



How to Make Rational Arguments Practical and Extractable

Matteo Campanelli¹ , Chaya Ganesh²  and Rosario Gennaro³ 

¹ Protocol Labs, United States

² Indian Institute of Science, India

³ The City University of New York, United States

Abstract. We investigate proof systems where security holds against *rational* parties instead of malicious ones. Our starting point is the notion of *rational arguments*, a variant of rational proofs (Azar and Micali, STOC 2012) where security holds against rational adversaries that are also *computationally bounded*.

Rational arguments are an interesting primitive because they generally allow for very efficient protocols, and in particular sublinear verification (i.e. where the Verifier does not have to read the entire input). In this paper, we aim at narrowing the gap between the literature on rational schemes and real world applications. Our contribution is two-fold.

We provide the first construction of rational arguments for the class of polynomial computations that is practical (i.e., it can be applied to real-world computations on reasonably common hardware) and with logarithmic communication. Techniques-wise, we obtain this result through a compiler from information-theoretic protocols and rational proofs for polynomial evaluation. The latter could be of independent interest. As a second contribution, we propose a new notion of *extractability* for rational arguments. Through this notion, we can obtain arguments where *knowledge* of a witness is incentivized (rather than incentivizing mere soundness). We show how our aforementioned compiler can also be applied to obtain efficient extractable rational arguments for NP.

1 Introduction

We study the problem of verifying that a statement is correct without having to re-execute the computation associated with it. In cryptographic literature a common primitive applied to this problem is a (cryptographic) *argument*. For deterministic (polynomial-time) computations, through an argument, we can convince a *weak* verifier, holding data D as input that $f(D) = y$ for some computation f . We can apply arguments to non-deterministic computations as well: a party holding a witness w can convince the verifier that it *knows* a witness such that $R(D, w) = 1$ for some NP relation R . This last type of argument is also called an *argument of knowledge* (or *extractable* argument, since their “knowledge” is technically defined in terms of the existence of an extractor).

Arguments have become more and more practical. There is a long line of work spanning now more than a decade that has produced schemes that are extremely efficient in several dimensions (e.g. [BCG⁺13, WTs⁺18, CFQ19, MBKM19, CHM⁺20, CFF⁺21, ABC⁺22, LSTW21, KPV22] for a partial list). The most notable of their improvements is *fast verification*. They are able to achieve this by obtaining a proof size substantially smaller than the witness (resp. the execution trace) of the relation (resp. computation).

E-mail: matteo.campanelli@gmail.com (Matteo Campanelli), chaya@iisc.ac.in (Chaya Ganesh), rosario@ccny.cuny.edu (Rosario Gennaro)



At the same time they feature very performant proving algorithms. From now on we will refer to the arguments from this line of work as *very succinct arguments* for short.

Our goal: in this work we are interested in improving on very succinct arguments by keeping the succinctness and the practicality of the prover and making the verifier *even faster*. In particular, we will aim for obtaining a *sublinear verifier* (the verifier does not even need to read the whole public input). Next, we shall motivate this problem and then describe the model in which we aim at achieving this.

MOTIVATING SUBLINEAR VERIFICATION.¹

- *Motivation 1: overhead in verification from public input preprocessing.* In very succinct arguments, the efficiency of the verifier is often a consequence of how simple checking the proof is. Let us consider for example Groth16, currently one of the schemes with the “most minimalistic” proof and verifier [Gro16]. Consider a data string D (the public input) on which we want to verify property P . No matter the size of D or the complexity of the property P , a Groth16 proof consists of three group elements (A, B, C) and the verifier checks an equation of the form

$$e(A, B) \stackrel{?}{=} Y_D \cdot e(C, Z)$$

where Z is a constant in the SRS (structured reference string) and Y_D is computed from the public input D . The verification check is as simple as it can get—at its core it involves two pairings only. However, for large public inputs, the bottleneck is from computing the element Y_D , which is defined as $\prod_i g_i^{D_i}$ for public generators g_i . For data as small as 2^{20} this computation can already be **three orders of magnitude** more expensive than the pairings themselves².

- *Motivation 2: lightweight access to data.* Processing a large public input may not be a bottleneck in all succinct arguments as it is in Groth16. Nonetheless, sublinear verification may still be advantageous in settings where it is impractical or strongly undesirable to access the whole public input. As a first example, a public input for a proof may be stored on the cloud, a fairly common practice for data in general. Downloading the whole input to verify a proof may be impractical when only a slow network is available (large parts of the world do not have access to fast internet). As another example, consider *light clients* in a blockchain. These are clients who do not intend to download the whole history of the chain but are still supposed to retrieve information about it. In order to verify statements about the chain they could instead access only parts of it rather than the chain in its entirety.

We have so far motivated designing proof schemes with a sublinear verifier. Nevertheless, it is apparent that achieving such schemes within the *standard* model of security of interactive proofs/arguments is infeasible³. Sublinear verification has been achieved in only two models. One is that of *proofs of proximity* [RVW13] where the prover only certifies a close approximation to the exact value of the computation. One drawback of this model, is that the efficiency of the protocol requires trading against correctness: this is no longer guaranteed to hold for instances that are only close to being in the language. For this reason and because we believe in the importance of modeling incentives explicitly, we will instead use the *rational* model of security described below.

¹See also further discussion on sublinear verifier in Section 1.3.

²These estimates refer to the commonly used curve BLS12-381 [ECK+23].

³Take the simple relation for parity where given a string $x \in \{0, 1\}^n$ we say it has parity 1 iff $\sum_i x_i \equiv 1 \pmod{2}$. Intuitively, we cannot convince a verifier of this without them reading the entire string. Why? Because by completeness of the proof system a string x_1 with parity 1 will have a high probability of being accepted by the verifier. But—and here is the absurd—so must be the accepting probability of an input x_0 identical to x_1 except in one position i where we flipped $x_1[i]$: since we are querying only a fraction of the input bits, the queries from x_0 and x_1 will be similarly distributed.

SECURITY AGAINST MALICIOUS RATIONAL PROVERS. The standard notion of security in interactive proofs considers a malicious adversary who should not be able to convince the verifier of a false statement. In [AM12] Azar and Micali propose a variant of interactive proof where the prover is not considered honest or malicious but simply *rational*—i.e., choosing the strategy that will maximize their reward. In a *rational* security model, a verifier may *in principle* accept proofs of false statements. However, at the same time, a “rationally secure” scheme is designed so that it would be irrational for a prover to do so.

This model is applicable, for example, in settings where the prover has incentives to act honestly. These settings include cloud computing—where a server executes a computation and is going to be rewarded for its services—and blockchains—where it is easy to embed incentives in different layers of the system.

Protocols within the rational framework tend to have some intriguing features, including: simplicity, efficiency and the (already mentioned) possibility of a verifier which needs to query only a sublinear number of input positions. As an example, the protocol in [CG15] obtains a protocol for NC^1 with a verifier running in logarithmic time and querying a constant number of input bits; the interaction transcript has logarithmic size. The prover has essentially no overhead on top of performing the computation (it just needs to send certain sub-portions of the circuit evaluation).

FROM PROOFS TO ARGUMENTS. Not only rational proofs can give us light-weight verification protocols, but they may achieve all that without requiring any cryptographic assumption—the aforementioned protocol in [CG15] works even if it turns out that $\text{P} = \text{NP}$. But by introducing assumptions we can achieve even more: a line of work starting in [GHRV14] and [GHRV16] explores the additional *benefits of applying cryptographic assumptions* to verifiable computation *in a rational model* (this will be our focus). For example, in [GHRV14] Guo et al. show how to obtain two-message rational *arguments* (as opposed to *proofs*, which use no cryptographic assumption) for NC^1 with polylogarithmic verification. Later work has improved this result obtaining schemes with two messages for polynomial-time computations [GHRV16]. This last work uses relatively strong assumptions (subexponential FHE) and is arguably not concretely practical.

In this work we extend rational arguments to the realm of concrete practicality considering both deterministic and non-deterministic computations.

1.1 Our Contributions

We continue the study of rational arguments. The scope of our contributions is (i) *practical*—we construct rational arguments for P more efficient than the state-of-the-art; (ii) *foundational*—we advance the theory of rational arguments by putting forth a new notion—*extractable* rational arguments—we build them for NP and argue their usefulness.

Rational Argument for P . We substantially improve the efficiency of rational arguments for polynomial computations. This work’s starting point is the question:

Can we obtain rational arguments with (1) a sublinear verifier; and (2) whose prover is as efficient as some of the best available SNARKs?

We answer this question in the positive and construct the first rational argument for the class of polynomial computations that is practical (i.e., it can be applied to real-world computations on reasonably common hardware) and with logarithmic communication (see also Table 1). Our results provide improvements compared to both the *rational* and *cryptographic* arguments literature.

In comparison to prior rational arguments we obtain the first rational argument for deterministic computation that has a concretely efficient prover, has a logarithmic verifier using only public coin and is provably secure from simple assumptions (see Table 1). The only prior work on rational arguments for P ([GHRV16]) uses expensive primitive under

Table 1: Comparison between this work and previous work on rational arguments for P. The parameter n is the instance size; T is the size of the computation. All asymptotic quantities implicitly contain a multiplicative factor polynomial in λ .

Scheme	Assumptions	Communication	P time	V time	Messages	Concretely efficient?	public coin?
[GHRV16]	subexp FHE	$O(\text{polylog}(n))$	$\text{poly}(T)$	$O(\text{polylog}(n))$	2	✗	✗
This work	SXDH ⁵	$O(\log(n))$	$O(T \log T)$	$O(\log(n))$	$O(\log(n))$	✓	✓

the hood, required a polynomial blow-up of the prover time and had a polylogarithmic verifier. We note that our results require a higher number of messages: logarithmic instead of constant as in [GHRV16]. From a technical standpoint, we obtain our results from a very different blueprint. We show how to leverage the efficient modular constructions of recent SNARKs modeled as Algebraic Holographic Proofs⁴ (or AHP [CHM⁺20]) and exploit techniques for rational arguments of polynomial evaluation, which may be of independent interest. The work in [GHRV16] instead closely follows the techniques from the delegation scheme of Kalai et al. [KRR14]: they construct δ -no-signaling *rational multi-prover proofs* (RMIPs). They then apply sub-exponentially secure Fully Homomorphic Encryption (FHE) scheme to transform no-signaling RMIPs into two-message rational arguments. Compared to other rational arguments for subclasses of P—e.g., the construction for NC circuits in [GHRV16]—ours removes the dependency of communication stemming from the depth of the circuit (logarithmic instead of polylogarithmic for a circuit of polylogarithmic depth).

