

MPC in the Head

CS 507, Topics in Cryptography: Secure Computation

David Heath

Fall 2026

Last time, we discussed connections between maliciously-secure protocols and *zero knowledge (ZK) proofs*. In particular, we examined the GMW compiler, which is a technique for converting semi-honest-secure protocols into maliciously-secure ones. Recall that the main idea was to insist that at each protocol step, the sending party proves in zero knowledge that their message was constructed as specified by the protocol. Thus, we can upgrade any semi-honest secure MPC protocol into a maliciously secure protocol.

Today we will explore quite a different connection between MPC and ZK proofs. In particular, we will see how MPC protocols can be used to construct ZK proofs for arbitrary statements. The particular paradigm we will study today is called *MPC in the Head* [IKOS07]. This paradigm has been explored in many works, and today we will study (a variant of) the IKOS construction. There are a few reasons for studying this paradigm:

- It demonstrates *one way* to construct ZK proofs for arbitrary statements. Note that ZK is an *incredibly* active area of research, and what we will see today is just one piece of an enormous landscape of ideas. Still...
- The paradigm is best-in-class for constructing proofs of particular statements. For instance, the FAEST signature scheme (<https://faest.info>) is based on MPC in the Head [BBD⁺23]. This signature scheme is a leading candidate for constructing signatures—a central primitive for secure communication—that remain secure in the presence of an adversary with access to a powerful quantum computer.
- The paradigm is a beautiful application of the techniques we have learned so far. It applies them in a non-obvious way, and many of the artifacts that we discuss simply as a means by which to define/prove security are used in a constructive way.

The protocol we will see today is constructible from a commitment scheme.

Zero Knowledge Proofs

Recall from last time that a ZK proof is an interaction between a prover P and a verifier V . The parties agree on some statement x , and P would like to convince V that the statement x is true (formally, x belongs to some language $\mathcal{L}$). To aid in this process, P has a *witness* w , which you can think of as P 's proof that x is true. We would like to achieve three properties:

- **Completeness.** If $x \in \mathcal{L}$, then V outputs 1 at the end of the protocol (with high probability).
- **Soundness.** If $x \notin \mathcal{L}$, then not even a malicious P can cause V to output 1 with high probability.
- **Zero Knowledge.** Not even a malicious V can learn anything about P 's witness w .

Today, we will concern ourselves with a *general-purpose* notion of ZK proofs. In particular, suppose we have some Boolean circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$; C will correspond to the proof statement. Namely, the statement we will consider is the following:

There exists some input $w \in \{0, 1\}^n$ such that $C(w) = 1$.

Of course, if one knows such a witness $w \in \{0,1\}^n$, it is easy to check that $C(w) = 1$ by simply running the circuit. This language (i.e., those circuits that are *satisfiable*) is *general-purpose* in the sense that any bounded-time computation outputting 0 or 1 can be converted to such a circuit, so we can essentially arrange that P can potentially prove to V that any (bounded) claim is true. Today we will build a generic ZK proof system for any such C .

Note that when we are constructing ZK proof techniques, we are typically concerned with three efficiency properties:

- **Proof size.** How many bits are transmitted between P and V .
- **Prover running time.** How long does it take P to generate the information it sends to V ?
- **Verifier running time.** How long does it take V to check P 's proof?

Note that in many modern proof systems there can be *huge* discrepancies between these costs. In particular, it is possible to construct proof systems where the verifier time and proof size are *extremely* low. Our focus today will be on simply achieving all three metrics of cost $O(\lambda \cdot |C|)$.

High Level Sketch of MPC-in-the-Head

Recall, P holds a witness w and both P and V hold a circuit C . P would like to convince V that it knows w such that $C(w) = 1$, while keeping w secret.

The main idea underlying MPC-in-the-head is this: P will emulate—in its head—the execution of a semi-honest secure MPC protocol that securely computes $C(w)$, writing down the view of each emulated party. P then *commits* to the view of each emulated MPC party; each party view is committed separately. Next, V challenges P to open a subset of the party views, and P opens them. V checks that the opened emulated parties sent messages consistent with their views, i.e. that they indeed followed the protocol as prescribed. If (1) all opened parties behaved honestly and (2) the output of the MPC protocol was indeed a 1, then V accepts, and otherwise V rejects.

