

Linear-time zero-knowledge SNARKs for R1CS

Jonathan Lee Srinath Setty Justin Thaler Riad Wahby
Nanotronics Imaging *Microsoft Research* *Georgetown University* *Stanford University*

Abstract

This paper studies zero-knowledge SNARKs for NP, where the prover incurs $O(N)$ finite field operations to prove the satisfiability of an N -sized R1CS instance. We observe that recent work of Bootle, Chiesa, and Groth (BCG, TCC 20) provides a polynomial commitment scheme that, when combined with the linear-time interactive proof system of Spartan (CRYPTO 20), yields linear-time SNARKs for R1CS. Specifically, for security parameter λ , and for an N -sized R1CS instance over a field of size $\exp(\lambda)$ and fixed $\epsilon > 0$, the prover incurs $O(N)$ finite field operations to produce a proof of size $O_\lambda(N^\epsilon)$ that can be verified in $O_\lambda(N^\epsilon)$ —after a one-time preprocessing step, which requires $O(N)$ finite field operations. This reestablishes the main result of BCG. Arguably, our approach is conceptually simpler and more direct. Additionally, the polynomial commitment scheme that we distill from BCG is of independent interest; it improves over the prior state of the art by offering the first scheme where the time to commit and to prove an evaluation of a committed polynomial are both $O(N)$ finite field operations for an N -sized polynomial.

We further observe that one can render the aforementioned SNARK zero knowledge and reduce the proof size and verifier time to polylogarithmic—while maintaining a linear-time prover—by outsourcing the verifier’s work via one layer of proof composition with an existing zkSNARK as the “outer” proof system. A similar result was recently obtained by Bootle, Chiesa, and Liu (ePrint 2020/1527).

1 Introduction

A zero-knowledge SNARK (zkSNARK) [Kil92, Mic94, GW11, BCCT12] is a cryptographic primitive that enables a *prover* to prove to a *verifier* the knowledge of a satisfying witness to an NP statement by producing a proof π such that: (1) π reveals no information beyond what is implied by the validity of the NP statement; and (2) the size of π and the cost to verify it are both sub-linear (ideally, at most polylogarithmic) in the size of the statement. Given their practical applications, constructing zkSNARKs with excellent asymptotics and concrete efficiency is a highly active area of research.

As with much of the literature on zkSNARKs, we focus on SNARKs for the following NP-complete problems: rank-1 constraint satisfiability (R1CS)¹ and arithmetic circuit satisfiability over a finite field \mathbb{F} . These problems are powerful “intermediate representations”, in that they are amenable to efficient checking via zkSNARKs and are highly expressive. For example, in theory, any non-deterministic random access machine running in time T can be transformed into an R1CS or an arithmetic-circuit-satisfiability instance of size “close” to T . In practice, there exist highly efficient transformations and compiler toolchains to transform applications of interest to these representations [SVP⁺12, PGHR13, BFR⁺13, BCGT13, WSR⁺15, SAGL18a, LNS20].

Our focus on this work is designing zkSNARKs for these problems with the fastest possible prover. We also wish to avoid a trusted setup, and desire a verifier that runs in at most polylogarithmic time. Since the verifier must at least read the statement that is being proven, we allow a one-time preprocessing phase for general, unstructured, R1CS or circuit-satisfiability instances, where the verifier computes a *computation commitment*, i.e., a cryptographic commitment to the structure of a circuit or R1CS matrices [Set20]. After the pre-processing phase, the verifier must run in polylogarithmic time. Furthermore, the pre-processing phase should be at least as efficient as the SNARK prover. Subsequent works refer to such preprocessing to achieve fast verification as leveraging *holography* [CHM⁺20, COS20, BCG20a].

¹R1CS is implicit in the QAPs of Gennaro et al. [GGPR13], but was made explicit in subsequent works [SBV⁺13, BCR⁺19].

Formalizing “fastest possible” provers. How fast can we hope for the prover in a zkSNARK to be? Letting N denote the size of the R1CS or arithmetic-circuit-satisfiability instance over an *arbitrary* finite field \mathbb{F} ,² a lower bound on the prover’s runtime is N operations in \mathbb{F} . This is because any prover that knows a witness w for the instance has to at least convince *itself* (much less the verifier) that w is valid. We refer to this procedure as *native evaluation* of the instance. So the natural goal, roughly speaking, is to achieve a zkSNARK prover that is only a constant-factor slower than native evaluation. Such a prover is said to run in *linear-time*. While our focus in this paper is largely theoretical (i.e., asymptotic), we anticipate that the development of linear-time provers is a prerequisite to deploying zkSNARKs on significantly larger statements than is currently feasible.

Achieving a linear-time prover may sound like a simple and well-defined goal, but it is in fact quite subtle to formalize, because one must be precise about what operations can be performed in one “time-step”, as well as the soundness error achieved and the choice of the finite field.

In known zkSNARKs, the bottleneck for the prover (both asymptotically and concretely) is typically one or more of the following operations (we discuss exceptions later): (1) Performing an FFT over a vector of length $O(N)$. (2) Building a Merkle-hash tree over a vector consisting of $O(N)$ elements of \mathbb{F} . (3) Performing a multiexponentiation of size $O(N)$ in a cryptographic group \mathbb{G} ; in this case, the field \mathbb{F} is of prime order p and \mathbb{G} is typically an elliptic curve group (or subgroup) of order p .

Should any of these operations count as “linear-time”?

FFTs. It is clear that an FFT of length $\Theta(N)$ over \mathbb{F} should *not* count as linear-time, because the fastest known algorithms require $\Theta(N \log N)$ operations over \mathbb{F} , which is a $\log N$ factor, rather than a constant factor, larger than native evaluation.

However, the remaining operations are somewhat trickier to render judgment upon, because they do not refer to field operations.

Merkle-hashing. Regarding (2), a first observation is that to build a Merkle tree over a vector of $O(N)$ elements of \mathbb{F} , computing $O(N)$ cryptographic hashes is necessary and sufficient, assuming the hash function takes as input $O(1)$ elements of \mathbb{F} . However, this is only “linear-time” if hashing $O(N)$ elements of \mathbb{F} can be done in time comparable to $O(N)$ operations over \mathbb{F} .

Bootle, Cerulli, Ghadafi, Groth, Hajiabadi and Jakobsen [BCG⁺17] observe that (assuming the intractability of certain lattice problems over \mathbf{F}_2 , specifically finding a low-Hamming vector in the kernel of a sparse matrix), a collision-resistant hash family of Applebaum, Haramaty, Ishai, Kushilevitz, and Vaikuntanathan [AHI⁺17] is capable of hashing strings consisting of $k \gg \lambda$ bits in $O(k)$ bit operations, with security parameter λ . This means that a vector of $O(N)$ elements of \mathbb{F} can be Merkle-hashed in $O(N \log |\mathbb{F}|)$ bit operations, which Bootle et al. [BCG⁺17] consider comparable to the cost of $O(N)$ operations in \mathbb{F} . Following their lead, we consider this to be “linear-time”.

Multiexponentiation. Pippenger’s algorithm can perform an $O(N)$ -sized multiexponentiation in a group \mathbb{G} of size $\sim 2^\lambda$ by performing $O(N \cdot \lambda / \log(N \cdot \lambda))$ group operations. Typically, one thinks of the security parameter λ as $\omega(\log N)$ (so that 2^λ is superpolynomial in N , ensuring the intractability of problems such as discrete logarithm in \mathbb{G}), and so $O(N \cdot \lambda / (\log N \cdot \lambda))$ group operations is considered $\omega(N)$ group operations. Each group operation is in practice at least as expensive (in fact, several times slower) than a field operation,³ and hence for purposes of this work, we do *not* consider this to be linear time.

Prior work. According to our taxonomy above, the problem of designing zkSNARKs with a linear-time prover remains open (the paragraphs ahead clarify this statement further). The following works have come closest to this goal: (1) Hyrax [WTS⁺18] and Libra [XZZ⁺19] for low-depth, uniform circuits; (2) Spartan [Set20] and Xiphos [SL20] for arbitrary, non-uniform statements represented with R1CS; (3) The recent work of Kothapalli et al. [KMP20] for a variant of R1CS; and (4) The recent work of Bootle et al. [BCG20a, BCL20] for R1CS.

²The size of an arithmetic-circuit-satisfiability instance is the number of gates in the circuit. The size of an R1CS instance of the form $Az \circ Bz = Cz$ is the number of non-zero entries in A, B, C , where \circ denotes the Hadamard (entry-wise) product.

³Typically, an operation in the elliptic-curve group \mathbb{G} requires performing a constant number of field operations within a field that is of similar size to, but different than, then prime-order field \mathbb{F} over which the circuit or R1CS instance is defined.

All these schemes combine the sum-check protocol [LFKN90, BCR⁺19] with cryptographic commitments (e.g., polynomial commitments).

In Hyrax and Libra, beyond performing $O(N)$ finite field operations,⁴ the prover performs $O(d \log N + W)$ exponentiations in a group, where W is the size of the witness to the circuit. As a result, these schemes achieve a linear-time prover so long as $d \log N + W \leq (N \cdot \log W)/\lambda$ (i.e., for circuits of sub-linear depth and with sub-linear sized witnesses, and when $|\mathbb{F}| = 2^{\Theta(\lambda)}$). Two downsides remain. First, their underlying interactive proof [GKR08, CMT12, WJB⁺17] has communication cost $\Theta(d \log N)$, which places a lower bound on the proof length (regardless of the polynomial commitment scheme used). Second, they require uniform circuits (e.g., data-parallel circuits) to achieve verifier’s costs that are sub-linear in the circuit size.