In comparison to prior cryptographic arguments our constructions have several advantages. First, we are able to achieve a sublinear verifier. Second, our construction concretely improves on the efficiency of the prover by not requiring any cryptographic operations on what are usually called “computation commitments”. A computation commitment is a commitment to a polynomial describing the computation. A significant amount of the prover’s time is used to prove evaluations of these polynomials. This may require a significant number of cryptographic operations (usually consisting of several multi-scalar exponentiations in the size of the computations) and for some proof schemes this may be a source of efficiency bottleneck (see, e.g., Section 7.5 in [CGG⁺23a]). See our instantiations in the full version [CGG23b] for more details. Finally, our constructions substantially simplify the “offline” stage (also called indexing) that is common in several cryptographic arguments (e.g., [CFF⁺21, CHM⁺20]). In this stage, the verifier is required to do a one-time preprocessing of the description of the computation. Intuitively, from the description of a computation C one can derive a key that can be to verify a proof on C efficiently. This process usually involves several multi-scalar exponentiations of a potentially large size (as large as the computation). This cost can be amortized over several proofs. However, it may still be infeasible for very weak devices wanting to verify large computations. Our construction allows the verifier not to have to perform any cryptographic operations on the description of the computation. Instead it just needs to be able to have RAM (or network) access to its encoding and query a few of its points at verification time.

Rational Argument for NP. *We show that rational proofs that are succinct with noticeable reward gap⁶ are impossible for NP.* This completes the landscape of results in succinct proofs for NP strengthening the impossibility results in [GHRV14] and [CG17].

The impossibility of succinct proofs motivates us to study succinct rational *arguments* for NP. *We put forth the notion of rational argument of knowledge, and construct a rational argument of knowledge for NP.* Our new notion models schemes where *knowledge* of a

⁴We note that our results could be easily adapted to the similar framework of PIOP [BFS20] and the more general framework of PHP [CFF⁺21].

⁵This assumption is used when we instantiate our construction with the Dory polynomial commitment [Lee21]. For a formal definition of SXDH, see Definition 1 in [Lee21].

⁶The reward gap bounds the difference between the reward of the honest and dishonest provers.

witness is incentivized (rather than incentivizing mere soundness). Our construction offers several advantages over succinct cryptographic arguments. First, the prover’s cryptographic operations grows only with the size of the witness, and not with the input. This is useful for proving computations over *large public* data. Second, since our construction is essentially the same as our construction for P, we also get a verifier that is sublinear in the public data, and only has oracle access to it.

Rational Proof for Polynomial Evaluation. As a tool of independent interest, we construct a rational proof for polynomial evaluation, which we think of as a rational analog of a cryptographic polynomial commitment scheme. This allows outsourcing polynomial evaluation to the prover where the verifier only has oracle access to the description of the polynomial (either as a vector of coefficients or as tuple of point-evaluation pairs). Thus, this rational proof for polynomial evaluation is a rational analogue of a cryptographic polynomial commitment scheme, where the verifier is sublinear in the degree of the polynomial by virtue of having oracle access to the polynomial (and querying few places). This is in contrast to a verifier being sublinear in the degree of the polynomial since the cryptographic commitment being succinct in the degree. Indeed, we use this rational proof in our construction in lieu of a polynomial commitment scheme, for polynomials that encode public values (description of the function to be computed, public input).

1.2 Technical Overview

Recent constructions of zkSNARKs [CHM⁺20, RZ21, CFF⁺21, GWC19] follow a modular approach where an information-theoretic protocol is constructed in an abstract model like Probabilistically Checkable Proof (PCP), Interactive Oracle Proof (IOP) etc., and then the information-theoretic protocol is compiled into an argument system via a cryptographic compiler. Our construction is modular: starting from an information-theoretic protocol, our compiler uses cryptographic and rational proofs in order to compile into a rational argument.

AHP to SNARK. We use the formalization of Algebraic Holographic Proofs (AHP) from [CHM⁺20]. In an AHP the prover receives as inputs a statement x and a witness w . In each round of interaction with the verifier, it sends oracle polynomials and the verifier responds with a random challenge. Then, in a query phase, the verifier queries the oracle polynomials at an evaluation point. For oracle p and evaluation query z , it obtains $v = p(z)$. Finally, based on the result of these evaluation queries the verifier outputs a bit indicating “accept” or “reject”. An AHP can be turned into an argument system by replacing the oracles and the query phase with a polynomial commitment scheme (PCS) [KZG10]. In the argument, the prover commits to the polynomials obtained from the information-theoretic prover, and then upon receiving an evaluation point z , the prover uses the evaluation proof of the PCS to convince the verifier that the claimed evaluation v is indeed $v = p(z)$ given a commitment to p .

Our main idea. Typically, the oracles sent by an AHP prover consist of the following: witness-carrying polynomials (WCP) (polynomials that encode the witness vector w) and index polynomials or computation commitments (which encode the description of the computation or the index i). The AHP verifier encodes the public input x into a polynomial and makes queries to the WCP and index polynomials. When compiled into an argument, the prover performs cryptographic work for the PCS to commit to WCP and index polynomials. Our first idea is to move the index polynomials from being cryptographically committed with a PCS to being available to the verifier as an “oracle” and *providing* rational proofs about evaluations. While typically, oracle access to the verifier is an “intermediate” model and oracles have to be realized in the real-world, our oracle model of index polynomials/computation commitments is limited to processing the index and storing a certain encoding in memory. This is part of the verifier’s preprocessing (which is anyway

done in cryptographic arguments too), but our oracle modeling is to capture the property that the verifier *does not read the entire preprocessed material*. Cryptographic arguments achieve sublinearity of the verifier in the index i by preprocessing the computation into a short cryptographic commitment, with respect to which the prover gives evaluation proofs of the PCS. We achieve sublinearity of the verifier in the index i by preprocessing the computation into an encoding; the verifier only queries this encoding at a few points. This makes the verifier sublinear in the size of the computation, but we also need to achieve sublinearity in the public input x . Here too, the verifier has oracle access to x and queries it a few points during verification. Our idea to achieve sublinearity in x is as follows. First, we observe that the computation that makes the verifier linear in x in the underlying AHP is in evaluating a polynomial encoding of x at a random point. We outsource this computation to the prover, so the prover provides this evaluation together with a rational proof of correct evaluation. Crucially, this rational proof can be verified given just oracle access to x . We note that x need not be preprocessed in any way, it is assumed to be stored in the cloud and the verifier accesses it by querying at certain positions. A key technical tool towards our construction is a rational proof for polynomial evaluation which we describe next.

Rational proofs for polynomial evaluation. Given a computation described as an arithmetic circuit of size n and depth d , the protocol of [CG15] works as follows. On input x , the prover sends the claimed output y to the verifier, and then they engage in a recursive protocol. The prover sends the two input values y_L, y_R to the output gate g , the verifier checks that $g(y_L, y_R) = y$, and then chooses one of y_L and y_R at random to recurse on. For uniform circuits, the verifier complexity is $O(d)$. We then describe polynomial evaluation as a parallel circuit that has depth that is logarithmic in the degree. Now, invoking the above protocol on this circuit, we obtain a rational proof with $O(\log d)$ verification where the verifier only needs oracle access to the polynomial, for example as a vector of d coefficients.

Defining rational arguments with extractable properties. In classical arguments of knowledge we require that “if the verifier accepts a proof with a reasonable probability, then we are able to (efficiently) extract a valid witness by interacting with the prover”. Translating this notion into the rational setting requires care.

First, it does not seem possible to have a meaningful notion for instances not in the language (this is not the problem in the standard notion). Here is why: Recall that the goal of designing a rational protocol is to *incentivize* parties towards the “right” choice, but for an input $x \notin L$ there is no way of incentivizing them in such a manner. This is because if $x \notin L$ then the only rational strategy of any prover is to just run the protocol (earning a possibly meager reward) even if it doesn’t know the witness (a witness the prover *cannot* possibly know). For this reason a meaningful definition can only make guarantees about instances $x \in L^7$, e.g., discrete logarithm.

Also, the notion should link the capability of the extractor of outputting a witness to the reward of the prover. Intuitively, we would like to guarantee that only a prover who knows the witness should obtain the highest reward. This notion is hard to capture directly. Our definition tries to instead capture this intuition:

A prover not knowing the witness is incentivized to search for it and only then run the protocol.

When is it reasonable to ask for the prover to search for a witness? Whenever the witness is *moderately* hard to find or hard to find for the verifier (but not for the prover). For example, maybe there is a (quasi-polynomial time or high-degree polynomial time) search algorithm that a prover with lots of computational power could run. Or, the witness

⁷The cryptographic literature has precedents for definitions with “one-sided” guarantees, i.e. for either yes or no instances only. Such an example is semantic security in witness encryption [GGSW13].

could be moderately hard to find in a non-complexity theoretic sense (it may require for example downloading and scanning the whole history of a blockchain or the Wikipedia corpus). Naturally, this is meaningful if the cost of the resources required for such provers to search for the witness is compensated by the received reward. See Section 1.3 for further discussion. We provide more details and intuitions on our definition for extractable rational arguments in Section 5.1.

Our final constructions. Our construction works as follows. We look at the polynomial oracles sent by the AHP prover that “encode” the witness, called the witness-carrying polynomials (WCP), and polynomial oracles arising from the indexer (preprocessing), that encode the computation (index). The prover in the compiled argument will treat the encodings of index and witness differently. The argumenting indexer runs the AHP indexer and outputs these index polynomials as oracles for the verifier. Then the argument prover commits to the WCPs sent by the AHP prover using a polynomial commitment scheme (PCS) and sends this commitment to the argument verifier. Now, (i) queries to WCP are answered by giving a PCS evaluation proof; (ii) queries to index polynomials are answered by giving a rational proof; (iii) the verifier’s linear step in the AHP decision algorithm is outsourced to the prover who provides a rational proof of correct evaluation of the encoding of the input. This construction (which is essentially a compiler) works for both P and NP. For NP, this also satisfies our notion of proof of knowledge. We extract WCPs from the knowledge soundness of the PCS and decode the witness from it. We show that this is a valid witness with respect to the oracle index and oracle input by relying on the guarantees of the rational proof. Assuming that PCS is knowledge sound, if extraction fails, then the prover gave an inconsistent answer in the rational proof. Now, if the rational proof for polynomial evaluation has a noticeable reward gap, then, we show that the probability of extraction failure affects the reward gap of the final prover.

1.3 Discussion: Alternative Sublinear Verifiers; Applications

1.3.1 On other approaches for large public inputs

There exist other techniques to let the verifier not read the whole public input in the context of arguments. They can be used as a way to achieve a sublinear verifier. These approaches, however, have some drawbacks which justify the alternative framework proposed in this paper. One folklore approach, for example, is to let the verifier preprocess the public input D and keep only its digest $h = H(D)$, where H is some collision-resistant hash function that is fast to compute, such as SHA256. Now, instead of providing a proof for just the statement “ $f(D) = y$ ” the prover will certify the augmented statement “ $f(D) = y \wedge H(D) = h$ ”.

