness indistinguishable, satisfies distributional weak zero-knowledge, strong WI and witness hiding against non-adaptive verifiers.

Let $\text{OT} = (\text{OT}_1, \text{OT}_2)$ denote a two-message bit oblivious transfer protocol according to Definition 2. Let $\text{OT}_1(b)$ denote the first message of the OT protocol with receiver input b , and let $\text{OT}_2(m_0, m_1)$ denote the second message of the OT protocol with sender input bits m_0, m_1 .

Let $\Sigma = (a, e, z)$ denote the three messages of a Σ -protocol. For most of this paper, we consider Σ -protocols that are a parallel composition of individual protocols with a single-bit challenge and constant soundness, i.e., the Σ -protocol contains three messages, denoted by (a, e, z) and that these messages can be parsed as $a = (a_1, \dots, a_\kappa)$, $e = (e_1, \dots, e_\kappa)$, and $z = (z_1, \dots, z_\kappa)$, where for each $i \in [\kappa]$, the triplet (a_i, e_i, z_i) are messages corresponding to an underlying Σ -protocol with a single-bit challenge (i.e., where $e_i \in \{0, 1\}$). We denote by f_1 and f_2 the functions that satisfy $a_i = f_1(x, w; r_i)$ and $z_i = f_2(x, w, r_i, e_i)$, where r_i is uniformly chosen randomness.

Examples of such Σ -protocols are the parallel Blum proof of Graph Hamiltonicity [13], and the Lapidot-Shamir [48] three round WI proof. By a Karp reduction to Graph Hamiltonicity, there exists such a Σ -protocol for all of NP.

<p>Witness Indistinguishable and Weak Distributional Zero-Knowledge Argument</p> <p>Prover Input: Instance $x \in L$, witness w such that $R_L(x, w) = 1$.</p> <p>Verifier Input: Instance x, language L.</p> <ul style="list-style-type: none"> – Verifier Message: The verifier picks challenge $e \xleftarrow{\\$} \{0, 1\}^\kappa$ for the Σ-protocol, and for $i \in [\kappa]$, sends $\text{OT}_{1,i}(e_i)$ in parallel. Each bit e_i is encrypted with a fresh OT instance. – Prover Message: For $i \in [\kappa]$, the prover sends $a_i, \text{OT}_{2,i}(z_i^0, z_i^1)$ in parallel. – Verifier Output: The verifier V recovers z_i as the output of OT_i for $i \in [\kappa]$, and outputs accept if for all $i \in [\kappa]$, $(a_i, e_i, z_i)_{i \in [\kappa]}$ is an accepting transcript of the underlying Σ-protocol.
--

Fig. 1. Two Round Argument System for NP

5.2 Adaptive Soundness

The protocol in Figure 1 compiles a three-round public coin proof to a two-round argument using oblivious transfer. Kalai-Raz [46] proved that such a compiler, applied to any public-coin proof system preserves soundness. Specifically, the following theorem in [46] proves (static) soundness of the above protocol, assuming sub-exponential oblivious transfer.

Imported Theorem 1 (*Rephrased*) Let $\Sigma = (a, e, z)$ denote a Σ -protocol, and let $\ell = \text{poly}(\kappa, s)$ be the size of z , where κ is the security parameter, and s is an upper bound on the length of allowed instances. Assuming the existence of an oblivious transfer protocol secure against probabilistic senders running in time at most 2^ℓ , the protocol in Figure 1 is sound.

We observe that the proof in Kalai-Raz [46] can be extended to prove adaptive soundness, i.e., soundness against malicious provers that can adaptively choose $x \notin L$ based on the verifier's input message.

Lemma 1. Let $\Sigma = (a, e, z)$ denote a Σ -protocol, and let ℓ be the size of z . Assuming the existence of an oblivious transfer protocol secure against probabilistic senders running in time at most 2^ℓ , the protocol in Figure 1 is adaptively sound.

