

Honest Majority Multi-Prover Interactive Arguments

Alexander R. Block*

Christina Garman*

May 6, 2022

Abstract

Interactive arguments, and their (succinct) non-interactive and zero-knowledge counterparts, have seen growing deployment in real world applications in recent years. Unfortunately, for large and complex statements, concrete proof generation costs can still be quite expensive. While recent work has sought to solve this problem by outsourcing proof computation to a group of workers in a privacy preserving manner, current solutions still require each worker to do work on roughly the same order as a single-prover solution.

We introduce the *Honest Majority Multi-Prover* (HMMP) model for interactive arguments. In these arguments, we distribute prover computation among M collaborating, but distrusting, provers. All provers receive the same inputs and have no private inputs, and we allow any $t < M/2$ provers to be statically corrupted before generation of public parameters, and all communication is done via an authenticated broadcast channel. In contrast with the recent works of Ozdemir and Boneh (USENIX '22) and Dayama et al. (PETS '22), our model targets prover efficiency over privacy.

We show that: (1) any interactive argument where the prover computation is suitably divisible into M sub-computations can be transformed into an interactive argument in the HMMP model; and (2) arguments that are obtained via compiling polynomial interactive oracle proofs with polynomial commitment schemes admit HMMP model constructions that experience a (roughly) $1/M$ speedup over a single-prover solution. The transformation of (1) preserves computational (knowledge) soundness, zero-knowledge, and can be made non-interactive via the Fiat-Shamir transformation. The constructions of (2) showcase that there are efficiency gains in proof distribution when privacy is not a concern.

1 Introduction

Interactive arguments [BCC88, GMR89] (or *arguments* in short) are interactive protocols that allow a computationally bounded (i.e., polynomial time) prover to certify an NP statement x with witness w to a computationally weak (i.e., sub-linear time) verifier. The (*perfect*) *correctness* condition states that an honest prover can always convince a verifier of a true statement, while the *soundness* condition states that any polynomial-time malicious prover cannot convince a verifier of a false statement (except with small probability). Arguments and their (succinct) non-interactive counterparts [Mic00, Val08, GW11, SG11, DFH12, BCCT12, GGPR13, BCCT13, BCI⁺13] remain relevant to both researchers and practitioners for their continued adoption in a variety of technologies, including privacy preserving smart contracts [KMS⁺16], e-cash [BSCG⁺14], decentralized private computation [BCG⁺20], and accountable anonymity [GGM16]. This continued study has advanced the state-of-the-art towards (nearly) optimal prover and verifier efficiency: for a witness of size N and security parameter λ , many arguments achieve $\tilde{O}_\lambda(N)^1$ prover time and $O_\lambda(\log(N))$ (or even $O_\lambda(1)$) verifier time under a variety of cryptographic assumptions, and additionally achieve other desirable properties such as zero-knowledge [GMR89] and various notions of knowledge-soundness [GI08, Lin03, GW11, WTS⁺18].

Though the asymptotic costs of proof generation are nearly optimal, in practice the concrete costs become prohibitive for all but the most powerful machines when statement and witness sizes become large (e.g.,

*Purdue University, {block9, c1g}@purdue.edu

¹ \tilde{O} omits $\text{polylog}(N)$ factors and subscript λ omits $\text{poly}(\lambda)$ factors.

$N \geq 2^{20}$). One solution for this problem is to outsource proof generation to more powerful machines: a simple solution is to give the worker machine the statement x and witness w and have it generate the proof. However, many applications wish to protect the privacy of the delegator; i.e., sharing w in the clear is undesirable. While it is impossible to avoid sending the true witness w with a single worker machine, recent work [OB22, DPP⁺22] overcomes this problem by constructing privacy preserving delegation schemes using multiple worker machines to collaboratively build a proof. Both these schemes can handle the scenario where there is an honest majority of workers, or even the case of a single honest worker by leveraging certain secure multi-party computation (MPC) protocols [Yao82, GMW87, Kil91, IKOS07].

The main drawback of this approach is that each worker runs in time proportional to a single-prover solution, with additional overheads to run specialized MPC protocols tailored for their applications. In essence, these delegation schemes share the witness w via a secret sharing scheme, and each worker produces a proof for their secret share. The key property of these schemes is that the proof produced via operating on a secret share of w is itself a secret share of a valid proof using witness w . Thus, the delegator can recover a complete proof by recombining the worker proofs. Therefore, the cost of computing a secret share proof is identical to the cost of a single prover producing a proof for w .

While preserving privacy is generally desirable, in some applications it is not necessary. For example, zero-knowledge rollups (zk-rollups) are an increasingly popular method for increasing the scalability of blockchains. In a zk-rollup, a node bundles together a number of (publicly known) transactions off-chain and then generates a zero-knowledge proof of validity. One significant challenge, though, is that these proofs are expensive for a single node to generate, which reduces the throughput gains that they are meant to achieve. Note that in this scenario, all proof inputs are publicly known to all nodes participating in the system [Eth, LNS20].

Of course, multiple provers are unnecessary if no privacy is required (i.e., just send w to one worker machine). Moreover, under the assumption that the majority of provers are honest, another strategy is to simply have each prover submit one bit certifying the validity of the statement x , with the delegator taking the majority and submitting this as a proof to the verifier. However, all provers in these solutions again run in roughly the same time as a single-prover solution. This seems wasteful:

if privacy is not a concern, is there a way for multiple distrusting provers to jointly compute a proof for a statement such that each participant performs less work than a single-prover solution?

1.1 Our Contribution

With the goal of leveraging collaboration to gain efficiency, we introduce the *Honest Majority Multi-Prover (HMMP) model* for interactive arguments. A HMMP model argument is an interactive argument where M mutually distrusting provers collaborate to convince a verifier of the validity of an NP statement x such that each prover is given the same witness w and has no other private inputs. Furthermore, up to $t < M/2$ provers can be statically corrupted before the generation of public parameters (or the first message of the verifier), and all provers communicate with each other and the verifier via an authenticated sender broadcast channel [PSL80, LSP82]. We require correctness to hold with high probability in the presence of $t < M/2$ static corruptions, and require soundness to hold when all M provers are corrupt.

We also define zero-knowledge in the HMMP model analogously to the single-prover definition of zero-knowledge. Informally, we say a HMMP model interactive argument is zero-knowledge if a malicious verifier learns nothing more than the validity of the statement x when interacting with M honest prover algorithms. Note this definition differs from the various zero-knowledge definitions of [OB22, DPP⁺22], but we see it as natural in our setting since all provers share the same witness and have no private inputs. See Section 2.1 for more discussion.

Our main result is showing that any argument where the prover computation is “suitably” partitionable can be transformed into an argument in the HMMP model. Suitably partitionable means that the prover computation can be divided into some number of sub-computations, then recombined into the same message that the single prover would have sent. Our transformation perfectly preserves soundness and zero-knowledge; in particular, it simply reduces to the soundness and zero-knowledge of the original argument. Moreover, if

the single prover argument is publicly verifiable or public coin,² then our transformed argument preserves these properties. However, we experience a (small) multiplicative increase in the correctness error over the original argument. For example, if the single-prover argument has perfect correctness, then our HMMP model argument has correctness with a small probability of error.

Our transformation leverages a simple “shuffle” and “audit” strategy, where all provers are shuffled during the generation of public parameters via a random permutation, and a committee of τ provers jointly check consistency of the next message. We stress that our transformation is *completely insecure* (i.e., not sound) if provers can be adaptively corrupt, or statically corrupt after the generation of public parameters (or the first message of the verifier). See [Section 2.1](#) for details.

We capture the key properties of our HMMP transformation in the following theorem.

Theorem 1.1. *Let $M, \tau, t \in \mathbb{N}$ be parameters such that $t < M/2$ and $\tau \leq t$. If Π_{arg} is an interactive argument such that the prover computation can be divided into M sub-computations, then there exists a HMMP model argument Π_{hmp} with M provers and the following parameters:*

- if Π_{arg} is publicly verifiable (resp., public coin), then Π_{hmp} is publicly verifiable (resp., public coin);
- if Π_{arg} has correctness error δ_{arg} , then Π_{hmp} has correctness error $\delta_{\text{hmp}} = 1 - (1 - \delta_{\text{arg}}) \cdot (1 - M \cdot (t/M)^\tau)$ when at most t provers are corrupt;
- if Π_{arg} has ε_{arg} (computational) soundness, then Π_{hmp} has $\varepsilon_{\text{hmp}} = \varepsilon_{\text{arg}}$ (computational) soundness; and
- if Π_{arg} has γ_{arg} zero-knowledge error, then Π_{hmp} has zero-knowledge error $\gamma_{\text{hmp}} = \gamma_{\text{arg}}$.

Furthermore, Π_{hmp} can be made non-interactive via the Fiat-Shamir transformation [[FS87](#)].

1.1.1 Efficiency of [Theorem 1.1](#).

Our HMMP model transformation is general in the sense that we only require the prover computation of a single-prover argument to be suitably partitionable; in particular, the M sub-computations need not be more efficient than the single-prover computation, and there may exist arguments with this property. This is a major hurdle towards our efficiency goals.

We overcome this limitation by directly working with a specific class of interactive arguments that admit efficient sub-division of prover computations. In particular, we consider the class of arguments that are obtained from compiling *polynomial interactive oracle proofs* (PIOPs) [[RRR16](#), [BSCS16](#)] with *polynomial commitment schemes* [[KZG10](#)], which has become one common paradigm for designing arguments [[MBKM19](#), [GWC19](#), [CHM⁺20](#), [BFS20](#), [BHR⁺20](#), [BHR⁺21](#)]. Briefly, a polynomial interactive oracle proof is an interactive proof (i.e., it is an argument where the prover is computationally unbounded) where large prover messages are encoded as polynomials and are sent as an oracle for the verifier to query, and all other communication between the prover and verifier is efficient. A polynomial commitment scheme is a special commitment scheme that allows a committer to commit to some polynomial to a receiver, and allows the committer to provably open (via an argument) to evaluations of the polynomial at a specific point specified by the receiver. The essential feature of a polynomial commitment scheme is that this “proof of evaluation” allows the committer to open to an evaluation point without opening the entire polynomial.

Arguments using this paradigm generally have the following structure: (1) a prover uses the polynomial commitment scheme to commit to the polynomial oracles for the verifier to query; (2) the efficient communication between the prover and the verifier is a sum-check protocol [[LFKN92](#)] over a polynomial that depends on the committed polynomial; and (3) any query made by the verifier to its oracle is completed using the evaluation proof of the polynomial commitment scheme.

We give an overview of this paradigm in [Section 2.2](#) and discuss how to obtain more efficient HMMP model provers by sub-division of computation, and capture these efficiencies in the following theorem.

²*Publicly verifiable* means any party with the transcript of the interaction (i.e., the proof) can verify its validity, and *public coin* means all messages sent by a verifier are uniformly and independently chosen at random.

Theorem 1.2. *Let $M, \tau, t, \lambda \in \mathbb{N}$ be parameters such that $t < M/2$ and $\tau \leq t$. For every NP relation with statements of size $\text{poly}(\lambda)$ and witnesses of size N , there exists a HMMP model interactive argument Π_{hmp} with M provers and the following efficiency properties:*

- *the individual prover computation is $\tilde{O}_\lambda(\tau \cdot (N/M))$;*
- *the communication complexity between provers is $\tilde{O}_\lambda(\tau \cdot (N/M))$ bits;*
- *the verifier computation is $O_\lambda(M + \log(N))$; and*
- *the verifier receives $O_\lambda(M \cdot \log(N))$ bits from all provers and sends $O_\lambda(\log(N))$ bits to all provers.*

Our efficiency guarantees in [Theorem 1.2](#) only hold with respect to t corrupt parties under the assumption that their goal is not to make the honest parties perform more computation. In particular, we assume that the goal of the corrupt parties is to try and force the verifier to reject even if the NP statement is true.

Corrupt parties could have the alternative goal of forcing honest parties to perform more computations. In this case, there is a simple attack: in each round of computation, one dishonest party intentionally submits an incorrect value to force an audit. This can happen $t < M/2$ times, which in the worst case (i.e., $t = \Theta(M)$) forces each honest prover to perform $\tilde{O}_\lambda(\tau \cdot N)$ computation, which is proportional to a single prover solution.

While not addressed in this work, this attack is unavoidable in our HMMP model transformation. However, one possible solution to this issue is to permanently kick parties and/or implement a reputation system to punish misbehaving provers; this could amortize the damage done by malicious provers.

1.2 Additional Related Work

Interactive arguments have a rich history of research over the past four decades [[BCC88](#), [GMR89](#), [Mic00](#), [Val08](#), [GW11](#), [SG11](#), [DFH12](#), [BCCT12](#), [GGPR13](#), [BCCT13](#), [BCI+13](#)]. The recent works of Ozdemir [[OB22](#)] and Dayama et al. [[DPP+22](#)] construct privacy preserving delegation schemes, and define their model in a similar way to ours. However, their goals are to preserve privacy, which is orthogonal to our goal of increasing efficiency while assuming no privacy. Another result loosely related to our work is due to Wu et al. [[WZC+18](#)]. [[WZC+18](#)] constructs DIZK, which works to distribute the creation of a zero-knowledge proof across multiple machines controlled by the prover, such as a compute cluster. This achieves something much closer to “centralized decentralization”, where a single prover distributes their computation to many (honest) trusted nodes. In contrast, our model supports any number of distinct provers that wish to work together to convince a verifier of a statement, with only an honest-majority assumption. Also somewhat related is the recent work of Naor, Parter, and Yogev [[NPY20](#)]. Here, the authors construct interactive proofs with *distributed verifiers*. In this setting, the verifiers are represented by an n node graph G defining their communication pattern, and a prover communicates with all nodes via short messages.

Polynomial commitment schemes are a relatively new cryptographic primitive introduced by Kate, Zaverucha, and Goldberg [[KZG10](#)]. Several variants of polynomial commitment schemes have been explored since their introduction, including privately verifiable schemes [[KZG10](#), [PST13](#)], publicly-verifiable schemes with trusted setup [[BFS20](#)], and zero-knowledge schemes [[WBT+17](#)]. Recent results have focused on obtaining publicly-verifiable schemes without a trusted setup [[BCC+16](#), [WBT+17](#), [BBB+18](#), [WTS+18](#), [BSBHR18](#), [KPV19](#), [BSGKS20](#), [BFS20](#), [Lee21](#), [SL20](#), [ZXZS20](#)]. More recently, Block et al. [[BHR+20](#), [BHR+21](#)] construct the first *space-efficient* polynomial commitment schemes, where both the prover and verifier space is poly-logarithmic in the size of the polynomial, and recent work [[BCHO22](#)] expands these results.

2 Technical Overview

We give an overview of [Theorems 1.1](#) and [1.2](#). At a high level, for [Theorem 1.1](#) we utilize a shuffle and audit strategy to ensure consistency among all provers. For [Theorem 1.2](#), we identify three core components of the polynomial IOP + polynomial commitment scheme paradigm and efficiently divide these components among M provers to get efficient arguments.

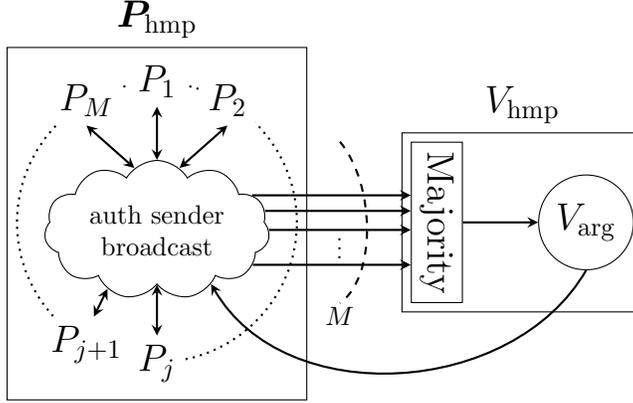


Figure 1: Diagram of the interaction between P_{hmp} and V_{hmp} .

2.1 Overview of Theorem 1.1

Fix an argument scheme $\Pi_{\text{arg}} = (\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$, where $\text{Setup}_{\text{arg}}$ is the setup algorithm and P_{arg} and V_{arg} are interactive algorithms representing the prover and verifier, respectively. Suppose that P_{arg} is divisible into M sub-computations, which we model as follows. Let f be the next message function of P_{arg} , let $S: [M] \times \{0, 1\}^* \rightarrow \{0, 1\}^*$ be a partition function, and let $f^*, f_1, \dots, f_M: \{0, 1\}^* \rightarrow \{0, 1\}^*$ be functions such that for all $z \in \{0, 1\}^*$, we have $f(z) = f^*(f_1(S(1, z)), \dots, f_M(S(M, z)))$. Intuitively, f_i is the i^{th} sub-computation and f^* reconstructs the next message of P_{arg} . Note that here $[M] := \{1, \dots, M\}$.

Given the above functions, one can envision the following straw man transformation from Π_{arg} to $\Pi_{\text{hmp}} = (\text{Setup}_{\text{hmp}}, P_{\text{hmp}}, V_{\text{hmp}})$, where $P_{\text{hmp}} = (P_1, \dots, P_M)$ are M collaborating provers.

1. Define $\text{Setup}_{\text{hmp}} := \text{Setup}_{\text{arg}}$.
2. Define P_{hmp} as follows: for any input z
 - (a) Prover P_i computes $y_i = f_i(S(i, z))$ and broadcasts y_i .
 - (b) All provers compute $y = f^*(y_1, \dots, y_M)$ and broadcast y to V_{hmp} .
3. Define V_{hmp} as follows: during any round of interaction, upon receiving M messages from P_{hmp} , take the majority and feed this value into V_{arg} . If V_{arg} returns a message, broadcast that message to P_{hmp} . Otherwise accept or reject as V_{arg} . Fig. 1 gives a pictorial representation of this interaction.

Clearly the above transformation is secure when *all* provers are honest; however, it is clear that even a single malicious prover \tilde{P}_j could inject arbitrary values into the computation. The key observation is that we do not take advantage of our honest majority assumption. To leverage this, we introduce a *permute and audit* mechanism to allow for cross-examination among provers.

2.1.1 Secure Transformation.

We modify the above straw man scheme in the following two ways. First, we introduce a τ *audit window* to the protocol: every f_i is computed by τ provers $P_i, P_{i+1}, \dots, P_{i+\tau-1}$, and their results are broadcast to all other provers. Then, all provers check to see if the τ submitted values are all the same. If they are not, then a *global audit* is initiated. In this global audit, *all* provers P_1, \dots, P_M compute the value f_i , broadcast this computation, and then take the majority of the answers. The computed majority value is then compared with the τ submitted values, and any prover that submitted an incorrect value is kicked from the protocol. By

our honest majority assumption, if a global audit is initiated then it is guaranteed to *only* remove dishonest parties from the protocol; this is an invariant property of our audit procedure (see [Lemma 4.7](#)).