One problem with this approach is that it substantially increases the proving time because hash functions commonly tend to be expensive when used proof systems (see, e.g., discussion in [WYX⁺21]). This problem can be mitigated using *Hash-and-Prove* schemes [FFG⁺16] but this requires strong assumptions on the hash function. The latter also needs to be “algebraic”, hence orders-of-magnitude slower to compute for the verifier.

Another limitation of the digest-based approach arises in settings where the data are dynamic. If the data change, then the digest h will change as well. Also, depending on the hash function used, the digest will plausibly need to be recomputed from scratch after each update. Rational arguments can arguably provide a simpler approach since they do not require processing again the input after an update.

1.3.2 Application scenarios for the rational model

Applications for rational arguments for deterministic computations:

- **Outsourced computations (cloud and volunteer computing):** In cloud computing businesses buy computing time from a service, rather than maintain their

own computing resources. In volunteer computing systems (such as SETI@Home or Folding@Home [KWA⁺01, BEJ⁺09]), where the rewards are non-fungible “points”. Rational arguments are suitable to both of these settings since the computing parties have an incentive through the payment (resp. reward points) they earn.

- **Smart contract execution:** Smart contracts are programs running on a blockchain⁸. Since it is costly to run them on chain, a common desideratum on contracts is for them to be as simple as possible. The sublinearity of the verification algorithm in rational arguments is an attractive feature because it may allow a smart contract not to even read its whole input to perform an action. As an example, consider a contract that releases some funds to a user if they provide y such that $y = f(x)$ where x is some value stored in the contract. The contract could be, e.g., enforcing the “incentives for an outsourced computation f ” as outlined above or representing a transition function of its own state through f (in this case the reward is tantamount to some form of mining). A party could now provide y to the contract together with a proof π ⁹. The contract would then query the required positions of x (a sublinear amount) and then release the funds depending on the reward function of the rational argument.

Applications for rational arguments of knowledge: in general, rational arguments of knowledge are a natural primitive for the following problem: party A is interested in rewarding someone (party B) if they know data with a specific feature. Party A does not necessarily need to see the data, it just need to be persuaded that they are “extractable” from someone. Some examples are: a foundation willing to reward someone for finding a proof of the Riemann Hypothesis, or private individuals willing to reward anyone storing specific parts of Wikipedia or the Internet Archive on a decentralized network for storage (see Section 7.2 in [CFK22] for more discussion of these and other examples).

We also provide one more application for extractable rational arguments that is specific to blockchains, in particular for the problem of *data availability*. In blockchains like Ethereum not all nodes have access to the full data of the chain. Nodes that are not expected to store the full data, *light nodes*, may verify that others do store them by downloading very small random chunks of such data (data availability sampling). The more assurance we want to have that someone is storing those data, the more light nodes need to query overall. Rational arguments of knowledge may provide a different approach to the problem where full nodes are provided incentives for storing the full data through a rational argument for the statement “I know (specific positions) of data D that is the opening of this public Merkle Tree root”.

Understanding the advantages the solutions above in the applications we mentioned (as well as others) may require a game-theoretic analysis of the specific setting, which is out of the scope of this paper. In the full version [CGG23b](Appendix C), however, we provide some general observations for when this approach may be viable, its limitations and heuristic mitigations.

1.4 Related Work

Rational Proofs and Arguments. Rational proofs, introduced by Azar and Micali [AM13] are a framework of interactive proofs where the prover is relaxed to be rational rather than malicious. In a rational proof, the prover is incentivized: the verifier pays the prover according to the quality of the result provided, and this reward is set up so that it is irrational for the prover to return the result of the computation incorrectly. Azar and

⁸See, e.g., <https://ethereum.org/en/developers/docs/smart-contracts>.

⁹Here we assume a non-interactive rational arguments. Our arguments are interactive, but public-coin. We leave it as an open problem to which extent rational arguments can be made non-interactive through Fiat-Shamir.

Micali also illustrated the power of rational proofs by constructing a single-round rational proof for all of $\#P$. The work of Campanelli and Gennaro [CG15, CG17] constructs rational proofs with composition properties (e.g., that are reusable for multiple executions) for bounded-depth circuits and bounded-space computations. The work of Guo, Hubacek, Rosen and Vald [GHRV14] restrict the rational prover to be computationally bounded, obtaining the notion of *rational arguments*. The advantage of rational arguments over their classical counterparts is that they allow for low communication and a *sublinear* verifier. [GHRV14] construct sublinear verifier rational arguments for class NC^1 . The work of [GHRV16] extends this result by constructing a single-round rational argument with sublinear verification for class P . Another line of work achieves verifiable computation against rational parties in an indirect manner through approaches based on *fine-grained* cryptographic primitives [CG18].

2 Preliminaries

Notation. We denote a finite field by \mathbb{F} . We denote by λ a security parameter. We consider interactive arguments for relations, where a prover P convinces the verifier that it knows a witness w such that for a public statement x , $(x, w) \in \mathcal{R}$. For a pair of PPT interactive algorithms P, V , we denote a protocol by (P, V) and by $(P(w), V)(x)$, the random variable representing the transcript between P and V when interacting where w is (an optional) P 's private input and x is a common input. We denote by $\text{out}((P(w), V)(x))$ the output of V after interacting with P on input x . We denote a *negligible* function, a function that vanishes faster than the inverse of any polynomial in λ by $\text{negl}(\lambda)$. A *noticeable* function is the inverse of a polynomial.

2.1 Modelling Access to Inputs

Given algorithm A , we write $A(x)$ (as it is standard) to denote that the string x is written on the input tape of A . This implies in particular that A will have to run at least in time $|x|$. Some of the algorithms in this paper will require only oracle access to inputs. We denote them through the following superscript notation as in the following: $A^{[u]}(x)$. Here A has oracle access to input u but “standard” (tape) access to input x . Algorithms may have oracle and standard access to more than one input: $A^{[u],[v]}(x, y)$ where the algorithm A has oracle access to inputs u, v and standard access to inputs x, y .

2.2 Rational Proofs and Arguments

We give the definition of Rational Proofs from [AM12]. Let $\text{rew}(\cdot)$ denote a randomized function computed by V on the transcript. The goal of a rational P is to maximize the expected value of reward rew , while the goal of V is to learn the correct evaluation of the function f on x .

Definition 1 (Rational Proof). A circuit $F : \{0, 1\}^n \rightarrow \{0, 1\}^*$ admits a rational proof if there exists an interactive proof (P, V) and a randomized reward function $\text{rew} : \{0, 1\}^* \rightarrow \mathbb{R}_{\geq 0}$ such that

1. For any input $x \in \{0, 1\}^n$, $\Pr[\text{out}((P(x), V^{[x]})) = F(x)] \geq 1 - \text{negl}(n)$.
2. For every prover \tilde{P} , and for any input $x \in \{0, 1\}^n$ there exists a $\delta_{\tilde{P}}(x) \geq 0$ such that $\mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] + \delta_{\tilde{P}}(x) \leq \mathbb{E}[\text{rew}((P(x), V^{[x]}))]$.

The expectations and probabilities are taken over the coins of both prover and verifier.

Rational arguments introduced in [GHRV14] capture a rational prover that is restricted to computationally bounded strategies.

Definition 2 (Rational Argument [GHRV14]). A circuit $F : \{0, 1\}^n \rightarrow \{0, 1\}^*$ admits a rational argument if there exists an interactive protocol (P, V) and a randomized reward function $\text{rew} : \{0, 1\}^* \rightarrow \mathbb{R}_{\geq 0}$ such that for any input $x \in \{0, 1\}^n$, any prover \tilde{P} of size $\leq 2^{\lambda(n)}$:

1. $\Pr[\text{out}((P, V)(x)) = F(x)] \geq 1 - \text{negl}(\lambda)$
2. $\mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] \leq \mathbb{E}[\text{rew}((P(x), V^{[x]}))] + \text{negl}(\lambda)$
3. If $\Pr[\text{out}((P, V)(x)) \neq F(x)] \geq 1/p(n)$ for some polynomial $p(\cdot)$, then there exists a polynomial $q(\cdot)$ such that $\mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] + 1/q(n) \leq \mathbb{E}[\text{rew}((P(x), V^{[x]}))]$

The expectations and the probabilities are taken over the random coins of the prover and verifier. The rational argument is efficient if the verifier runs in time $o(n)$.

Property 1 corresponds to the notion of completeness, Property 2 guarantees that the gain attained by deviating from the prescribed strategy in a computationally bounded way is at most negligible. Property 3 guarantees that not reporting the correct output with noticeable probability results in a noticeable loss in reward.

Reward gap measures how big the loss of a prover that always reports $F(x)$ incorrectly is. A noticeable gap in expectation between such a prover and the prescribed behavior guarantees that it is beneficial for the prover to act honestly to significantly increase its utility, allowing us to argue for rationality in the presence of computational cost.

Definition 3 (Reward Gap [GHRV14]). Let (P, V) be a rational argument for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$ with reward function rew . Let $\epsilon_{\tilde{P}} = \Pr[\text{out}((\tilde{P}, V)(x)) \neq F(x)]$, and $\delta_{P^*}(x) = \mathbb{E}[\text{rew}((P(x), V^{[x]}))] - \mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))]$. The reward gap is a function $\Delta : \mathbb{N} \rightarrow \mathbb{R}$ such that for every $n \in \mathbb{N}$,

$$\Delta(n) = \min_x \min_{P^* : \epsilon_{P^*} = 1} [\delta_{P^*}(x)]$$

For an arbitrary prover \tilde{P} we have $\delta_{\tilde{P}}(x) \geq \epsilon_{\tilde{P}} \cdot \Delta(x)$.