Why does this work? We can argue each ZK property in turn:

- **Completeness:** If $C(w) = 1$ and P simply behaves honestly, then the emulated MPC protocol will indeed output 1, due to correctness of the underlying MPC scheme.
- **Zero Knowledge:** V learns nothing about w due to the security of the MPC protocol. Indeed, V does not open *all* party views.
- **Soundness:** If P wishes to cheat, the insight is that she must “corrupt” at least one of the emulated MPC parties. Indeed, if all emulated parties behave honestly and $C(w) = 0$, then V will not accept, so cheating P must force some emulated party to behave dishonestly. But V opens party views randomly and checks consistency. So if we are careful in our design, V will catch the cheating party with some probability (scaling in the number of emulated parties).

In a bit more detail, a cheating P must corrupt at least one *communication channel* between parties. Namely, the protocol must instruct some party p_i to send a message m to p_j , yet p_i instead sends some different message $m' \neq m$. If V opens both p_i and p_j , it will detect this discrepancy.

In more detail...

Recall that the view of party P_i consists of its input (concretely, its share of the witness w_i), its random tape, and every message it receives. Each honest party's outgoing messages are a function of its view, so a verifier holding view_i can locally compute every message party P_i should send.

Let's assume that we have a p -party MPC protocol Π that is perfectly correct and secure against any $p - 1$ semi-honest corruptions (e.g. the GMW protocol). For such a Π , the MPC-in-the-head paradigm proceeds as follows:

- P secret shares its witness w into p shares $w_0, \dots, w_{p-1}$ and gives one share to each of the p emulated parties.

- P emulates execution of an MPC protocol Π that securely computes $C(w_0 \oplus \dots \oplus w_{p-1})$ and records the view of each party. P commits to each view $c_i = \text{Commit}(\text{view}_i; r_i)$ (r_i is commitment randomness) and sends each c_i to V .
- V randomly samples a size $p - 1$ subset of parties and sends their ids to P .
- P opens all such committed views to V .
- V computes each opened party's outgoing messages from its view and checks that these messages are consistent with those recorded in the other opened views. V also computes the output of each opened party and checks it is 1. If any check fails, V rejects, and otherwise V accepts.

In a p -party protocol, there are $p(p - 1)$ total communication channels, and V opens $(p - 1)(p - 2)$ of those channels. Again, to cheat it is necessary for cheating prover to corrupt at least *one* of these channels, so the cheating prover will be caught with *at least* the following probability:

$$\Pr_{\text{catch}} \geq \frac{(p - 1)(p - 2)}{p(p - 1)} = \frac{p - 2}{p}$$

Setting $p = 3$, we see that V catches a cheating P with probability at least $1/3$.

Exercise 1. Above, we assumed the MPC protocol Π is perfectly correct (i.e. the semi-honest MPC protocol outputs the correct value with probability 1), and indeed GMW has this property. Suppose instead that the semi-honest parties need only output the correct value with overwhelming probability, say probability $1 - \text{negl}(\lambda)$. Show that there exists a Π with such a correctness property such that the MPC in the head transformation is unsound.

Exercise 2. For $p = 3$, compute the exact probability that V catches a prover who corrupts exactly one channel, and exactly two. Which is better for the cheating prover?

Soundness amplification. Of course, we might want *much* better soundness than this. For instance, we might want to achieve a soundness error that is negligible in the security parameter. This can be achieved by simply repeating the protocol many times. This is a standard trick in interactive proof systems.

Improving Efficiency

So far, the protocol we have described is incredibly inefficient. Indeed, if we recall how the GMW protocol works, it involves pairwise utilization of Oblivious Transfer (OT). Therefore, if we plug GMW into the above schematic as is, we implicitly make non-black-box use of OT: the prover P must write down all messages sent/received as part of the OT protocol. Note that this implies non-black-box usage of some cryptographic primitive, e.g. DDH-friendly groups.