However, even introducing the above audit window is not sufficient for security unless $\tau > t$. In particular, if $\tau \leq t$, then an adversary can simply corrupt τ provers in quick succession and is still able to attack the protocol. Further, setting $\tau > t$ is undesirable; if $t = c \cdot M$ for some constant $c < 1/2$, then every prover has a $\tau = \Omega(M)$ blow-up in computational complexity, and at that point it would be simpler for all provers to compute all values f_i , defeating our efficiency goals for the HMMP model.

To remedy this, we introduce a uniformly random permutation π on the set $[M]$ and modify the above audit procedure as follows. For computing function f_i , the parties $P_{\pi(i)}, \dots, P_{\pi(i+\tau-1)}$ all compute f_i , and then follow the above audit procedure as prescribed. So long as the permutation is sampled *after* the adversary performs t corruptions, then with good probability every audit window will contain at least one honest prover. So long as at least one honest prover is in every audit window, then any incorrect behavior from corrupt provers will be caught by a global audit.

This leads to the following secure transformation of Π_{arg} to Π_{hmp} .

1. Define $\text{Setup}_{\text{hmp}}$ as follows: sample $\text{pp}' \leftarrow \text{Setup}_{\text{arg}}$ and sample uniformly random permutation $\pi: [M] \rightarrow [M]$. Output $\text{pp} = (\text{pp}', \pi)$.
2. Define \mathbf{P}_{hmp} as follows: for any input z
 - (a) For $j \in \{i, i+1, \dots, i+\tau-1\}$, prover $P_{\pi(j)}$ computes $y_i = f_i(S(i, z))$ and broadcasts y_i .
 - (b) All provers audit the value y_i as necessary.
 - (c) After all values f_i are computed, all provers compute $y = f^*(y_1, \dots, y_M)$ and broadcast y to V_{hmp} .
3. V_{hmp} is defined analogously as before.

[Fig. 2](#) presents this transformation pictorially (on the right), along with the insecure transformation (on the left).

By definition, our HMMP model requires an adversary to perform t static corruptions *before* the generation of public parameters by $\text{Setup}_{\text{hmp}}$. Thus $\text{Setup}_{\text{hmp}}$ samples the permutation π and we obtain a secure transformation.³

Correctness. Suppose that Π_{arg} has correctness error δ_{arg} ; that is, the probability that V_{arg} rejects a correct proof from an honest prover P_{arg} is at most δ_{arg} . Then our transformation guarantees correctness error δ_{arg} so long as every τ audit window contains at least one honest prover.

Proposition 2.1. *Let Π_{arg} be an argument with δ_{arg} correctness error such that [Theorem 1.1](#) can be applied to Π_{arg} . Let $\mathcal{H} \subseteq [M]$ be the indices of the honest provers and let π be a permutation over $[M]$. For every $i \in [M]$, if $\mathcal{H} \cap \pi(\mathcal{J}) \neq \emptyset$ for $\mathcal{J} = \{i, i+1, \dots, i+\tau-1\}/(M\mathbb{Z})$ then the argument Π_{hmp} has correctness error δ_{arg} .*

Clearly, if the permutation π is fixed ahead of time before the corruption of t provers, then there is always a way for the adversary to construct a set \mathcal{J} such that $\mathcal{H} \cap \pi(\mathcal{J}) = \emptyset$. Hence, we enforce the corruption of t provers before a uniformly random permutation π is sampled. Then we can guarantee with suitably high probability that [Proposition 2.1](#) holds. The following lemma characterizes the probability that [Proposition 2.1](#) holds over a uniformly random permutation.

Lemma 2.2. *Let $\tau, t, M \in \mathbb{N}$ be integers such that $\tau \leq t$ and $t < M$, and let $T \subset \{1, \dots, M\}$ be a subset of size t . For permutation π over $[M]$, let $S_{(i,j)} = \{S_{\ell, \pi}: \ell = (i+k \bmod M) \forall k \in \{0, \dots, j-1\}\}$, where $S_{\ell, \pi} = \pi(\ell)$ for all $\ell \in \{1, \dots, M\}$. Then $\Pr[\exists i \in [M]: S_{(i, \tau)} \subseteq T] \leq M \cdot (t/M)^\tau$, where the probability is taken over uniformly random permutation π .*

³Alternatively, the verifier can sample π and send it as its first message, after t static corruptions have occurred.

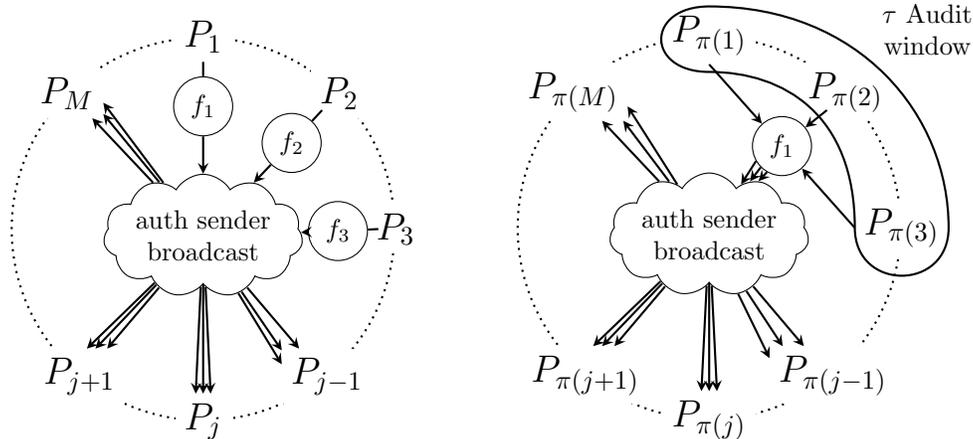


Figure 2: Demonstration of the insecure straw man implementation of \mathbf{P}_{hmp} (pictured left) and the secure implementation of \mathbf{P}_{hmp} (pictured right) with the τ audit window for computing f_1 , for $\tau = 3$. Here π is a uniformly random permutation on the set $\{1, \dots, M\}$.

[Lemma 2.2](#) states that the probability there exists an audit window with no honest provers is at most $M \cdot (t/M)^\tau$; i.e., with probability at least $1 - M \cdot (t/M)^\tau$, \mathbf{P}_{hmp} behaves identically to the honest prover algorithm P_{arg} . Note τ is a parameter for the argument Π_{hmp} and can be chosen such that $M \cdot (t/M)^\tau = \text{negl}(M)$, e.g., $\tau = \text{polylog}(M)$. [Proposition 2.1](#) and [Lemma 2.2](#) together guarantee that correctness holds with probability at least $(1 - \delta_{\text{arg}}) \cdot (1 - M \cdot (t/M)^\tau)$, thus our correctness error is $\delta_{\text{hmp}} = 1 - (1 - \delta_{\text{arg}}) \cdot (1 - M \cdot (t/M)^\tau)$ as desired.

Soundness. As stated in [Section 1.1](#), we require soundness to hold with respect to all M provers being corrupt. Informally, we say that a HMMP model argument Π_{hmp} is ε -computationally sound if the probability M malicious collaborating provers convince a verifier of a false statement with probability at most ε . We similarly extend our soundness definition to *knowledge soundness* [[BG92](#)]. Informally, we say that a HMMP model argument Π_{hmp} has knowledge soundness if whenever M malicious collaborating provers cause a verifier to accept, then a valid (NP) witness can be efficiently extracted via an extractor by interacting with and rewinding the M malicious provers. Our definitions mirror both [[OB22, DPP+22](#)].

Under these definitions, it is not difficult to see that both soundness and knowledge soundness directly reduce to the single-prover argument case (this observation is similarly made in both [[OB22, DPP+22](#)]). In particular, any M malicious prover strategy that is successful in breaking soundness is itself a single-prover strategy for breaking soundness: the single prover simply simulates the interaction “in its head”. For knowledge soundness, from an extractor perspective, any property that holds for a malicious prover also holds for a collection of malicious provers. See [Section 4](#) and [Proposition 4.10](#) for more details.

Defining and Obtaining Zero-Knowledge. We define zero-knowledge analogously to [[DPP+22](#)]. In particular, we say that a HMMP model argument is zero-knowledge with error γ if any malicious verifier interacting with M honest collaborating prover algorithms learns nothing other than the validity of the statement, except with probability at most γ . In our setting this definition of zero-knowledge is the most natural since all M provers are given the same witness and have no other private inputs to protect.

If collaborating provers each have some private input to protect, then one can define zero-knowledge with respect to an honest verifier and t malicious provers by stipulating the malicious provers learn nothing by interacting with the honest provers; [[OB22](#)] defines t -zero knowledge in this way, and [[DPP+22](#)] calls this witness confidentiality. Moreover, [[DPP+22](#)] defines another zero-knowledge variant with respect to t

malicious provers *and* a malicious verifier, which they call witness confidentiality with collusion. While both of these notions of zero-knowledge are interesting, they do not apply (nor make sense) in our case since all provers have the same inputs.

As stated in [Theorem 1.1](#), our transformation preserves zero-knowledge, which can be seen via the following abstraction. Given a zero-knowledge argument Π_{arg} , the prover next message function can be viewed as two functions: the non-zero knowledge function f and the zero-knowledge function g . Given sufficiently many random coins f and input z , if $f(z)$ is sent by the prover in the non-zero knowledge setting, then $g(f(z); r)$ is sent in the zero-knowledge setting. This abstraction allows us to preserve zero-knowledge by giving \mathcal{P}_{hmp} a shared random string r as input in addition to the witness. Each prover then uses the same random string r and sends $y' = g(y; r)$, where y is computed as in [Item 2c](#). In place of receiving r as input, one can assume a randomness beacon [[Rab83](#), [SJK⁺17](#)] to broadcast r , or leverage honest-majority MPC to agree upon r [[BOGW88](#), [RBO89](#)].

Obtaining Non-Interactivity. Briefly, a non-interactive argument is an argument Π_{arg} where only a single message is sent from the prover to the verifier, and there is no other interaction. Similarly, we say a HMMP model argument Π_{hmp} is non-interactive if M collaborating and communicating provers generate a single message (i.e., the proof string), send this message to the verifier, and there is no other interaction.

We obtain non-interactivity directly by applying the Fiat-Shamir transformation [[FS87](#)] to any public coin HMMP model argument. All provers are assumed to have access to a shared random oracle H , and the provers each use H to generate the next message of the verifier. In particular, instead of sending M messages to the verifier, each prover takes the majority of these messages, appends this new message to the transcript so far, then applies the random oracle H to the updated transcript and uses the output as the verifier message. Security now holds information theoretically in the random oracle model: if any dishonest prover does not follow this procedure, then with high probability over the random oracle, the verifier challenges generated in a dishonest way will be incorrect, which guarantees that with high probability the dishonest parties are caught in subsequent rounds.

2.2 Overview of [Theorem 1.2](#)

Recall that [Theorem 1.1](#) does not guarantee efficiency gains for collaborating provers, which is the goal of our model. We overcome this limitation of [Theorem 1.1](#) by directly applying our HMMP model to a class of arguments for NP that follow the paradigm of compiling polynomial interactive oracle proofs with polynomial commitment schemes. Arguments obtained via this paradigm have algebraic structure that we can take advantage of during the collaborative proof generation. We first give an overview of the paradigm, then discuss how to efficiently distribute the computation of each sub-component, then describe how to compose these components to obtain [Theorem 1.2](#).

2.2.1 Paradigm Overview.

As mentioned before, this paradigm compiles any polynomial interactive oracle proof (polynomial IOP, PIOP) for NP with a suitable polynomial commitment scheme to obtain interactive arguments for NP. We begin by giving an overview of polynomial commitment schemes, or polynomial commitments in short. Briefly, polynomial commitments consist of a tuple of algorithms (Setup, Com, Open, Eval). This tuple allows a sender to (non-interactively) produce a commitment $C = \text{Com}(f)$ for polynomial f , open a commitment C to a polynomial f using Open, and additionally supports “proofs of evaluation” using Eval. The Eval algorithm is an interactive protocol between the sender and receiver which does the following: on input commitment C , evaluation point x from the receiver, and value y from the sender, Eval certifies that the sender knows a polynomial f such that $f = \text{Open}(C)$ and $y = f(x)$. Intuitively, Eval allows the sender to open the commitment C at specific evaluation points without revealing the entire polynomial. Crucially, this proof of evaluation requires far less computation than just opening the entire polynomial f .

Given the outline of a polynomial commitment scheme, we turn to polynomial IOPs. We focus on PIOPs where the prover sends a single oracle, and the rest of the communication between the prover and verifier is

succinct (i.e., poly-logarithmic in the size of the witness). These PIOPs (roughly) operate as follows:

1. The first message sent by the prover is an oracle proof string consisting of all evaluations of a (low-degree) polynomial, say $f \in \mathbb{F}[X_1, \dots, X_n]$ for finite field \mathbb{F} . The verifier has query access to these evaluations.
2. The prover and verifier engage in a sum-check protocol [LFKN92] over some auxiliary (low-degree) polynomial, say $p \in \mathbb{F}[X_1, \dots, X_n]$, that depends on the f .
3. At the end of the sum-check, the verifier has some claim of the form “ $p(r) = y$ ” for $r \in \mathbb{F}^n$ sampled uniformly at random. The verifier then makes a constant number of queries to the oracle f to verify this claim, and accepts or rejects.

To obtain arguments, we compile the above PIOP with a polynomial commitment. This is done by (1) having the prover commit to the polynomial f in [Item 1](#) using `Com`; and (2) for every query to f made by the verifier in [Item 3](#), the prover and verifier instead run the `Eval` protocol to obtain these queries. The remainder of the PIOP remains the same; i.e., the prover and verifier still run a sum-check protocol. Note this extends to PIOPs that send multiple oracles: simply commit to each oracle and run `Eval` whenever a query is made.

This paradigm yields interactive arguments for all of NP [BHR⁺20, BHR⁺21, BCHO22]. Therefore, if we extend the above paradigm to the HMMP model with the efficiency gains proportional to $1/M$ for M provers, then we obtain [Theorem 1.2](#). There are three tasks to implement in the HMMP model: (1) generating the polynomial f in [Item 1](#); (2) computing the sum-check in [Item 2](#); and (3) computing the commitment for [Item 1](#) and the proof of evaluation for [Item 3](#); i.e., computing the polynomial commitment. We concretely consider the following problems in the HMMP model to address these tasks. First, we give a HMMP model protocol for circuit satisfiability. Second, we give a HMMP model protocol for the sum-check protocol. Third, we instantiate the well-known Bulletproofs polynomial commitment scheme in the HMMP model. In what follows, we give an overview of these protocols along with formal propositions outlining their efficiency guarantees. All together, this yields [Theorem 1.2](#).

We first fix some notation. We focus on constructing an argument for the NP complete language of (arithmetic) circuit satisfiability. Let \mathbb{F} be a finite field and let $C: \{0, 1\}^* \rightarrow \mathbb{F}$ be an arithmetic circuit of size N ; i.e., C has N wires (including input and output wires), and gates in C perform addition and multiplication over \mathbb{F} . Let $M \leq N$ be the number of provers, let $t < M/2$ be the number of corrupt parties, and let $\tau \leq t$ be the audit window parameter.

2.2.2 Circuit SAT.

Generally, the polynomial f sent in [Item 1](#) of a PIOP is the encoding of a wire transcript (i.e., values of wires) of an arithmetic circuit C of size N on input x , which corresponds to proving circuit satisfiability. For arithmetic circuits, this is proving that $C(x) = y$ for some $y \in \mathbb{F}$. Our first step is to give an efficient HMMP model protocol for producing this transcript.

For simplicity, assume that N/M is an integer. Partition the circuit C into M partitions C_1, \dots, C_M each of size N/M such that for any $j \in \{1, \dots, M\}$, partition C_j only depends on wires from partitions C_i for $i \leq j$. Then the M provers (P_1, \dots, P_M) produce a wire transcript of C as follows. Let x be the input to the circuit C . For $j = 1, 2, \dots, M$, audit window $P_{\pi(j)}, P_{\pi(j+1)}, \dots, P_{\pi(j+\tau-1)}$ compute the wire transcript of partition C_j , using input x and previously obtained transcripts for circuits C_i for all $i < j$. The audit window then broadcasts their computed transcript, invoking a global audit and re-balance as necessary.

At the end of the protocol, each party combines all partial transcripts in order to obtain the final wire transcript for C . Since each C_j consists of N/M gates, each audit window performs $O(N/M)$ computations to generate the wire transcript for C_j , yielding $O(\tau \cdot (N/M))$ arithmetic gate evaluations for each prover.⁴ For simplicity, we assumed that N/M was an integer. However, this is not necessary: we can instead let $k = \lfloor N/M \rfloor$ and partition C into at most $2M$ partitions C_1, \dots, C_{2M} each of size at most k with the same

⁴We do not count the work required to write down the entire transcript, as this is minor in comparison to generating the actual transcript.

dependency properties as before. This increases individual prover complexity by a factor of at most 2, yielding the same asymptotic results. The following proposition formally captures our HMMP model protocol for circuit satisfiability.

Proposition 2.3. *There exists an interactive proof in the HMMP model for arithmetic circuit satisfiability with the following properties: for circuits of size N over finite field \mathbb{F} and parameters τ, M such that $M \leq N$,*

- *each prover computes $O(\tau \cdot (N/M))$ arithmetic gate evaluations; and*
- *each prover broadcasts $O(\tau \cdot (N/M))$ field elements.*

2.2.3 Sum-Check Protocol.

Let $f \in \mathbb{F}[X_1, \dots, X_n]$ be a polynomial of constant individual degree (i.e., $\deg_i(f) = \Theta(1)$ for all i). The sum-check protocol [LFKN92] is an interactive proof certifying that $\sum_{x \in \{0,1\}^n} f(x) = y$ for some $y \in \mathbb{F}$. The protocol proceeds for n rounds, where in the first round the prover sends y and for every round $j = 1, \dots, n$, the prover computes a polynomial $g_j(X_j) := \sum_{x \in \{0,1\}^{n-j}} g(r, X_j, x)$, where $r = (r_1, \dots, r_{j-1})$ are random field elements sent by the verifier in each prior round.

During every round $j \in \{1, \dots, n\}$ of the protocol, the prover evaluates g at 2^{n-j} points to compute the constant-degree polynomial $g_j(X_j)$. Because computing g_j is linear, we can sub-divide this computation into M different sums. For simplicity, assume that $2^{n-j}/M$ is an integer. Then we can partition $\{0, 1\}^{n-j}$ into subsets $S_1, \dots, S_M \subset \{0, 1\}^{n-j}$ each of size $2^{n-j}/M$. For every $i \in \{1, \dots, M\}$, the audit window $P_{\pi(i)}, \dots, P_{\pi(i+\tau-1)}$ computes and broadcasts the polynomial $g_{j,i}(X_j) := \sum_{x \in S_i} g(r, X_j, x)$. Note that g has constant individual degree, so each audit window broadcasts $O(1)$ field elements. After broadcasting all $g_{j,i}$, each prover reconstructs the polynomial $g_j := \sum_i g_{j,i}$, sends it to the verifier, and proceeds to the next round.