EXAMPLES OF RATIONAL PROOFS. For concreteness here we show the protocol for a single threshold gate (readers are referred to [AM12, AM13, GHRV14] for more examples). Let $G_{n,k}(x_1, \dots, x_n)$ be a threshold gate with n Boolean inputs, that evaluates to 1 if at least k of the input bits are 1. The protocol in [AM13] to evaluate this gate goes as follows. The Prover announces the number \tilde{m} of input bits equal to 1, which allows the Verifier to compute $G_{n,k}(x_1, \dots, x_n)$. The Verifier select a random index $i \in [1..n]$ and looks at input bit $b = x_i$ and rewards the Prover using Brier's Rule $BSR(\tilde{p}, b)$ where $\tilde{p} = \tilde{m}/n$ i.e. the probability claimed by the Prover that a randomly selected input bit be 1. Then

$$BSR(\tilde{p}, 1) = 2\tilde{p} - \tilde{p}^2 - (1 - \tilde{p})^2 + 1 = 2\tilde{p}(2 - \tilde{p})$$

$$BSR(\tilde{p}, 0) = 2(1 - \tilde{p}) - \tilde{p}^2 - (1 - \tilde{p})^2 + 1 = 2(1 - \tilde{p}^2)$$

Let m be the true number of input bits equal to 1, and $p = m/n$ the corresponding probability, then the expected reward of the Prover is

$$pBSR(\tilde{p}, 1) + (1 - p)BSR(\tilde{p}, 0) \tag{1}$$

which is easily seen to be maximized for $p = \tilde{p}$ i.e. when the Prover announces the correct result. Moreover one can see that when the Prover announces a wrong \tilde{m} his reward goes down by $2(p - \tilde{p})^2 \geq 2/n^2$. In other words for all n -bit input x , we have $\Delta(x) = 2/n^2$ and if a dishonest Prover \tilde{P} cheats with probability $\epsilon_{\tilde{P}}$ then $\delta_{\tilde{P}} > 2\epsilon_{\tilde{P}}/n^2$.

2.3 Interactive Proofs

Indexed Relations. In order to achieve a succinct verifier, we model relations as triples instead of pairs; this allows us to split the verifier’s input into one part for the offline phase and one part for the online phase. The offline input, called the index, is encoded by an indexer algorithm, and this encoding can be reused for proofs over different instances for the same index. Additionally, the *encoded* index allows verifier efficiency – in cryptographic protocols the encoding is a cryptographic digest and reading it is more efficient than reading the index; in rational protocols, the encoding once generated and stored, is accessed as an oracle at very few locations.

Definition 4 (Indexed relation [CHM⁺20]). An *indexed relation* \mathcal{R} is given by a set of triples (i, x, w) where i is the index, x is the instance, and w is the witness. The corresponding *indexed language* $\mathcal{L}_{\mathcal{R}}$ is the set of pairs (i, x) for which there exists a witness w such that $(i, x, w) \in \mathcal{R}$.

Rank-1 Constraint Systems (R1CS). *Rank-1 Constrained Systems (R1CS)* are a popular way to encode computations to be proven via a SNARK. R1CS were implicitly defined as *Quadratic Arithmetic Programs* in [GGPR13] where it is proven that they are NP-complete (and therefore can express any arbitrary non-deterministic polynomial computation).

An R1CS instance is a tuple $(\mathbb{F}, A, B, C, x, N, m)$ where x denotes the public input of the instance, $A, B, C \in \mathbb{F}^{N \times N}$ are matrices defined over a field \mathbb{F} , with $N \geq |x| + 1$, and there are at most m non-zero entries in each matrix.

An R1CS instance $(\mathbb{F}, A, B, C, x, N, m)$ is satisfiable if there exists a witness $w \in \mathbb{F}^{N-|x|-1}$ such that

$$(A \cdot z) \circ (B \cdot z) = (C \cdot z)$$

where $z = (x, 1, w)$, \cdot is the matrix-vector product, and \circ is the Hadamard product.

R1CS for deterministic computations. The above Indexed relation definition is for NP, can be used to capture P as well. A tuple (i, x, w) corresponds to $(f, (x, y), z)$ where $f(x) = y$ and z consists of all intermediate values in computing $f(x)$. Now, R1CS representation for this tuple is well-defined.

A proof/argument system for a language \mathcal{L} allows a prover P to convince a verifier V that $x \in L$ for a common input x . A proof of knowledge intuitively captures not only the truth of a statement $x \in L$, but also that the prover is in “possession” of a witness w .

2.4 Polynomial Commitment Scheme

A polynomial commitment scheme [KZG10] allows a prover to open evaluations of the committed polynomial succinctly. A polynomial commitment scheme over \mathbb{F} is given by a tuple $\text{PC} = (\text{setup}, \text{commit}, \text{open}, \text{eval})$ where:

- $\text{setup}(1^\lambda, D) \rightarrow \text{pp}_{\text{pc}}$. On input security parameter λ , and an upper bound $D \in \mathbb{N}$ on the degree, setup generates public parameters pp_{pc} .
- $\text{commit}(\text{pp}_{\text{pc}}, f(X), d) \rightarrow (C, \tilde{c})$. On input the public parameters pp_{pc} , and a univariate polynomial $f(X) \in \mathbb{F}[X]$ with degree at most $d \leq D$, commit outputs a commitment to the polynomial C , and additionally an opening hint \tilde{c} .
- $\text{open}(\text{pp}_{\text{pc}}, f(X), d, C, \tilde{c}) \rightarrow b$. On input the public parameters pp_{pc} , the commitment C and the opening hint \tilde{c} , a polynomial $f(X)$ of degree $d \leq D$, open outputs a bit indicating accept or reject.
- $\text{eval}(\text{pp}_{\text{pc}}, C, d, x, v; f(X)) \rightarrow b$. A public coin interactive protocol $\langle P_{\text{eval}}(f(X)), V_{\text{eval}} \rangle(\text{pp}_{\text{pc}}, C, d, z, v)$ between a PPT prover and a PPT verifier. The parties have as common input public parameters pp_{pc} , commitment C , degree d ,

evaluation point x , and claimed evaluation v . The prover has, in addition, the opening $f(X)$ of C , with $\deg(f) \leq d$. At the end of the protocol, the verifier outputs 1 indicating accepting the proof that $f(x) = v$, or outputs 0 indicating rejection.

A polynomial commitment scheme must satisfy completeness, commitment binding, evaluation binding. It may also satisfy stronger properties like extractability and hiding.

Definition 5 (Completeness). For all polynomials $f(X) \in \mathbb{F}[X]$ of degree $d \leq D$, for all $x \in \mathbb{F}$,

$$\Pr \left(b = 1 : \begin{array}{l} \text{pp}_{\text{pc}} \leftarrow \text{setup}(1^\lambda, D) \\ (C, \tilde{\mathbf{c}}) \leftarrow \text{commit}(\text{pp}_{\text{pc}}, f(X), d) \\ v \leftarrow f(x) \\ b \leftarrow \text{eval}(\text{pp}_{\text{pc}}, C, d, x, v; f(X)) \end{array} \right) = 1.$$

Definition 6 (Commitment Binding). A polynomial commitment scheme PC is binding if for all PPT \mathcal{A} , the following probability is negligible in λ :

$$\Pr \left(\begin{array}{l} \text{open}(\text{pp}_{\text{pc}}, f_0, d, C, \tilde{\mathbf{c}}_0) = 1 \wedge \\ \text{open}(\text{pp}_{\text{pc}}, f_1, d, C, \tilde{\mathbf{c}}_1) = 1 \wedge \\ f_0 \neq f_1 \end{array} : \begin{array}{l} \text{pp}_{\text{pc}} \leftarrow \text{setup}(1^\lambda, D) \\ (C, f_0, f_1, \tilde{\mathbf{c}}_0, \tilde{\mathbf{c}}_1, d) \leftarrow \mathcal{A}(\text{pp}_{\text{pc}}) \end{array} \right).$$

Definition 7 (Extractability). For any PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a PPT algorithm \mathcal{E} such that the following probability is negligible in λ :

$$\Pr \left(b = 1 \wedge \mathcal{R}_{\text{eval}}(\text{pp}_{\text{pc}}, C, x, v; \tilde{f}, \tilde{\mathbf{c}}) = 0 : \begin{array}{l} \text{pp}_{\text{pc}} \leftarrow \text{setup}(1^\lambda, D) \\ (C, d, x, v, \text{st}) \leftarrow \mathcal{A}_1(\text{pp}_{\text{pc}}) \\ (\tilde{f}, \tilde{\mathbf{c}}) \leftarrow \mathcal{E}^{\mathcal{A}_2}(\text{pp}_{\text{pc}}) \\ b \leftarrow \langle \mathcal{A}_2(\text{st}), V_{\text{eval}} \rangle(\text{pp}_{\text{pc}}, C, d, x, v) \end{array} \right).$$

where the relation $\mathcal{R}_{\text{eval}}$ is defined as follows:

$$\mathcal{R}_{\text{eval}} = \{((\text{pp}_{\text{pc}}, C \in \mathbb{G}, x \in \mathbb{F}, v \in \mathbb{F}); (f(X), \tilde{\mathbf{c}})) : (\text{open}(\text{pp}_{\text{pc}}, f, d, C, \tilde{\mathbf{c}}) = 1) \wedge v = f(x)\}$$

Definition 8 (Succinctness). The scheme is *proof succinct* if the commitments and the evaluation proofs are of size independent of the degree of the polynomial, $|C|$ is $\text{poly}(\lambda)$, $|\pi|$ is $\text{poly}(\lambda)$ where π is the transcript obtained by applying FS to eval . Additionally, the scheme is *verifier succinct* if eval runs in time $\text{poly}(\lambda) \cdot \log(d)$ for the verifier.

2.5 Algebraic Holographic Proof

Definition 9 (AHP [CHM⁺20]). An *Algebraic Holographic Proof* (AHP) over a field family \mathcal{F} for an indexed relation \mathcal{R} is given by the following tuple:

$$\text{AHP} = (\mathbf{k}, \mathbf{s}, \mathbf{d}, \mathcal{I}, \mathcal{P}, \mathcal{V})$$

where $\mathbf{k}, \mathbf{s}, \mathbf{d} : \{0, 1\}^* \rightarrow \mathbb{N}$ are polynomial-time computable functions; $\mathcal{I}, \mathcal{P}, \mathcal{V}$ are the *indexer*, *prover*, and *verifier* algorithms; \mathbf{k} denotes the number of rounds, \mathbf{s} denotes the number of polynomials in each round, and \mathbf{d} specifies degree bounds on these polynomials. The protocol proceeds as follows:

- **Indexing phase** The indexer \mathcal{I} receives as input a field $\mathbb{F} \in \mathcal{F}$, index i for \mathcal{R} , and outputs $\mathbf{s}(0)$ polynomials $p_{0,1}, \dots, p_{0,\mathbf{s}(0)} \in \mathbb{F}[X]$ of degrees at most $d(|i|, 0, 1), \dots, d(|i|, 0, \mathbf{s}(0))$ respectively. This phase does not depend on the public input or witness and simply consists of encoding the given index i .
- **Online phase** The prover \mathcal{P} receives¹⁰ $(i, \mathbf{x}, \mathbf{w})$, for an instance \mathbf{x} and witness \mathbf{w} such that $(i, \mathbf{x}, \mathbf{w}) \in \mathcal{R}$. The verifier \mathcal{V} receives \mathbf{x} and oracle access to the polynomials

¹⁰ \mathbb{F} is an implicit input to all algorithms, and omitted for brevity.

output by $\mathcal{I}(\mathbb{F}, i)$. The prover \mathcal{P} and the verifier \mathcal{V} interact over $k = k(|i|)$ rounds. In the i -th round, $i \in [k]$, the verifier \mathcal{V} sends a message $\rho_i \in \mathbb{F}^*$ to the prover \mathcal{P} ; the prover \mathcal{P} responds with $s(i)$ oracle polynomials $p_{i,1}, \dots, p_{i,s(i)} \in \mathbb{F}[X]$. At the end of k rounds, the verifier outputs additional randomness $\rho_{k+1} \in \mathbb{F}^*$ which is an auxiliary input to \mathcal{V} in subsequent phases.