In contrast, Spartan and Xiphos apply to arbitrary R1CS and when using appropriate polynomial commitment schemes (such as Dory [Lee20]), their proof length is $O(\log N)$ group elements. Furthermore, they achieve $O_\lambda(\log N)$ verification times, after a one-time preprocessing step to create a commitment to R1CS matrices.⁵ However, after performing $O(N)$ operations over \mathbb{F} , Spartan’s and Xiphos’s provers perform an $O(N)$ -sized multiexponentiation, which stems from their use of a polynomial commitment schemes applied to a multilinear polynomial with $O(N)$ coefficients.⁶ As a result, according to the aforementioned accounting, they do not achieve a linear-time prover even though concretely these zkSNARKs have a prover time that is in fact the state of the art in most cases.

In a recent theoretical work, Kothapalli et al [KMP20] propose a zkSNARK that achieves $O_\lambda(1)$ proof sizes and verification times for a variant of R1CS, but their prover performs an $O(N)$ -sized multiexponentiation for an N -sized R1CS instance, so it does not achieve a linear-time prover according to the above accounting. Additionally, their scheme requires a one-time trusted setup.

Another recent theoretical work by Bootle, Chiesa, and Groth [BCG20a] proposes a SNARK, where the prover’s work is $O(N)$ finite field operations for an N -sized R1CS instance over any finite field of size $\Omega(N)$ and with a multiplicative subgroup of size N . However, to achieve a soundness error that is negligible in the security parameter λ , one must either sacrifice the linear-time prover, or restrict to R1CS instances over a finite field where $|\mathbb{F}| = 2^{\Theta(\lambda)}$. Furthermore, their construction does not achieve zero-knowledge nor polylogarithmic proofs and verification times (the proof sizes and verification times are $O_\lambda(N^{1/t})$, where t is a constant, and not $O_\lambda(\log N)$ or $O_\lambda(1)$).

Concurrent with our work, Bootle, Chiesa, and Liu [BCL20] address the latter issue by both achieving zero-knowledge as well as polylogarithmic proof sizes and verification times via a host of new techniques. A more detailed discussion of the relationship between our results and those of [BCL20] is in Section 1.3.

In summary, existing works leave open the following problems: (1) designing a zkSNARK for R1CS where the prover incurs only a constant overhead over a native evaluation of the R1CS instance defined over an arbitrary finite field; and (2) designing a zkSNARK for R1CS, which in addition to incurring only a constant overhead to the prover, achieves concrete prover performance better than existing high-speed zkSNARKs [WTS⁺18, XZZ⁺19, Set20, SL20] for “reasonable” instance sizes (e.g., $N \leq 2^{30}$).

1.1 Our results

This work does *not* fully address the aforementioned open problems. We however make some progress toward the second open problem by building on existing high-speed zkSNARKs.

Our first result is a non-interactive argument for R1CS that is unconditionally secure in the random oracle model and achieves the following efficiency. For security parameter λ , and for an N -sized R1CS instance over a field of size $\exp(\lambda)$ and any fixed $\epsilon > 0$, the prover incurs $O(N)$ finite field operations to produce a proof of

⁴Hyrax employs Giraffe [WJB⁺17] that shows how to implement the prover via $O(N)$ finite field operations for *data-parallel* circuits. Libra showed how to achieve this runtime bound for arbitrary circuits.

⁵Similar preprocessing applies to Hyrax and Libra to achieve an $O_\lambda(d \log N)$ -time verifier even for non-uniform, depth- d circuits; see Figure 1.

⁶There exist polynomial commitment schemes, e.g., [ZXZS20], that do not require a multiexponentiation, but they require the prover to perform an FFT over a vector of length N , which as discussed is not linear-time, and in practice is slower than an $O(N)$ -sized multiexponentiation for large enough N .

size $O_\lambda(N^\epsilon)$ that can be verified in $O_\lambda(N^\epsilon)$ —after a one-time preprocessing step, which requires $O(N)$ finite field operations. This reestablishes the main result of Bootle, Chiesa, and Groth [BCG20a], but we believe that our approach is simpler and more direct. It also reuses an existing high-speed, linear-time interactive proof system for R1CS as well as its approach to achieve sub-linear verification costs [Set20].

Second, we observe that one can render the aforementioned SNARK zero knowledge and reduce the proof size and verifier time to at most polylogarithmic—while maintaining a linear-time prover—by outsourcing the verifier’s work via one layer of proof composition with an existing zkSNARK as the “outer” proof system.

Along the way, we distill from [BCG20a] a polynomial commitment scheme with the following efficiency characteristics (this commitment scheme is little more than a rephrasing of the results in [BCG20a]). For security parameter λ , and for a $\log N$ -variate multilinear polynomial over a field of size $\exp(\lambda)$ and any fixed $t > 0$, the time to commit to a polynomial and prove an evaluation of a committed polynomial are both $O(N)$ finite field operations, the size of a commitment is $O_\lambda(1)$, and the time to verify a proof of an evaluation and the size of the proof are both $O_\lambda(N^{1/t})$. This improves over the prior state-of-the-art polynomial commitment schemes [KZG10, ZGK⁺17, WTS⁺18, ZXZS20, BFS20, SL20, Lee20] by offering the first polynomial commitment scheme where the time to commit and to prove an evaluation of a committed polynomial are both linear in the size of the polynomial. Our stated result is for multilinear polynomials, but the scheme generalizes to arbitrary polynomials such as univariate polynomials (see e.g., [Lee20]).

For completeness, we provide a self-contained description and an analysis of the distilled polynomial commitment scheme for the case of $t = 2$. The scheme and its analysis is far simpler than the scheme for general integers $t > 0$ in [BCG20a]. Moreover, even the $t = 2$ case suffices to achieve all of our results that make use of one layer of recursive composition (as to achieve these results, it is sufficient for the proof length and verifier time of the “inner SNARK” to be smaller than N by any factor that dominates $\lambda/\log N$). The distilled polynomial commitment scheme for the $t = 2$ case is already implicit in the argument system (for the arithmetic circuit satisfiability problem) of [BCG⁺17], an ancestor of [BCG20a].

We have not implemented our schemes, but we believe our approach will eventually lead to concretely-efficient zkSNARKs, as most of the machinery in our schemes is inherited from existing high-speed zkSNARKs [Set20, SL20]. Figure 1 depicts the efficiency of our schemes and compares it with prior work.

1.2 Our approach

Distilling a linear-time polynomial commitment scheme. We observe that a core technique in the work of Bootle, Chiesa, and Groth [BCG20a] can be used, in a straightforward manner, to build a non-interactive argument of knowledge for a specific type of inner product relations between two vectors over \mathbb{F} , where one of the vectors is committed with a binding commitment and the other is the tensor product of a constant number of smaller vectors. Specifically, an untrusted prover commits to a vector $z \in \mathbb{F}^N$ by producing an $O_\lambda(1)$ -sized commitment, and then later proves that $v = \langle q, z \rangle$, where $q \in \mathbb{F}^N$ is any vector that is a tensor product of t vectors of length $N^{1/t}$ (for a constant t) and $v \in \mathbb{F}$. For N -sized vectors, the cost to commit and to later prove an inner product relation are both $O(N)$ operations over \mathbb{F} ; the size of an evaluation proof and the time to verify are both $O_\lambda(N^{1/t})$ where λ is the security parameter (the verification time is sub-linear because it exploits the tensor structure in q).

We further observe that the above non-interactive proof system implies a polynomial commitment scheme for multilinear polynomials with efficiency characteristics stated earlier. In a nutshell, one can express the evaluations of a multilinear polynomial using the special inner product operation (§4).⁷

Constructing linear-time SNARKs. Recall that an interactive oracle proof (IOP) [BCS16, RRR16] is a generalization of an interactive proof, where in each round, the prover sends a string as an oracle, and the verifier may read one or more entries in the oracle. Recent works have considered several generalizations of IOPs, such as “polynomial IOPs” [BFS20] (in which in each round the prover may send the verifier a polynomial, and the verifier may query the polynomial to obtain its evaluation at any desired point) and

⁷A similar polynomial commitment scheme is implicit in earlier work of Rothblum and Ron-Zewi [RR20], albeit with a worse dependence of polynomial evaluation proof lengths on both the security parameter λ and the constant parameter t .

“tensor IOPs” [BCG20a] (in which each round the prover may send a vector to the verifier and the verifier can pose “tensor queries” to the vector, meaning the verifier request the inner product of the vector with any vector possessing “tensor structure”). These formalisms are meant to capture the most general type of protocols that we know how to transform into SNARKs. For example, Bootle, Chiesa, and Groth [BCG20a] give a transformation from any tensor IOP to a standard IOP, in a manner that preserves linear runtime for the prover. The latter is in turn known to be transformable into SNARKs [Mic94, Val08, BCS16] (also in a manner that preserves linear runtime for the prover). Similarly, Bünz, et al [BFS20] provide a compiler from polynomial IOPs to SNARKs using an (extractable) polynomial commitment scheme.

To construct SNARKs with a linear-time prover, we take two separate routes (both produce the same result).

1. We observe that the prover’s work in Spartan [Set20], excluding its work as part of a polynomial commitment scheme, is $O(N)$ finite field operations for an N -sized R1CS instance. Thus, by employing the linear-time polynomial commitment scheme that we distill from [BCG20a] in Spartan, we can obtain SNARKs for R1CS with a linear-time prover. To formalize this, we first observe that Spartan [Set20] gives a polynomial IOP for R1CS with a linear-time prover (§3). We then apply the compiler of [BFS20] using the linear-time polynomial commitment scheme.
2. We observe that any polynomial IOP is also a tensor IOP (§7). This observation is in fact implicit or explicit in many prior works, e.g., [AHIV17, WTS⁺18, RR20, Lee20, Tha20]. Hence, applying the transformation of [BCG20a] to Spartan’s polynomial IOP yields a SNARK with linear-time prover.

Finally, our recursively-composed zkSNARKs are obtained via a standard application of (one level of) recursive composition with known zkSNARKs. Prior work [CFH⁺15, BCG⁺20b, KMP20] also applies one level of proof composition to achieve similar properties as us, but in different contexts.