Each audit window performs $2^{n-j}/M$ evaluations of g and $O(2^{n-j}/M)$ field additions to compute $g_{j,i}$. Additionally, each prover performs $O(M)$ field additions to reconstruct g_j since g_j has constant degree. Thus each prover performs $O(\tau \cdot (2^{n-j}/M))$ evaluations of g and $O(\tau \cdot (2^{n-j}/M) + M)$ field additions in round j . As with the circuit SAT algorithm, if $2^{n-j}/M \geq 1$ and is not an integer, we can set $k = \lfloor 2^{n-j}/M \rfloor$ and partition $\{0, 1\}^{n-j}$ into S_1, \dots, S_{2M} sets of size at most k . This increases the prover complexity by a factor of at most 2, yielding the same asymptotic complexity. Thus in the worst case, the total prover computation over all rounds $j = 1, \dots, n$ is $O(\tau \cdot (2^n/M))$ evaluations of g and $O(\tau \cdot (2^n/M) + n \cdot M)$ field additions, and each prover broadcasts $O(\tau \cdot n)$ field elements.

We conveniently assumed that $2^{n-j}/M \geq 1$, but it is possible that in some round j we have $2^{n-j} < M$. To handle this case, consider j^* as the unique index such that $2^{n-j^*} < M \leq 2^{n-j^*+1}$. Then we let the first 2^{n-j^*} audit windows evaluate g at a single point to construct g_{j^*} ; i.e., they compute $g_{j^*,i} = g(r, X_{j^*}, i)$ where $i \in \{0, \dots, 2^{n-j^*} - 1\}$ is interpreted in binary in the natural way. Then the next 2^{n-j^*-1} audit windows evaluate g at a single point and construct g_{j^*+1} as above. This continues until a single audit window constructs the polynomial g_n .

By dividing up the computation this way, we ensure that every prover computes $O(\tau)$ evaluations of g and $O(M)$ field additions and broadcasts $O(\tau)$ field elements over *all* rounds j^*, \dots, n . To see this, notice that $2^{n-j^*} < M \leq 2^{n-j^*+1} < 2M$; moreover $\sum_{j=j^*}^n 2^{n-j} < 2M$. Thus, our above division of computation only requires each audit window to evaluate g at at most 2 points, giving the stated complexity $O(\tau)$. Additionally, each round $j = j^*, \dots, n$ requires 2^{n-j} additions to reconstruct the polynomial g_j , giving $O(M)$ additions overall. We capture our sum-check protocol in the following corollary.

Corollary 2.4. *For any n -variate polynomial g over finite field \mathbb{F} of constant individual degree, there exists a HMMP model protocol for computing the sum-check protocol of f on $\{0, 1\}^n$ such that for parameters τ and $M \leq 2^n$,*

- *each prover computes $O(\tau \cdot (2^n/M))$ evaluations of g and $O(\tau \cdot (2^n/M) + n \cdot M)$ field additions; and*
- *each prover broadcasts $O(\tau \cdot n)$ field elements.*

We actually prove a more general result about the sum-check protocol for polynomials of individual degree d being summed over a set H^n , where $H \subset \mathbb{F}$. See [Proposition 5.1](#) and [Section 5.2](#) for details.

2.2.4 Polynomial Commitment Scheme.

Recall that polynomial commitments allow a sender to commit to a polynomial and additionally allow the sender to open the polynomial at specific evaluations specified by the receiver without opening the entire polynomial. Such schemes are the continued subject of recent research and have many attractive features such as succinct proofs, small commitments, and their use in the PIOP + polynomial commitment paradigm to obtain succinct arguments. As such, there are many different schemes from a variety of cryptographic assumptions, including discrete-log [BCC⁺16, BBB⁺18, WTS⁺18, BHR⁺20], groups of unknown order [BFS20, BHR⁺21], and bilinear pairings [KZG10, Lee21, OB22, BCHO22].

In this work, we do not tackle implementing every polynomial commitment scheme in the HMMP model. We instead focus on the well-known and studied Bulletproofs scheme due to Bootle et al. [BCC⁺16] and Bünz et al. [BBB⁺18]. We give the complete details of how we implement Bulletproofs in the HMMP model in Section 5.3. Here, we give an overview of the core components of Bulletproofs and how to implement them efficiently in the HMMP model.

The Bulletproofs polynomial commitment scheme utilizes a vector-commitment and an inner product. Let $Y \in \mathbb{F}^N$ be a vector given to a prover and let $X \in \mathbb{F}^N$ be a public vector, and for simplicity assume $N = 2^n$. Then the prover commits to Y using a Pedersen commitment [Ped91] over a prime-order group \mathbb{G} . For generators g_1, \dots, g_N , the commitment to Y is $C := \prod_i g_i^{Y_i}$. Computation of C is easily divisible among M provers: let N/M be an integer and partition the set $\{1, \dots, N\}$ into subsets S_1, \dots, S_M each of size N/M . Then audit window $P_{\pi(i)}, \dots, P_{\pi(i+\tau-1)}$ computes and broadcasts the partial commitment $C_i = \prod_{j \in S_i} g_j^{Y_j}$. Once all C_i are obtained, each prover locally computes the commitment $C = \prod_i C_i$. Therefore each prover performs $O(\tau \cdot (N/M))$ group exponents and $O(\tau \cdot (N/M) + M)$ group multiplications, and each prover broadcasts $O(\tau)$ group elements. As before, if N/M is not an integer, we can partition $\{1, \dots, N\}$ into subsets S_1, \dots, S_{2M} each of size at most $\lfloor N/M \rfloor$, and each prover does at most two times the computation as the case where N/M is an integer.

The inner-product argument uses a halving protocol to reduce the inner-product claim $\gamma = \langle X, Y \rangle$ into another inner-product claim $\gamma' = \langle X', Y' \rangle$, where $X', Y' \in \mathbb{F}^{N/2}$. This is done by splitting X and Y into (X_L, X_R) and (Y_L, Y_R) , where X_L, Y_L are the first $N/2$ symbols of X, Y , respectively, and X_R, Y_R are the last $N/2$ symbols of X, Y , respectively. The prover computes “cross inner-products” $\gamma_0 = \langle X_R, Y_L \rangle$ and $\gamma_1 = \langle X_L, Y_R \rangle$. Then, given a random verifier challenge α , new vectors X', Y' are computed as $\alpha X_L + \alpha^{-1} X_R$ and $\alpha^{-1} Y_L + \alpha Y_R$; moreover, γ' is defined using $\alpha, \gamma_0, \gamma_1$, and γ such that $\gamma' = \langle X', Y' \rangle$. This halving protocol continues until the vector Y' is of length 1.

We divide the above computation among M provers, then generalize it to any round $j = 1, \dots, n$ of the halving protocol. Suppose that $N/(2^j M)$ is an integer. Partition X_L, X_R and Y_L, Y_R into M vectors $X_{L,i}, X_{R,i}$ and $Y_{L,i}, Y_{R,i}$ each of length $N/(2^j M)$ for all $i \in \{1, \dots, M\}$. Then each audit window $P_{\pi(i)}, \dots, P_{\pi(i+\tau-1)}$ computes and broadcasts partial sums $\gamma_{0,i} = \langle X_{R,i}, Y_{L,i} \rangle$ and $\gamma_{1,i} = \langle X_{L,i}, Y_{R,i} \rangle$. Once all $\gamma_{0,i}, \gamma_{1,i}$ are obtained, each prover locally computes $\gamma_0 = \sum_i \gamma_{0,i}$ and $\gamma_1 = \sum_i \gamma_{1,i}$. Upon receiving α from the verifier, each audit window now computes partial parts of X', Y' . That is, audit window $P_{\pi(i)}, \dots, P_{\pi(i+\tau-1)}$ computes and broadcasts vectors $X'_i = \alpha X_{L,i} + \alpha^{-1} X_{R,i}$ and $Y'_i = \alpha^{-1} Y_{L,i} + \alpha Y_{R,i}$. Once all X'_i, Y'_i are obtained, each prover locally computes $X' = (X'_1, \dots, X'_M)$ and $Y' = (Y'_1, \dots, Y'_M)$.

For general round $j = 1, \dots, n$, if $2^{n-j-1}/M \geq 1$ is an integer, then each audit window performs $O(2^{n-j}/M)$ field multiplications and additions to compute values $\gamma_{0,i}, \gamma_{1,i}$ and vectors X'_i, Y'_i . Each audit window also broadcasts $O(2^{n-j}/M)$ field elements, and each prover performs $O(M)$ additions to reconstruct γ_0, γ_1 and $O(\tau \cdot (2^{n-j}/M))$ additions to reconstruct vectors X', Y' . Thus each prover performs $O(\tau \cdot (2^{n-j}/M))$ field multiplications and $O(\tau \cdot (2^{n-j}/M) + M)$ additions, and broadcasts $O(\tau \cdot (2^{n-j}/M))$ field elements. This gives that over all rounds $j = 1, \dots, n$, each prover performs $O(\tau \cdot (N/M))$ field multiplications and $O(\tau \cdot (N/M) + n \cdot M)$ additions, and broadcasts $O(\tau \cdot n \cdot (N/M))$ field elements. Again we note that if $2^{n-j-1}/M \geq 1$ is not an integer, we obtain the same asymptotic complexities. By the same ideas as the division of the sum-check, if j^* is the unique round where $2^{n-j^*-1} < M \leq 2^{n-j^*}$, we can divide up the remaining computation such that each prover only performs $O(\tau)$ field multiplications, $O(M)$ additions, and broadcasts $O(\tau)$ field elements over all rounds $j = j^*, \dots, n$. Thus our total prover complexity is $O(\tau \cdot (N/M))$ field multiplications and $O(\tau \cdot (N/M) + n \cdot M)$ field additions, and each prover broadcasts $O(\tau \cdot n \cdot (N/M))$

field elements.

In the context of Bulletproofs, each prover also computes “folded” Pedersen commitments. In fact, we divide up computation of these commitments in the same way as we divide up the inner products above. This results in similar complexities for group exponents and multiplications.

Proposition 2.5. *Assuming the hardness of discrete-log in prime-order groups, there exists a HMMP model inner-product argument with the following properties: for vectors of length $N = 2^n$ over a finite field and parameters τ, M such that $M \leq N$,*

- *each prover computes $O(\tau \cdot (N/M))$ field multiplications and group exponents, $O(\tau \cdot (N/M + n \cdot M))$ field additions and group multiplications; and*
- *each prover broadcasts $O(\tau \cdot n \cdot (N/M))$ field and group elements.*

Efficiency of the Bulletproofs Verifier. The original verifier from the Bulletproofs inner-product argument runs in time $\tilde{O}(N)$ because it also must compute foldings of the public values. A nice side-effect of implementing the Bulletproofs prover in the HMMP model is that we can take advantage of our honest majority assumption to get an exponential speedup in our verifier. In particular, we can modify the verifier to only sample the required elements α needed to shuffle the inner products, and at the end of the protocol all provers can send the “folded” values needed for the verifier to perform the final verification. This results in a $O(\log(N))$ verifier for our HMMP model Bulletproofs inner-product argument. See [Section 5.3](#) for details.

Corollary 2.6. *The verifier of our HMMP model Bulletproofs inner-product argument samples $O(\log(N))$ field elements and performs $\Theta(1)$ group exponents and field multiplications.*

2.2.5 Efficient HMMP Model Arguments for NP.

Given the above building blocks, we now combine them all to obtain [Theorem 1.2](#). Let \mathbb{F} be a finite field and \mathbb{G} be a prime-order group for some λ -bit prime. Let \mathbf{P}_{hmp} denote our HMMP model prover algorithm

Using our Circuit SAT algorithm ([Proposition 2.3](#)), \mathbf{P}_{hmp} obtains a wire transcript. We view this wire transcript as a function $f: \{0, 1\}^n \rightarrow \mathbb{F}$, where $N = 2^n$ is the size of the arithmetic circuit. While f may not be a polynomial, it uniquely defines a multi-linear extension $\tilde{f}: \mathbb{F}^n \rightarrow \mathbb{F}$ (see [Fact 5.4](#)). A multi-linear extension is a polynomial with individual degree at most 1 such that $\tilde{f}(x) = f(x)$ for all $x \in \{0, 1\}^n$.

Using our Bulletproofs algorithm, \mathbf{P}_{hmp} commits to \tilde{f} using a Pedersen commitment to commit to the vector $Y \in \mathbb{F}^N$ defined as $Y_i = f(i-1)$, where $i-1$ is viewed as an n -bit string in the natural way. The prover and verifier then engage in a sum-check over some auxiliary polynomial g of constant individual degree that depends on the polynomial \tilde{f} and some other efficiently computable constant individual degree polynomials; in particular, these other polynomials are computable in $O(\log(N))$ time [[CMT12](#), [Tha13](#), [XZZ⁺19](#)].

After engaging in the sum-check protocol, the verifier invokes the polynomial commitment scheme to obtain some constant number of evaluations of the polynomial \tilde{f} . Note that the evaluation of \tilde{f} is exactly given by an inner product $\langle Y, Z \rangle$, where $Z \in \mathbb{F}^N$ is a vector that depends on an evaluation point x specified by the verifier, and some interpolation polynomials ([Fact 5.4](#)); moreover, this vector is computable using $\tilde{O}(N)$ field multiplications and can be divided up between \mathbf{P}_{hmp} to obtain $\tilde{O}(\tau \cdot (N/M))$ field multiplications for each prover ([Fact 5.4](#)). Thus for any evaluation of \tilde{f} at point x , \mathbf{P}_{hmp} computes the vector Z and engages in the Bulletproofs inner-product argument with the verifier. This happens a constant number of times. Noting that we assume \mathbb{F} and \mathbb{G} operations take $\text{poly}(\lambda)$ time, and that we do not count the cost of writing down field/group elements, [Propositions 2.3](#) and [2.5](#) and [Corollaries 2.4](#) and [2.6](#) all together yield the desired efficiencies given in [Theorem 1.2](#).

3 Preliminaries

We let $\lambda \in \mathbb{Z}^+$ denote the security parameter. For natural number $n \in \mathbb{N}$, we let $N = 2^n$. A function $\mu: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is called *negligible* if $1/\mu(n) = o(p(n))$ for any positive polynomial p . For integers $a \leq b$, we let

$[a, b]$ denote the set $\{a, a + 1, \dots, b\}$, and for positive integer k , we let $[k]$ denote the set $[1, k]$. For integer M and integers $a \leq b$, we let $[a, b]_M$ denote the set $\{a \pmod{M}, (a + 1) \pmod{M}, \dots, b \pmod{M}\} \subset \mathbb{Z}_M$. For a finite, non-empty set S , we let $x \stackrel{\$}{\leftarrow} S$ denote sampling x uniformly at random from S . We let \mathbb{F} denote a finite field and \mathbb{G} denote a finite group.

3.1 Interactive Arguments

We give formal definitions of standard interactive arguments. In [Section 4](#), we complement these definitions with our HMMP model definitions.

Definition 3.1 (Witness Relation Ensemble). *A witness relation ensemble (or relation ensemble) is a binary relation $\mathcal{R}_{\mathcal{L}}$ that is polynomially bounded, polynomial-time recognizable, and defines the language $\mathcal{L} = \{x : \exists w, (x, w) \in \mathcal{R}_{\mathcal{L}}\}$.*

We now define interactive arguments for relation ensembles.

Definition 3.2 (Interactive Arguments). *Let \mathcal{R} be a relation ensemble. Let (Setup, P, V) denote three algorithms such that P and V are a pair of PPT interactive algorithms and Setup is a non-interactive algorithm that outputs public parameters pp on input 1^λ and aux , for any auxiliary input aux . Let $\langle P(w), V(z) \rangle(\text{pp}, x)$ denote the output of V after interacting with P on common inputs pp and x , where P additionally receives w such that $(x, w) \in \mathcal{R}$ and V receives auxiliary input $z \in \{0, 1\}^*$. We say that (Setup, P, V) is an interactive argument for relation \mathcal{R} if the following hold:*

1. **δ -Correctness.** *For all $(x, w) \in \mathcal{R}$ and every $z \in \{0, 1\}^*$ we have*

$$\Pr[\langle P(w), V(z) \rangle(\text{pp}, x) = 1] \geq 1 - \delta(\lambda),$$

where the probability is taken over $\text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux})$ and the random coins of V .

2. **ε -Computational Soundness.** *For every x and w such that $(x, w) \notin \mathcal{R}$, every non-uniform interactive polynomial-sized adversary P^* , and every $z \in \{0, 1\}^*$, we have*

$$\Pr[\langle P^*(w), V(z) \rangle(\text{pp}, x) = 1] < \varepsilon(\lambda),$$

where the probability is taken over $\text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux})$ and the random coins of V .

If the above soundness definition holds with respect to unbounded cheating provers, then we say that (Setup, P, V) is an interactive proof for \mathcal{R} .

Two desirable properties of interactive arguments include public verifiability and public coin, which we define here.

Definition 3.3 (Publicly Verifiable). *An interactive argument (Setup, P, V) is publicly verifiable if any party viewing the transcript $T = \langle P(w), V(z) \rangle(x)$ can certify the validity of the interaction.*

Definition 3.4 (Public Coin). *An interactive argument (Setup, P, V) is called public coin if all messages sent from V to P are chosen uniformly and independently at random, and independently of the prover's messages.*

Zero-knowledge is often desirable as well. We recall the standard definition of zero-knowledge with auxiliary input.

Definition 3.5 (Zero-knowledge Arguments). *A public-coin interactive argument (Setup, P, V) for relation ensemble \mathcal{R} is said to have computational zero-knowledge with respect to an auxiliary input if for every PPT interactive algorithm V^* there exists a PPT simulator \mathcal{S} that runs in polynomial time with respect to its first input such that for every $(x, w) \in \mathcal{R}$ and any $z \in \{0, 1\}^*$: $\text{View}(\langle P(w), V^*(z) \rangle(x)) \approx_c \mathcal{S}(x, z)$. Here $\text{View}(\langle P(w), V^*(z) \rangle(x))$ denotes the distribution of the transcript of the interaction between P and V^* and \approx_c denotes computational indistinguishability. If the statistical distance between the two distributions is negligible then the interactive argument is said to have statistical zero-knowledge. If the simulator is allowed to abort with probability at most $1/2$, but the distribution of its output conditioned on no abort is identical to $\text{View}(\langle P(w), V^*(z) \rangle(x))$, then the interactive argument is said to have perfect zero-knowledge.*

Arguments often need some form of knowledge soundness for security. We recall one such notion: witness-extended emulation.