Proof. We will use a prover that breaks soundness to break sub-exponential receiver security of the underlying oblivious transfer. The reduction samples two random challenge strings e_0, e_1 and reduction sends them to an external OT challenger. The external OT challenger picks $b \xleftarrow{\$} \{0, 1\}$, and outputs $\text{OT}_1(e_{i,b})$ for $i \in [\kappa]$, which the reduction forwards to the cheating prover P^* .

P^* outputs $x \notin L$, together with messages $a_i, \text{OT}_2(z_i^0, z_i^1)$ for $i \in [\kappa]$. Next, the reduction R does a brute-force search over all possible values of z , checking whether (a, e_0, z) is an accepting transcript for any $z \in \{0, 1\}^\ell$ and whether (a, e_1, z') is an accepting transcript for any $z' \in \{0, 1\}^\ell$.

Suppose a cheating prover breaks soundness with probability $p = \frac{1}{\text{poly}(\kappa)}$ over the randomness of the experiment. Since the reduction chooses prover messages e_0, e_1 uniformly at random, with probability p , the prover P^* outputs $a_i^*, \text{OT}_2(z_i^0, z_i^1)$ for $i \in [\kappa]$ that cause the verifier to accept.

Thus, with probability p , R finds at least one z such that (a^*, e_b, z) is an accepting transcript.

Since $e_{\bar{b}}$ was picked uniformly at random and independent of e_b , we argue that with at most $\text{negl}(\kappa)$ probability, R finds one or more z' such that $(a^*, e_{\bar{b}}, z')$ is an accepting transcript. Note that with probability $1 - 2^{-\kappa}$, we have that $e_b \neq e_{\bar{b}}$. By special-soundness of the underlying Σ -protocol, if there exists z' such that $(a^*, e_{\bar{b}}, z')$ is an accepting transcript, conditioned on $e_b \neq e_{\bar{b}}$, this would allow obtaining a witness w from (a, e_b, z) and $(a, e_{\bar{b}}, z')$, which is a contradiction since $x \notin L$.

Therefore, if R finds z such that (a^*, e_b, z) is an accepting transcript, R outputs e_b as its guess for the first OT message, and this guess is correct with probability at least $p - \text{negl}(\kappa)$. Since R runs in time 2^ℓ and guesses the OT message with non-negligible probability, this is a contradiction to the security of OT against 2^ℓ -time malicious senders.

Observing the Verifier's output. The protocol is not sound when the prover is allowed to generate a-priori unbounded arguments using the same verifier message, as an adaptive function of the *verifier's accept/reject outputs on prior*

arguments. Looking ahead, such a prover can use the simulation strategy from Section 5.4 to explicitly break soundness.

However, the protocol is sound when the prover is only allowed to generate an *a-priori bounded* arguments that adaptively depend on the verifier’s accept/reject outputs on prior arguments. This can be ensured via simply having the verifier output a longer challenge string – to obtain adaptive soundness for B executions, the protocol requires the verifier to generate $e \xleftarrow{\$} \{0, 1\}^{\kappa \cdot B}$, and encrypt it using $\kappa \cdot B$ OT instances. The prover uses the first κ instances for the first argument, the second set of κ instances for the second, and so forth. It is easy to see then that the argument of Lemma 1 easily extends to the bounded execution case.

5.3 Witness Indistinguishability

We have the following theorem, the proof of which can be found in the full version of the paper.

Theorem 6. *Assuming two-round oblivious transfer (OT) secure against malicious PPT receivers, the two-round protocol in Figure 1 is witness-indistinguishable against PPT verifiers.*

Recall that witness indistinguishability (WI) is closed under parallel composition [29], therefore it suffices to prove WI for a single repetition (i.e., for some $i \in [\kappa]$) of the protocol in Figure 1. Our proof proceeds via a sequence of hybrid arguments, where, in an intermediate hybrid, we construct a distinguisher-dependent simulator, that learns (using both witnesses w_1 and w_2), an approximation for the verifier’s challenge e . Upon learning the challenge, the simulator uses the honest-verifier ZK property to generate a simulated proof, without using any of the witnesses.