Bootle, Chiesa, and Groth [BCG20a] remark that Spartan [Set20] combined with a subset of their techniques could potentially establish their main result, but suggest that this could be a complicated endeavor as one must distill a tensor IOP from Spartan. We believe that our work shows otherwise. Additionally, our first route offers a more direct way of achieving the same result. This route makes no explicit reference to the tensor IOP formalism, and by extension may render the core technique of [BCG20a] more easily applicable in diverse settings, as polynomial commitment schemes are now well-understood primitives with many applications [BGH19, ZXZS20, BFS20, BCMS20, BDFG20].

1.3 Additional comparisons

On cryptographic assumptions and detailed comparison to [BCL20]. Bootle, Chiesa, and Liu [BCL20] obtain a zkSNARK that is unconditionally knowledge sound in the random oracle model. However, the SNARK must ultimately be instantiated in the plain model, with the random oracle replaced with a hash function supporting linear-time evaluation. The hash function used [AHI⁺17] is collision-resistant based on certain lattice assumptions.

Our first (non-zero-knowledge) SNARK satisfies the same security property as BCL. However, our other SNARKs are not unconditionally sound in the random oracle model, owing to their use of (one layer of) recursive composition. Rather, they are knowledge-sound assuming our first SNARK remains knowledge sound in the plain model when the random oracle is instantiated by the appropriate concrete hash function. This is a well-known issue that arises whenever one recursively composes two SNARKs, where the “inner” SNARK makes use of random oracles [Val08, BCCT13, COS20, BCMS20]. The random oracle used by the inner SNARK must be instantiated before recursive composition can occur, and knowledge soundness of the composed SNARK requires knowledge soundness of the inner SNARK.

From a practical perspective, this difference in security (unconditional security in the random oracle model, vs. plain model security when the random oracle is instantiated with an appropriate hash function) is unimportant for the reasons indicated above. Namely, security of any real-world linear-time instantiation of the SNARKs that are unconditionally sound in the random oracle model (e.g., [BCL20], or our first SNARK) will ultimately need to assume knowledge soundness when the random oracle is instantiated with a hash function that can be evaluated in linear time, just as our recursively-composed SNARKs do.

	prover time	encoder time	proof size/ verifier time	ZK?	assumptions	computational model
Hyrax [WTS ⁺ 18]	$O(W + d \log N)$ \mathbb{G} -exp $O(N)$ \mathbb{F} -ops	N/A	$O_\lambda(\sqrt{W} + d \log N)$	✓	DLOG, RO	data-parallel circuits with low-depth d
Libra [XZZ ⁺ 19] [†]	$O(W + d \log N)$ \mathbb{G} -exp $O(N)$ \mathbb{F} -ops	N/A	$O_\lambda(d \log N)$	✓	q-type, RO	uniform circuits with low-depth d
Virgo [ZXZS20]	$O(N + W \log W)$ \mathbb{F} -ops $O(W)$ hashes	N/A	$O_\lambda(d \log N + \log^2 W)$	✓	RO	uniform circuits with low-depth d
Spartan [Set20]	$O(N)$ \mathbb{G} -exp	$O(N)$ \mathbb{G} -exp	$O_\lambda(\sqrt{N})$	✓	DLOG, RO	R1CS
Spartan++ [SL20]	$O(N)$ \mathbb{G} -exp	$O(N)$ \mathbb{F} -ops	$O_\lambda(\sqrt{N})$	✓	DLOG, RO	R1CS
Xiphos [SL20]	$O(N)$ \mathbb{G} -exp	$O(N)$ \mathbb{F} -ops	$O_\lambda(\log N)$	✓	SXDH, RO	R1CS
[KMP20] [†]	$O(N)$ \mathbb{G} -exp	$O(N)$ \mathbb{G} -exp	$O_\lambda(1)$	✓	q-type, RO*	ACS
Hyrax*	$O(W + d \log N)$ \mathbb{G} -exp $O(N)$ \mathbb{F} -ops	$O(N)$ \mathbb{F} -ops	$O_\lambda(d \log N)$	✓	SXDH, RO	non-uniform circuits with low-depth d
[BCG20a]	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O_\lambda(N^\epsilon)$	✗	RO	R1CS
[BCL20]	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O_\lambda(\log^c(N))$	✓	RO	R1CS
This work	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O(N)$ \mathbb{F} -ops	$O_\lambda(N^\epsilon)$	✗	RO	R1CS
This work	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O(N)$ \mathbb{F} -ops	$O_\lambda(\log^2 N)$	✓	RO*	R1CS
This work	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O(N)$ \mathbb{F} -ops	$O_\lambda(\log N)$	✓	RO*, SXDH	R1CS
This work [†]	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O(N)$ \mathbb{F} -ops	$O_\lambda(1)$	✓	RO*, q-type	R1CS

[†] Requires universal trusted setup

Figure 1: Asymptotic efficiency of prior zkSNARKs (we borrow style from [Set20, SL20, BCG20a]). The depicted costs are for an NP statement of size N over a finite field \mathbb{F} , where $|\mathbb{F}| = \exp(\Theta(\lambda))$ and $\lambda \geq \omega(\log N)$ is the security parameter. \mathbb{F} -ops refers to field multiplications or additions; \mathbb{G} -exp refers to an exponentiation in a group \mathbb{G} whose scalar field is \mathbb{F} . We focus on zkSNARKs with a linear number of operations of a certain type. We do not depict schemes that do not achieve sub-linear verification costs [BBB⁺18, BCG⁺17, WYKW20, BMRS20, DIO20]. For Hyrax [WTS⁺18] and Libra [XZZ⁺19], we assume a layered arithmetic circuit, with N gates, depth d , and a non-deterministic witness of size W . The parameters $\epsilon \in (0, 1)$ and $c > 0$ are constants. ACS refers to a specialization of R1CS [KMP20]. Spartan++ and Xiphos use an untrusted assistant (e.g., the prover) to accelerate the encoder, thereby avoiding $O(N)$ \mathbb{G} -exp operations [SL20]. Hyrax* refers to Hyrax with the following modifications from subsequent works: (1) linear-time sum-checks for non-uniform circuits from Libra [XZZ⁺19]; (2) computation commitments from Spartan [Set20] to preprocess the structure of a non-uniform circuit; and (3) polynomial commitment scheme from Dory [Lee20], which is also used in Xiphos [SL20]. Our first scheme is Theorem 6. The latter three schemes are obtained by applying one-level of proof composition to our first scheme using one of three existing zkSNARKs as the “outer” proof system: Spartan_{RO} [Set20], Xiphos [SL20], and Groth16 [Gro16]. RO in the “assumptions” column refers to the random oracle model (rows for which RO is the only assumption yield an unconditionally knowledge sound protocol in the random oracle model, and *interactive* protocols in the plain model that are knowledge sound assuming CRHFs. Rows for which RO* appears in the assumption column require assuming plain-model security when the random oracle is instantiated with a concrete hash function). Finally, to achieve a linear-time prover by using a linear-time hash function for Merkle hashes, our schemes and the schemes of Bootle et al. [BCG20a, BCL20] require assuming the hardness of certain lattice problems, which are not listed in the “assumptions” column for brevity (§1).

On a technical level, the reason Bootle et al. [BCL20] obtain a zkSNARK in the random oracle model while we do not is that they apply proof composition to interactive oracle proofs (IOPs) that are only later transformed into SNARKs via Merkle hashing and the Fiat-Shamir transformation [Mic94, Val08, BCS16]. In contrast, we apply Merkle hashing and the Fiat-Shamir transformation to obtain a SNARK, *before* performing recursive composition.

Additional related work. GGPR-based zkSNARKs [GGPR13, PGHR13, Gro16], which build on earlier work [IKO07, Gro10, Lip12, SMBW12], offer the best proof sizes and verification times in practice, but they require the prover to perform an FFT of length $\Theta(N)$ and also require a trusted setup (a trusted authority, or a set of authorities of which at least one is honest, that creates public parameters using a secret trapdoor that they later forget). Other approaches to zero-knowledge arguments not discussed above either require the prover to devote $\Theta(N \log N)$ field operations to FFTs [AHIV17, BCR⁺19, GWC19, CHM⁺20], or they require a trusted setup [GWC19, CHM⁺20], or they do not achieve sub-linear verification costs [BBB⁺18, BCG⁺17]. Some works [GMO16, CDG⁺17, WYKW20] achieve a linear-time prover analogous to [BCG20a], but they do not achieve sub-linear verification costs nor proof sizes; many of these works do not achieve non-interactivity nor publicly-verifiable proofs.

1.4 Roadmap

Section 3 describes a linear-time polynomial IOP for R1CS from Spartan [Set20]. Section 4 describes a linear-time polynomial commitment scheme distilled from [BCG20a]; Section 5 extends it to handle sparse polynomials efficiently using techniques from Spartan [Set20]. Section 6 describes our first route to construct linear-time SNARKs, and Section 7 describes our second route to construct linear-time SNARKs. Appendix A explicitly describes the linear-time polynomial commitment scheme when $t = 2$.

2 Preliminaries

We use \mathbb{F} to denote a finite field, λ to denote the security parameter, and $\text{negl}(\lambda)$ to denote a negligible function in λ . Unless we specify otherwise, $|\mathbb{F}| = 2^{\Theta(\lambda)}$.

Polynomials. We recall a few basic facts about polynomials. Detailed treatment of these facts can be found elsewhere [Tha20].

- A polynomial g over \mathbb{F} is an expression consisting of a sum of *monomials* where each monomial is the product of a constant (from \mathbb{F}) and powers of one or more variables (which take values from \mathbb{F}); all arithmetic is performed over \mathbb{F} .
- The degree of a monomial is the sum of the exponents of variables in the monomial; the (total) degree of a polynomial g is the maximum degree of any monomial in g . Furthermore, the degree of a polynomial g in a particular variable x_i is the maximum exponent that x_i takes in any of the monomials in g .
- A *multivariate* polynomial is a polynomial with more than one variable; otherwise it is called a *univariate* polynomial. A multivariate polynomial is called a *multilinear* polynomial if the degree of the polynomial in each variable is at most one.
- A multivariate polynomial g over a finite field \mathbb{F} is called *low-degree* if the degree of g in each variable is bounded above by a constant.