Definition 3.6 (Witness-extended Emulation [Lin03]). *Given a public-coin interactive argument (Setup, P, V) and arbitrary prover algorithm P^* , let $\text{Rec}(P^*, \text{pp}, x, st)$ denote the message transcript between P^* and V on shared input x , initial prover state st , and pp generated by Setup . Furthermore, let $\mathbf{E}^{\text{Rec}(P^*, \text{pp}, x, st)}$ denote a machine \mathbf{E} with a transcript oracle for this interaction that can be rewound to any round and run again on fresh verifier randomness. The tuple (Setup, P, V) has witness-extended emulation if for every deterministic polynomial-time P^* there exists an expected polynomial-time emulator \mathbf{E} such that for all non-uniform polynomial-time adversaries A the following holds:*

$$\Pr \left[\begin{array}{l} A(tr) = 1: \quad \text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux}), (x, st) \leftarrow A(\text{pp}), \\ \quad \quad \quad tr \leftarrow \text{Rec}(P^*, \text{pp}, x, st) \end{array} \right] \approx \quad (1)$$

$$\Pr \left[\begin{array}{l} A(tr) = 1 \wedge \\ tr \text{ accepting} \implies (x, w) \in \mathcal{R} : \quad \text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux}), (x, st) \leftarrow A(\text{pp}), \\ \quad \quad \quad (tr, w) \leftarrow \mathbf{E}^{\text{Rec}(P^*, \text{pp}, x, st)}(\text{pp}, x) \end{array} \right].$$

4 The Honest Majority Multi-Prover Model

In this section, we give our formal definition of the Honest Majority Multi-Prover Model for interactive arguments and prove [Theorem 1.1](#). We begin with the definition of the HMMP model.

Definition 4.1. *The Honest Majority Multi-Prover (HMMP) Model is an interactive protocol framework where M provers P_1, \dots, P_M collectively convince a verifier V of the validity of a statement. During any round of communication, the provers P_1, \dots, P_M are allowed to communicate and collectively send a message to V . Communication is done via an authenticated sender broadcast channel.*

Given [Definition 4.1](#), we define interactive arguments in this model.

Definition 4.2 (HMMP Model Interactive Arguments). *Let \mathcal{R} be a relation ensemble and let $M \in \mathbb{Z}^+$. Let $\mathbf{P} = (P_1, \dots, P_M)$, let (\mathbf{P}, V) denote a pair of PPT interactive HMMP model algorithms, and let Setup denote a non-interactive setup algorithm that outputs public parameters pp given security parameter 1^λ . Let $\langle \mathbf{P}(\text{pp}, x, w), V(\text{pp}, x) \rangle$ denote the output of V 's interaction with \mathbf{P} on common inputs pp and statement x , where \mathbf{P} additionally has witness w . The triple $(\text{Setup}, \mathbf{P}, V)$ is an M -prover interactive argument for \mathcal{R} in the HMMP model if*

1. **δ -Correctness.** *For every $(x, w) \in \mathcal{R}$, for every non-uniform polynomial-sized adversary A that statically corrupts $t < M/2$ provers before the selection of public parameters, and for all $z \in \{0, 1\}^*$, we have $\Pr[\langle \tilde{\mathbf{P}}(w), V(z) \rangle(\text{pp}, x) = 1] \geq 1 - \delta(\lambda)$, where $\tilde{\mathbf{P}}$ denotes all honest and corrupt provers, and the probability is taken over $\text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux})$ and the random coins of V . If $\delta = 0$, then we say the argument has perfect correctness.*
2. **ε -Computational Soundness.** *For every x and w such that $(x, w) \notin \mathcal{R}$, all non-uniform polynomial-sized adversaries $\tilde{\mathbf{P}} = (\tilde{P}_1, \dots, \tilde{P}_M)$, and every $z \in \{0, 1\}^*$, we have $\Pr[\langle \tilde{\mathbf{P}}(w), V(z) \rangle(\text{pp}, x) = 1] \leq \varepsilon(\lambda)$, where the probability is taken over $\text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux})$ and the random coins of V .*

If P_1, \dots, P_M are computationally unbounded algorithms, then we replace all adversaries above with another computationally unbounded algorithm and obtain an interactive proof for \mathcal{R} in the HMMP model.

Our HMMP model definitions public verifiability, public coin, zero-knowledge, and non-interactivity are identical to the standard definitions.

Definition 4.3 (HMMP Model Public Verifiability). *A HMMP model interactive argument $(\text{Setup}, \mathbf{P}, V)$ is publicly verifiable if any party viewing the transcript $T = \langle \mathbf{P}(w), V(z) \rangle(x)$ can certify the validity of the interaction. Importantly, the interaction among provers \mathbf{P} is not part of the transcript T .*

Definition 4.4 (HMMP Model Public Coin). A HMMP model interactive argument $(\text{Setup}, \mathbf{P}, V)$ is called public coin if all messages sent from V to \mathbf{P} are chosen uniformly and independently at random, and independently of the prover’s messages.

Definition 4.5 (HMMP Model Zero-Knowledge). A public-coin interactive HMMP model argument $(\text{Setup}, \mathbf{P}, V)$ for a relation ensemble \mathcal{R} is said to have γ -computational zero-knowledge with respect to an auxiliary input if for every PPT interactive algorithm V^* there exists a PPT simulator \mathcal{S} running in polynomial time with respect to its first input such that for every $(x, w) \in \mathcal{R}$ and any $z \in \{0, 1\}^*$, we have $\text{View}(\langle \mathbf{P}(w), V^*(z) \rangle(x)) \approx_\gamma S(x, z)$, where all provers in \mathbf{P} are honest. Here, View denotes the distribution of the transcript of the interaction between \mathbf{P} and V^* and \approx_γ denotes γ -computational indistinguishability. If the statistical distance between the above two distributions is negligible, then the argument has statistical zero-knowledge. If the simulator is allowed to abort with probability at most $1/2$, but the distribution of its output conditioned on no abort is identical to $\text{View}(\langle \mathbf{P}(w), V^*(z) \rangle(x))$, then the argument has perfect zero-knowledge.

We formally define HMMP model witness-extended emulation as well.

Definition 4.6 (HMMP Model Witness-extended Emulation). Given a public-coin HMMP model interactive argument $(\text{Setup}, \mathbf{P}, V)$ and adversary A_{hmp} against the argument, let $\text{Rec}(\tilde{\mathbf{P}}, \text{pp}, x, st)$ denote the message transcript between $\tilde{\mathbf{P}}$ and V on shared input x , initial prover state st , and pp generated by Setup , where $\tilde{\mathbf{P}} = (\tilde{P}_1, \dots, \tilde{P}_M)$ are M corrupt HMMP model prover algorithms. Furthermore, let $\mathbf{E}^{\text{Rec}(\tilde{\mathbf{P}}, \text{pp}, x, st)}$ denote a machine \mathbf{E} with a transcript oracle for this interaction that can be rewound to any round and run again on fresh verifier randomness. The tuple $(\text{Setup}, \mathbf{P}, V)$ has witness-extended emulation if for every deterministic polynomial-time A_{hmp} there exists an expected polynomial-time emulator \mathbf{E} such that for all non-uniform polynomial-time adversaries A the following holds:

$$\Pr \left[\begin{array}{l} A(tr) = 1 : \text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux}), (x, st) \leftarrow A(\text{pp}), \\ tr \leftarrow \text{Rec}(\tilde{\mathbf{P}}, \text{pp}, x, st) \end{array} \right] \approx \tag{2}$$

$$\Pr \left[\begin{array}{l} A(tr) = 1 \wedge \\ tr \text{ accepting} \implies (x, w) \in \mathcal{R} : \text{pp} \leftarrow \text{Setup}(1^\lambda, \text{aux}), (x, st) \leftarrow A(\text{pp}), \\ (tr, w) \leftarrow \mathbf{E}^{\text{Rec}(\tilde{\mathbf{P}}, \text{pp}, x, st)}(\text{pp}, x) \end{array} \right].$$

4.1 The HMMP Model Transformation

Given any interactive argument such that the prover has a suitably partitionable next-message function, we can transform this interactive argument to one that is secure in the HMMP model. This transformation preserves zero-knowledge and non-interactivity. Our construction makes explicit use of the partitioned function and employs a “permute and audit” strategy where provers are shuffled and must perform computations in audit windows to cross-check other provers. The properties of this transformation are given in [Theorem 1.1](#).

4.1.1 Construction.

Let $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$ be an interactive argument where P_{arg} is suitably partitionable. Then we construct a HMMP model interactive argument $(\text{Setup}_{\text{hmp}}, \mathbf{P}_{\text{hmp}}, V_{\text{hmp}})$ with the desired properties.

The Setup Algorithm. The algorithm $\text{Setup}_{\text{hmp}}$ is a modification of $\text{Setup}_{\text{arg}}$ that additionally outputs a random permutation π on the set $[M]$. Suppose $\text{Setup}_{\text{arg}}$ takes as input 1^λ for security parameter λ and additional input aux . Then we define $\text{Setup}_{\text{hmp}}$ as follows: on input $1^\lambda, \text{aux}, 1^M, 1^\tau$ for security parameter λ , auxiliary input aux , number of provers M , and audit parameter τ , first sample $\text{pp}_{\text{arg}} \leftarrow \text{Setup}_{\text{arg}}(1^\lambda, \text{aux})$. Next sample uniformly random permutation $\pi: [M] \rightarrow [M]$. Finally output $\text{pp} := (\text{pp}_{\text{arg}}, \pi, M, \tau)$.

As stated previously, we include the random permutation π to shuffle the (potentially corrupt) provers to ensure soundness. We remark that π can also be the first message sent by the verifier V_{hmp} in the definition of the interaction, but we choose to include this in the setup phase.

Protocol 1: P_{hmp} Implementation

Parties: M provers $\mathbf{P}_{\text{hmp}} = (P_1, \dots, P_M)$.

Input : $\text{pp} \leftarrow \text{Setup}_{\text{hmp}}(1^\lambda, \text{aux}, 1^M, 1^\tau)$, $(x, w) \in \mathcal{R}_{\mathcal{L}}$, partial transcript Π with all messages sent between \mathbf{P}_{hmp} and V_{hmp} so far.

Internal State: Each P_i has tuple $\mathcal{P} \in [M]^M$ of the current provers, initially set as $\mathcal{P} = (1, 2, \dots, M)$.

```
1 foreach  $i \in [M]$  do
2   foreach  $j \in [i, i + \tau - 1]$  do
3      $P_{\pi(j)}$  computes and broadcasts  $y_j^{(i)} = f_i(S(i, \text{pp} \circ x \circ w \circ \Pi))$ .
4   if  $y_j^{(i)}$  not all equal for  $j \in [i, i + \tau - 1]$  then
5      $\lfloor$  All provers set  $y^{(i)} = \text{Audit}(\text{pp}, x, w, \Pi, i, \{y_j^{(i)}\}_j)$ .
6   else All provers set  $y^{(i)} = y_i^{(i)}$ .
7 For all  $i \in [M]$ ,  $P_i$  sets  $\beta^{(i)} = f^*(y^{(1)}, \dots, y^{(M)})$  and broadcasts  $\beta^{(i)}$  to  $V_{\text{hmp}}$ .
```

The Multi-Prover Algorithm. We present \mathbf{P}_{hmp} in Protocol 1 and describe it in detail. Let f be the next message function of P_{arg} , and let f_1, \dots, f_M, f^* be functions and S be a partition function such that for all $z \in \{0, 1\}^*$ that are valid prover inputs $f(z) = f^*(f_1(S(1, z)), \dots, f_M(S(M, z)))$. To implement P_{arg} using M provers, we utilize the functions f_i, f^* , and S . Given a permutation π and parameter τ , function f_i is computed by all provers $P_{\pi(j)}$ for $j \in [i, i + \tau - 1]_M$, where $[i, i + \tau - 1]_M$ denotes the set $\{i, i + 1, \dots, i + \tau - 1\}$ modulo M . Then each $P_{\pi(j)}$ broadcasts their computation of the function f_i , and if any of the τ messages differ then the sub-routine **Audit** is invoked.

The sub-routine **Audit** (presented in Protocol 2) is a *global* audit of the current message being computed. That is, if function f_i is being computed and **Audit** is called, all M provers engage in a global audit of this computation. In this sub-routine, all provers compute the value f_i and compare it to the values broadcast by the provers $P_{\pi(j)}$ for $j \in [i, i + \tau - 1]_M$. Then for every j such that the value presented by $P_{\pi(j)}$ differs from the globally computed value, the index j is added to a set A , and then the sub-routine **Rebalance** (presented in Protocol 3) is called.

The sub-routine **Rebalance** removes prover $P_{\pi(j)}$ from the protocol and replaces it with the most recent prover $P_{\pi(k)}$ for $k < j$ that was not removed from the protocol, for every $j \in A$. This ensures that over the course of audits if malicious parties submit incorrect values then they are removed from the protocol for future computations.

Our **Audit** and **Rebalance** sub-routines also allow us to handle aborting provers. Since all provers have the same inputs, if any prover $P_{\pi(i)}$ aborts during the execution of the protocol, the remaining provers simply treat the value that was to be computed by this prover as a null value (e.g., \perp) and proceed as normal. Then when this value is given, the provers engage in **Audit** and **Rebalance** to replace the prover that aborted.

The Verifier Algorithm. We present the algorithm V_{hmp} in Protocol 4. The algorithm is identical to V_{arg} , except that additionally V_{hmp} takes the majority of the prover messages it receives, and then uses this majority value as input to V_{arg} . It then forwards the message from V_{arg} to the broadcast channel, sending this message to all M provers, and accepts or rejects as V_{arg} . Additionally, in the case that **Audit** is called by \mathbf{P}_{hmp} , V_{hmp} receives updated prover tuples \mathcal{P}_i for every $i \in [M]$, computes the majority of these tuples and updates its state \mathcal{P} with this majority.

4.1.2 Correctness and Soundness.

We show that the defined interactive argument $(\text{Setup}_{\text{hmp}}, \mathbf{P}_{\text{hmp}}, V_{\text{hmp}})$ is an argument in the HMMP model, as per Definition 4.1. To show correctness and soundness, we first show some key properties about the algorithm \mathbf{P}_{hmp} .

Protocol 2: Audit sub-routine	Protocol 3: Rebalance sub-routine
<p>Input : Public parameters pp, statement x, audit parameter τ, partial transcript Π, current index i, claimed values $\{y_j^{(i)}\}_{j \in [i, i+\tau-1]_M}$.</p> <p>Output : Value $y^{(i)}$.</p> <ol style="list-style-type: none"> 1 foreach $k \in [M]$ do 2 P_k computes and broadcasts $\tilde{y}_k^{(i)} = f_i(S(i, \text{pp} \circ x \circ w \circ \Pi))$. 3 foreach $k \in [M]$ do 4 P_k computes $y^{(i)} = \text{majority}(\tilde{y}_1^{(i)}, \dots, \tilde{y}_M^{(i)})$. 5 P_k computes and broadcasts $A_k = \{j \in [i, i+\tau-1]_M : y_j^{(i)} \neq y^{(i)}\}$. 6 foreach $k \in [M]$ do 7 P_k computes $A = \text{majority}(A_1, \dots, A_M)$. 8 Run $\text{Rebalance}(A)$. 9 return $y^{(i)}$. 	<p>Input : Set $A \subset [M]$.</p> <ol style="list-style-type: none"> 1 foreach $i \in A$ do 2 foreach $j = i, i-1, \dots, 1, M, \dots, i+1$ do 3 Find first j such that $k = (i-j) \bmod M \notin A$. 4 Remove party $P_{\pi(i)}$ from the protocol and set $P_{\pi(i)} := P_{\pi(k)}$. 5 foreach $j \in [M]$ do 6 P_j updates internal state \mathcal{P}_j by replacing $\pi(i)$ with $\pi(k)$. 7 P_j broadcasts \mathcal{P}_j to V_{hmp}.

Fix the number of provers $M \in \mathbb{N}$, the number of corrupted parties $t \in \mathbb{N}$ such that $t < M/2$, and the audit window size $\tau \in \mathbb{N}$ such that $\tau \leq t$. We first show that the algorithm **Audit** never evicts an honest prover.

Lemma 4.7. *Let $\text{pp}, x, \tau, \Pi, i, \{y_j^{(i)} : j \in [i, i+\tau-1]_M\}$ be an input to **Audit**, and let $\mathcal{H} \subseteq [M]$ be the indices of the honest provers. Let $A \subset [M]$ be the set computed in [Protocol 2](#) via **Audit**($\text{pp}, x, \tau, \Pi, i, \{y_j^{(i)}\}_j$). Then $\mathcal{H} \cap \pi(A) = \emptyset$, where $\pi(A) := \{\pi(i) : i \in A\}$.*

Proof. The proof follows directly from the honest majority assumption. That is, in [Line 7](#) of [Protocol 2](#), since there is an honest majority of provers, the set A contains all indices j such that $P_{\pi(j)}$ did not compute the correct value of f_i . Since any honest prover computes the correct value of f_i , this implies that the computed A only contains indices j such that $P_{\pi(j)}$ submitted an incorrect value, which implies that $\mathcal{H} \cap \pi(A) = \emptyset$. \square

The next property we observe is that if every window of τ provers contains at least one honest party, then this property is invariant under **Audit**. That is, running **Audit** never introduces a window of τ dishonest provers.

Lemma 4.8. *Let $\mathcal{H} \subseteq [M]$ be the indices of the honest provers and fix $\text{pp} \leftarrow \text{Setup}_{\text{hmp}}(1^\lambda, \text{aux}, 1^M, 1^\tau)$. If for every $i \in [M]$ it holds that $\mathcal{H} \cap \pi(\mathcal{J}) \neq \emptyset$ for $\mathcal{J} = [i, i+\tau-1]_M$, then after executing **Audit**($\text{pp}, x, \tau, \Pi, i, \{y_j^{(i)}\}_j$) for any $i \in [M]$, it holds that $\mathcal{H} \cap \pi(\mathcal{J}) \neq \emptyset$.*

Proof. Note that **Audit** is only called during round $i \in [M]$ of the computation if not all messages computed during this round by the window of τ provers is equal; thus without loss of generality, **Audit** is only called when there exists at least one dishonest prover. By assumption, there exists at least one honest prover in every window of τ provers. This implies that every honest prover is at most distance τ from another honest prover; that is, for every $i, j \in \mathcal{H}$, it holds that $|\pi(i) - \pi(j)| \leq \tau$. By [Lemma 4.7](#), honest provers are never evicted from their windows during an audit, which implies that after any **Audit** it holds that $|\pi(i) - \pi(j)| \leq \tau$ for every $i, j \in \mathcal{H}$, which in turn applies that there exists at least one honest prover in every window of τ provers. \square

Now given [Lemmas 4.7](#) and [4.8](#), we recall and prove [Proposition 2.1](#).