5.4 Distributional Weak Zero Knowledge

In this section, we have the following theorem:

Theorem 7. *Assuming oblivious transfer (OT) secure against malicious PPT receivers, the protocol in Figure 1 is distributional weak zero-knowledge against non-adaptive verifiers.*

Proof. (Overview) The proof of weak zero-knowledge is more involved than WI, because weak ZK is not closed under parallel composition. We develop an inductive analysis and a simulation strategy that learns the receiver’s challenge bit-by-bit.

Fix any PPT V^* , any distinguisher \mathcal{D} , any distribution $(\mathcal{X}, \mathcal{W}, \mathcal{Z})$, and any $\epsilon > 0$. We construct a simulator Sim_ϵ that obtains non-uniform advice z , $p_\epsilon = \text{poly}(1/\epsilon)$ random instance-witness samples $(x_1^*, w_1^*), (x_2^*, w_2^*), \dots (x_{p_\epsilon}^*, w_{p_\epsilon}^*)$ from the distribution $(\mathcal{X}, \mathcal{W})$. Or, if the distribution $(\mathcal{X}, \mathcal{W})$ is efficiently samplable, Sim_ϵ samples $(x_1^*, w_1^*), (x_2^*, w_2^*), \dots (x_{p_\epsilon}^*, w_{p_\epsilon}^*)$ these on its own.

At a high level, the simulator uses these instances to approximately-learn the verifier's challenge string e (call this approximation e_{approx}), and then generates a transcript corresponding to a random $x \sim \mathcal{X}$, by using the honest-verifier ZK simulation strategy of the underlying Σ -protocol, corresponding to verifier challenge e_{approx} . Our simulation strategy can be found, together with the complete proof, in the full version of the paper.

5.5 Strong Witness Indistinguishability

We note that the simulator's learning is monotone for two distributions, i.e., given two distributions $\mathcal{X}_1, \mathcal{X}_2$, then the view generated by a simulator Sim_ϵ that learns using samples from both distributions, $\mathcal{X}_1 \cup \mathcal{X}_2$, but outputs the simulation for a sample from \mathcal{X}_1 , is indistinguishable from the view generated by a simulator Sim_ϵ that learns using samples from only \mathcal{X}_1 and then outputs the simulation for a sample from \mathcal{X}_1 .

In other words, learning using additional distributions can only provide “more” information to the simulator. This observation coupled with the proof of weak ZK, directly implies strong witness indistinguishability, when the instances are sampled either from distribution \mathcal{X}_1 or from (an indistinguishable) distribution \mathcal{X}_2 . This is because, the simulator can learn (in all hybrids) using instances from $\mathcal{X}_1 \cup \mathcal{X}_2$, and use these to simulate external samples generated according to either \mathcal{X}_1 or \mathcal{X}_2 .

Corollary 8 *Assuming oblivious transfer (OT) secure against malicious PPT receivers, the protocol in Figure 1 is strong witness-indistinguishable against non-adaptive verifiers.*

5.6 Witness Hiding

It is easy to see that distributional weak zero-knowledge implies witness hiding. Suppose there exists a distribution \mathcal{X}_κ and a PPT verifier V^* with auxiliary input z , that interacts with prover P . P samples random $X \sim \mathcal{X}_\kappa$ together with some witness $W(X)$ and generates a proof for V^* – such that V^* outputs a witness for $X \in \mathcal{X}$ with probability $\gamma = \frac{1}{\text{poly}(\kappa)}$ for some polynomial $\text{poly}(\cdot)$. Then, by the distributional weak zero-knowledge property, there exists a non-uniform simulator Sim_ϵ that uses V^* to output a witness for $X \sim \mathcal{X}$ with probability at least $\gamma - \epsilon$. Setting $\epsilon = \frac{\gamma}{2}$, we obtain a non-uniform polynomial size circuit $(\text{Sim}_\epsilon, V^*)$ that outputs a witness for $X \sim \mathcal{X}$ with probability at least $\gamma/2$, which is a contradiction to the assumption in Definition 7. This implies the following corollary.

Corollary 9 *Assuming two-message oblivious transfer (OT) secure against malicious PPT receivers, the protocol in Figure 1 is witness-hiding against non-adaptive verifiers.*

5.7 Extensions

In this section, we sketch some simple extensions of our main results.