- **Query phase** Let $\mathbf{p} = (p_{i,j})_{i \in [k], j \in [s(i)]}$ be a vector consisting of all the polynomials sent by the prover \mathcal{P} . The verifier \mathcal{V} executes a subroutine $\mathbf{Q}_{\mathcal{V}}$ that receives $(x; \rho_1, \dots, \rho_{k+1})$ and outputs a query set Q consisting of tuples $((i, j), z)$ that are interpreted as “query $p_{i,j}$ at $z \in \mathbb{F}^n$ ”. We denote a vector consisting of query answers by $\mathbf{p}(Q)$.
- **Decision phase** The verifier outputs **accept** or **reject** based on the answers received to the queries and its randomness. That is, \mathcal{V} executes a subroutine $\mathbf{D}_{\mathcal{V}}$ that receives $(x, \mathbf{p}(Q); \rho_1, \dots, \rho_{k+1})$ as input, and outputs a decision bit.

The function d determines what kind of provers are considered for the completeness and soundness properties of the proof system. A (potentially malicious) prover $\tilde{\mathcal{P}}$ is considered *admissible* for AHP if, in an interaction with the verifier \mathcal{V} , it holds that for every round $i \in [k]$ and oracle index $j \in [s(i)]$ we have $\deg(p_{i,j}) \leq d(|i|, i, j)$. The honest prover \mathcal{P} is required to be admissible under this definition. An AHP satisfies completeness and soundness as defined below.

- **Completeness:** An AHP is complete if for all $\mathbb{F} \in \mathcal{F}$ and any $(i, x, w) \in \mathcal{R}$, the decision bit returned by $\mathcal{V}^{\mathcal{I}(\mathbb{F}, i)}(x)$ after interacting with an honest $\mathcal{P}(i, x, w)$ is 1.
- **Soundness:** An AHP is ϵ -sound if for any prover \mathcal{P}^* , field $\mathbb{F} \in \mathcal{F}$, for every $(i, x) \notin \mathcal{L}_{\mathcal{R}}$, and auxiliary input z :

$$\Pr[\langle \mathcal{P}^*(i, x, z), \mathcal{V}^{\mathcal{I}(\mathbb{F}, i)}(x) \rangle = 1] \leq \epsilon$$

Using the compilation process in [BFS20, CHM+20], AHP = $(\mathcal{I}, \mathcal{P}, \mathcal{V})$ for \mathcal{R} can be turned into a *preprocessing argument*, denoted by AoK = (S, I, P, V) for \mathcal{R} .

3 Definitions and Building Blocks

3.1 Rational Arguments with Indexing

We now formalize preprocessing the function so that a rational verifier can be sublinear in the size of the function representation (for instance, sublinear in the size of the circuit computing f). Once the preprocessing is done, the verifier has oracle access to preprocessed information and can then engage in multiple proofs for the same function by using the same oracles. We call this model rational argument with *indexing* to indicate that the function is preprocessed or indexed and the verifier only needs this indexed material in order to compute the reward. Crucially, accessing this indexed material is via oracle queries, and the verifier only makes sublinear number of queries.

Definition 10 (Rational Argument with Indexing). A Rational Argument with Indexing for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$ is a tuple of four algorithms $(\text{Setup}, \text{idx}, \text{Prv}, V, \text{rew})$. **Setup** is a probabilistic polynomial-time setup algorithm that samples public parameters pp_{rat} . The indexer algorithm idx is a deterministic algorithm that takes $\text{pp}_{\text{rat}}, f$, and outputs $\phi \in \{0, 1\}^*$ that is used as an oracle by V . $(\text{Prv}, V^{[\phi]}, \text{rew})$ is a rational argument for f where the reward function $\text{rew}^{[\phi]}(\text{pp}_{\text{rat}}, x, (\text{Prv}, V)(x)) \rightarrow \mathbb{R}_{\geq 0}$ can be computed with oracle access to ϕ .

In our constructions, the verifier needs to do no preprocessing for a function being indexed except saving a representation of the function f . In particular, when the function

is polynomial evaluation, the verifier just needs to store the polynomial (for example, as a vector of coefficients) and make few oracle accesses to it during verification.

A note on input access. The standard verifier in interactive proofs needs to read the whole public input (i.e., the input is written on an input tape and read sequentially by a Turing machine). Works on rational proofs [AM12] model the verifier’s access to input differently. As in proofs of proximity [RVW13], we do not require the verifier to read all of it, but only to query specific indices (usually randomly sampled). In this work, we exploit this fact plus more: we explicitly model the fact that the verifier may also have oracle access to (some description of) the function and oracle access to the input. This provides efficiency of the verifier when combined with the fact that, in the online stage, the verifier needs to query very few points in the representation of f . $(\text{Prv}, V^{[\phi],[x]}, \text{rew})$ is a rational argument for f with *both function and input indexing* where the reward function $\text{rew}^{[\phi],[x]}(\text{pp}_{\text{rat}}, x, (\text{Prv}, V)(x)) \rightarrow \mathbb{R}_{\geq 0}$ can be computed with oracle access to ϕ and x .¹¹

Definition 11 (Rational Argument with Full Indexing). A Rational Argument with Full Indexing for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$ is a tuple of four algorithms $(\text{Setup}, \text{idx}, \text{Prv}, V, \text{rew})$. Setup is a probabilistic polynomial-time setup algorithm that samples public parameters pp_{rat} . The indexer algorithm idx is a deterministic algorithm that takes $\text{pp}_{\text{rat}}, f$, and outputs $\phi \in \{0, 1\}^*$ that is used as an oracle by V . $(\text{Prv}, V^{[\phi],[x]}, \text{rew})$ is a rational argument for f where the reward function $\text{rew}^{[\phi],[x]}(\text{pp}_{\text{rat}}, x, (\text{Prv}, V)(x)) \rightarrow \mathbb{R}_{\geq 0}$ can be computed with oracle access to ϕ and x .

3.2 Efficient Rational Proofs for Polynomial Evaluation

Let $f \in \mathbb{F}[X]$ be a univariate polynomial of degree d , $f(X) = \sum_{i=0}^d a_i X^i$. We now consider rational proofs for polynomial evaluation. That is for the function

$$F_{\text{poly},f}(t) := f(t)$$

In order to construct *efficient* rational proofs for polynomial evaluation we exploit parallel circuits as much as possible proceed as follows:

- We describe a highly parallel circuit for $F_{\text{poly},f}$ (see `PolyEvalPreproc` in Fig. 1) which works by requiring a small preprocessing by the verifier. This preprocessing requires logarithmic steps in the degree of the polynomial. Its output is described as the vector \mathbf{u} in Fig. 2.
- We then apply a rational proof for highly parallel arithmetic circuits from [CG15] (see also Lemma 2) to `PolyEvalPreproc`.

Theorem 1. *There exists a rational proof for polynomial evaluation where:*

- *the verifier runs in time $O(\log(d))$;*
- *the query and the round complexity is $O(\log(d))$;*
- *the prover runs in time $O(d)$;*
- *the reward gap is noticeable.*

Proof. Completeness and rationality properties of the protocol follow by applying Lemma 2 and Lemma 1 to the protocol in Fig. 2.

The efficiency of the protocol follows from these observations: In order to evaluate f at point t , the sum consists of $d + 1$ terms where each term can be computed using $O(d)$ multiplications in the circuits after a preprocessing. This results into an arithmetic circuit of depth $O(\log d)$ as described in Fig. 1. We can thus obtain a rational proof that runs in logarithmic rounds and communication and with a logarithmic verifier (Theorem 1). \square

¹¹See also Section 2.1.

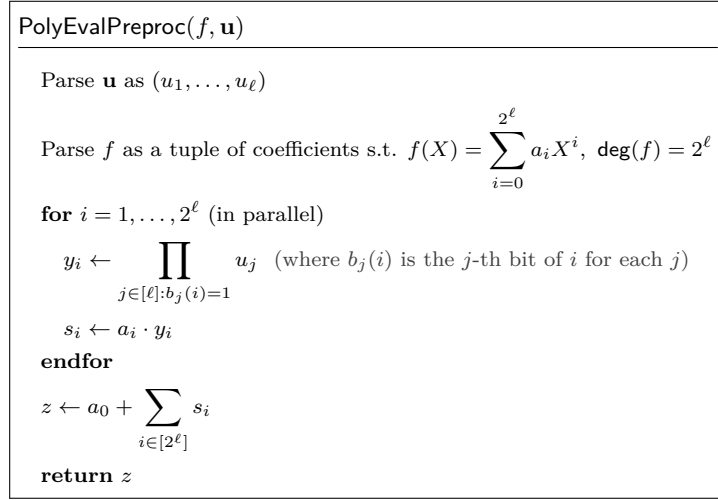


Figure 1: An arithmetic circuit with logarithmic depth for polynomial evaluation.

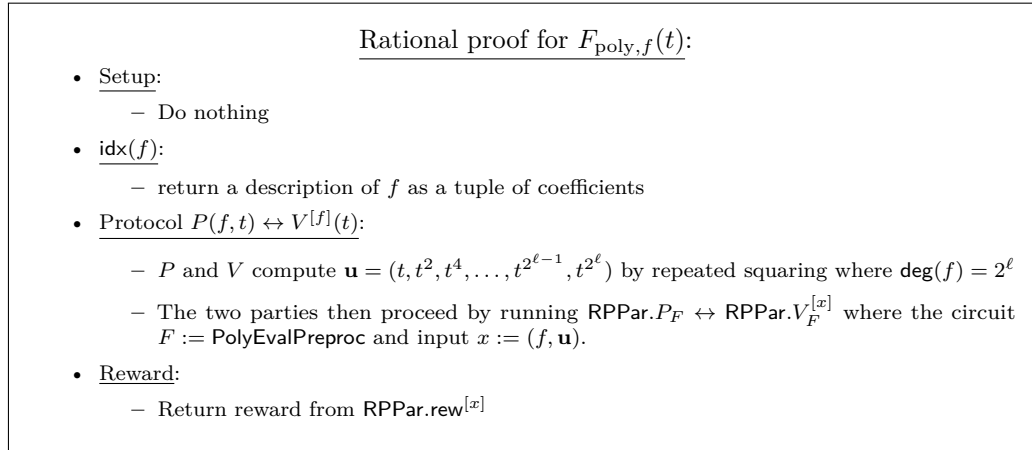


Figure 2: Rational proof for polynomial evaluation. The protocol RPPar is from the statement of Lemma 2.