Protocol 4: V_{hmp} Implementation

Input : $\text{pp} \leftarrow \text{Setup}_{\text{hmp}}(1^\lambda, \text{aux}, 1^M, 1^\tau)$, $x \in \{0, 1\}^*$, partial transcript Π with all messages sent between \mathbf{P}_{hmp} and V_{hmp} so far.

Output : Accept or Reject.

Internal State: Tuple $\mathcal{P} \in [M]^M$ of the current provers, initially set as $\mathcal{P} = (1, 2, \dots, M)$.

- 1 **if** Audit and Rebalance have been invoked **then**
 - 2 \lfloor Update $\mathcal{P} = \text{majority}(\mathcal{P}_1, \dots, \mathcal{P}_M)$, where \mathcal{P}_i are broadcast from Rebalance.
 - 3 Set $\beta = \text{majority}\{\{\beta^{(i)}\}_{i \in \mathcal{P}}\}$.
 - 4 Compute $\alpha \leftarrow V_{\text{arg}}(\text{pp}, x, (\Pi, \beta))$.
 - 5 **if** α is accept or reject **then return** α .
 - 6 Broadcast α to \mathbf{P}_{hmp} .
-

Proposition 2.1. Let Π_{arg} be an argument with δ_{arg} correctness error such that [Theorem 1.1](#) can be applied to Π_{arg} . Let $\mathcal{H} \subseteq [M]$ be the indices of the honest provers and let π be a permutation over $[M]$. For every $i \in [M]$, if $\mathcal{H} \cap \pi(\mathcal{J}) \neq \emptyset$ for $\mathcal{J} = \{i, i+1, \dots, i+\tau-1\}/(M\mathbb{Z})$ then the argument Π_{hmp} has correctness error δ_{arg} .

Proof. By [Lemmas 4.7](#) and [4.8](#), we have that if every window of τ provers contains at least one honest party, then the output of \mathbf{P}_{hmp} with $t < M/2$ corruptions is identical to the output of \mathbf{P}_{hmp} with $t = 0$ corruptions. This shows the result. \square

Finally, the following proposition follows directly from applying [Lemma 2.2](#). This proposition states that the probability of sampling a bad permutation π (i.e., a permutation where no honest parties exist in some window of τ provers) is small (depending of the choice of τ).

Proposition 4.9. Let $\mathcal{H} \subseteq [M]$ be the indices of the honest provers. Then $\Pr[\exists i \in [M]: \mathcal{H} \cap \pi(\mathcal{J}_i) = \emptyset] \leq M \cdot (t/M)^\tau$, where t is the number of corrupt provers and the probability is taken over uniformly random permutation $\pi: [M] \rightarrow [M]$.

For completeness, we recall and prove [Lemma 2.2](#).

Lemma 2.2. Let $\tau, t, M \in \mathbb{N}$ be integers such that $\tau \leq t$ and $t < M$, and let $T \subset \{1, \dots, M\}$ be a subset of size t . For permutation π over $[M]$, let $S_{(i,j)} = \{S_{\ell,\pi}: \ell = (i+k \bmod M) \forall k \in \{0, \dots, j-1\}\}$, where $S_{\ell,\pi} = \pi(\ell)$ for all $\ell \in \{1, \dots, M\}$. Then $\Pr[\exists i \in [M]: S_{(i,\tau)} \subseteq T] \leq M \cdot (t/M)^\tau$, where the probability is taken over uniformly random permutation π .

Proof. Fix $i \in [M]$. Note that $S_{(i,\tau)} \subseteq T$ if and only if for every $j \in [0, \tau-1]$, we have $S_{i+j \bmod M} \in T$. For a fixed j we have that

$$\Pr_{\pi} [S_{i+j \bmod M} \in T] = \Pr[\exists k \in T: \pi(i+j \bmod M) = k] \leq (t/M).$$

Note that all events $S_{i+j \bmod M} \in T$ are independent for all j . Therefore we have

$$\Pr_{\pi} [S_{(i,\tau)} \subseteq T] = \prod_{j=0}^{\tau-1} \Pr_{\pi} [S_{i+j \bmod M} \in T] \leq (t/M)^\tau.$$

Applying Union Bound over indices $i \in [M]$ yields the desired probability. \square

Correctness. By [Proposition 2.1](#), as long as the permutation sampled by $\text{Setup}_{\text{hmp}}$ is not a bad permutation, then the interaction between \mathbf{P}_{hmp} and V_{hmp} is identical to the interaction between \mathbf{P}_{arg} and V_{arg} . In particular, if Π_{arg} has δ_{arg} -correctness, then Π_{hmp} has δ_{arg} -correctness conditioned on the permutation being good. By [Proposition 4.9](#), the probability that a good permutation is sampled is at least $1 - M \cdot (t/M)^\tau$. Let $p := p(M, t, \tau) = 1 - M \cdot (t/M)^\tau$. Then the probability that $\langle \mathbf{P}_{\text{hmp}}(w), V_{\text{hmp}} \rangle(\text{pp}, x) = 1$ is at least $(1 - \delta_{\text{arg}}) \cdot p$. This gives our stated correctness error $\delta_{\text{hmp}} = 1 - (1 - \delta_{\text{arg}}) \cdot (1 - M \cdot (t/M)^\tau)$.

Soundness. As discussed in [Section 2.1](#), the soundness of our HMMP model protocol directly reduces to the soundness of the underlying interactive argument. This is because by our definition, we require soundness to hold with respect to M corrupt provers $\tilde{P}_1, \dots, \tilde{P}_M$. Notice that any set of M corrupt provers that can break the soundness of our HMMP model argument can directly be used to break the soundness of the underlying argument with identical probability. The reduction is simple: first run the corrupt provers and wait for them to broadcast their answers. Take the M values broadcast by the provers, take the majority, and directly forward this to the verifier for the argument. By our construction of V_{hmp} , it is easy to see that if $\tilde{P}_1, \dots, \tilde{P}_M$ break soundness with V_{hmp} , then our constructed adversary breaks soundness with V_{arg} .

Zero-Knowledge. Given a zero-knowledge argument $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$, we modify our algorithm \mathbf{P}_{hmp} to obtain zero-knowledge for the transformed HMMP model argument. For this transformation, we first observe that the next message function of P_{arg} can be abstracted as two functions $f: \{0, 1\}^* \rightarrow \{0, 1\}^*$ and $g: \{0, 1\}^* \rightarrow \{0, 1\}^*$ such that for any input $x \in \{0, 1\}^*$: (1) $f(x)$ computes the desired functionality; and (2) $g(f(x); r)$ gives the zero-knowledge property, given uniformly random string $r \xleftarrow{\$} \{0, 1\}^*$. Intuitively, we achieve zero-knowledge for $(\text{Setup}_{\text{hmp}}, \mathbf{P}_{\text{hmp}}, V_{\text{hmp}})$ by additionally requiring that

1. All provers P_i engage in a protocol to generate a shared (almost) uniformly random string $r \xleftarrow{\$} \{0, 1\}^*$ (or simply use a randomness beacon); and
2. Compute $\beta^{(i)} = g(f^*(y^{(1)}, \dots, y^{(M)}))$ in the final step of \mathbf{P}_{hmp} in [Protocol 1](#).

Thus with high probability over the choice of permutation π , so long as the honest majority of provers all receive the same random string r , the verifier V_{hmp} receives the correct zero-knowledge message. The random string r can be generated in every round either by assuming access to a randomness beacon [[Rab83, SJK⁺17](#)] that broadcasts random bits to all provers, or by any MPC protocol for collective coin tossing with an honest majority of provers [[BOGW88, RBO89](#)]. In the MPC case, since we are in the honest majority setting, adversaries can only bias the output of the coins by a negligible amount [[BOGW88, RBO89](#)]. Further, in the case of abort, parties run an **Audit** with the set A containing the index of the aborting party, then re-run the randomness generation protocol.

Non-Interactive Arguments. To obtain non-interactivity via the Fiat-Shamir transform [[FS87](#)], we assume all provers in \mathbf{P}_{hmp} have access to the same random oracle H . Let $\mathbf{pp}, x, w, \tau, \Pi$ be the input to \mathbf{P}_{hmp} during any round of the computation, where the public parameters are generated as $\mathbf{pp} \leftarrow \text{Setup}_{\text{hmp}}(1^\lambda, \text{aux}, 1^M, 1^\tau)$, $(x, w) \in \mathcal{R}_{\mathcal{L}}$, and Π is a partial transcript containing all prior messages sent by \mathbf{P}_{hmp} and messages generated by H (i.e., the simulated messages of V_{arg}). Then the provers \mathbf{P}_{hmp} generate the next message α of V_{hmp} as follows:

1. upon the broadcast of all $\beta^{(i)}$ for $i \in [M]$, every prover P_j computes $\beta = \text{majority}(\beta^{(1)}, \dots, \beta^{(M)})$; and
2. every P_j computes $\alpha \leftarrow H(\mathbf{pp}, x, (\Pi, \beta))$ and uses this α as the verifier message.

It is clear that this transformation is identical to the case of transforming $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$ into a non-interactive argument. Note that if any malicious prover of the protocol \mathbf{P}_{hmp} deviates from defining α as described above, with high probability over the choice of random oracle H these malicious provers are caught in the proceeding rounds of the protocol.

Witness-Extended Emulation. We show that the transformation of [Theorem 1.1](#) also preserves witness-extended emulation. We capture this via the following proposition.

Proposition 4.10. *If $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$ is an interactive argument with witness-extended emulation, then $(\text{Setup}_{\text{hmp}}, \mathbf{P}_{\text{hmp}}, V_{\text{hmp}})$ obtained by applying [Theorem 1.1](#) has witness-extended emulation.*

Proof. Suppose that $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$ is an interactive argument with witness-extended emulation. Let $(\text{Setup}_{\text{hmp}}, \mathbf{P}_{\text{hmp}}, V_{\text{hmp}})$ be the HMMP model interactive argument obtained by applying [Theorem 1.1](#). We argue that this HMMP model argument has witness-extended emulation (as per [Definition 4.6](#)). We show this by constructing an emulator E_{hmp} such that [Eq. \(2\)](#) holds. Suppose E_{hmp} has oracle access to $\text{Rec}(\tilde{\mathbf{P}}_{\text{hmp}}, \text{pp}, x, st)$ where $\text{pp} \leftarrow \text{Setup}_{\text{hmp}}(1^\lambda, \text{aux})$ and $(x, st) \leftarrow A(\text{pp})$. Then given pp, x as input, E_{hmp} does the following

1. Parse pp as $(\text{pp}', \pi, M, \tau)$, where $\text{pp}' \leftarrow \text{Setup}_{\text{arg}}(1^\lambda, \text{aux})$.
2. E_{hmp} initializes an empty transcript tr .
3. Let E be the emulator for argument $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$. E_{hmp} runs $E(\text{pp}', x)$ and simulates oracle access to $\text{Rec}(P^*, \text{pp}', x, st)$, where P^* is some arbitrary prover strategy. In particular, E_{hmp} simulates this oracle access by directly forwarding the transcript derived from its oracle $\text{Rec}(\tilde{\mathbf{P}}_{\text{hmp}}, \text{pp}, x, st)$ with the following modifications:
 - Whenever $\tilde{\mathbf{P}}_{\text{hmp}}$ sends M messages to V_{hmp} , E_{hmp} replaces these M messages in the transcript with the majority of these messages.
 - Any message from V_{hmp} that is broadcast to the M provers of $\tilde{\mathbf{P}}_{\text{hmp}}$ is instead replaced with a single message.
 - Any messages generated due to a call to `Audit` are not given as part of the interaction.

Note that this simulation is otherwise unchanged; whenever E rewinds to round i , then E_{hmp} similarly rewinds to round i . Note that E_{hmp} stores the unmodified transcript in tr .

4. Upon receiving $(tr', w) \leftarrow E(\text{pp}, x)$, output (tr, w) .

We now argue that [Eq. \(2\)](#) holds. In particular, we show that adversary A on input tr cannot distinguish between whether tr was output by $\text{Rec}(\tilde{\mathbf{P}}_{\text{hmp}}, \text{pp}, x, st)$ or E_{hmp} . Further, if tr output by E_{hmp} is accepting then $(x, w) \in \mathcal{R}$. First note that A cannot distinguish whether tr is generated by Rec or E_{hmp} since the transcript generated by E_{hmp} is identical to a transcript generated by Rec , except for that E_{hmp} might rewind and resample *verifier* randomness. Since the verifier is public-coin, as long as E_{hmp} samples new coins the same way the verifier does, the transcripts are indistinguishable. Next suppose that tr is accepting. Note that a valid arbitrary prover strategy P^* is to simulate $\tilde{\mathbf{P}}_{\text{hmp}}$ honestly, send the majority of M messages received from $\tilde{\mathbf{P}}_{\text{hmp}}$ to V_{arg} . This implies that if tr is accepting then tr' is also accepting. Further, by [Theorem 1.1](#), if $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$ is an interactive argument for \mathcal{R} , then so is $(\text{Setup}_{\text{hmp}}, \mathbf{P}_{\text{hmp}}, V_{\text{hmp}})$. By witness-extended emulation of $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$, since tr' is accepting we have that $(x, w) \in \mathcal{R}$; this implies that tr is accepting implies that $(x, w) \in \mathcal{R}$. Finally, we establish indistinguishability by observing that if [Eq. \(2\)](#) does not hold, then we immediately have a prover P^* such that [Eq. \(1\)](#) does not hold, breaking witness-extended emulation of $(\text{Setup}_{\text{arg}}, P_{\text{arg}}, V_{\text{arg}})$. \square

5 Efficient Interactive Arguments for NP in the HMMP Model

We dedicate this section to proving [Theorem 1.2](#), which we recall here.

Theorem 1.2. *Let $M, \tau, t, \lambda \in \mathbb{N}$ be parameters such that $t < M/2$ and $\tau \leq t$. For every NP relation with statements of size $\text{poly}(\lambda)$ and witnesses of size N , there exists a HMMP model interactive argument Π_{hmp} with M provers and the following efficiency properties:*

- the individual prover computation is $\tilde{O}_\lambda(\tau \cdot (N/M))$;
- the communication complexity between provers is $\tilde{O}_\lambda(\tau \cdot (N/M))$ bits;
- the verifier computation is $O_\lambda(M + \log(N))$; and

- the verifier receives $O_\lambda(M \cdot \log(N))$ bits from all provers and sends $O_\lambda(\log(N))$ bits to all provers.

To show [Theorem 1.2](#), it suffices to prove [Propositions 5.1, 2.3](#) and [2.5](#). Then by our discussion in the overview, we obtain [Theorem 1.2](#).

5.1 HMMP Model Circuit SAT

In [Section 2.2](#), we outlined a HMMP model protocol for circuit satisfiability. We give the formal protocol here and add some discussion.

One key condition is the need for a circuit C to be partitionable into C_1, \dots, C_M circuits each of size N/M such that for any $j \in [M]$, partition C_j depends only on wires from C_i for any $i \leq j$ and the input x . We remark that a suitable topological sorting will define a sequence of individual gates with this property, after which the partitions C_1, \dots, C_M can be suitably defined.

Given the above discussion, the discussion in [Section 2.2](#), and [Protocol 5, Proposition 2.3](#) follows.

Protocol 5: HMMP Model Circuit SAT	
<hr/>	
Parties :	M provers $\mathbf{P}_{\text{hmp}} = (P_1, \dots, P_M)$.
Input :	An arithmetic circuit $C: \{0, 1\}^n \rightarrow \mathbb{F}$ of size N , an input to the circuit x , an output y , and a partition of C into C_1, \dots, C_M .
Output :	A wire transcript w certifying that $C(x) = y$.
1	Each prover sets wire transcript $w_i = \emptyset$ for $i \in [M]$.
2	foreach $i \in [M]$ do
3	foreach $j \in [i, i + \tau - 1]_M$ do
4	$P_{\pi(j)}$ uses w_k for all $k < i$ and x to compute and broadcast a wire transcript $w_{i,j}$ for C_i .
5	if all $w_{i,j}$ are not equal for $j \in [i, i + \tau - 1]_M$ then
6	All provers set $w_i = \text{Audit}(C, x, \{w_k\}_{k < i}, \{w_{i,j}\}_{j \in [i, i + \tau - 1]_M})$.
7	else
8	All provers set $w_i := w_{i,i}$.
9	Each prover defines $w = (w_1, \dots, w_M)$.
10	return w .

5.2 HMMP Model Sum-Check

We give the formal description of our HMMP model sum-check protocol and prove the following proposition.

Proposition 5.1. *There exists a HMMP model protocol for the sum-check protocol with the following properties: for a finite field \mathbb{F} , an n -variate polynomial $f \in \mathbb{F}[X_1, \dots, X_n]$ of individual degree at most d , a subset $H \subset \mathbb{F}$, and parameters τ, M such that $M \leq |H|^n$,*

- the individual prover computation is $O(\tau \cdot ((n \cdot |H|^n)/M + |H|))$ polynomial evaluations and $O(\tau \cdot ((n \cdot |H|^n)/M + d \cdot |H|))$ field additions; and
- each prover broadcasts $O(\tau \cdot d \cdot (n + |H|))$ field elements.

Note that setting $H = \{0, 1\}$ and $d = \Theta(1)$ in [Proposition 5.1](#) directly yields [Corollary 2.4](#).