Two Round WI and Distributional WZK from *any* Σ -Protocol. So far, we assumed that the Σ -protocol contains three messages, denoted by (a, e, z) and that these messages can be parsed as $a = (a_1, \dots, a_\kappa)$, $e = (e_1, \dots, e_\kappa)$, and $z = (z_1, \dots, z_\kappa)$, where for each $i \in [\kappa]$, the triplet (a_i, e_i, z_i) are messages corresponding to an underlying Σ -protocol with a single-bit challenge (i.e., where $e_i \in \{0, 1\}$). We denote by f_1 and f_2 the functions that satisfy $a_i = f_1(x, w; r_i)$ and $z_i = f_2(x, w, r_i, e_i)$, where r_i is uniformly chosen randomness.

However, there is a large class of Σ -protocols that do not have this special structure. In Figure 2, we describe how any Σ -protocol can be compiled into 2-message WI and 2-message distributional weak ZK, assuming 2-message malicious-secure OT and garbled circuits. Our protocol is described in Figure 2.

Witness Indistinguishable and Distributional Weak Zero-Knowledge Argument

Prover Input: Instance $x \in L$, witness w such that $R_L(x, w) = 1$.

Verifier Input: Instance x , language L .

- **Verifier Message:** The verifier picks challenge $e \xleftarrow{\$} \{0, 1\}^\kappa$ for the Σ -protocol, and for $i \in [\kappa]$, sends $\text{OT}_{1,i}(e_i)$ in parallel. Each e_i is encrypted with a fresh OT instance.
- **Prover Message:** The prover samples a , and then constructs a garbled circuit $\text{GC}(a, \cdot)$ for a function that on input e (the verifier challenge), outputs the corresponding message z of the underlying Σ -protocol. Let $(\text{label}_i^0, \text{label}_i^1)_{i \in [\kappa]}$ denote the labels of the garbled circuit. The prover sends $a, \text{GC}(a, \cdot)$, together with $\text{OT}_{2,i}(\text{label}_i^0, \text{label}_i^1)$ for all $i \in [\kappa]$.
- **Verifier Output:** The verifier V recovers z as the output of the garbled circuit on the labels obtained via OT, and outputs **accept** if (a, e, z) is an accepting transcript of the underlying Σ -protocol.

Fig. 2. Two Round Argument System for NP from any Σ -Protocol

Three Round Protocols from Polynomial Assumptions. Our three round protocol from polynomial assumptions is described in Figure 3. We denote the three messages of a Σ -protocol by (a, e, z) , and assume that the Σ -protocol is a parallel repetition of protocols with a single bit receiver challenge. We further assume that a consists of a string of commitments, and z contains decommitment information for some of these commitments. We denote the i^{th} set of commitments (in the i^{th} parallel repetition of the Σ -protocol) by $a_i = \text{commit}(h_i)$. We will implement this commitment differently in our protocol in Figure 3. We let com denote a non-interactive statistically binding commitment scheme, and let $\text{wi} = (\text{wi}_1, \text{wi}_2, \text{wi}_3)$ denote the messages of a 3-message delayed-input WI argument for NP. We also assume the existence of dense cryptosystems which are known based on DDH, QR, RSA, etc.

Theorem 10. *There exists a 3-message argument that satisfies distributional weak zero-knowledge, strong witness indistinguishability, witness hiding and witness indistinguishability against non-adaptive malicious verifiers, assuming either polynomially-hard DDH, N^{th} -residuosity or Quadratic Residuosity.*

The proof of soundness, and privacy against malicious verifiers, of the scheme in Figure 3 is deferred to the full version of the paper.