Definition 2.1. Suppose $f : \{0, 1\}^\ell \rightarrow \mathbb{F}$ is a function that maps ℓ -bit strings to an element of \mathbb{F} . A *polynomial extension* of f is a low-degree ℓ -variate polynomial $\tilde{f}(\cdot)$ such that $\tilde{f}(x) = f(x)$ for all $x \in \{0, 1\}^\ell$.

Definition 2.2. A *multilinear extension (MLE)* of a function $f : \{0, 1\}^\ell \rightarrow \mathbb{F}$ is a low-degree polynomial extension where the extension is a multilinear polynomial.

It is well-known that every function $f : \{0, 1\}^\ell \rightarrow \mathbb{F}$ has a unique multilinear extension, and similarly every ℓ -variate multilinear polynomial of \mathbb{F} extends a unique function mapping $\{0, 1\}^\ell \rightarrow \mathbb{F}$. In the rest of the document, for a function f , we use \tilde{f} to denote the unique MLE of f .

For an ℓ -variate multilinear polynomial \tilde{f} extending a function f , let $v_{\tilde{f}}$ denote the 2^ℓ -dimensional vector containing all 2^ℓ evaluations of \tilde{f} . We refer to $v_{\tilde{f}}$ as the representation of \tilde{f} in the Lagrange basis, because the entries of $v_{\tilde{f}}$ are the (unique) set of coefficients of \tilde{f} when \tilde{f} is expressed as a linear combination of Lagrange basis polynomials.

The sum-check protocol. Recall that the sum-check protocol [LFKN90] is an interactive proof system for proving claims of the form: $T = \sum_{x \in \{0,1\}^\ell} G(x)$, where G is ℓ -variate polynomial over \mathbb{F} with the degree in each variable at most μ , and $T \in \mathbb{F}$. In the sum-check protocol, the verifier \mathcal{V}_{SC} interacts with the prover \mathcal{P}_{SC} over a sequence of ℓ rounds. At the end of this interaction, \mathcal{V}_{SC} must evaluate $G(r)$ where $r \in_R \mathbb{F}^\ell$ is a vector of random field elements chosen by \mathcal{V}_{SC} . Other than evaluating $G(r)$, \mathcal{V}_{SC} performs just $O(\ell \cdot \mu)$ field operations. As in prior work, it is natural to view the sum-check protocol as a mechanism to transform claims of the form $\sum_{x \in \{0,1\}^m} G(x) \stackrel{?}{=} T$ to the claim $G(r) \stackrel{?}{=} e$, where $e \in \mathbb{F}$. This is because in most cases, the verifier uses an auxiliary protocol to verify the latter claim.

Rank-1 constraint satisfiability (R1CS). R1CS refers to the following problem.

Definition 2.3. An R1CS instance is a tuple $(\mathbb{F}, A, B, C, M, N, \text{io})$, where $A, B, C \in \mathbb{F}^{M \times M}$, $M \geq |\text{io}| + 1$, io denotes the public input and output, and there are at most $N = \Omega(M)$ non-zero entries in each matrix.

We denote the set of R1CS (instance, witness) pairs as:

$$\mathcal{R}_{\text{R1CS}} = \{ \langle (\mathbb{F}, A, B, C, \text{io}, m, n), w \rangle : A \cdot (w, 1, \text{io}) \circ B \cdot (w, 1, \text{io}) = C \cdot (w, 1, \text{io}) \}$$

In the rest of the manuscript, WLOG, we assume that M and N are powers of 2, and that $M = |\text{io}| + 1$. Throughout this manuscript, all logarithms are to base 2.

3 A linear-time polynomial IOP for R1CS from Spartan

This section recapitulates the results of Spartan [Set20] using a subsequent formalism, a polynomial IOP [BFS20]. This is a variant of IOPs [BCS16, RRR16] where in each round, the prover sends a polynomial as an oracle, and the verifier query may request an evaluation of the polynomial at a point in its domain.

The following theorem formalizes the polynomial IOP at the core of Spartan.

For an R1CS instance, $\mathbb{X} = (\mathbb{F}, A, B, C, M, N, \text{io})$, we interpret the matrices A, B, C as functions mapping domain $\{0, 1\}^{\log M} \times \{0, 1\}^{\log M}$ to \mathbb{F} in the natural way. That is, an input in $\{0, 1\}^{\log M} \times \{0, 1\}^{\log M}$ is interpreted as the binary representation of an index $(i, j) \in [M] \times [M]$, where $[M] := \{1, \dots, M\}$ and the function outputs the (i, j) 'th entry of the matrix.

Theorem 1 ([Set20]). *For any finite field \mathbb{F} , there exists a polynomial IOP for $\mathcal{R}_{\text{R1CS}}$, with the following parameters, where M denotes the dimension of the R1CS coefficient matrices, and N denotes the number of non-zero entries in the matrices:*

- soundness error is $O(\log M)/|\mathbb{F}|$
- round complexity is $O(\log M)$;
- at the start of the protocol, the prover sends a single $(\log M - 1)$ -variate multilinear polynomial \widetilde{W} , and the verifier has a query access to three additional $2 \log M$ -variate multilinear polynomials $\widetilde{A}, \widetilde{B}$, and \widetilde{C} ;
- the verifier makes a single evaluation query to each of the four polynomials $\widetilde{W}, \widetilde{A}, \widetilde{B}$, and \widetilde{C} , and otherwise performs $O(\log M)$ operations over \mathbb{F} ;
- the prescribed prover performs $O(N)$ operations over \mathbb{F} to compute its messages over the course of the polynomial IOP (and to compute answers to the verifier's four queries to $\widetilde{W}, \widetilde{A}, \widetilde{B}$, and \widetilde{C}).

Proof. Let $s = \log M$. For an R1CS instance, $\mathbb{X} = (\mathbb{F}, A, B, C, M, N, \text{io})$ and a purported witness W , let $Z = (W, 1, \text{io})$. As explained prior to the theorem statement, we can interpret A, B, C as functions mapping $\{0, 1\}^s \times \{0, 1\}^s$ to \mathbb{F} , and similarly we interpret Z and $(1, \text{io})$ as functions with the following respective signatures in the same manner: $\{0, 1\}^s \rightarrow \mathbb{F}$ and $\{0, 1\}^{s-1} \rightarrow \mathbb{F}$. It is easy to check that the MLE \widetilde{Z} of Z satisfies

$$\widetilde{Z}(X_1, \dots, X_{\log M}) = (1 - X_1) \cdot \widetilde{W}(X_2, \dots, X_{\log M}) + X_1 \cdot \widetilde{(1, \text{io})}(X_2, \dots, X_{\log M}) \quad (1)$$

Indeed, the right hand side of Equation (1) is a multilinear polynomial, and it is easily checked that $\tilde{Z}(x_1, \dots, x_{\log M}) = Z(x_1, \dots, x_{\log M})$ for all $x_1, \dots, x_{\log M}$ (since the first half of the evaluations of Z are given by W and the second half are given by the vector $(1, \text{io})$). Hence, the right hand side of Equation (1) must be the unique multilinear extension of Z .

From [Set20, Theorem 4.1], checking if $(\mathbb{X}, W) \in \mathcal{R}_{\text{RICS}}$ is equivalent, except for a soundness error of $\log M/|\mathbb{F}|$ over the choice of $\tau \in \mathbb{F}^s$, to checking if the following identity holds:

$$0 \stackrel{?}{=} \left(\sum_{x \in \{0,1\}^s} \tilde{e}q(\tau, x) \cdot \left(\left(\sum_{y \in \{0,1\}^s} \tilde{A}(x, y) \cdot \tilde{Z}(y) \right) \cdot \left(\sum_{y \in \{0,1\}^s} \tilde{B}(x, y) \cdot \tilde{Z}(y) \right) - \sum_{y \in \{0,1\}^s} \tilde{C}(x, y) \cdot \tilde{Z}(y) \right) \right), \quad (2)$$

where $\tilde{e}q$ is the MLE of $eq : \{0, 1\}^s \times \{0, 1\}^s \rightarrow \mathbb{F}$:

$$eq(x, e) = \begin{cases} 1 & \text{if } x = e \\ 0 & \text{otherwise.} \end{cases}$$

That is, if $(\mathbb{X}, W) \in \mathcal{R}_{\text{RICS}}$, then Equation (2) holds with probability 1 over the choice of τ , and if $(\mathbb{X}, W) \notin \mathcal{R}_{\text{RICS}}$, then Equation (2) holds with probability at most $O(\log M/|\mathbb{F}|)$ over the random choice of τ .