5.2.1 Single-Prover Sum-Check Protocol

We first recall the single-prover sum-check protocol. We present it in [Protocol 6](#), and give an overview here. For a finite field \mathbb{F} , polynomial $g \in \mathbb{F}[X_1, \dots, X_n]$ of individual degree at most d , and subset $H \subset \mathbb{F}$, the sum-check protocol [[LFKN92](#)] is an interactive proof certifying that $\sum_{x \in H^n} g(x) = y$ for some $y \in \mathbb{F}$. The

Protocol 6: Sum-Check Protocol

- Parties** : Prover P and verifier V .
- Prover Input** : A multi-variate polynomial $g \in \mathbb{F}[X_1, \dots, X_n]$ of individual degree at most d , a subset $H \subset \mathbb{F}$, and a value $y \in \mathbb{F}$.
- Verifier Input** : Individual degree upper bound d , a subset $H \subset \mathbb{F}$, and value $y \in \mathbb{F}$.
- Verifier Output** : Accept or Reject.
- Oracle** : The verifier V has oracle access to evaluations of the polynomial g .
- 1 P computes and sends to V the univariate polynomial $g_1(X_1) := \sum_{x \in H^{n-1}} g(X_1, \mathbf{x})$.
 - 2 V checks that g_1 is univariate with degree at most d and that $y = \sum_{x \in H} g_1(x)$, aborting with Reject if not.
 - 3 V samples and sends $r_1 \xleftarrow{\$} \mathbb{F}$ to P .
 - 4 **foreach** $1 < j < n$ **do**
 - 5 P computes and sends to V the univariate polynomial $g_j(X_j) := \sum_{x \in H^{n-j}} g(r, X_j, x)$, where $r := (r_1, \dots, r_{j-1})$.
 - 6 V checks that g_j is univariate with degree at most d and that $g_{j-1}(r_{j-1}) = \sum_{x \in H} g_j(x)$, aborting with Reject if not.
 - 7 V samples and sends $r_j \xleftarrow{\$} \mathbb{F}$ to P .
 - 8 P computes and sends to V the univariate polynomial $g_n(X_n) := g(r_1, \dots, r_{n-1}, X_n)$.
 - 9 V checks that g_n is univariate with degree at most d and that $g_{n-1}(r_{n-1}) = \sum_{x \in H} g_n(x)$, aborting with Reject if not.
 - 10 V samples $r_n \xleftarrow{\$} \mathbb{F}$ and queries the oracle at $g(r_1, \dots, r_n)$. V checks that $g_n(r_n) = g(r_1, \dots, r_n)$, aborting with Reject if not.
 - 11 V outputs Accept.
-

protocol assumes that the verifier has oracle access to evaluations of the polynomial f and proceeds in n rounds: in the first round, the prover sends y and the polynomial $g_1(X_1) := \sum_{x \in H^{n-1}} g(X_1, x)$. The verifier checks that (1) g_1 has degree at most d ; and (2) $y = \sum_{b \in H} g_1(b)$, rejecting if either do not hold. The verifier then samples $r_1 \in \mathbb{F}$ uniformly at random and sends it to the prover, and additionally sets $y' = g_1(r_1)$. Then, for every $j = 2, \dots, n$

- the prover computes and sends polynomial $g_j(X_j) := \sum_{x \in H^{n-j}} g(r, X_j, x)$ to the verifier, where $r = (r_1, \dots, r_{j-1})$;
- the verifier checks that g_j has degree at most d and $y' = \sum_{b \in H} g_j(b)$, rejecting if either do not hold; and
- the verifier samples $r_j \in \mathbb{F}$ uniformly at random, sets $y' = g_j(r_j)$, and sends r_j to the prover. If $j = n$, then the verifier instead queries $g(r_1, \dots, r_n)$ and checks if it is equal to $g_n(r_n)$, rejecting if not.

In the sum-check protocol, generating the polynomial g_j for $j = 1, \dots, n$ requires evaluating the polynomial g at $|H|^{n-j}$ points (keeping the variable X_j free) and computing the same number of polynomial additions over \mathbb{F} .

5.2.2 HMMP Model Sum-Check Protocol

We outline how to obtain [Proposition 5.1](#) given the single-prover sum-check protocol and present our HMMP model sum-check in [Protocol 7](#). Since g_j is a polynomial and is a linear function, we can easily sub-divide the computation between multiple parties, then reconstruct the complete polynomial via addition. Suppose we are in round j of the computation and for simplicity assume that $|H|^{n-j}/M$ is an integer. Partition the set H^{n-j} into M subsets H'_1, \dots, H'_M each of size $|H|^{n-j}/M$. Then the computation of g_j is divided as follows. For $i = 1, \dots, M$, audit window $P_{\pi(i)}, P_{\pi(i+1)}, \dots, P_{\pi(i+\tau-1)}$ computes the polynomial $g_{j,i}(X_j) := \sum_{x \in H'_i} g(r_1, \dots, r_{j-1}, X_j, x)$. The audit window then broadcasts $g_{j,i}(X_j)$ and invokes a

global audit and re-balance as necessary. At the end of this round of computation, each party obtains $g_j(X_j) := \sum_{i=1}^M g_{j,i}(X_j)$ and sends this polynomial to the verifier.

Thus during this round j of the computation, each audit window performs $O(|H|^{n-j}/M)$ evaluations of g and $O(|H|^{n-j}/M+M)$ additions over \mathbb{F} , yielding $O(\tau \cdot (|H|^{n-j}/M))$ evaluations and $O(\tau \cdot (|H|^{n-j}/M+M))$ additions for each prover. Moreover, each prover broadcasts $O(d \cdot \tau)$ field elements since each $g_{j,i}$ is a univariate polynomial of degree at most d . As before with the circuit satisfiability algorithm, if $|H|^{n-j}/M$ is not an integer (but is at least 1), we can let $k = \lfloor |H|^{n-j}/M \rfloor$ and partition the computation of g_j into at most $2M$ polynomials $g_{j,1}, \dots, g_{j,2M}$ by partitioning H^{n-j} into at most $2M$ subsets H'_1, \dots, H'_{2M} each of size at most k . Then each audit window computes at most 2 polynomials in round j , yielding the same asymptotic complexity.

Things are slightly more complicated if $M > |H|^{n-j}$. Let j^* be the index such that $|H|^{n-j^*} < M \leq |H|^{n-j^*+1}$. From this same inequality, we know that $|H|^{n-j^*+1} < |H| \cdot M$. Moreover, it holds that $\sum_{j=j^*}^n |H|^{n-j} = (|H|^{n-j^*+1} - 1)/(|H| - 1) < |H| \cdot M$. Thus we can sub-divide the computation of the final $n - j^*$ polynomials with only a $O(|H|)$ overhead for each prover as follows. In a round-robin fashion, the first $|H|^{n-j^*}$ audit windows compute the polynomial $g_{j^*,i}$ for $i = 1, \dots, |H|^{n-j^*}$, followed by the next $|H|^{n-j^*-1}$ audit windows each computing the polynomial $g_{j^*+1,i}$ for $i = 1, \dots, |H|^{n-j^*-1}$, and so on until the final round. Each audit window participates in the round-robin computation at most $O(|H|)$ times, which gives our $O(|H|)$ overhead and $O(\tau \cdot |H|)$ evaluations and $O(\tau \cdot |H| \cdot d)$ field additions in total over all rounds $j = j^*, \dots, n$.

As before, let $j^* \in \{1, \dots, n\}$ be the round such that $|H|^{n-j^*} < M \leq |H|^{n-j^*+1}$. By our previous discussion, the total computation done by each prover in all rounds $j = j^*, j^* + 1, \dots, n$ is $O(\tau \cdot |H|)$ evaluations and $O(\tau \cdot |H| \cdot d)$ field additions. Moreover, for any round $j = 1, \dots, j^* - 1$, each prover performs $O(\tau \cdot (|H|^{n-j}/M))$ evaluations and $O(d \cdot \tau \cdot (|H|^{n-j}/M + M))$ field additions. This gives $O(n \cdot \tau \cdot (|H|^n/M))$ evaluations and $O(n \cdot \tau \cdot (|H|^n/M + |H|))$ field additions over all rounds $j = 1, \dots, j^* - 1$. Thus we have obtained [Proposition 5.1](#). desired.

5.3 HMMP Model Polynomial Commitment Schemes

We formally show [Proposition 2.5](#) in this section. For our purposes, we utilized the Bulletproofs-based Block et al. [\[BHR⁺20\]](#) tailored for multi-linear polynomials.

5.3.1 Polynomial Commitment Scheme Preliminaries

Much of this section is borrowed from Block et al. [\[BHR⁺20\]](#). We let $\mathbb{F} := \mathbb{F}_p$ denote a finite field of prime order p . When clear from context we omit p above. For finite cyclic group \mathbb{G} , we assume that \mathbb{G} is a multiplicative group. For $n \in \mathbb{N}$ and bit-string $b = (b_n, \dots, b_1) \in \{0, 1\}^n$, we assume that b_n is the most significant bit and b_1 is the least significant bit. We let “ \circ ” denote the string concatenation operator; that is, for $b \in \{0, 1\}^n$ and $c \in \{0, 1\}$, we have $b \circ c := (b_n, \dots, b_1, c) \in \{0, 1\}^{n+1}$. Often, we index a vector and/or sequence using a binary string. That is, for $N = 2^n$ and $Y \in \mathbb{F}^N$, for $b \in \{0, 1\}^n$ we have that Y_b is the b -th element of Y , where we interpret b naturally as an integer in the range $\{0, 1, \dots, N - 1\}$.

Discrete-log Assumption. Let GGen be a randomized algorithm that on input 1^λ for security parameter λ returns (\mathbb{G}, p, g) such that \mathbb{G} is the description of a finite cyclic group of prime order p , where p is a λ -bit prime, and g is a generator of \mathbb{G} .

Assumption 5.2 (Discrete-log Assumption). *The discrete-log assumption holds for GGen if for all PPT adversaries A there exists a negligible function μ such that*

$$\Pr \left[\alpha' = \alpha : (\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\lambda), \alpha \xleftarrow{\$} \mathbb{Z}_p, \alpha' \leftarrow A(\mathbb{G}, g, g^\alpha) \right] \leq \mu(\lambda).$$

We use the following variant of the discrete-log assumption, which is equivalent to the assumption above.

Assumption 5.3 (Discrete-log Relation Assumption [BCC⁺16]). *The discrete-log relation assumption holds for GGen if for all PPT adversaries A and for all $n \geq 2$ there exists a negligible function μ such that*

$$\Pr \left[\exists \alpha_i \neq 0 \wedge \prod_{i=1}^n g_i^{\alpha_i} = 1 : \begin{array}{l} (\mathbb{G}, g, p) \leftarrow \text{GGen}(1^\lambda), g_1, \dots, g_n \xleftarrow{\$} \mathbb{G}, \\ (\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_p^n \leftarrow A(\mathbb{G}, g_1, \dots, g_n). \end{array} \right] \leq \mu(\lambda).$$

For α_i not all equal to 0, we say that $\prod_i g_i^{\alpha_i}$ is a non-trivial discrete-log relation between group elements g_1, \dots, g_n . The discrete-log relation assumption states that any PPT adversary cannot find a non-trivial relation between randomly chosen group elements.

Multi-linear Polynomials. An n -variate polynomial $f \in \mathbb{F}[X_1, \dots, X_n]$ is a *multi-linear* polynomial if every variable has individual degree at most 1. Any multi-linear polynomial is uniquely represented by its evaluations over the Boolean hypercube. We capture this in [Fact 5.4](#).

Fact 5.4. *For n -variate multi-linear polynomial f , for any $\zeta \in \mathbb{F}^n$ we have*

$$f(\zeta) = \sum_{b \in \{0,1\}^n} f(b) \cdot \beta(\zeta, b), \quad \beta(\zeta, b) := \prod_{i=1}^n \zeta_i \cdot b_i + (1 - \zeta_i) \cdot (1 - b_i). \quad (3)$$

Note that any sequence $Y \in \mathbb{F}^N$ for $N = 2^n$ uniquely defines a n -variate multi-linear polynomial [Eq. \(3\)](#). For sequence $Y \in \mathbb{F}^N$ and evaluation point $\zeta \in \mathbb{F}^n$, we define $\text{MLE}(Y, \zeta) := \sum_b Y_b \cdot \beta(\zeta, b)$.

Definition 5.5 ([BHR⁺20]). *A tuple of protocols (Setup, Com, Open, Eval) is a multi-linear polynomial commitment scheme if it satisfies the following*

1. $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^N)$ takes as input the unary representations of security parameter $\lambda \in \mathbb{N}$ and size parameter $N = 2^n$ corresponding to $n \in \mathbb{N}$, and produces public parameter pp . We allow pp to contain the description of the field \mathbb{F} over which the multi-linear polynomials will be defined.
2. $(C; d) \leftarrow \text{Com}(\text{pp}, Y)$ takes as input public parameter pp and sequence $Y = (Y_b)_{b \in \{0,1\}^n} \in \mathbb{F}^N$ that defines the multi-linear polynomial to be committed, and outputs public commitment C and secret decommitment d .
3. $b \leftarrow \text{Open}(\text{pp}, C, Y, d)$ takes as input pp , a commitment C , sequence committed Y and a decommitment d and returns a decision bit $b \in \{0, 1\}$.
4. $\text{Eval}(\text{pp}, C, \zeta, \gamma; Y, d)$ is a public-coin interactive protocol between a prover P and a verifier V with common inputs—public parameter pp , commitment C , evaluation point $\zeta \in \mathbb{F}^n$ and claimed evaluation $\gamma \in \mathbb{F}$, and prover has secret inputs Y and d . The prover then engages with the verifier in an interactive argument system for the relation

$$\mathcal{R}_{\text{mle}}(\text{pp}) = \{(C, \zeta, \gamma; Y, d) : \text{Open}(\text{pp}, C, Y, d) = 1 \wedge \gamma = \text{MLE}(Y, \zeta)\}. \quad (4)$$

The output of V is the output of Eval protocol.

We require the following three properties.

1. Computational Binding. For all PPT adversaries A and $n \in \mathbb{N}$

$$\Pr \left[b_0 = b_1 \neq 0 \wedge \mathcal{Y}_0 \neq \mathcal{Y}_1 : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda, 1^N) \\ (C, \mathcal{Y}_0, \mathcal{Y}_1, d_0, d_1) \leftarrow A(\text{pp}) \\ b_0 \leftarrow \text{Open}(\text{pp}, C, Y^{(0)}, d_0) \\ b_1 \leftarrow \text{Open}(\text{pp}, C, Y^{(1)}, d_1) \end{array} \right] \leq \text{negl}(\lambda).$$

2. Perfect Correctness. For all $n, \lambda \in \mathbb{N}$ and all $Y \in \mathbb{F}^N$ and $\zeta \in \mathbb{F}^n$,

$$\Pr \left[\text{Eval}(\text{pp}, C, Z, \gamma; Y, d) = 1 : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda, 1^N), \\ (C; d) \leftarrow \text{Com}(\text{pp}, Y), \gamma = \text{MLE}(Y, \zeta) \end{array} \right] = 1.$$

3. Witness-extended Emulation. We say that the polynomial commitment scheme has witness-extended emulation if Eval has a witness-extended emulation as an interactive argument for the relation ensemble $\{\mathcal{R}_{\text{mle}}(\text{pp})\}_{\text{pp}}$ (Eq. (4)) except with negligible probability over the coins of $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^N)$.

5.3.2 Multi-Linear Polynomial Commitment Scheme in the HMMP Model

We dedicate this section to proving Proposition 2.5. Since the prover of a polynomial commitment scheme is responsible for computing three functionalities—the commitment Com , the opening Open , and the evaluation Eval —we shall define the above functions for each of these phases for simplicity. We first recall the multi-linear polynomial commitment scheme of Block et al. [BHR⁺20], which we refer to as the BHRRS polynomial commitment scheme.

BHRRS Polynomial Commitment Scheme. We recall the BHRRS polynomial commitment scheme. Note that the BHRRS polynomial commitment scheme focuses on *space-efficiency* of the prover; that is, the prover uses space which is poly-logarithmic in the description size of the multi-linear polynomial. They achieve this space-efficiency in the random oracle model by assuming multi-pass streaming access to the description of the multi-linear polynomial that is being committed. For simplicity, we present the BHRRS polynomial commitment scheme without its space-efficient implementation.

We now specify the BHRRS polynomial commitment scheme.

1. $\text{Setup}(1^\lambda, 1^N)$: on input security parameter 1^λ and size parameter 1^N for $N = 2^n$, where $n \in \mathbb{N}$, sample $(\mathbb{G}, p, \mathbf{g} \in \mathbb{G}^N)$, set $\mathbb{F} := \mathbb{F}_P$, and return $\text{pp} := (\mathbb{G}, \mathbb{F}, N, p, \mathbf{g})$. Here, \mathbf{g} is a vector of N distinct generators of \mathbb{G} .
2. $\text{Com}(\text{pp}, Y)$: return commitment $C = \prod_{b \in \{0,1\}^n} \mathbf{g}_b^{Y_b}$.
3. $\text{Open}(\text{pp}, C, Y)$: return 1 if and only if $C = \text{Com}(\text{pp}, Y)$.
4. $\text{Eval}(\text{pp}, C, \zeta, \gamma; Y)$ is an interactive protocol $\langle P, V \rangle$ which begins with V sending $g_V \xleftarrow{\$} \mathbb{G}$. Then both P and V compute $C_\gamma = C \cdot g_V^\zeta$ to bind claimed evaluation γ . P and V then engage in EvalReduce on input $(C_\gamma, Z, \mathbf{g}, g_V, \gamma; Y)$ where P proves knowledge of Y such that

$$C_\gamma = \text{Com}(\text{pp}, Y) \cdot g_V^\zeta \wedge \langle Y, Z \rangle = \gamma,$$

where $\zeta \in \mathbb{F}^n$ and $Z := (z_b = \beta(\zeta, b))_{b \in \{0,1\}^n}$ for function β defined in Eq. (3). The protocol is presented in Protocol 8

BHRRS Polynomial Commitment Scheme in the HMMP Model. We specify how to construct each function Com , Open , and Eval in our HMMP model. Excluding the Open function, the specification and analysis of Com and Eval directly yield Proposition 2.5.

HMMP Model Com. The algorithm Com receives $\text{pp} \leftarrow \text{Setup}_{\text{hmp}}(1^\lambda, 1^N, 1^M, 1^\tau)$ as input along with sequence of evaluations $Y \in \mathbb{F}^N$ defining a multi-linear polynomial. Here $\text{pp} = (\text{pp}', \pi, M, \tau)$ where $\text{pp}' \xleftarrow{\$} \text{Setup}(1^\lambda, 1^N)$ for Setup algorithm of the BHRRS polynomial commitment scheme. By assumption $N = 2^n$ and $M = 2^m$ for $m \leq n$.

Let g_0, \dots, g_{N-1} be generators of a prime-order group \mathbb{G} . Then the commitment to f is $C = \prod_{i=0}^{N-1} g_i^{f(i)}$, where i is interpreted as an element in $\{0, 1\}^n$ in the natural way when evaluating $f(i)$. To implement this

commitment in the HMMP model, first partition the set $\{0, 1, \dots, N-1\}$ into M subsets S_1, \dots, S_M each of size N/M . Then for every $i \in \{1, \dots, M\}$, the audit window $P_{\pi(i)}, \dots, P_{\pi(i+\tau-1)}$ computes the partial commitment $C_i = \prod_{j \in S_i} g_j^{f(j)}$, and broadcasts C_i , invoking a global audit and re-balance as necessary. After the computation of C_M , each party combines all commitments into the final commitment $C = \prod_i C_i$ and broadcast this to the verifier. Here, each prover performs $O(\tau \cdot (N/M))$ group multiplications and exponents, with at most a 2 fold increase if N/M is not an integer, and broadcasts $O(\tau)$ group elements.

HMMP Model Open. Each prover in the **Open** algorithm must send its entire witness Y . In particular, this is identical to the single-prover case, with the modification that V now needs to take majority of all these vectors.