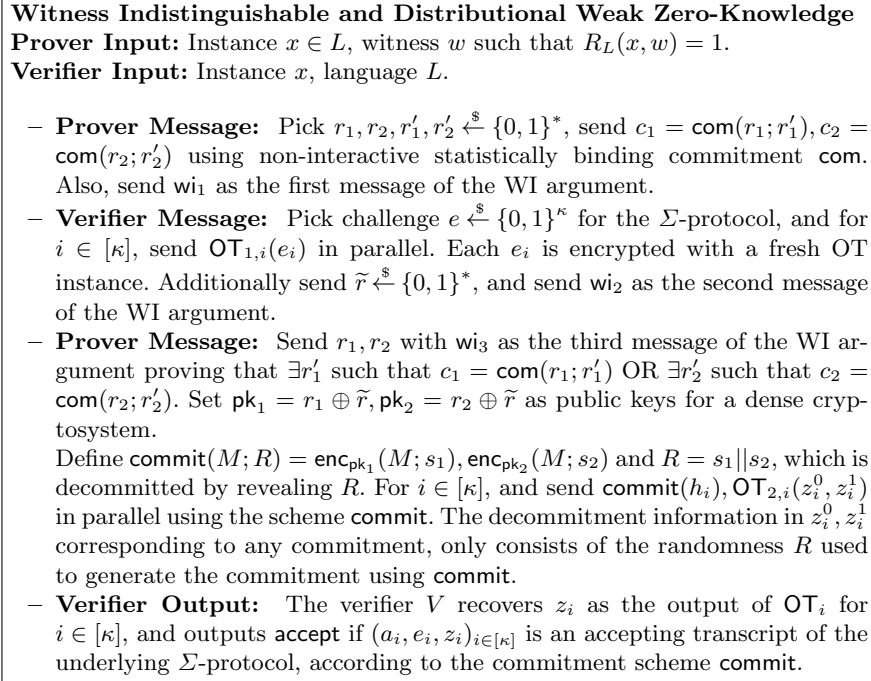


Fig. 3. Three Round Argument System for NP

References

1. Aiello, W., Bhatt, S.N., Ostrovsky, R., Rajagopalan, S.: Fast verification of any remote procedure call: Short witness-indistinguishable one-round proofs for NP. In: Automata, Languages and Programming, 27th International Colloquium, ICALP 2000, Geneva, Switzerland, July 9-15, 2000, Proceedings. pp. 463–474 (2000)
2. Badrinarayanan, S., Garg, S., Ishai, Y., Sahai, A., Wadia, A.: Two-message witness indistinguishability and secure computation in the plain model from new assumptions. IACR Cryptology ePrint Archive 2017, 433 (2017), <http://eprint.iacr.org/2017/433>

3. Barak, B.: How to go beyond the black-box simulation barrier. In: FOCS. pp. 106–115 (2001)
4. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings. pp. 1–18 (2001)
5. Bellare, M., Palacio, A.: The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In: Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings. pp. 273–289 (2004)
6. Bellare, M., Stepanovs, I., Tessaro, S.: Contention in cryptoland: Obfuscation, leakage and UCE. In: Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II. pp. 542–564 (2016)
7. Bitansky, N., Brakerski, Z., Kalai, Y.T., Paneth, O., Vaikuntanathan, V.: 3-message zero knowledge against human ignorance. In: TCC-B (2016)
8. Bitansky, N., Canetti, R.: On strong simulation and composable point obfuscation. In: Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings. pp. 520–537 (2010)
9. Bitansky, N., Canetti, R., Paneth, O., Rosen, A.: On the existence of extractable one-way functions. In: Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014. pp. 505–514 (2014)
10. Bitansky, N., Paneth, O.: Point obfuscation and 3-round zero-knowledge. In: Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings. pp. 190–208 (2012)
11. Bitansky, N., Paneth, O.: Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In: Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II. pp. 401–427 (2015)
12. Bitansky, N., Paneth, O., Wichs, D.: Perfect structure on the edge of chaos - trapdoor permutations from indistinguishability obfuscation. In: Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I. pp. 474–502 (2016)
13. Blum, M.: How to prove a theorem so no one else can claim it. In: Proceedings of the International Congress of Mathematicians. pp. 1444–1451 (1987)
14. Brzuska, C., Mittelbach, A.: Indistinguishability obfuscation versus multi-bit point obfuscation with auxiliary input. In: Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II. pp. 142–161 (2014)
15. Canetti, R.: Towards realizing random oracles: Hash functions that hide all partial information. In: Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings. pp. 455–469 (1997)
16. Canetti, R., Goldreich, O., Goldwasser, S., Micali, S.: Resettable zero-knowledge (extended abstract). In: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA. pp. 235–244 (2000)
17. Chor, B., Goldreich, O., Kushilevitz, E., Sudan, M.: Private information retrieval. In: 36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, 23-25 October 1995. pp. 41–50 (1995)