Consider computing the right hand side of Equation (2) by applying the sum-check protocol to the polynomial

$$g(x) := \tilde{e}q(\tau, x) \cdot \left(\left(\sum_{y \in \{0,1\}^s} \tilde{A}(x, y) \cdot \tilde{Z}(y) \right) \cdot \left(\sum_{y \in \{0,1\}^s} \tilde{B}(x, y) \cdot \tilde{Z}(y) \right) - \sum_{y \in \{0,1\}^s} \tilde{C}(x, y) \cdot \tilde{Z}(y) \right).$$

From the verifier's perspective, this reduces the task of computing the right hand side of Equation (2) to the task of evaluating g at a random input $r_x \in \mathbb{F}^s$. Note that the verifier can evaluate $\tilde{e}q(\tau, r_x)$ unassisted in $O(\log M)$ field operations, as it is easily checked that $\tilde{e}q(\tau, r_x) = \prod_{i=1}^s (\tau_i r_{x,i} + (1 - \tau_i)(1 - r_{x,i}))$. With $\tilde{e}q(\tau, r_x)$ in hand, $g(r_x)$ can be computed in $O(1)$ time given the three quantities

$$\sum_{y \in \{0,1\}^s} \tilde{A}(r_x, y) \cdot \tilde{Z}(y),$$

$$\sum_{y \in \{0,1\}^s} \tilde{B}(r_x, y) \cdot \tilde{Z}(y),$$

and

$$\sum_{y \in \{0,1\}^s} \tilde{C}(r_x, y) \cdot \tilde{Z}(y).$$

These three quantities can be computed by applying the sum-check protocol three more times in parallel, once to each of the following three polynomials (using the same random vector of field elements, $r_y \in \mathbb{F}^s$, in each of the three invocations):

$$\tilde{A}(r_x, y) \cdot \tilde{Z}(y),$$

$$\tilde{B}(r_x, y) \cdot \tilde{Z}(y),$$

$$\tilde{C}(r_x, y) \cdot \tilde{Z}(y).$$

To perform the verifier's final check in each of these three invocations of the sum-check protocol, it suffices for the verifier to evaluate each of the above 3 polynomials at the random vector r_y , which means it suffices for the verifier to evaluate $\tilde{A}(r_x, r_y)$, $\tilde{B}(r_x, r_y)$, $\tilde{C}(r_x, r_y)$, and $\tilde{Z}(r_y)$. The first three evaluations can be obtained via the verifier's assumed query access to \tilde{A} , \tilde{B} , and \tilde{C} . $\tilde{Z}(r_y)$ can be obtained from one query to \tilde{W} and one query to $(1, \text{io})$ via Equation (1).

In summary, we have the following polynomial IOP:

1. $\mathcal{P} \rightarrow \mathcal{V}$: a $(\log M - 1)$ -variate multilinear polynomial \widetilde{W} as an oracle.
2. $\mathcal{V} \rightarrow \mathcal{P}$: $\tau \in_R \mathbb{F}^s$
3. $\mathcal{V} \leftrightarrow \mathcal{P}$: run the sum-check reduction to reduce the check in Equation (2) to checking if the following hold, where r_x, r_y are vectors in \mathbb{F}^s chosen at random by the verifier over the course of the sum-check protocol:
 - $\widetilde{A}(r_x, r_y) \stackrel{?}{=} v_A$, $\widetilde{B}(r_x, r_y) \stackrel{?}{=} v_B$, and $\widetilde{C}(r_x, r_y) \stackrel{?}{=} v_C$; and
 - $\widetilde{Z}(r_y) \stackrel{?}{=} v_Z$.
4. \mathcal{V} :
 - check if $\widetilde{A}(r_x, r_y) \stackrel{?}{=} v_A$, $\widetilde{B}(r_x, r_y) \stackrel{?}{=} v_B$, and $\widetilde{C}(r_x, r_y) \stackrel{?}{=} v_C$, with one query to each of $\widetilde{A}, \widetilde{B}, \widetilde{C}$;
 - check if $\widetilde{Z}(r_y) \stackrel{?}{=} v_Z$ by checking if: $v_Z = (1 - r_y[1]) \cdot v_W + r_y[1] \cdot \widetilde{W}(r_y[2..])$, where $r_y[2..]$ refers to a slice of r_y without the first element of r_y , and $v_W \leftarrow \widetilde{W}(r_y[2..])$ via an oracle query (see Equation (1)).

Completeness. Perfect completeness follows from perfect completeness of the sum-check protocol and the fact that Equation (2) holds with probability 1 over the choice of τ if $(\mathbb{X}, W) \in \mathcal{R}_{\text{RICS}}$.

Soundness. Applying a standard union bound to the soundness error introduced by probabilistic check in Equation (2) with the soundness error of the sum-check protocol [LFKN90], we conclude that the soundness error for the depicted polynomial IOP is at most $O(\log M)/|\mathbb{F}|$.

Round and communication complexity. The sum-check protocol is applied 4 times (although 3 of the invocations occur in parallel and in practice combined into one [Set20]). In each invocation, the polynomial to which the sum-check protocol is applied has degree at most 3 in each variable, and the number of variables is $s = \log M$. Hence, the round complexity of the polynomial IOP is $O(\log M)$. Since each polynomial has degree at most 3 in each variable, the total communication cost is $O(\log M)$ field elements.

Verifier time. The asserted bounds on the verifier’s runtime are immediate from the verifier’s runtime in the sum-check protocol, and the fact that $\widetilde{e}q$ can be evaluated at any input $(\tau, r_x) \in \mathbb{F}^{2s}$ in $O(\log M)$ field operations.

Prover Time. [Set20] shows how to implement the prover’s computation in the polynomial IOP in $O(N)$ \mathbb{F} -ops using prior techniques for linear-time sum-checks [Tha13, XZZ⁺19] (see also [Tha20, Section 7.5.2] for an exposition). This includes the time required to compute $\widetilde{A}(r_x, r_y)$, $\widetilde{B}(r_x, r_y)$, $\widetilde{C}(r_x, r_y)$, and $\widetilde{Z}(r_y)$ (i.e., to compute answers to the verifier’s queries to the polynomials \widetilde{A} , \widetilde{B} , \widetilde{C} , and \widetilde{Z}). \square

4 Linear-time commitments for multilinear polynomials

This section describes a polynomial commitment scheme from multilinear polynomials (see [Set20, BFS20] for definitions). This is essentially a direct application of the “tensor IOP to point-query IOP” transformation in Bootle et al. [BCG20a] to the case of multilinear polynomial evaluations. Figure 2 compares the asymptotics of this scheme with prior polynomial commitment schemes.

For ease of exposition, instead of first constructing an argument system for special inner products and then using it for polynomial commitments (§1), we directly prove the existence of a linear-time polynomial commitment scheme, using results from [BCG20a]. Furthermore, our focus here is on multilinear polynomials, but the scheme described here generalizes to other polynomials such as univariate polynomials (see e.g., [Lee20]).

Background on tensor IOPs. Recall that an interactive oracle proof (IOP) [BCS16, RRR16] is a generalization of an interactive proof (IP), in which in each round the i prover sends a (possibly long) message string Π_i , and the verifier is given query access to Π_i . [BCG20a] define a generalization of IOPs called *tensor IOPs*. To distinguish tensor IOPs from the standard notion of IOPs that they generalize, [BCG20a] refers to standard IOPs as *point-query IOPs*.

Definition 4.1 ([BCG20a]). A (\mathbb{F}, k, t) -tensor IOP is an IOP except for the following modifications: (1) the prover's message in each round i consists of a message m_i that the verifier reads in full, followed optionally by a vector $\Pi_i \in \mathbb{F}^{k^t}$; and (2) a verifier query to Π_i may request the value $\langle q_1 \otimes q_2 \otimes \dots \otimes q_t, \Pi_i \rangle$ for a chosen round i and chosen vectors $q_1, \dots, q_t \in \mathbb{F}^k$.

Polynomial evaluation as a tensor query to the coefficient vector. For an ℓ -variate multilinear polynomial g represented in the Lagrange basis via a vector $z \in \mathbb{F}^n$ (where $n = 2^\ell$), given an evaluation point $r \in \mathbb{F}^\ell$, $g(r)$ can be evaluated using the following tensor product identity:

$$g(r) = \langle ((r_1, 1 - r_1) \otimes (r_2, 1 - r_2) \otimes \dots \otimes (r_\ell, 1 - r_\ell)), z \rangle.$$

The above equality expresses $g(r)$ as the inner product of the coefficient vector z of the polynomial with another vector described as a tensor product of dimensionality ℓ . However, the identity generalizes to any dimensionality $t \leq \ell$ (for simplicity of presentation, we assume henceforth that $t|\ell$). Specifically, given $r \in \mathbb{F}^\ell$ and $t \in (1, \ell)$, there always exist vectors $q_1, \dots, q_t \in \mathbb{F}^{n^{1/t}}$ (which can be computed from r using $O(n^{1/t})$ operations over \mathbb{F}) such that the following holds:

$$(r_1, 1 - r_1) \otimes (r_2, 1 - r_2) \otimes \dots \otimes (r_\ell, 1 - r_\ell) = (q_1 \otimes q_2 \otimes \dots \otimes q_t). \quad (3)$$

Equation (3) implies:

$$g(r) = \langle (q_1 \otimes q_2 \otimes \dots \otimes q_t), z \rangle.$$

Indexed relations. The completeness and soundness requirements of IPs and IOPs are typically defined with respect to a language \mathcal{L} . For languages, completeness requires that for every input $x \in \mathcal{L}$, there exists a prover strategy that causes the verifier to output 1 with high probability. Soundness requires that for every input $x \notin \mathcal{L}$ and for every prover strategy, the verifier outputs 0 with high probability.

Recall that a relation \mathcal{R} is a set of tuples (\mathbb{X}, W) , where \mathbb{X} is the instance and W is the witness. Moreover, for any relation \mathcal{R} , there is a corresponding language $\mathcal{L}_{\mathcal{R}}$, which is the set of \mathbb{X} for which there exists a witness W such that $(\mathbb{X}, W) \in \mathcal{R}$. In this section, we consider a generalized notion of relations called *indexed relations*, which are implicit in [Set20] but formalized in [COS20]. An indexed relation \mathcal{R} is a set of triples $(\mathbb{I}, \mathbb{X}, W)$, where \mathbb{I} is the index, \mathbb{X} is the instance, and W is the witness. Naturally, there is an *indexed language* $\mathcal{L}_{\mathcal{R}}$ associated with an indexed relation, namely the set of tuples (\mathbb{I}, \mathbb{X}) for which there exists a witness W such that $(\mathbb{I}, \mathbb{X}, W) \in \mathcal{R}$.

Tensor IOPs for multilinear polynomial evaluation. Consider the following indexed relation.