HMMP Model Eval. To produce proofs of evaluations, the BHRRS polynomial commitment scheme utilizes the Bulletproofs inner-product argument in conjunction with a Pedersen commitment. The inner-product argument certifies that (1) $\gamma = \langle X, Y \rangle$, where $X \in \mathbb{F}^N$ is a public vector and $Y \in \mathbb{F}^N$ is a private vector held by the prover; and (2) the vector Y used to compute $\langle X, Y \rangle$ is consistent with some Pedersen commitment C . The inner-product argument employs a “half-and-fold” technique to reduce the claim $\gamma = \langle X, Y \rangle$ to a claim $\gamma' = \langle X', Y' \rangle$, where $X', Y' \in \mathbb{F}^{N/2}$.

We outline the reduction from size N to size $N/2$ vectors. In a bit more detail, let C be a Pedersen commitment to the vector Y using generators $G = (g_0, \dots, g_{N-1})$ and suppose γ is the claimed inner-product $\langle X, Y \rangle$. The verifier first samples and sends a random group generator g , and both the prover and verifier update the commitment C to $C \cdot g^\gamma$. Let $X_o, X_e \in \mathbb{F}^{N/2}$, $Y_o, Y_e \in \mathbb{F}^{N/2}$, and $G_o, G_e \in \mathbb{G}^N$ denote the odd and even halves of the vectors X, Y , and G , respectively. That is, $X_o := (X_{2i})_{i \in [0, N-1]}$ and $X_e := (X_{2i+1})_{i \in [0, N-1]}$, with Y_o, Y_e and G_o, G_e being defined analogously. The prover computes $\gamma_0 = \langle X_o, Y_e \rangle$ and $\gamma_1 = \langle X_e, Y_o \rangle$, and additionally computes two commitments C_0 and C_1 , where C_0 is g^{γ_0} times the Pedersen commitment to Y_e using generators G_o and C_1 is g^{γ_1} times the Pedersen commitment to Y_o using generators G_e . Upon receiving C_0, C_1 , the verifier samples and sends uniformly random challenge $\alpha \in \mathbb{F}$ to the prover. The prover then sends $\gamma' = \alpha^2 \cdot \gamma_0 + \gamma + \alpha^{-2} \cdot \gamma_1$ to the verifier, and both parties compute new commitment $C' = (C_0)^{\alpha^2} \cdot C \cdot (C_1)^{\alpha^{-2}}$, new public vector $X' = \alpha^{-1} \cdot X_e + \alpha \cdot X_o$, and new generator vector $G' = (G_e)^{\alpha^{-1}} * (G_o)^\alpha$, where “ $*$ ” denotes entry-wise multiplication and the exponents are computed for every entry. Additionally, the prover computes $Y' = \alpha \cdot Y_e + \alpha^{-1} \cdot Y_o$. This reduction continues until the vectors X', Y', G' all have length 1, at which point the prover simply sends Y' to the verifier, yielding $n = \log(N)$ rounds for the protocol.

Like the sum-check protocol, we take advantage of the above reduction’s structure to give an efficient HMMP model protocol. Suppose that $N/(2M)$ is an integer. For $i \in \{o, e\}$, partition vectors X_i, Y_i, G_i into vectors $X_{i,j}, Y_{i,j}, G_{i,j}$ each of size $N/(2M)$ for $j \in \{1, \dots, M\}$. Then each audit window $P_{\pi(j)}, \dots, P_{\pi(j+\tau-1)}$ computes and broadcasts the values

$$\begin{aligned} \gamma_{0,j} &= \langle X_{o,j}, Y_{e,j} \rangle & C_{0,j} &= \text{Com}(G_{o,j}, Y_{e,j}) \\ \gamma_{1,j} &= \langle X_{e,j}, Y_{o,j} \rangle & C_{1,j} &= \text{Com}(G_{e,j}, Y_{o,j}) \end{aligned}$$

where $\text{Com}(u, v) := \prod_i u_i^{v_i}$ (i.e., a Pedersen commitment). After all audit windows have broadcast their values, every prover computes

$$\begin{aligned} \gamma_0 &= \sum_{j=1}^M \gamma_{0,j} & C_0 &= g^{\gamma_0} \cdot \prod_{j=1}^M C_{0,j} \\ \gamma_1 &= \sum_{j=1}^M \gamma_{1,j} & C_1 &= g^{\gamma_1} \cdot \prod_{j=1}^M C_{1,j} \end{aligned}$$

and broadcasts (C_0, C_1) to the verifier.

Upon receiving α from the verifier, each prover locally computes the values γ' and C' as specified. The provers then jointly compute the vectors $X', Y',$ and G' as follows. Each audit window $P_{\pi(j)}, \dots, P_{\pi(j+\tau-1)}$

computes and broadcasts

$$X'_j = \alpha^{-1}X_{e,j} + \alpha X_{o,j} \quad Y'_j = \alpha Y_{e,j} + \alpha^{-1}Y_{o,j} \quad G'_j = (G_{e,j})^{\alpha^{-1}} * (G_{o,j})^\alpha.$$

Once all provers have broadcast their values, each prover obtains X', Y', G' by concatenating X'_j, Y'_j, G'_j in order.

For any round $r \in \{1, \dots, n\}$ such that $2^{n-r-1}/M \geq 1$, every prover performs $O(\tau \cdot (2^{n-r}/M))$ field multiplications and group exponents, and performs $O(\tau \cdot (2^{n-r}/M + M))$ field additions and group multiplications. Additionally, each prover broadcasts $O(\tau \cdot (2^{n-r}/M))$ field and group elements. Moreover, for any round r such that $2^{n-r-1} < M$, we can divide up the remaining computation the same way we did for the sum-check protocol. If r^* is the round such that $2^{n-r^*-1} < M \leq 2^{n-r^*}$, then over all rounds $r = r^*, \dots, n$ each prover performs $O(\tau)$ field and group operations, and a total of $O(\tau \cdot M)$ field and group elements are broadcast. Additionally, the complexity of each prover over all rounds $r = 1, \dots, r^* - 1$ is $O(\tau \cdot (N/M))$ field multiplications and group exponents and $O(\tau \cdot (N/M + n \cdot M))$ field additions and group multiplications. This yields total prover complexity $O(\tau \cdot (N/M))$ field multiplications and group exponents and $O(\tau \cdot (N/M + n \cdot M))$ field additions and group additions. Moreover, each prover broadcasts a total of $O(\tau \cdot n \cdot (N/M))$ field elements. This gives [Proposition 2.5](#).

The HMMP Model Verifier. We note that the HMMP Model verifier V_{hmp} for our polynomial commitment scheme is constructed exactly as given in [Theorem 1.1](#), with one exception. In the BRRS polynomial commitment scheme, the verifier must compute the new vectors \mathbf{g}' and Z' during every recursive round. Instead, we leverage the honest majority assumption and have \mathbf{P}_{hmp} send the final folded vectors $\mathbf{g} \in \mathbb{G}$ and $Z \in \mathbb{F}$ in addition to $Y \in \mathbb{F}$, then V_{hmp} performs the final check that V would perform in [Protocol 8](#). This gives [Corollary 2.6](#).

6 Extensions of the Honest Majority Multi-Prover Model

One can easily re-imagine the HMMP model and interactive arguments in the HMMP model by tweaking any of the moving parts. We discuss some potential modifications below, and leave it as future work to achieve HMMP model constructions with such modifications.

Adaptive Corruption. It is crucial for our construction of \mathbf{P}_{hmp} in [Protocol 1](#) that adversarial corruption is static and occurs before the generation of the permutation π . In our construction, *any* adaptive corruption break both correctness and soundness. Though we believe our static corruption is reasonable to assume in the real world, giving a HMMP model argument that is secure against adaptive corruptions would greatly increase the applicability of this model, as well as give interesting new transformations and protocols in this setting.

Communication Between Provers. For simplicity, we assume an authenticated sender broadcast channel with which the provers use for communication. This allows provers to easily identify who sent which messages and easily perform audits as necessary. One could also change this communication assumption (e.g., secure point-to-point, asynchronous sending, etc.) and define the HMMP model with respect to this communication channel. Of course, this could simplify some aspects of the HMMP model as well as complicate other aspects.

Communication with the Verifier. Again for simplicity, we assume that the provers communicate with the verifier via the same authenticated broadcast channel. This allows the verifier to easily send its message to all provers simultaneously, as well as allowing the prover to identify which provers are sending which messages. The second point above helps a verifier ignore multiple messages sent by the same prover during any round (e.g., a malicious prover trying to flood the verifier with messages).

7 Future Work

We have demonstrated the efficiency gains of collaborative proof generation when privacy is not a concern. Even for small values of M , our HMMP model directly yields concrete efficiency gains; e.g., for $M = 2^3$ and $\tau = 2^2$, our provers experience a roughly $2\times$ decrease in overall computational costs, at the cost of (a modest amount of) communication. For another setting of $M = 2^5$ and $\tau = 2^3$, we obtain a $4\times$ decrease.

Our HMMP model can easily be tweaked by modifying different parts of it. Some potential modifications, left as future work, include modifying the model to handle adaptive corruptions, modifying the communication channels used by the provers, and modifying the communication channels used between the provers and the verifier. Additionally, one can imagine adjusting the adversarial model to trying to protect against adversaries with the goal of forcing other provers to perform computation proportional to a single prover. Our framework is modular and lends itself well to different aspects being tweaked to handle many situations.

8 Acknowledgements

Alexander R. Block was supported in part by NSF CCF #1910659. Christina Garman was supported in part by Intelligence Advanced Research Projects Activity (IARPA) contract #2019-19020700004.

References

- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *IEEE Symposium on Security and Privacy*, 2018.
- [BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 1988.
- [BCC⁺16] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In *Advances in Cryptology – EUROCRYPT*, 2016.
- [BCCT12] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *Innovations in Theoretical Computer Science Conference*, 2012.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for snarks and proof-carrying data. In *ACM Symposium on Theory of Computing*, 2013.
- [BCG⁺20] Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. Zexe: Enabling decentralized private computation. In *IEEE Symposium on Security and Privacy*, 2020.
- [BCHO22] Jonathan Bootle, Alessandro Chiesa, Yuncong Hu, and Michele Orrù. Gemini: elastic snarks for diverse environments. In *Advances in Cryptology – EUROCRYPT (to appear)*, 2022.
- [BCI⁺13] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Omer Paneth, and Rafail Ostrovsky. Succinct non-interactive arguments via linear interactive proofs. In *Theory of Cryptography*, 2013.
- [BFS20] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent snarks from dark compilers. In *Advances in Cryptology – EUROCRYPT*, 2020.
- [BG92] Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In *Advances in Cryptology – CRYPTO*, 1992.

- [BHR⁺20] Alexander R. Block, Justin Holmgren, Alon Rosen, Ron D. Rothblum, and Pratik Soni. Public-coin zero-knowledge arguments with (almost) minimal time and space overheads. In *Theory of Cryptography*, 2020.
- [BHR⁺21] Alexander R. Block, Justin Holmgren, Alon Rosen, Ron D. Rothblum, and Pratik Soni. Time- and space-efficient arguments from groups of unknown order. In *Advances in Cryptology – CRYPTO*, 2021.
- [BOGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *ACM Symposium on Theory of Computing*, 1988.
- [BSBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast Reed-Solomon Interactive Oracle Proofs of Proximity. In *International Colloquium on Automata, Languages, and Programming*, 2018.
- [BSCG⁺14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *IEEE Symposium on Security and Privacy*, 2014.
- [BSCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In *Theory of Cryptography*, 2016.
- [BSGKS20] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: Sampling Outside the Box Improves Soundness. In *Theoretical Computer Science Conference*, 2020.
- [CHM⁺20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In *Advances in Cryptology – EUROCRYPT*, 2020.
- [CMT12] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In *Innovations in Theoretical Computer Science Conference*, 2012.
- [DFH12] Ivan Damgård, Sebastian Faust, and Carmit Hazay. Secure two-party computation with low communication. In *Theory of Cryptography*, 2012.
- [DPP⁺22] Pankaj Dayama, Arpita Patra, Protik Paul, Nitin Singh, and Dhinakaran Vinayagamurthy. How to prove any NP statement jointly? efficient distributed-prover zero-knowledge protocols. Privacy Enhancing Technologies Symposium, 2022.
- [Eth] Ethereum Foundation. Zk-rollups. <https://docs.ethhub.io/ethereum-roadmap/layer-2-scaling/zk-rollups/>.
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology – CRYPTO*, 1987.
- [GGM16] Christina Garman, Matthew Green, and Ian Miers. Accountable privacy for decentralized anonymous payments. In *International Conference on Financial Cryptography and Data Security*, 2016.
- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *Advances in Cryptology – EUROCRYPT*, 2013.
- [GI08] Jens Groth and Yuval Ishai. Sub-linear zero-knowledge argument for correctness of a shuffle. In *Advances in Cryptology – EUROCRYPT*, 2008.

- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 1989.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *ACM Symposium on Theory of Computing*, 1987.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *ACM Symposium on Theory of Computing*, 2011.
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019.
- [IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *ACM Symposium on Theory of Computing*, 2007.
- [Kil91] Joe Kilian. A general completeness theorem for two party games. In *ACM Symposium on Theory of Computing*, 1991.
- [KMS⁺16] Ahmed Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *IEEE Symposium on Security and Privacy*, 2016.
- [KPV19] Assimakis Kattis, Konstantin Panarin, and Alexander Vlasov. Redshift: Transparent snarks from list polynomial commitment iops. Cryptology ePrint Archive, Report 2019/1400, 2019.
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *Advances in Cryptology - ASIACRYPT*, 2010.
- [Lee21] Jonathan Lee. Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. In *Theory of Cryptography*, 2021.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 1992.
- [Lin03] Y. Lindell. Parallel coin-tossing and constant-round secure two-party computation. *Journal of Cryptology*, 2003.
- [LNS20] Jonathan Lee, Kirill Nikitin, and Srinath Setty. Replicated state machines without replicated execution. In *IEEE Symposium on Security and Privacy*, 2020.
- [LSP82] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 1982.
- [MBKM19] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge snarks from linear-size universal and updatable structured reference strings. In *ACM SIGSAC Conference on Computer and Communications Security*, 2019.
- [Mic00] Silvio Micali. Computationally sound proofs. *SIAM J. Comput.*, 2000.
- [NPY20] Moni Naor, Merav Parter, and Eylon Yogev. *The Power of Distributed Verifiers in Interactive Proofs*. 2020.
- [OB22] Alex Ozdemir and Dan Boneh. Experimenting with collaborative zk-snarks: Zero-knowledge proofs for distributed secrets. USENIX Security Symposium, 2022.
- [Ped91] Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Advances in Cryptology — CRYPTO*, 1991.

- [PSL80] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 1980.
- [PST13] Charalampos Papamanthou, Elaine Shi, and Roberto Tamassia. Signatures of correct computation. In *Theory of Cryptography*, 2013.
- [Rab83] Michael O. Rabin. Transaction protection by beacons. *Journal of Computer and System Sciences*, 1983.
- [RBO89] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *ACM Symposium on Theory of Computing*, 1989.
- [RRR16] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In *ACM Symposium on Theory of Computing*, 2016.
- [SG11] Aviad Rubinfeld, Shafi Goldwasser, Huijia Lin. Delegation of computation without rejection problem from designated verifier cs-proofs. Cryptology ePrint Archive, Report 2011/456, 2011.
- [SJK⁺17] Ewa Syta, Philipp Jovanovic, Eleftherios Kogias, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Michael J. Fischer, and Bryan Ford. Scalable bias-resistant distributed randomness. In *IEEE Symposium on Security and Privacy*, 2017.
- [SL20] Srinath Setty and Jonathan Lee. Quarks: Quadruple-efficient transparent zkSNARKs. Cryptology ePrint Archive, Report 2020/1275, 2020.
- [Tha13] Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In *Advances in Cryptology – CRYPTO*, 2013.
- [Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *Theory of Cryptography*, 2008.
- [WBT⁺17] Primal Wijesekera, Arjun Baokar, Lynn Tsai, Joel Reardon, Serge Egelman, David Wagner, and Konstantin Beznosov. The feasibility of dynamically granted permissions: Aligning mobile privacy with user preferences. In *IEEE Symposium on Security and Privacy*, 2017.
- [WTS⁺18] Riad S. Wahby, Ioanna Tzialla, Abhi Shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *IEEE Symposium on Security and Privacy*, 2018.
- [WZC⁺18] Howard Wu, Wenting Zheng, Alessandro Chiesa, Raluca Ada Popa, and Ion Stoica. DIZK: A distributed zero knowledge proof system. In *USENIX Security Symposium*, 2018.
- [XZZ⁺19] Tiacheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. Libra: Succinct zero-knowledge proofs with optimal prover computation. In *Advances in Cryptology – CRYPTO*, 2019.
- [Yao82] Andrew C. Yao. Protocols for secure computations. In *Symposium on Foundations of Computer Science*, 1982.
- [ZXZS20] Jiaheng Zhang, Tiancheng Xie, Yupeng Zhang, and Dawn Song. Transparent polynomial delegation and its applications to zero knowledge proof. In *IEEE Symposium on Security and Privacy*, 2020.

Protocol 7: HMMP Model Sum-Check Protocol

Parties : M provers \mathbf{P} and verifier V .

Prover Input : A multi-variate polynomial $g \in \mathbb{F}[X_1, \dots, X_n]$ of individual degree at most d , a subset $H \subset \mathbb{F}$, and a value $y \in \mathbb{F}$.

Verifier Input : Individual degree upper bound d , a subset $H \subset \mathbb{F}$, and value $y \in \mathbb{F}$.

Verifier Output : Accept or Reject.

Oracle : The verifier V has oracle access to evaluations of the polynomial g .

- 1 \mathbf{P} initializes $r = \emptyset$ and $\text{count} = 1$.
- 2 V initializes $y' = y$.
- 3 **foreach** $j \in [n]$ **do**
- 4 **if** $|H|^{n-j} \geq M$ **then**
- 5 \mathbf{P} partitions H^{n-j} into M subsets H_1, \dots, H_M each of size H^{n-j}/M .
- 6 **foreach** $i \in [M]$ **do**
- 7 **foreach** $k \in [i, i + \tau - 1]_M$ **do**
- 8 $P_{\pi(k)}$ computes and broadcasts $g_{j,i,k}(X_j) := \sum_{x \in H_k^{n-j}} g(r, X_j, x)$.
- 9 **else**
- 10 **foreach** $i = 1, \dots, |H|^{n-j}$ **do**
- 11 **foreach** $k \in [\text{count}, \text{count} + \tau - 1]_M$ **do**
- 12 $P_{\pi(k)}$ computes and broadcasts $g_{j,i,k}(X_j) := g(r, X_j, h_i)$, where h_i is the i^{th} element of H^{n-j} .
- 13 \mathbf{P} updates $\text{count} = \text{count} + 1 \pmod{M}$.
- 14 **if not all** $g_{j,i,k}$ **are equal** **then**
- 15 All provers set $g_{j,i}(X_j) := \text{Audit}(i, \{g_{j,i,k}\}_k)$.
- 16 **else** All provers set $g_{j,i}(X_j) := g_{j,i,i}(X_j)$.
- 17 **foreach** $i \in [M]$ **do** P_i computes $g_j^{(i)}(X_j) := \sum_k g_{j,i,k}(X_j)$ and broadcasts it to V .
- 18 V sets $g_j(X_j) = \text{majority}(g_j^{(1)}(X_j), \dots, g_j^{(M)}(X_j))$.
- 19 V checks that $g_j(X_j)$ has degree at most d and $y' = \sum_{x \in H} g_j(x)$, outputting reject if either do not hold.
- 20 V samples $r_j \xleftarrow{\$} \mathbb{F}$.
- 21 **if** $j < n$ **then**
- 22 V sets $y' = g_j(r_j)$, and broadcasts r_j to \mathbf{P} .
- 23 \mathbf{P} updates $r = (r, r_j)$.
- 24 **else** V queries the oracle and checks $g_n(r_n) = g(r_1, \dots, r_n)$, aborting with Reject if not.
- 25 V outputs Accept.