18. Chung, K., Kalai, Y.T., Vadhan, S.P.: Improved delegation of computation using fully homomorphic encryption. In: *Advances in Cryptology - CRYPTO 2010*, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings. pp. 483–501 (2010)
19. Chung, K., Lui, E., Pass, R.: From weak to strong zero-knowledge and applications. In: *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015*, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I. pp. 66–92 (2015)
20. Ciampi, M., Ostrovsky, R., Siniscalchi, L., Visconti, I.: Concurrent non-malleable commitments (and more) in 3 rounds. In: *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III. pp. 270–299 (2016)
21. Ciampi, M., Persiano, G., Scafuro, A., Siniscalchi, L., Visconti, I.: Improved or-composition of sigma-protocols. In: *Theory of Cryptography - 13th International Conference, TCC 2016-A*, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II. pp. 112–141 (2016)
22. Ciampi, M., Persiano, G., Scafuro, A., Siniscalchi, L., Visconti, I.: Online/offline OR composition of sigma protocols. In: *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Vienna, Austria, May 8-12, 2016, Proceedings, Part II. pp. 63–92 (2016)
23. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: *Advances in Cryptology - CRYPTO '94*, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings. pp. 174–187 (1994)
24. Crescenzo, G.D., Persiano, G., Visconti, I.: Constant-round resettable zero knowledge with concurrent soundness in the bare public-key model. In: *Advances in Cryptology - CRYPTO 2004*, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings. pp. 237–253 (2004)
25. Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography (extended abstract). In: *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, May 5-8, 1991, New Orleans, Louisiana, USA. pp. 542–552 (1991)
26. Döttling, N., Fleischhacker, N., Krupp, J., Schröder, D.: Two-message, oblivious evaluation of cryptographic functionalities. In: *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III. pp. 619–648 (2016)
27. Dwork, C., Naor, M.: Zaps and their applications. In: *41st Annual Symposium on Foundations of Computer Science, FOCS 2000*, 12-14 November 2000, Redondo Beach, California, USA. pp. 283–293 (2000)
28. Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.J.: Magic functions. In: *40th Annual Symposium on Foundations of Computer Science, FOCS '99*, 17-18 October, 1999, New York, NY, USA. pp. 523–534 (1999)
29. Feige, U., Shamir, A.: Witness indistinguishable and witness hiding protocols. In: *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, May 13-17, 1990, Baltimore, Maryland, USA. pp. 416–426 (1990)
30. Garg, S., Goyal, V., Jain, A., Sahai, A.: Concurrently secure computation in constant rounds. In: *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Cambridge, UK, April 15-19, 2012. Proceedings. pp. 99–116 (2012)