One crucial and elegant observation is that due to the structure of the above proof, we can replace invocations of the OT protocol by its ideal functionality. Indeed, V already checks communication channels between parties, so it can just as easily check that the ideal OT functionality is properly executed between that pair of parties.

With this crucial optimization and setting $p = 3$, the total communication cost per repetition is only $O(|C| + \lambda)$ bits. Essentially, the whole proof involves (1) transmitting commitments, and (2) transmitting views of size $O(|C|)$. Of course, achieving negligible soundness error requires repeating the protocol $O(\lambda)$ times, for $O(\lambda \cdot |C|)$ total cost.

One major improvement of subsequent works in the area of MPC-in-the-head was to (1) decrease the number of needed repetitions and (2) achieve *sublinear* scaling in the circuit size [AHIV17]. A key technique for achieving these improvements is to change the underlying MPC protocol to something more suited to the setting.

Non-interactive ZK Proofs (NIZKs)

So far, the scheme we have described is *interactive*. Namely, the prover and verifier interact, and that interaction convinces only that verifier of the truth of the statement. If P wishes to convince some other verifier, she will need to run the proof again.

In many cases, it might be desirable for P to prove the desired statement “once and for all”. Namely, P would like to one time write down a proof π , then distribute π to verifiers, or even post π online. The document π achieves all the desired guarantees of an interactive proof: each verifier learns the proved statement must be true, while learning nothing about P ’s witness. Such an object is called a *non-interactive* zero-knowledge proof.

The MPC-in-the-head proof system sketched above can be used to achieve NIZKs, using the *Fiat-Shamir* heuristic [FS87]. The key observation about our ZK protocol is that it is a so-called *public coin* protocol. This means that the verifier V has no private randomness. Indeed, V does have randomness (the choice of which parties it opens), but these choices are not “secret”. It is as if, once P commits to party views, V can flip its coins in plain view of the prover P .

Public-coin protocols can be made non-interactive via a cryptographic hash function. The idea is simply this: rather than having V flip random coins, we will have the hash function flip the coins. This is the Fiat-Shamir heuristic.

Formally, Fiat-Shamir is proven secure in the *random oracle model*.

Definition 1 (Random oracle model). *In the random oracle model, we assume all parties have access to a hash function $H : \{0,1\}^* \rightarrow \{0,1\}^\lambda$. We model the behavior of H as a trusted third party that implements a truly random function.*

It’s worth briefly stating that random oracles (ROs) are objects of some contention. Indeed, we know that random oracles are *uninstantiable* in some situations [CGH98]. Still, RO is an incredibly useful heuristic for designing new cryptography [BR93], and in this course we will often use it unapologetically.

In particular, our ZK protocol involves three rounds:

- **Commit:** P sends commitments of party views to V .
- **Challenge:** V sends back a public-coin random challenge.
- **Response:** P opens party views according to the challenge.

(Note, this style of three-round protocol is called a *sigma protocol*.) In the non-interactive version of the protocol, P does not send its commitments to V (indeed, there is no V), but rather inputs the commitments (together with the statement x) to a random oracle. The output of the RO is treated as the random challenge. P ’s proof π is the concatenation of the first-round commitments with the third-round openings. A verifier recomputes the challenge from x and the commitments, then runs the same checks that the interactive verifier would have run.

Thus, the MPC-in-the-head paradigm can yield simple, reasonably efficient proof systems, and they can even achieve non-interactivity.

Next Time

We have now seen our first techniques for achieving maliciously-secure protocols, and we have introduced the idea of emulating MPC players. From here, we will investigate more concretely-efficient techniques for achieving malicious security.

References

- [AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. Liger: Lightweight sublinear arguments without a trusted setup. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 2087–2104. ACM Press, October / November 2017.

- [BBD⁺23] Carsten Baum, Lennart Braun, Cyprien Delpéch de Saint Guilhem, Michael Klooß, Emmanuela Orsini, Lawrence Roy, and Peter Scholl. Publicly verifiable zero-knowledge and post-quantum signatures from VOLE-in-the-head. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 581–615. Springer, Cham, August 2023.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218. ACM Press, May 1998.
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Berlin, Heidelberg, August 1987.
- [IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30. ACM Press, June 2007.