Definition 4.2 (Multilinear evaluation indexed relation). The indexed relation \mathcal{R}_{MLE} is the set of all triples

$$(\mathbb{I}, \mathbb{X}, W) = ((\mathbb{F}, Z), (\ell, r, e), \perp),$$

where \mathbb{F} is a finite field, $Z \in \mathbb{F}^N$ for $n = 2^\ell$, $r \in \mathbb{F}^\ell$, $v \in \mathbb{F}$, such that

$$e = \langle Z, (r_1, 1 - r_1) \otimes (r_2, 1 - r_2) \otimes \dots \otimes (r_\ell, 1 - r_\ell) \rangle.$$

Theorem 2. For a finite field \mathbb{F} and positive integers k, t , there exists a (\mathbb{F}, k, t) -tensor IOP for \mathcal{R}_{MLE} with the following parameters, where $N = 2^\ell = k^t$ and ℓ is a parameter in an instance:

- soundness error is 0;

	commit time	commit size	$\mathcal{P}_{\text{Eval}}$	π_{Eval}	$\mathcal{V}_{\text{Eval}}$	assumptions
vSQL-VPD [ZGK ⁺ 17] [†]	$O(N)$ \mathbb{G} -exp	$O_\lambda(1)$	$O(N)$ \mathbb{G} -exp	$O_\lambda(\log N)$	$O_\lambda(\log N)$	q-type
BP-PC [BBB ⁺ 18, BGH19]	$O(N)$ \mathbb{G} -exp	$O_\lambda(1)$	$O(N)$ \mathbb{G} -exp	$O_\lambda(\log N)$	$O_\lambda(N)$	DLOG
Hyrax-PC [WTS ⁺ 18]	$O(N)$ \mathbb{G} -exp	$O_\lambda(\sqrt{N})$	$O(N)$ \mathbb{F} -ops	$O_\lambda(\log N)$	$O_\lambda(\sqrt{N})$	DLOG
Virgo-VPD [ZXZS20]	$O(N \log N)$ \mathbb{F} -ops	$O_\lambda(1)$	$O(N \log N)$ \mathbb{F} -ops	$O_\lambda(\log^2 N)$	$O_\lambda(\log^2 N)$	CRHF
Kopis-PC [SL20]	$O(N)$ \mathbb{G} -exp	$O_\lambda(1)$	$O(N)$ \mathbb{F} -ops	$O_\lambda(\log N)$	$O_\lambda(\sqrt{N})$	SXDH
Dory-PC [Lee20]	$O(N)$ \mathbb{G} -exp	$O_\lambda(1)$	$O(N)$ \mathbb{F} -ops	$O_\lambda(\log N)$	$O_\lambda(\log N)$	SXDH
Scheme in Section 4 (distilled from [BCG20a])	$O(N)$ \mathbb{F} -ops $O(N)$ hashes	$O_\lambda(1)$	$O(N)$ \mathbb{F} -ops	$O_\lambda(N^{1/t})$	$O_\lambda(N^{1/t})$	CRHF

[†] Requires trusted setup

Figure 2: Asymptotic efficiency of prior polynomial commitment schemes. The depicted costs are for an ℓ -variate multilinear polynomial ($N = 2^\ell$) over a finite field \mathbb{F} , where $|\mathbb{F}| = \exp(\Theta(\lambda))$ and $\lambda \geq \omega(\log N)$ is the security parameter. \mathbb{F} -ops refers to field multiplications or additions; \mathbb{G} -exp refers to an exponentiation in a group \mathbb{G} whose scalar field is \mathbb{F} . The parameter t refers to a constant. For non-interactivity, all schemes except vSQL-VPD [ZGK⁺17] assume the random oracle model; vSQL-VPD assumes q-type, knowledge of exponent assumptions. Furthermore, to achieve a linear-time commit time by using a linear-time hash function for Merkle hashes, our scheme requires assuming the hardness of certain lattice problems, which are not listed in the “assumptions” column for brevity (§1). Finally, one can reduce the verifier time and proof size in our scheme by using one layer of proof composition with an existing zkSNARK [Gro16, Set20, SL20], but we do not depict those variants here.

- round complexity is $O(1)$;
- index length is $O(N)$ elements in \mathbb{F} ;
- query complexity is $O(1)$;
- the prover performs $O(N)$ operations over \mathbb{F} ; and
- the verifier performs $O(N^{1/t})$ operations over \mathbb{F} , given a tensor-query access to the index.

Proof. Suppose that the evaluation point is $r \in \mathbb{F}^\ell$. Consider the following tensor IOP.

<ol style="list-style-type: none"> 1. \mathbb{I}: a vector Π specifying the coefficients of the ℓ-variate multilinear polynomial in the Lagrange basis. 2. $\mathcal{V} \rightarrow \mathcal{P}$: a tensor query (q_1, \dots, q_t) to \mathcal{P}, where $(r_1, 1 - r_1) \otimes (r_2, 1 - r_2) \otimes \dots \otimes (r_\ell, 1 - r_\ell) = (q_1 \otimes q_2 \otimes \dots \otimes q_t)$. This query is answered with an evaluation $e \in \mathbb{F}$ equal to $\langle (q_1, \dots, q_t), \Pi \rangle$.
--

The index consists of $O(N)$ elements in \mathbb{F} . The round complexity and the query complexity of the depicted tensor IOP is $O(1)$; soundness error is 0. In terms of time complexity, the verifier starts with an evaluation point $r \in \mathbb{F}^\ell$, which it transforms to a set of vectors (q_1, \dots, q_t) , where each $q_i \in \mathbb{F}^{N^{1/t}}$; this costs $O(N^{1/t})$ operations over \mathbb{F} to the verifier. The prover performs $O(N)$ operations over \mathbb{F} to compute a response to the tensor query. \square

The theorem below requires a linear code, and to achieve the stated asymptotics, the linear code must support linear-time encoding. As noted in prior work, explicit constructions of such codes are known [Spi96, DI14]. One can also use Reed-Solomon codes, but it introduces a superlinear encoding time and hence a superlinear prover. In fact, in the case of $t = 2$ and using Reed-Solomon codes, the tensor IOP (Theorem 3) and resulting polynomial commitment scheme (Theorem 4) given below is implicit in Ligerio [AHIV17] and explicitly distilled in [Tha20, Section 9.6]. Meanwhile, the case of $t = 2$ using general linear codes is implicitly considered in [BCG⁺17].

Theorem 3. *For security parameter λ , given an (\mathbb{F}, k, t) -tensor IOP for \mathcal{R}_{MLE} with $N = 2^\ell = k^t$ (where ℓ is the size parameter) and fixed value of t , and a linear code over \mathbb{F} with rate $\rho = k/n$, relative distance $\delta = d/n$, and encoding time k , there exists a point-query IOP for \mathcal{R}_{MLE} with the following parameters:*

- soundness error is $O(d^t/|\mathbb{F}| + (1 - \frac{\delta^t}{2})^\lambda)$;
- round complexity is $O(1)$;
- index length is $O(N)$ elements in \mathbb{F} ;
- query complexity is $O(N^{1/t})$;
- the indexer performs $O(N)$ operations over \mathbb{F} ;
- the prover performs $O(N)$ operations over \mathbb{F} ; and
- the verifier performs $O(N^{1/t})$ operations over \mathbb{F} .

Proof. Applying [BCG20a, Theorem 3] to the tensor IOP from Theorem 2 provides the desired result. \square

Theorem 4. *For security parameter λ and a positive integer t , given a hash function that can compute a Merkle-hash of N elements of \mathbb{F} with the same time complexity as $O(N)$ \mathbb{F} -ops, there exists a linear-time polynomial commitment scheme for multilinear polynomials. Specifically, there exists an algorithm that, given as input the coefficient vector of an ℓ -variate multilinear polynomial over \mathbb{F} over the Lagrange basis, with $N = 2^\ell$, commits to the polynomial, where:*

- the size of the commitment is $O_\lambda(1)$; and
- the running time of the commit algorithm is $O(N)$ operations over \mathbb{F} .

Furthermore, there exists a non-interactive argument of knowledge in the random oracle model to prove the correct evaluation of a committed polynomial with the following parameters:

- the running time of the prover is $O(N)$ operations over \mathbb{F} ;
- the running time of the verifier is $O_\lambda(N^{1/t})$ operations over \mathbb{F} ; and
- the proof size is $O_\lambda(N^{1/t})$.

Proof. The desired commit algorithm and its claimed efficiency follows from applying the BCS transform [BCS16] (with a linear-time hash function [AHI⁺17, BCG⁺17]) to the indexer in the point-query IOP from Theorem 3 (obtained by using any linear-time encodable linear error-correcting code of constant rate and relative distance). Similarly, the non-interactive argument of knowledge along with its claimed efficiency follows from applying the BCS transform [BCS16] (with a linear-time hash function [AHI⁺17, BCG⁺17]) to the point-query IOP from Theorem 3. \square

Theorem 4 obtains the claimed polynomial commitment scheme using the tensor-IOP-to-point-IOP transformation of [BCG20a] as a black box, and hence provides the reader with essentially no information about how the polynomial commitment scheme actually operates. For concreteness, we provide a self-contained description of the polynomial commitment scheme and its analysis for the case of $t = 2$ in Appendix A. The analysis for $t = 2$ is far simpler than the more general case of arbitrary t considered in [BCG20a].

5 Linear-time commitments for sparse multilinear polynomials

To construct SNARKs under both routes (§1), we need a commitment scheme for multilinear polynomials \tilde{A} , \tilde{B} , and \tilde{C} capturing the R1CS instance. Specifically, we desire a polynomial commitment scheme that when applied to these polynomials, the time to commit and prove an evaluation are both $O(N)$.

The previous section gave a commitment scheme for ℓ -variate multilinear polynomials where the prover runs in time $O(2^\ell)$. However, \tilde{A} , \tilde{B} , and \tilde{C} are each defined over $\ell = 2 \log M$ variables. So if the commitment scheme is applied naively to these polynomials, the runtime of the committer is $O(M^2)$, which is superlinear in the size of the R1CS instance (unless a constant fraction of the entries of the matrices A, B, C are non-zero). We call this an efficient commitment scheme for “dense” polynomials, since the prover’s runtime is linear in the number of coefficients of the polynomial, irrespective of how many of those coefficients are non-zero.