31. Garg, S., Mukherjee, P., Pandey, O., Polychroniadou, A.: The exact round complexity of secure computation. In: *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Vienna, Austria, May 8-12, 2016, Proceedings, Part II. pp. 448–476 (2016)
32. Garg, S., Pandey, O., Srinivasan, A., Zhandry, M.: Breaking the sub-exponential barrier in obfuscation. In: *CRYPTO* (2017)
33. Gennaro, R., Gentry, C., Parno, B.: Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In: *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference*, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings. pp. 465–482 (2010)
34. Goldreich, O.: A uniform-complexity treatment of encryption and zero-knowledge. *J. Cryptology* 6(1), 21–53 (1993)
35. Goldreich, O., Krawczyk, H.: On the composition of zero-knowledge proof systems. *SIAM J. Comput.* 25(1), 169–192 (1996)
36. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game. In: *STOC* (1987)
37. Goldreich, O., Oren, Y.: Definitions and properties of zero-knowledge proof systems. *J. Cryptology* 7(1), 1–32 (1994)
38. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems. In: *STOC*. pp. 291–304 (1985)
39. Goyal, V., Pandey, O., Richelson, S.: Textbook non-malleable commitments. In: Wicks, D., Mansour, Y. (eds.) *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*, Cambridge, MA, USA, June 18-21, 2016. pp. 1128–1141. ACM (2016), <http://doi.acm.org/10.1145/2897518.2897657>
40. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for NIZK. In: *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference*, Santa Barbara, California, USA, August 20-24, 2006, Proceedings. pp. 97–111 (2006)
41. Hada, S., Tanaka, T.: On the existence of 3-round zero-knowledge protocols. In: *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference*, Santa Barbara, California, USA, August 23-27, 1998, Proceedings. pp. 408–423 (1998)
42. Haitner, I., Rosen, A., Shaltiel, R.: On the (im)possibility of arthur-merlin witness hiding protocols. In: *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009*, San Francisco, CA, USA, March 15-17, 2009. Proceedings. pp. 220–237 (2009)
43. Halevi, S., Kalai, Y.T.: Smooth projective hashing and two-message oblivious transfer. *J. Cryptology* 25(1), 158–193 (2012)
44. Hazay, C., Venkatasubramanian, M.: On the power of secure two-party computation. In: *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II. pp. 397–429 (2016)
45. Kalai, Y.T.: Smooth projective hashing and two-message oblivious transfer. In: *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Aarhus, Denmark, May 22-26, 2005, Proceedings. pp. 78–95 (2005)
46. Kalai, Y.T., Raz, R.: Probabilistically checkable arguments. In: *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings. pp. 143–159 (2009)

47. Katz, J., Ostrovsky, R.: Round-optimal secure two-party computation. In: Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology-Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings. pp. 335–354 (2004)
48. Lapidot, D., Shamir, A.: Publicly verifiable non-interactive zero-knowledge proofs. In: Menezes, A., Vanstone, S.A. (eds.) Advances in Cryptology - CRYPTO '90, 10th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1990, Proceedings. Lecture Notes in Computer Science, vol. 537, pp. 353–365. Springer (1990), http://dx.doi.org/10.1007/3-540-38424-3_26
49. Micali, S., Pass, R., Rosen, A.: Input-indistinguishable computation. In: 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06). pp. 367–378 (Oct 2006)
50. Mittelbach, A., Venturi, D.: Fiat-shamir for highly sound protocols is instantiable. In: Security and Cryptography for Networks - 10th International Conference, SCN 2016, Amalfi, Italy, August 31 - September 2, 2016, Proceedings. pp. 198–215 (2016)
51. Naor, M., Pinkas, B.: Efficient oblivious transfer protocols. In: Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA. pp. 448–457 (2001)
52. Ostrovsky, R., Persiano, G., Visconti, I.: On input indistinguishable proof systems. In: Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I. pp. 895–906 (2014)
53. Pass, R.: Simulation in quasi-polynomial time, and its application to protocol composition. In: Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings. pp. 160–176 (2003)
54. Pass, R.: Limits of provable security from standard assumptions. In: Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011. pp. 109–118 (2011)
55. Prabhakaran, M., Rosen, A., Sahai, A.: Concurrent zero knowledge with logarithmic round-complexity. In: 43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings. pp. 366–375 (2002)
56. Rosen, A.: A note on constant-round zero-knowledge proofs for NP. In: Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings. pp. 191–202 (2004)
57. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014. pp. 475–484. ACM (2014), <http://doi.acm.org/10.1145/2591796.2591825>
58. Wee, H.: Black-box, round-efficient secure computation via non-malleability amplification. In: 51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA. pp. 531–540 (2010)
59. Yao, A.C.: How to generate and exchange secrets (extended abstract). In: 27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986. pp. 162–167 (1986)
60. Yung, M., Zhao, Y.: Generic and practical resettable zero-knowledge in the bare public-key model. In: Advances in Cryptology - EUROCRYPT 2007, 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Barcelona, Spain, May 20-24, 2007, Proceedings. pp. 129–147 (2007)