In the previous theorem we used the following lemma, which we prove in the full version [CGG23b].

Lemma 1. *The following properties hold for the circuit PolyEvalPreproc in Fig. 1:*

- if f is a polynomial and $t \in \mathbb{F}$ then $\text{PolyEvalPreproc}\left(f, (t, t^2, t^4, \dots, t^{2^{\ell-1}}, t^{2^\ell})\right)$ outputs $f(t)$;
- the circuit wiring can be decided in time logarithmic in the degree of the polynomial.

Lemma 2 ([CG15]). *Let \mathbb{F} be a field and F be a logtime-uniform arithmetic circuit of logarithmic depth then there exists a rational proof RPPar for F . For an input x of size n , the verifier runs in logarithmic time in n , the query and the round complexity is logarithmic in n . Its reward gap is noticeable.*

4 Our Compiler: AHP to Rational Argument for P

Consider the indexed relation \mathcal{R} given by triples (i, x, w) that represent $(f, (x, y), w)$ such that $f(x) = y$ and w consists of all intermediate values in computing $f(x)$. We now construct a compiler that transforms an AHP into a rational argument for deterministic computations.

The high-level idea behind our compiler is along the lines of the AHP to succinct argument compiler [BFS20, CHM⁺20]. In an AHP, the prover and the verifier interact where the prover sends polynomial oracles, and the verifier sends random challenges. Then, the verifier queries the polynomials at some challenge points for evaluation, and finally accepts or rejects. Existing compilers use a polynomial commitment scheme as a cryptographic object to realize the polynomial oracle. The prover in the compiled argument commits to the polynomials and then provides evaluations together with proofs of correct evaluation. The argument verifier accepts if both the PCS verifier and the AHP verifier accept. Our key idea is to separate the queries, and have the prover answer some of them with a cryptographic proof and the rest with a rational proof.

- Queries to index polynomials $p_{0,j}$: on query z , prover responds with y and a rational proof for the statement “ $p_{0,j}(z) = y$ ”.
- Queries to witness polynomials: $p_{i,j}$: on query z , prover responds with y together with an evaluation proof of a polynomial commitment scheme for “ $p_{i,j}(z) = y$ AND $C_{i,j} = \text{Commit}(p_{i,j})$ ”

Our compiler removes cryptographic operations in the indexing step and the prover performs cryptographic operations that grows only with the input and not the circuit. The resulting argument is a rational argument. Now, in order to make the verifier sublinear in the public input, we outsource the computation involving the public input to the prover, and have the prover provide a rational proof.

Additional properties of AHP. We use the definition of witness-carrying polynomials (WCPs) [ABC⁺22] that identify a minimum set of polynomials containing enough information about the whole witness, with which auxiliary commitments are shown to be consistent. We also restrict the AHP extractor to be straightline and deterministic so that it is essentially a witness decoding algorithm that works for both honest and malicious provers. Additionally, we assume that the public input access in the computation of the AHP verifier is restricted to querying a polynomial encoding of it. We note that the AHPs of PLONK, Marlin and Sonic already satisfy these properties.

4.1 A Formal Description of our Compiler from AHP to Rational Argument

Let $\text{AHP} = (\mathcal{I}, \mathcal{P}, \mathcal{V})$ be an AHP for \mathcal{R} . Let $\text{PC} = (\text{setup}, \text{commit}, \text{open}, \text{eval})$ be a polynomial commitment scheme. and $\text{rat-pc} = (\text{idx}, P \leftrightarrow V, \text{rew})$ be a rational argument system for polynomial evaluations. We assume that \mathcal{V} has a query algorithm $Q_{\mathcal{V}}$ and a decision algorithm $D_{\mathcal{V}}$. Furthermore, whenever $D_{\mathcal{V}}$ uses the statement x , it is for the computation of a polynomial encoding \hat{x} and therefore giving oracle access to \hat{x} suffices to correctly run the decision algorithm $D_{\mathcal{V}}$. We therefore think of $D_{\mathcal{V}}(x, \mathbf{v}; \rho_1, \dots, \rho_{k+1})$ as $D_{\mathcal{V}}^{\hat{x}}(\mathbf{v}; \rho_1, \dots, \rho_{k+1})$. This is without loss of generality for AHPs of interest. The rational argument with full indexing $(\text{Setup}, \text{idx}, P, V, \text{rew})$ for \mathcal{R} is described in Fig. 3.

Theorem 2. *Let \mathcal{R} be the relation (i, x, w) that represents $(f, (x, y), w)$ Let AHP be an AHP for \mathcal{R} with soundness error ϵ_1 , let PC be a polynomial commitment scheme with*

¹²Here, we assume that the AHP encodes the input as a polynomial in monomial basis. The compiler extends to AHPs that encode input as a polynomial in Lagrange basis as well, by simply using a rational proof for evaluating a polynomial given point-evaluation pairs. We show this in the full version [CGG23b].

Let $\text{AHP} = (\mathcal{I}, \mathcal{P}, \mathcal{V})$ be an AHP for \mathcal{R} . Let $\text{PC} = (\text{setup}, \text{commit}, \text{open}, \text{eval})$ be a polynomial commitment scheme. Let $\text{rat-pc} = (\text{idx}, P \leftrightarrow V, \text{rew})$ be a rational argument system for polynomial evaluations.

- **Setup** $\text{Setup}(1^\lambda)$: Run $\text{pp}_{\text{pc}} \leftarrow \text{setup}(1^\lambda)$, $\text{pp}_{\text{rat}} \leftarrow \text{Setup}(N)$. Output $\text{srs} := (\text{pp}_{\text{pc}}, \text{pp}_{\text{rat}})$.
- **Preprocessing idx**: On input i :
 1. Run the AHP indexer \mathcal{I} on i to obtain polynomials $(p_{0,j})_{j=1}^{s(0)} \in \mathbb{F}[X]$.
 2. For $j \in [s(0)]$, invoke the rational argument indexer to preprocess $p_{0,j}$; $\phi_j \leftarrow \text{idx}(p_{0,j})$.
 3. Input indexing: Run the rational argument indexer idx on x ; $\phi_v \leftarrow \text{idx}(x)$.

Output $\text{ipk} := (\phi_v, (\phi_j)_{j=1}^{s(0)})$.

- **Prover P and Verifier V**: In every round $i \in [k]$, P and V simulate the interaction between the AHP prover $\mathcal{P}(i, x, w)$ and verifier $\mathcal{V}(x)$. In round i :
 1. V interacts with P and internally runs \mathcal{V} . V receives $\rho_i \in \mathbb{F}$ from \mathcal{V} , and forwards it to P.
 2. P interacts with V and internally runs \mathcal{P} . P forwards the received ρ_i to \mathcal{P} , and receives $s(i)$ polynomials $p_{i,1}, \dots, p_{i,s(i)} \in \mathbb{F}[X]$. P invokes the polynomial commitment scheme to commit to each of these polynomials.

$$C_{i,j} = \text{commit}(p_{i,j}), \text{ for } j = 1, \dots, s(i)$$

P sends the vector of commitments $\mathbf{C} = \{C_{i,j}\}$ to V.

- P and V simulate the query phase of the AHP.
 1. Let $\mathbf{p} = (p_{i,j})_{i \in [k], j \in [s(i)]}$ denote the vector consisting of all the polynomials sent by \mathcal{P} , and \mathbf{C} the vector of commitments to \mathbf{p} . \mathcal{V} executes $\mathcal{Q}_{\mathcal{V}}^x(\rho_1, \dots, \rho_{k+1})$ and outputs a query set Q consisting of tuples $(p_{i,j}, z)$. V separates the set of query points $(p_{i,j}, z)$ into rational queries Q_r (that are to index polynomials, that is, $i = 0$), and cryptographic queries Q_c ($i \neq 0$) and forwards (Q_c, Q_r) to P.
 - *Rational queries* ($i = 0$): For $(p_{0,j}, z_j) \in Q_r$, P and V run a rational argument for polynomial evaluation. For each $(f_j, z_j) \in Q_r$, execute rational proof for $F_{\text{poly}, f_j}(z_j)$: $P(f_j, z_j) \leftrightarrow V^{[\phi_j]}(z_j)$.
 - *Cryptographic queries* ($i > 0$): For $(p_{i,j}, z) \in Q_c$, P and V run $\text{eval}(\text{pp}_{\text{pc}}, \mathbf{C}, d, \mathbf{z}, \mathbf{v}; \mathbf{p})$ where \mathbf{z} is the vector of all query points in Q_c and \mathbf{v} is vector of all claimed evaluations. In the above, eval is the *batched* evaluation protocol of the polynomial commitment scheme that proves the evaluation of multiple polynomial commitments.
- **Reward phase**:
 1. V runs \mathcal{V} 's decision algorithm $\mathcal{D}_{\mathcal{V}}^{\hat{x}}(\mathbf{v}; \rho_1, \dots, \rho_{k+1})$. Whenever $\mathcal{D}_{\mathcal{V}}$ queries \hat{x} on α , P and V run a rational argument for polynomial evaluation. For each (\hat{x}, α_j) , execute rational proof for $F_{\text{poly}, \hat{x}}(\alpha_j)$: $P(\hat{x}, \alpha_j) \leftrightarrow V^{[\hat{x}]}(\alpha_j)^{12}$. V receives a decision bit b_1 from $\mathcal{D}_{\mathcal{V}}$.
 2. Let b_2 be the decision bit of V from the execution of eval .
 3. If $(b_1 = 1 \wedge b_2 = 1)$, for each $(p_{0,j}, z_j) \in Q_r$, V computes $\text{rew}_j \leftarrow \text{rew}^{[\phi_j]}(\text{pp}_{\text{rat}}, z_j, (P, V)(z_j))$, and sets $\text{rew} \leftarrow \min_j \{\text{rew}_j\}$.

Figure 3: Our compiler from AHP to rational arguments.

extractability error ϵ_2 , and `rat-pc` be a rational proof for polynomial evaluation with reward gap Δ . The, the compiled system is a rational argument for \mathcal{R} with reward gap $\Delta(1 - \epsilon_1 - \epsilon_2)$. The verifier runs in time $O(|y| + \log |f| + \log |x|)$ given oracle access to x , where $|f|$ is the R1CS complexity of f .