Fortunately, for each of these three polynomials, only $O(N)$ of coefficients over the Lagrange basis are non-zero. If $N \ll M^2$, we call such polynomials *sparse*. What we require is a commitment scheme for sparse polynomials in which the committer runs in time proportional to the number of non-zero coefficients in both the commitment phase and the evaluation phase. Spartan [Set20] gave a technique that transforms any efficient polynomial commitment scheme for dense polynomials into an efficient polynomial commitment scheme for sparse polynomials. By applying Spartan’s construction to the polynomial commitment scheme of the previous section, we obtain our desired sparse polynomial commitment scheme.

We first state the result from Spartan [Set20, Lemma 7.6] in a more general form; below, provide a detailed proof of this result for completeness. An alternate perspective on this result is in [Tha20, §13].

Theorem 5 ([Set20]). *Given a polynomial commitment scheme for $\log M$ -variate multilinear polynomials with the following parameters (where M is a positive integer and WLOG a power of 2):*

- the size of the commitment is $\mathbf{c}(M)$;
- the running time of the commit algorithm is $\mathbf{tc}(M)$;
- the running time of the prover to prove a polynomial evaluation is $\mathbf{tp}(M)$;
- the running time of the verifier to verify a polynomial evaluation is $\mathbf{tv}(M)$; and
- the proof size is $\mathbf{p}(M)$,

there exists a polynomial commitment scheme for $2 \log M$ -variate multilinear polynomials that evaluate to a non-zero value at at most $N = \Omega(M)$ locations over the Boolean hypercube $\{0, 1\}^{2 \log M}$, with the following parameters, assuming that the commit algorithm is run by an honest entity:

- the size of the commitment is $O(\mathbf{c}(N))$;
- the running time of the commit algorithm is $O(\mathbf{tc}(N))$;
- the running time of the prover to prove a polynomial evaluation is $O(\mathbf{tp}(N))$;
- the running time of the verifier to verify a polynomial evaluation is $O(\mathbf{tv}(N))$; and
- the proof size is $O(\mathbf{p}(N))$.

The following corollary follows from applying the above theorem to the polynomial commitment scheme for $\log M$ -variate multilinear polynomials (Theorem 4).

Corollary 1. *Given a polynomial commitment scheme for $\log M$ -variate multilinear polynomials with the following parameters (where M is a positive integer and WLOG a power of 2, t is a constant, and λ is the security parameter):*

- the size of the commitment is $O_\lambda(1)$;
- the running time of the commit algorithm is $O(M)$ operations over \mathbb{F} ;
- the running time of the prover to prove a polynomial evaluation is $O(M)$ operations over \mathbb{F} ;
- the running time of the verifier to verify a polynomial evaluation is $O_\lambda(M^{1/t})$ operations over \mathbb{F} ; and
- the proof size is $O_\lambda(M^{1/t})$,

there exists a polynomial commitment scheme for $2 \log M$ -variate multilinear polynomials that evaluate to a non-zero value at at most $N = \Omega(M)$ locations over the Boolean hypercube $\{0, 1\}^{2 \log M}$, with the following parameters, assuming that the commit algorithm is run by an honest entity:

- the size of the commitment is $O_\lambda(1)$;
- the running time of the commit algorithm is $O(N)$ operations over \mathbb{F} ;
- the running time of the prover to prove a polynomial evaluation is $O(N)$ operations over \mathbb{F} ;
- the running time of the verifier to verify a polynomial evaluation is $O_\lambda(N^{1/t})$ operations over \mathbb{F} ; and

– the proof size is $O_\lambda(N^{1/t})$.

Representing sparse polynomials with dense polynomials. Let D denote a $2 \log M$ -variate multilinear polynomial that evaluates to a non-zero value at at most $N = \Omega(M)$ locations over $\{0, 1\}^{2 \log M}$. For any $r \in \mathbb{F}^{2 \log M}$, we can express the evaluation of $D(r)$ as follows. Interpret $r \in \mathbb{F}^{2 \log M}$ as a tuple (r_x, r_y) in a natural manner, where $r_x, r_y \in \mathbb{F}^{\log M}$.

$$D(r_x, r_y) = \sum_{(i,j) \in \{0,1\}^{\log M} \times \{0,1\}^{\log M} : D(i,j) \neq 0} D(i, j) \cdot \tilde{e}q(i, r_x) \cdot \tilde{e}q(j, r_y). \quad (4)$$

Claim 1. *Given a $2 \log M$ -variate multilinear polynomial D that evaluates to a non-zero value at at most N locations over $\{0, 1\}^{2 \log M}$, there exist three $\log N$ -variate multilinear polynomials $\text{row}, \text{col}, \text{val}$ such that the following holds for all $r_x, r_y \in \mathbb{F}^s$:*

$$D(r_x, r_y) = \sum_{k \in \{0,1\}^{\log N}} \text{val}(k) \cdot \tilde{e}q(\text{row}(k), r_x) \cdot \tilde{e}q(\text{col}(k), r_y). \quad (5)$$

Moreover, the polynomials' coefficients in the Lagrange basis can be computed in $O(N)$ time.

Proof. Let $R, C, V \in \mathbb{F}^n$. Since D evaluates to a non-zero value at at most N locations over $\{0, 1\}^{2 \log M}$, D can be represented uniquely with N tuples of the form $(i, j, D(i, j)) \in (\{0, 1\}^{\log M}, \{0, 1\}^{\log M}, \mathbb{F})$. By using a natural injection from $\{0, 1\}^{\log M}$ to \mathbb{F} , we can view entries in these tuples as elements of \mathbb{F} . Furthermore, these tuples can be represented with three N -sized vectors R, C, V , where tuple k (for all $k \in [N]$) is stored across the three vectors at the k th location in the vector, i.e., the first entry in the tuple is stored in R , the second entry in C , and the third entry in V . Take row as the unique MLE of R viewed as a function $\{0, 1\}^{\log N} \rightarrow \mathbb{F}$. Similarly, col is the unique MLE of C , and val is the unique MLE of V . The claim holds by inspection since Equations (4) and (5) are both multilinear polynomials in r_x and r_y and agree with each other at every pair $r_x, r_y \in \{0, 1\}^{\log M}$. \square

A first attempt at the commit phase. Here is a first attempt at designing the commit phase. To commit to D , the committer can send commitments to the three $\log N$ -variate multilinear polynomials $\text{row}, \text{col}, \text{val}$ from Claim 1. Using the provided polynomial commitment scheme, this costs $O(N)$ finite field operations, and the size of the commitment to D is $O_\lambda(1)$. As we will see, to aid in the evaluation phase, we will ultimately have to extend the commit phase to include commitments to several additional polynomials.

A first attempt at the evaluation phase. Given $r_x, r_y \in \mathbb{F}^{\log M}$, to prove an evaluation of a committed polynomial, i.e., to prove that $D(r_x, r_y) = v$ for a purported evaluation $v \in \mathbb{F}$, consider the following polynomial IOP, where assume that the verifier has oracle access to the three $\log N$ -variate multilinear polynomial oracles that encode D ($\text{row}, \text{col}, \text{val}$):

1. $\mathcal{P} \rightarrow \mathcal{V}$: two $\log N$ -variate multilinear polynomials E_{r_x} and E_{r_y} as oracles. These polynomials are purported to respectively equal $\tilde{e}q(\text{row}(k), r_x)$ and $\tilde{e}q(\text{row}(k), r_y)$.
2. $\mathcal{V} \leftrightarrow \mathcal{P}$: run the sum-check reduction to reduce the check that

$$v = \sum_{k \in \{0,1\}^{\log N}} \text{val}(k) \cdot E_{r_x}(k) \cdot E_{r_y}(k)$$

to checking if the following hold, where $r_z \in \mathbb{F}^{\log N}$ is chosen at random by the verifier over the course of the sum-check protocol:

- $\text{val}(r_z) \stackrel{?}{=} v_{\text{val}}$;

- $E_{rx}(r_z) \stackrel{?}{=} v_{E_{rx}}$ and $E_{ry}(r_z) \stackrel{?}{=} v_{E_{ry}}$. Here, v_{val} , $v_{E_{rx}}$, and $v_{E_{ry}}$ are values provided by the prover at the end of the sum-check protocol.

3. \mathcal{V} : check if the three equalities hold with an oracle query to each of val , E_{rx} , E_{ry} .

If the prover is honest, it is easy to see that it can convince the verifier about the correct of evaluations of D . Unfortunately, the two oracles that the prover sends in the first step of the depicted polynomial IOP can be completely arbitrary. To fix, this, \mathcal{V} must *additionally* check that the following two conditions hold.

- $\forall k \in \{0, 1\}^{\log N}$, $E_{rx}(k) = \tilde{e}q(\text{row}(k), r_x)$; and
- $\forall k \in \{0, 1\}^{\log N}$, $E_{ry}(k) = \tilde{e}q(\text{col}(k), r_y)$.

A core insight of Spartan [Set20] is to check these two conditions using memory-checking techniques [BEG⁺91]. In particular, the RHS in each of the above conditions can be viewed as N memory lookups over an M -sized memory, initialized to contain the values $\{\tilde{e}q(\text{row}(k), r_x) : k \in \{0, 1\}^{\log N}\}$ and $\{\tilde{e}q(\text{col}(k), r_y) : k \in \{0, 1\}^{\log N}\}$ (note that all of these values can be computed in $O(N)$ total time using standard techniques).

Specifically, focusing on the RHS of the first condition, the M -sized memory mem_{rx} is initialized to satisfy $\text{mem}_{rx}[i] = \tilde{e}q(i, r_x)$ for all $i \in \{0, 1\}^{\log M}$, and for memory lookup $k \in [N]$, the memory address that is read is $\text{row}(k)$. Similarly, in the second condition, the M -sized memory mem_{ry} is the evaluations of $\tilde{e}q(j, r_y)$ for all $j \in \{0, 1\}^{\log M}$, and for memory lookup $k \in [N]$, the memory address read is $\text{col}(k)$.