Protocol 8: Eval protocol

Parties : Prover P and verifier V .

Prover Input : Public parameters pp , commitment C , $\zeta \in \mathbb{F}^n$, $\gamma \in \mathbb{F}$, and $Y \in \mathbb{F}^N$

Verifier Input : Public parameters pp , commitment C , $\zeta \in \mathbb{F}^n$, and $\gamma \in \mathbb{F}$

Verifier Output : Accept or Reject.

1 V samples and sends $g_V \xleftarrow{\$} \mathbb{G}$.

2 P and V define $C_\gamma := C \cdot g_V^\gamma \in \mathbb{G}$.

3 P and V define sequence $Z = (Z_b = \beta(b, \zeta))_{b \in \{0,1\}^n} \in \mathbb{F}^N$.

4 P and V engage in $\text{EvalReduce}(C_\gamma, Z, \gamma, \mathbf{g}, g_V, N; Y)$.

5 **Procedure** $\text{EvalReduce}(C_\gamma, Z, \gamma, \mathbf{g}, g_V, N; Y)$

6 **if** $N = 1$ **then**

7 P sends $Y \in \mathbb{F}$ to V .

8 V outputs Accept if and only if $C_\gamma = \mathbf{g}^Y \cdot g_V^{\langle Y, Z \rangle}$ and $\gamma = \langle Y, Z \rangle$, where $\mathbf{g} \in \mathbb{G}$ and $Z \in \mathbb{F}$; else outputs Reject.

9 Set $n = \log(N)$.

10 P computes γ_0, γ_1 as

$$\gamma_0 = \sum_{b \in \{0,1\}^{n-1}} Y_{b00} \cdot Z_{b01} \qquad \gamma_1 = \sum_{b \in \{0,1\}^{n-1}} Y_{b01} \cdot Z_{b00}.$$

11 P computes and sends C_0, C_1 to V , where

$$C_0 = g_V^{\gamma_0} \cdot \prod_{b \in \{0,1\}^{n-1}} (\mathbf{g}_{b01})^{Y_{b00}} \qquad C_1 = g_V^{\gamma_1} \cdot \prod_{b \in \{0,1\}^{n-1}} (\mathbf{g}_{b00})^{Y_{b01}}.$$

12 V samples and sends $\alpha \xleftarrow{\$} \mathbb{F}$ to P .

13 P computes and sends $\gamma' = \alpha^2 \cdot \gamma_0 + \gamma + \alpha^{-2} \cdot \gamma_1$.

14 P and V both compute

$$\begin{aligned} C'_{\gamma'} &= (C_0)^{\alpha^2} \cdot C_\gamma \cdot (C_1)^{\alpha^{-2}} \\ Z' &= (Z'_b = \alpha^{-1} \cdot Z_{b00} + \alpha \cdot Z_{b01})_{b \in \{0,1\}^{n-1}} \\ \mathbf{g}' &= (\mathbf{g}'_b = (\mathbf{g}_{b00})^{\alpha^{-1}} \cdot (\mathbf{g}_{b01})^\alpha)_{b \in \{0,1\}^{n-1}}. \end{aligned}$$

15 P computes $Y' = (Y'_b = \alpha \cdot Y_{b00} + \alpha^{-1} \cdot Y_{b01})_{b \in \{0,1\}^{n-1}}$.

16 P and V engage in $\text{EvalReduce}(C'_{\gamma'}, Z', \gamma', \mathbf{g}', g_V, N/2; Y')$

Protocol 9: HMMP Model Eval Protocol

Parties : M provers \mathbf{P} and verifier V .
Prover Input : Public parameters \mathbf{pp} , commitment C , $\zeta \in \mathbb{F}^n$, $\gamma \in \mathbb{F}$, and $Y \in \mathbb{F}^N$
Verifier Input : Public parameters \mathbf{pp} , commitment C , $\zeta \in \mathbb{F}^n$, and $\gamma \in \mathbb{F}$.
Verifier Output : Accept or Reject.

- 1 V samples and sends $g_V \xleftarrow{\$} \mathbb{G}$.
- 2 \mathbf{P} and V define $C_\gamma := C \cdot g_V^\gamma \in \mathbb{G}$. \mathbf{P} additionally initializes $\text{count} = 1$.
- 3 \mathbf{P} partitions $\{0, 1\}^n$ into M subsets S_1, \dots, S_M each of size N/M .
- 4 **foreach** $i \in [M]$ **do**
- 5 **foreach** $j \in [i, i + \tau - 1]_M$ **do**
- 6 $P_{\pi(j)}$ computes and broadcasts $Z_{i,j} = (\beta(b, \gamma))_{b \in S_i} \in \mathbb{F}^{N/M}$.
- 7 **if not all** $Z_{i,j}$ **are equal then**
- 8 All provers set $Z_i := \text{Audit}(i, \{Z_{i,j}\}_j)$.
- 9 **else** All provers set $Z_i := Z_{i,i}$.
- 10 **while** $N > 1$ **do**
- 11 Set $n = \log(N)$.
- 12 **if** $N/2 \geq M$ **then**
- 13 Run [Protocol 10](#) with inputs $(N/2, n, Y, Z, \mathbf{g})$ and obtain $\{\gamma_{0,i}\}_i$, $\{\gamma_{1,i}\}_i$, $\{C_{0,i}\}_i$, and $\{C_{1,i}\}_i$.
- 14 **else**
- 15 Run [Protocol 11](#) with inputs $(N/2, Y, Z, \mathbf{g}, \text{count})$ and obtain $\{\gamma_{0,i}\}_i$, $\{\gamma_{1,i}\}_i$, $\{C_{0,i}\}_i$, and $\{C_{1,i}\}_i$.
- 16 **foreach** $i \in [M]$ **do**
- 17 P_i computes and broadcasts $\gamma_0^{(i)} := \sum_i \gamma_{0,i}$, $\gamma_1^{(i)} := \sum_i \gamma_{1,i}$, $C_0^{(i)} := g_V^{\gamma_0} \cdot \prod_i C_{0,i}$, and $C_1^{(i)} := g_V^{\gamma_1} \cdot \prod_i C_{0,i}$ to V .
- 18 V samples and broadcasts $\alpha \xleftarrow{\$} \mathbb{F}$ to \mathbf{P}
- 19 **foreach** $i \in [M]$ **do**
- 20 P_i computes and broadcasts $\gamma' = \alpha^2 \cdot \gamma_0 + \gamma + \alpha^{-2} \cdot \gamma_1$.
- 21 V and all provers locally compute $C_\gamma := (C_0)^{\alpha^2} \cdot C_\gamma \cdot (C_1)^{\alpha^{-2}}$.
- 22 **if** $N/2 \geq M$ **then**
- 23 Run [Protocol 12](#) with inputs $(N/2, n, Y, Z, \mathbf{g}, \alpha)$ and obtain $Y, Z \in \mathbb{F}^{N/2}$ and $\mathbf{g} \in \mathbb{G}^{N/2}$.
- 24 **else**
- 25 Run [Protocol 13](#) with inputs $(N/2, Y, Z, \mathbf{g}, \alpha, \text{count})$ and obtain $Y, Z \in \mathbb{F}^{N/2}$ and $\mathbf{g} \in \mathbb{G}^{N/2}$.
- 26 Update $N \leftarrow N/2$.
- 27 Each P_i broadcasts $Y \in \mathbb{F}$, $\mathbf{g} \in \mathbb{G}$, and $Z \in \mathbb{F}$ to V .
- 28 V outputs Accept if and only if $C_\gamma = \mathbf{g}^Y \cdot g_V^{(Y \cdot Z)}$; else outputs Reject.
- 29 \mathbf{P} and V both compute

$$\begin{aligned} C'_{\gamma'} &= (C_0)^{\alpha^2} \cdot C_\gamma \cdot (C_1)^{\alpha^{-2}} \\ Z' &= (Z'_b = \alpha^{-1} \cdot Z_{b00} + \alpha \cdot Z_{b01})_{b \in \{0,1\}^{n-1}} \\ \mathbf{g}' &= (\mathbf{g}'_b = (\mathbf{g}_{b00})^{\alpha^{-1}} \cdot (\mathbf{g}_{b01})^\alpha)_{b \in \{0,1\}^{n-1}} \end{aligned}$$

30 \mathbf{P} computes $Y' = (Y'_b = \alpha \cdot Y_{b00} + \alpha^{-1} \cdot Y_{b01})_{b \in \{0,1\}^{n-1}}$.

Protocol 10: Polycom Compute Partial Evals and Commitments 1

Input : Integers $N/2$ and n , $Y, Z \in \mathbb{F}^N$, and $\mathbf{g} \in \mathbb{G}^N$
Output : Values $\{\gamma_{0,i}\}_i$, $\{\gamma_{1,i}\}_i$, $\{C_{0,i}\}_i$, and $\{C_{1,i}\}_i$

- 1 P partitions $\{0, 1\}^{n-1}$ into M subsets B_1, \dots, B_M each of size $N/(2M)$.
- 2 **foreach** $i \in [M]$ **do**
- 3 **foreach** $j \in [i, i + \tau - 1]_M$ **do**
- 4 | $P_{\pi(j)}$ computes and broadcasts
- $$\gamma_{0,i,j} = \sum_{b \in B_i} Y_{b \circ 0} \cdot Z_{b \circ 1} \qquad C_{0,i,j} = \prod_{b \in B_i} (\mathbf{g}_{b \circ 1})^{Y_{b \circ 0}}$$
- $$\gamma_{1,i,j} = \sum_{b \in B_i} Y_{b \circ 1} \cdot Z_{b \circ 0} \qquad C_{1,i,j} = \prod_{b \in B_i} (\mathbf{g}_{b \circ 0})^{Y_{b \circ 1}}.$$
- 5 **if not all** $\{\gamma_{0,i,j}\}_j$ **are equal then** All provers set $\gamma_{0,i} = \text{Audit}(i, \{\gamma_{0,i,j}\}_j)$.
- 6 **else** All provers set $\gamma_{0,i} = \gamma_{0,i,i}$.
- 7 **if not all** $\{\gamma_{1,i,j}\}_j$ **are equal then** All provers set $\gamma_{1,i} = \text{Audit}(i, \{\gamma_{1,i,j}\}_j)$.
- 8 **else** All provers set $\gamma_{1,i} = \gamma_{1,i,i}$.
- 9 **if not all** $\{C_{0,i,j}\}_j$ **are equal then** All provers set $C_{0,i} = \text{Audit}(i, \{C_{0,i,j}\}_j)$.
- 10 **else** All provers set $C_{0,i} = C_{0,i,i}$.
- 11 **if not all** $\{C_{1,i,j}\}_j$ **are equal then** All provers set $C_{1,i} = \text{Audit}(i, \{C_{1,i,j}\}_j)$.
- 12 **else** All provers set $C_{1,i} = C_{1,i,i}$.
- 13 **return** $\{\gamma_{0,i}\}_i$, $\{\gamma_{1,i}\}_i$, $\{C_{0,i}\}_i$, and $\{C_{1,i}\}_i$.

Protocol 11: Polycom Compute Partial Evals and Commitments 2

Input : Integer $N/2$, $Y, Z \in \mathbb{F}^N$, $\mathbf{g} \in \mathbb{G}^N$, and integer count.
Output : Values $\{\gamma_{0,i}\}_i$, $\{\gamma_{1,i}\}_i$, $\{C_{0,i}\}_i$, and $\{C_{1,i}\}_i$

- 1 **foreach** $i = 0, \dots, N/2 - 1$ **do**
- 2 **foreach** $j \in [\text{count}, \text{count} + \tau - 1]_M$ **do**
- 3 | $P_{\pi(j)}$ computes and broadcasts $\gamma_{0,i,j} = Y_{i \circ 0} \cdot Z_{i \circ 1}$, $\gamma_{1,i,j} = Y_{i \circ 1} \cdot Z_{i \circ 0}$, $C_{0,i,j} = (\mathbf{g}_{i \circ 1})^{Y_{i \circ 0}}$,
- $C_{1,i,j} = (\mathbf{g}_{i \circ 0})^{Y_{i \circ 1}}$.
- 4 **if not all** $\{\gamma_{0,i,j}\}_j$ **are equal then** All provers set $\gamma_{0,i} = \text{Audit}(i, \{\gamma_{0,i,j}\}_j)$.
- 5 **else** All provers set $\gamma_{0,i} = \gamma_{0,i,i}$.
- 6 **if not all** $\{\gamma_{1,i,j}\}_j$ **are equal then** All provers set $\gamma_{1,i} = \text{Audit}(i, \{\gamma_{1,i,j}\}_j)$.
- 7 **else** All provers set $\gamma_{1,i} = \gamma_{1,i,i}$.
- 8 **if not all** $\{C_{0,i,j}\}_j$ **are equal then** All provers set $C_{0,i} = \text{Audit}(i, \{C_{0,i,j}\}_j)$.
- 9 **else** All provers set $C_{0,i} = C_{0,i,i}$.
- 10 **if not all** $\{C_{1,i,j}\}_j$ **are equal then** All provers set $C_{1,i} = \text{Audit}(i, \{C_{1,i,j}\}_j)$.
- 11 **else** All provers set $C_{1,i} = C_{1,i,i}$.
- 12 count = count + 1 (mod M).
- 13 **return** $\{\gamma_{0,i}\}_i$, $\{\gamma_{1,i}\}_i$, $\{C_{0,i}\}_i$, and $\{C_{1,i}\}_i$.

Protocol 12: Polycom Compute Partial Sequences 1

Input : Integers $N/2$ and n , $Y, Z \in \mathbb{F}^N$, $\mathbf{g} \in \mathbb{G}^N$, and $\alpha \in \mathbb{F}$.

Output : Sequences $Y, Z \in \mathbb{F}^{N/2}$, $\mathbf{g} \in \mathbb{G}^{N/2}$.

```
1 P partitions  $\{0, 1\}^{n-1}$  into  $M$  subsets  $B_1, \dots, B_M$  each of size  $N/(2M)$ .
2 foreach  $i \in [M]$  do
3   foreach  $j \in [i, i + \tau - 1]_M$  do
4      $P_{\pi(j)}$  computes and broadcasts
        
$$Z'_{i,j} = (\alpha^{-1} \cdot Z_{b_{o0}} + \alpha \cdot Z_{b_{o1}})_{b \in B_i}$$

        
$$Y'_{i,j} = (\alpha \cdot Y_{b_{o0}} + \alpha^{-1} \cdot Y_{b_{o1}})_{b \in B_i}$$

        
$$\mathbf{g}'_{i,j} = ((\mathbf{g}_{b_{o0}})^{\alpha^{-1}} \cdot (\mathbf{g}_{b_{o1}})^\alpha)_{b \in B_i}$$

5   if not all  $\{Z'_{i,j}\}_j$  are equal then All provers set  $Z'_i = \text{Audit}(i, \{Z'_{i,j}\}_j)$ .
6   else All provers set  $Z'_i = Z'_{i,i}$ .
7   if not all  $\{Y'_{i,j}\}_j$  are equal then All provers set  $Y'_i = \text{Audit}(i, \{Y'_{i,j}\}_j)$ .
8   else All provers set  $Y'_i = Y'_{i,i}$ .
9   if not all  $\{\mathbf{g}'_{i,j}\}_j$  are equal then All provers set  $\mathbf{g}'_i = \text{Audit}(i, \{\mathbf{g}'_{i,j}\}_j)$ .
10  else All provers set  $\mathbf{g}'_i = \mathbf{g}'_{i,i}$ .
11 return  $Y = (Y'_1, \dots, Y'_M)$ ,  $Z = (Z'_1, \dots, Z'_M)$ , and  $\mathbf{g} = (\mathbf{g}'_1, \dots, \mathbf{g}'_M)$ .
```

Protocol 13: Polycom Compute Partial Sequences 2

Input : Integers $N/2$ and n , $Y, Z \in \mathbb{F}^N$, $\mathbf{g} \in \mathbb{G}^N$, $\alpha \in \mathbb{F}$, and integer count.

Output : Sequences $Y, Z \in \mathbb{F}^{N/2}$, $\mathbf{g} \in \mathbb{G}^{N/2}$.

```
1 foreach  $i = 0, \dots, N/2 - 1$  do
2   foreach  $j \in [\text{count}, \text{count} + \tau - 1]_M$  do
3      $P_{\pi(j)}$  computes and broadcasts  $Z'_{i,j} = \alpha^{-1} \cdot Z_{i_{o0}} + \alpha \cdot Z_{i_{o1}}$ ,  $Y'_{i,j} = \alpha \cdot Y_{i_{o0}} + \alpha^{-1} \cdot Y_{i_{o1}}$ , and
        
$$\mathbf{g}'_{i,j} = (\mathbf{g}_{i_{o0}})^{\alpha^{-1}} \cdot (\mathbf{g}_{i_{o1}})^\alpha.$$

4   if not all  $\{Z'_{i,j}\}_j$  are equal then All provers set  $Z'_i = \text{Audit}(i, \{Z'_{i,j}\}_j)$ .
5   else All provers set  $Z'_i = Z'_{i,i}$ .
6   if not all  $\{Y'_{i,j}\}_j$  are equal then All provers set  $Y'_i = \text{Audit}(i, \{Y'_{i,j}\}_j)$ .
7   else All provers set  $Y'_i = Y'_{i,i}$ .
8   if not all  $\{\mathbf{g}'_{i,j}\}_j$  are equal then All provers set  $\mathbf{g}'_i = \text{Audit}(i, \{\mathbf{g}'_{i,j}\}_j)$ .
9   else All provers set  $\mathbf{g}'_i = \mathbf{g}'_{i,i}$ .
10  count = count + 1(mod  $M$ ).
11 return  $Y = (Y'_1, \dots, Y'_M)$ ,  $Z = (Z'_1, \dots, Z'_M)$ , and  $\mathbf{g} = (\mathbf{g}'_1, \dots, \mathbf{g}'_M)$ .
```