Proof. Let $\text{AHP}^{\{p_{i,j}\}}$ be the starting AHP with polynomial oracles. We first consider a hybrid intermediate oracle argument system $\text{AHP}^{\{p_{0,j}\}}$ where the polynomial oracles of $\{p_{i,j}\}$, for $i \neq 0$ are realized via a cryptographic PCS PC, and the polynomials $\{p_{0,j}\}$ are still available as oracles. Suppose prover \tilde{P} of $\text{AHP}^{\{p_{0,j}\}}$ convinces V of instance (f, x, \tilde{y}) in $L_{\mathcal{R}}$. We construct an adversary \mathcal{A}_{PC} against the extractability game of PC. \mathcal{A}_{PC} receives the commitment key ck , and random coins and internally invokes \tilde{P} to obtain a set of commitments $\{C_{i,j}\}_{i \in [k], j \in [s(i)]}$. We now invoke the extractor \mathcal{E}_{PC} , which given the same input as \mathcal{A}_{PC} outputs a set of polynomials $\tilde{p} = \{p_{i,j}\}_{i \in [k], j \in [s(i)]}$. If \mathcal{E}_{PC} fails, that is the extracted polynomials are inconsistent with the alleged evaluations ($\tilde{p}(\mathbf{z}) \neq \mathbf{v}$), then \mathcal{A}_{PC} wins the extractability game, which happens with probability ϵ_2 . Hence, except with probability ϵ_2 , $\tilde{p}(\mathbf{z}) = \mathbf{v}$. Given that the evaluations are valid w.r.t the set of polynomials \tilde{p} , V accepts when AHP verifier accepts. Since the soundness error of the AHP is ϵ_1 , the probability that V accepts $(f, x, \tilde{y}) \notin L_{\mathcal{R}}$ is $\epsilon_1 + \epsilon_2$.

Now, we replace oracle calls to $\{p_{0,j}\}$ polynomials with rational proofs for polynomial evaluation. The hybrid protocol $\text{AHP}^{\{p_{0,j}\}}$ is compiled into a rational argument Π where $p_{0,j}$ and \hat{x} are realized via a rational proof system. To show the reward gap of Π , we consider the following events:

- Event A : V accepts, that is, $b_1 = 1 \wedge b_2 = 1$, and the rational prover correctly answers all queries to $p_{0,j}$, \hat{x} , that is, for all j , $\text{out}((P^*, V)(z)) = p_{0,j}(z)$, for every query z to $p_{0,j}$, and $\text{out}((P^*, V)(z)) = \hat{x}(z)$, for every query z to \hat{x} .
- Event B : V accepts, that is, $b_1 = 1 \wedge b_2 = 1$, and the rational prover incorrectly answers queries, that is, for some j , $\text{out}((P^*, V)(z)) \neq p_{0,j}(z)$.
- Event C : V rejects, that is $b_1 = 0 \vee b_2 = 0$

Now, the expected reward is

$$\begin{aligned} \mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] &= \Pr[A] \cdot \mathbb{E}[\text{rew}((\tilde{P}, V)(x_A))|A] + \\ &\Pr[B] \cdot \mathbb{E}[\text{rew}((\tilde{P}, V)(x_B))|B] + \Pr[C] \cdot \mathbb{E}[\text{rew}((\tilde{P}, V)(x_C))|C] \end{aligned} \quad (2)$$

where x_E denotes the distribution of the instance for the rational argument given the occurrence of event E . When event C occurs, by construction, the reward is 0, therefore,

$$\mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] = \Pr[A] \cdot \mathbb{E}[\text{rew}((\tilde{P}, V)(x_A))|A] + \Pr[B] \cdot \mathbb{E}[\text{rew}((\tilde{P}, V)(x_B))|B]$$

Now, we bound the probability of event A .

$$\Pr[A] \leq \epsilon_1 + \epsilon_2$$

This is because in the hybrid world where only rational proofs are oracles and others are cryptographically realized, the soundness error is $\epsilon_1 + \epsilon_2$, and when event A occurs, the prover gives correct answers to all rational queries (and therefore behave ideally). Moreover, the prover in the compiled argument does not learn anything more from the query itself since in a public-coin AHP, the queries are independent of the rest of the communication. Now, we have

$$\mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] \leq \Pr[A] \cdot \mathbb{E}[\text{rew}((\tilde{P}, V)(x_A))|A] + (1 - \Pr[A]) \cdot (\mathbb{E}[\text{rew}((P, V)(x_B))] - \Delta)$$

In a public-coin AHP, the verifier queries to the statement polynomials are from tossing public coins, the distribution of queries to polynomials of rational proof is independent of the transcript and instance. Thus, x_A and x_B are distributed identically.

$$\begin{aligned} \mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] &\leq \Pr[A] \cdot \mathbb{E}[\text{rew}((P, V)(x))] + (1 - \Pr[A]) \cdot (\mathbb{E}[\text{rew}((P, V)(x))] - \Delta) \\ \mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] &\leq \mathbb{E}[\text{rew}((P, V)(x))] - (1 - \Pr[A]) \cdot \Delta \\ \mathbb{E}[\text{rew}((\tilde{P}(x), V^{[x]}))] &\leq \mathbb{E}[\text{rew}((P, V)(x))] - (1 - \epsilon_1 - \epsilon_2) \cdot \Delta \end{aligned}$$

□

4.2 Instantiating Our Compiler

Instantiating our compiler requires selecting an AHP (with a decision algorithm with specific properties), a polynomial commitment scheme and a rational argument system for polynomial evaluation. The last can be instantiated through the construction described in Section 3.2. Below we discuss AHP and polynomial commitment candidates. These instantiations are also suitable for our construction in Section 5.2.

In the full version, we show how to make the decision algorithm D_V sublinear in the statement x by realizing $D_V^x(\mathbf{v}; \rho_1, \dots, \rho_{k+1})$ for candidate AHPs: Marlin, PLONK and Sonic.

For polynomial commitments, example candidate constructions are Dory [Lee21] (provably extractable under the SXDH assumption, with transparent setup, logarithmic communication/rounds/verification) and KZG [KZG10] (non-interactive, with trusted setup, constant communication and verification but with extraction provable in the AGM [FKL18]).

5 Rational Arguments of Knowledge for NP

The work of [GHRV14] showed how to downscale rational proofs for $\#P$ of Azar and Micali to NP. They give a succinct (what they call extremely laconic) rational proof for any language in NP where the communication is one bit. However, there are two downsides: (i) the verifier is not succinct, and (ii) the reward gap is negligible. We begin by exploring the possibility of succinct rational proofs for NP with noticeable reward gap.

Succinct Rational Proofs for NP. [GHRV14] also shows that such succinct rational proofs for NP cannot have non-negligible reward gap when the proof is public coin and non-interactive. We show that we cannot have rational proofs for NP that have communication complexity logarithmic in the size of the witness, even when interaction is allowed.¹³

The proof of the following theorem is in the full version [CGG23b].

Theorem 3 (Impossibility of *succinct* rational proofs for NP). *Let (P, V) be a public coin interactive rational proof for a language L with perfect completeness. Let rew be the reward function, $\log(n)$ and $\text{poly}(n)$ be the communication complexity and verification complexity respectively. If the reward gap Δ is a noticeable function, then L is decidable $\text{poly}(n)$ time.*

This theorem shows that even in the rational setting (as in the classical setting), we cannot have succinctness by increasing the number of rounds if we want a *proof* system. We therefore relax the soundness requirement from proof to *argument* just as in the classical setting to achieve succinctness. Most languages that show up in practice need the guarantee of knowledge of a witness: for instance proof of discrete logarithm of a public value or preimage of a compressing function do not guarantee much since they always exist. Instead, what is meaningful is if the prover actually knows the discrete logarithm or the preimage. In the classical setting, this is captured by proofs/arguments of knowledge. We now define *rational arguments of knowledge* for NP.

¹³Theorem 16 of [GHRV14] shows that rational proofs with *one round* and logarithmic communication for NP are impossible, assuming $\text{NP} \not\subseteq \text{BPP}$. The theorem we show extends the impossibility even for rational proofs with logarithmically many rounds. Theorem 5 of [CG17] shows that rational proofs with polylog rounds and communication complexity for NP are impossible, assuming $\text{NP} \not\subseteq \text{BPQP}$. Our theorem is under the weaker assumption that $\text{NP} \not\subseteq \text{BPP}$.

5.1 Succinct Rational Arguments of Knowledge for NP

We now define rational arguments of knowledge where probability of the extractor outputting a witness is tied to the reward of the prover (see also discussion in Section 1.2).

5.1.1 Intuitions on our notion

Recall from Section 1.2 that our notion is silent on the case $x \notin L$. An informal version of our definition is as follows:

For every $x \in L$:

1. the honest prover (running the protocol with some valid witness as input) will earn the highest reward;
2. a prover not knowing the witness should earn a “low” amount.

The two requirements together capture the intuition discussed in Section 1.2, i.e. that a prover not knowing a witness will be incentivized to search for it before protocol execution:

- Requirement 2 implies that not knowing the witness (as in “acting in a way that makes the extractor fail”) gives a low reward.
- By requirement 1, there is at least one strategy that provides the high reward (the honest prover’s strategy).
- Note that requirement 1 and 2 together prevent the extractor from being trivial (never returning valid witnesses for statements in the language).

The two requirements are formalized in Definition 12.

Definition 12 (Rational AoK). A Rational AoK Π for $\mathcal{L} \in NP$ is given by a tuple $(\text{Setup}, \text{idx}, P, V, \text{rew})$: **Setup** is a probabilistic polynomial-time setup algorithm that samples public parameters pp_{rat} . The indexer algorithm idx is a deterministic algorithm that takes $\text{pp}_{\text{rat}}, \mathcal{R}$, and outputs $\phi \in \{0, 1\}^*$ that is used as an oracle by V , where $\mathcal{R}(x, \cdot) = 1$ iff $x \in \mathcal{L}$. The reward function $\text{rew}^{[\phi], [x]}(\text{pp}_{\text{rat}}, x, (P, V)(x)) \rightarrow \mathbb{R}_{\geq 0}$ can be computed with oracle access to ϕ and x . Π is said to have reward gap Δ if for any PPT prover \tilde{P} , any input $x \in \mathcal{L}$, the following hold:

1. (Completeness) $\text{rew}(P, V)(x) > 0$
2. (Knowledge soundness) There exists a PPT extractor \mathcal{E} such that if $\Pr[w \leftarrow \mathcal{E}^{\tilde{P}}(x) : \mathcal{R}(x, w) = 0] \geq 1/q(n)$ for some polynomial $q(\cdot)$, then there exists a polynomial $p(\cdot)$ such that $\mathbb{E}[\text{rew}(\tilde{P}(x), V^{[x]})] + \Delta \leq \mathbb{E}[\text{rew}(P(x), V^{[x]})]$, for $\Delta = 1/p(n)$.

5.2 Our Compiler: AHP to Rational AoK for NP

Consider the indexed relation \mathcal{R} given by triples (i, x, w) where the circuit for the relation is given by the index i , x is the instance, and w is the witness. The corresponding *indexed language* $\mathcal{L}(\mathcal{R})$ is the set of pairs (i, x) for which there exists a witness w s.t. $(i, x, w) \in \mathcal{R}$. We construct a compiler that transforms an AHP into a rational argument for NP.

Our compiler is the same as the compiler for P in Section 4. Here, we show that the compiler yields an argument of knowledge for NP, by showing an extractor that satisfies Definition 12. The resulting rational AoK does not perform any cryptographic operations in the indexing step and the prover performs cryptographic operations that grows only with the witness and not the circuit.

Theorem 4. *Let \mathcal{R} be an NP relation. Let AHP be an AHP for \mathcal{R} with soundness error ϵ_1 , let PC be a polynomial commitment scheme with extractability error ϵ_2 , and rat-pc be a rational proof for polynomial evaluation with reward gap Δ . Then, the compiled system is a rational argument of knowledge for \mathcal{R} with reward gap $\Delta(1 - \epsilon_1 - \epsilon_2)$.*

Proof of the above theorem is deferred to the full version [CGG23b].

References

- [ABC⁺22] Diego F. Aranha, Emil Madsen Bennedsen, Matteo Campanelli, Chaya Ganesh, Claudio Orlandi, and Akira Takahashi. ECLIPSE: Enhanced compiling method for pedersen-committed zkSNARK engines. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part I*, volume 13177 of *LNCS*, pages 584–614. Springer, Heidelberg, March 2022. doi:[10.1007/978-3-030-97121-2_21](https://doi.org/10.1007/978-3-030-97121-2_21).
- [AM12] Pablo Daniel Azar and Silvio Micali. Rational proofs. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 1017–1028. ACM, 2012. doi:[10.1145/2213977.2214069](https://doi.org/10.1145/2213977.2214069).
- [AM13] Pablo Daniel Azar and Silvio Micali. Super-efficient rational proofs. In Michael J. Kearns, R. Preston McAfee, and Éva Tardos, editors, *Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, pages 29–30. ACM, 2013. doi:[10.1145/2492002.2482561](https://doi.org/10.1145/2492002.2482561).
- [BCG⁺13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, August 2013. doi:[10.1007/978-3-642-40084-1_6](https://doi.org/10.1007/978-3-642-40084-1_6).
- [BEJ⁺09] Adam L. Beberg, Daniel L. Ensign, Guha Jayachandran, Siraj Khaliq, and Vijay S. Pande. Folding@home: Lessons from eight years of volunteer distributed computing. In *23rd IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2009, Rome, Italy, May 23-29, 2009*, pages 1–8. IEEE, 2009. doi:[10.1109/IPDPS.2009.5160922](https://doi.org/10.1109/IPDPS.2009.5160922).
- [BFS20] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 677–706. Springer, Heidelberg, May 2020. doi:[10.1007/978-3-030-45721-1_24](https://doi.org/10.1007/978-3-030-45721-1_24).
- [CFF⁺21] Matteo Campanelli, Antonio Faonio, Dario Fiore, Anaïs Querol, and Hadrián Rodríguez. Lunar: A toolbox for more efficient universal and updatable zk-SNARKs and commit-and-prove extensions. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 3–33. Springer, Heidelberg, December 2021. doi:[10.1007/978-3-030-92078-4_1](https://doi.org/10.1007/978-3-030-92078-4_1).
- [CFK22] Matteo Campanelli, Dario Fiore, and Hamidreza Khoshakhlagh. Witness encryption for succinct functional commitments and applications. Cryptology ePrint Archive, Paper 2022/1510, 2022. <https://eprint.iacr.org/2022/1510>. URL: <https://eprint.iacr.org/2022/1510>.
- [CFQ19] Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2075–2092. ACM Press, November 2019. doi:[10.1145/3319535.3339820](https://doi.org/10.1145/3319535.3339820).

- [CG15] Matteo Campanelli and Rosario Gennaro. Sequentially composable rational proofs. In M. H. R. Khouzani, Emmanouil A. Panaousis, and George Theodorakopoulos, editors, *Decision and Game Theory for Security - 6th International Conference, GameSec 2015, London, UK, November 4-5, 2015, Proceedings*, volume 9406 of *Lecture Notes in Computer Science*, pages 270–288. Springer, 2015. doi:10.1007/978-3-319-25594-1_15.
- [CG17] Matteo Campanelli and Rosario Gennaro. Efficient rational proofs for space bounded computations. In Stefan Rass, Bo An, Christopher Kiekintveld, Fei Fang, and Stefan Schauer, editors, *Decision and Game Theory for Security - 8th International Conference, GameSec 2017, Vienna, Austria, October 23-25, 2017, Proceedings*, volume 10575 of *Lecture Notes in Computer Science*, pages 53–73. Springer, 2017. doi:10.1007/978-3-319-68711-7_4.
- [CG18] Matteo Campanelli and Rosario Gennaro. Fine-grained secure computation. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 66–97. Springer, Heidelberg, November 2018. doi:10.1007/978-3-030-03810-6_3.
- [CGG+23a] Matteo Campanelli, Nicolas Gailly, Rosario Gennaro, Philipp Jovanovic, Mara Mihali, and Justin Thaler. sftestudo: Linear time prover snarks with constant size proofs and square root size universal setup. In Abdelrahman Aly and Mehdi Tibouchi, editors, *Progress in Cryptology - LATINCRYPT 2023 - 8th International Conference on Cryptology and Information Security in Latin America, LATINCRYPT 2023, Quito, Ecuador, October 3-6, 2023, Proceedings*, volume 14168 of *Lecture Notes in Computer Science*, pages 331–351. Springer, 2023. doi:10.1007/978-3-031-44469-2_17.
- [CGG23b] Matteo Campanelli, Chaya Ganesh, and Rosario Gennaro. How to make rational arguments practical and extractable. Cryptology ePrint Archive, Paper 2023/1966, 2023. <https://eprint.iacr.org/2023/1966>. URL: <https://eprint.iacr.org/2023/1966>.
- [CHM+20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 738–768. Springer, Heidelberg, May 2020. doi:10.1007/978-3-030-45721-1_26.
- [ECK+23] Jens Ernstberger, Stefanos Chaliasos, George Kadianakis, Sebastian Steinhorst, Philipp Jovanovic, Arthur Gervais, Benjamin Livshits, and Michele Orrù. zk-bench: A toolset for comparative evaluation and performance benchmarking of snarks. Cryptology ePrint Archive, Paper 2023/1503, 2023. <https://eprint.iacr.org/2023/1503>. URL: <https://eprint.iacr.org/2023/1503>.
- [FFG+16] Dario Fiore, Cédric Fournet, Esha Ghosh, Markulf Kohlweiss, Olga Ohrimenko, and Bryan Parno. Hash first, argue later: Adaptive verifiable computations on outsourced data. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1304–1316. ACM Press, October 2016. doi:10.1145/2976749.2978368.
- [FKL18] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, August 2018. doi:10.1007/978-3-319-96881-0_2.

- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013. doi:[10.1007/978-3-642-38348-9_37](https://doi.org/10.1007/978-3-642-38348-9_37).
- [GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476. ACM Press, June 2013. doi:[10.1145/2488608.2488667](https://doi.org/10.1145/2488608.2488667).
- [GHRV14] Siyao Guo, Pavel Hubáček, Alon Rosen, and Margarita Vald. Rational arguments: single round delegation with sublinear verification. In Moni Naor, editor, *ITCS 2014*, pages 523–540. ACM, January 2014. doi:[10.1145/2554797.2554845](https://doi.org/10.1145/2554797.2554845).
- [GHRV16] Siyao Guo, Pavel Hubáček, Alon Rosen, and Margarita Vald. Rational sumchecks. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 319–351. Springer, Heidelberg, January 2016. doi:[10.1007/978-3-662-49099-0_12](https://doi.org/10.1007/978-3-662-49099-0_12).
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016. doi:[10.1007/978-3-662-49896-5_11](https://doi.org/10.1007/978-3-662-49896-5_11).
- [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. <https://eprint.iacr.org/2019/953>.
- [KPV22] Assimakis A. Kattis, Konstantin Panarin, and Alexander Vlasov. RedShift: Transparent SNARKs from list polynomial commitments. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 1725–1737. ACM Press, November 2022. doi:[10.1145/3548606.3560657](https://doi.org/10.1145/3548606.3560657).
- [KRR14] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. How to delegate computations: the power of no-signaling proofs. In David B. Shmoys, editor, *46th ACM STOC*, pages 485–494. ACM Press, May / June 2014. doi:[10.1145/2591796.2591809](https://doi.org/10.1145/2591796.2591809).
- [KWA⁺01] Eric Korpela, Dan Werthimer, David P. Anderson, Jeff Cobb, and Matt Lebofsky. Seti@home-massively distributed computing for SETI. *Comput. Sci. Eng.*, 3(1):78–83, 2001. doi:[10.1109/5992.895191](https://doi.org/10.1109/5992.895191).
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, December 2010. doi:[10.1007/978-3-642-17373-8_11](https://doi.org/10.1007/978-3-642-17373-8_11).
- [Lee21] Jonathan Lee. Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part II*, volume 13043 of *LNCS*, pages 1–34. Springer, Heidelberg, November 2021. doi:[10.1007/978-3-030-90453-1_1](https://doi.org/10.1007/978-3-030-90453-1_1).

- [LSTW21] Jonathan Lee, Srinath Setty, Justin Thaler, and Riad Wahby. Linear-time and post-quantum zero-knowledge SNARKs for R1CS. Cryptology ePrint Archive, Report 2021/030, 2021. <https://eprint.iacr.org/2021/030>.
- [MBKM19] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2111–2128. ACM Press, November 2019. doi:10.1145/3319535.3339817.
- [RVW13] Guy N. Rothblum, Salil P. Vadhan, and Avi Wigderson. Interactive proofs of proximity: delegating computation in sublinear time. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 793–802. ACM Press, June 2013. doi:10.1145/2488608.2488709.
- [RZ21] Carla Ràfols and Arantxa Zapico. An algebraic framework for universal and updatable SNARKs. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 774–804, Virtual Event, August 2021. Springer, Heidelberg. doi:10.1007/978-3-030-84242-0_27.
- [WTs⁺18] Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy*, pages 926–943. IEEE Computer Society Press, May 2018. doi:10.1109/SP.2018.00060.
- [WYX⁺21] Chenkai Weng, Kang Yang, Xiang Xie, Jonathan Katz, and Xiao Wang. Mystique: Efficient conversions for zero-knowledge proofs with applications to machine learning. In Michael D. Bailey and Rachel Greenstadt, editors, *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, pages 501–518. USENIX Association, 2021. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/weng>.