We take a detour to introduce prior results that we rely on here.

Detour: offline memory checking. Recall that in the offline memory checking algorithm of [BEG⁺91], a *trusted checker* issues operations to an untrusted memory. For our purposes, it suffices to consider only operation sequences in which each memory address is initialized to a certain value, and all subsequent operations are read operations. To enable efficient checking using set-fingerprinting techniques, the memory is modified so that in addition to storing a value at each address, the memory also stores a timestamp with each address. Moreover, each read operation is followed by a write operation that updates the timestamp associated with that address (but not the value stored there). This is captured in the codebox below.

Local state of the checker:

- a timestamp counter ts initialized to 0;
- Two sets: RS and WS , which are initialized as follows.⁸ $RS = \{\}$, and for an M -sized memory, WS is initialized to the following set of tuples: for all $i \in [M]$, the tuple $(i, v_i, 0)$ is included in WS , where v_i is the value stored at address i , and the third entry in the tuple, 0, is an “initial timestamp” associated with the value (intuitively capturing the notion that v_i was written to address i at time step 0, i.e., at initialization).

Read operations and an invariant. For a read operation at address a , suppose that the untrusted memory responds with a value-timestamp pair (v, t) . Then the checker updates its local state as follows:

1. $RS \leftarrow RS \cup \{(a, v, t)\}$;
2. $ts \leftarrow \max(ts, t) + 1$;
3. store (v, ts) at address a in the untrusted memory; and
4. $WS \leftarrow WS \cup \{(a, v, ts)\}$.

The following claim captures the invariant maintained on the sets of the checker:

⁸The checker in [BEG⁺91] maintains a fingerprint of these sets, but for our exposition, we let the checker maintain full sets.

Claim 2. *There exists a set S with cardinality M consisting of tuples of the form (k, v_k, t_k) for all $k \in [M]$ such that $WS = RS \cup S$ if and only if for every read operation the untrusted memory returns the tuple last written to that location.*

Proof. A proof of a more general claim is at [SAGL18b, Lemma C.1]. □

Observe that if the untrusted memory is honest, S can be constructed trivially from the current state of the memory. It is simply the current state of the memory viewed as a set of address-value-timestamp tuples.

Timestamp polynomials. To aid the polynomial evaluation proof of the sparse polynomial commitment scheme, the commit algorithm of the sparse polynomial commitment commits to additional multilinear polynomials, beyond `val`, `row`, and `col`. We now describe these additional polynomials and how they are constructed.

Observe that given the size of memory M and a list of N addresses involved in read operations, one can locally compute three vectors $T_r \in \mathbb{F}^N, T_w \in \mathbb{F}^N, T_f \in \mathbb{F}^M$ defined as follows. For $k \in [N]$, $T_r[k]$ stores the timestamp that would have been returned by the untrusted memory if it were honest during the k th read operation. Let $T_w[k]$ store the timestamp that the checker stores in step (3) of its specification when processing the k th read operation. Similarly, for $k \in [M]$, let $T_f[k]$ store the final timestamp stored at memory location k of the untrusted memory (if the untrusted memory were honest) at the termination of the N read operations. Computing these three vectors requires computation comparable to $O(N)$ operations over \mathbb{F} .

Let $\text{read} = \widetilde{T}_r, \text{write} = \widetilde{T}_w, \text{final} = \widetilde{T}_f$. We refer these polynomials as *timestamp polynomials*, which are unique for a given memory size M and a list of N addresses involved in read operations.

The actual commit algorithm. Given a $2 \log M$ -variate multilinear polynomial D that evaluates to a non-zero value at at most N locations over $\{0, 1\}^{2 \log M}$, the commitment algorithm commits to D by committing to seven $\log N$ -variate multilinear polynomials (`row`, `col`, `val`, `readrow`, `writerow`, `readcol`, `writecol`), two $\log M$ -variate multilinear polynomials (`finalrow`, `finalcol`), where `row`, `col`, `val` are described in Claim 1, and (`readrow`, `writerow`, `finalrow`) and (`readcol`, `writecol`, `finalcol`) are respectively the timestamp polynomials for the N addresses specified by `row` and `col` over a memory of size M .

There are two crucial subtleties unique to the setting of holography that we are exploiting in the above commitment procedure. First, that the timestamp polynomials (`readrow`, `writerow`, `finalrow`, `readcol`, `writecol`, `finalcol`) depend only on the sparse polynomial being committed, and in particular not on the evaluation point (r_x, r_y) . Second, in the holography context, the commit algorithm is run by an honest entity. This means that the commit phase does not need to include any proof that the committed timestamp polynomials actually equal the polynomials $\widetilde{T}_r, \widetilde{T}_w$, and \widetilde{T}_f described above. This appears to be essential for avoiding superlinear operations such as sorting.⁹

In total, using the provided polynomial commitment, the commit algorithm incurs $O(N)$ finite field operations, and the commitment size is $O_\lambda(1)$.

The actual evaluation procedure. To prove the correct evaluation of a $2 \log M$ -variate multilinear polynomial D , in addition to performing the polynomial IOP depicted earlier in the proof, the core idea is to check if the two oracles sent by the prover satisfy the conditions identified earlier using Claim 2.

Claim 3. *Given a $2 \log M$ -variate multilinear polynomial, suppose that (`row`, `col`, `val`, `readrow`, `writerow`, `finalrow`, `readcol`, `writecol`, `finalcol`) denote multilinear polynomials committed by the commit algorithm. Then, for any $r_x \in \mathbb{F}^{\log M}$, checking that $\forall k \in \{0, 1\}^{\log N}, E_{rx}(k) = \tilde{e}q(\text{row}(k), r_x)$ is equivalent to checking $WS = RS \cup S$, where*

⁹If desired in other contexts, such a proof can be produced with $O(N \log N)$ operations over \mathbb{F} . We are unaware of a mechanism to produce a proof faster than this. Straightforward approaches either require breaking field elements into their binary representations (which introduces $\log N$ factor to the number of coefficients of the polynomials being committed) or require the prover to sort a vector of timestamps, which is a superlinear-time operation. The challenging issue is to ensure after every read operation, the timestamp associated with the memory location gets updated to the maximum of the returned timestamp t_s and the current timestamp t .

- $WS = \{(i, \tilde{e}q(i, r_x), 0) : i \in [M]\} \cup \{(\text{row}(k), E_{rx}(k), \text{write}_{\text{row}}(k)) : k \in [N]\};$
- $RS = \{(\text{row}(k), E_{rx}(k), \text{read}_{\text{row}}(k)) : k \in [N]\};$ and
- $S = \{(i, \tilde{e}q(i, r_x), \text{final}_{\text{row}}(i)) : i \in [M]\}.$

Similarly, for any $r_y \in \mathbb{F}^{\log M}$, checking that $\forall k \in \{0, 1\}^{\log N}$, $E_{ry}(k) = \tilde{e}q(\text{col}(k), r_y)$ is equivalent to checking $WS' = RS' \cup S'$, where

- $WS' = \{(j, \tilde{e}q(j, r_y), 0) : j \in [M]\} \cup \{(\text{col}(k), E_{ry}(k), \text{write}_{\text{col}}(k)) : k \in [N]\};$
- $RS' = \{(\text{col}(k), E_{ry}(k), \text{read}_{\text{col}}(k)) : k \in [N]\};$ and
- $S' = \{(j, \tilde{e}q(j, r_y), \text{final}_{\text{col}}(j)) : j \in [M]\}.$

Proof. The desired result follows from a straightforward application of the invariant in Claim 2. \square

There is no direct way to prove that the checks on sets in Claim 3 hold. Instead, we rely on public-coin, multiset hash functions to compress RS , WS , and S into a single element of \mathbb{F} each. Specifically:

Claim 4 ([Set20]). *Given two multisets A, B where each element is from \mathbb{F}^3 , checking that $A = B$ is equivalent to checking the following, except for a soundness error of $O(|A| + |B|)/|\mathbb{F}|$ over the choice of γ, τ : $\mathcal{H}_{\tau, \gamma}(A) = \mathcal{H}_{\tau, \gamma}(B)$, where $\mathcal{H}_{\tau, \gamma}(A) = \prod_{(a, v, t) \in A} (h_\gamma(a, v, t) - \tau)$, and $h_\gamma(a, v, t) = a \cdot \gamma^2 + v \cdot \gamma + t$. That is, if $A = B$, $\mathcal{H}_{\tau, \gamma}(A) = \mathcal{H}_{\tau, \gamma}(B)$ with probability 1 over randomly chosen values τ and γ in \mathbb{F} , while if $A \neq B$, then $\mathcal{H}_{\tau, \gamma}(A) = \mathcal{H}_{\tau, \gamma}(B)$ with probability at most $O(|A| + |B|)/|\mathbb{F}|$.*

We are now ready to depict a polynomial IOP for proving evaluations of a committed sparse multilinear polynomial. Given $r_x, r_y \in \mathbb{F}^{\log M}$, to prove that $D(r_x, r_y) = v$ for a purported evaluation $v \in \mathbb{F}$, consider the following polynomial IOP, where assume that the verifier has an oracle access to multilinear polynomial oracles that encode D (namely, row , col , val , read_{row} , $\text{write}_{\text{row}}$, $\text{final}_{\text{row}}$, read_{col} , $\text{write}_{\text{col}}$, $\text{final}_{\text{col}}$).

1. $\mathcal{P} \rightarrow \mathcal{V}$: two $\log N$ -variate multilinear polynomials E_{rx} and E_{ry} as oracles.
2. $\mathcal{V} \leftrightarrow \mathcal{P}$: run the sum-check reduction to reduce the check that

$$v = \sum_{k \in \{0, 1\}^{\log N}} \text{val}(k) \cdot E_{rx}(k) \cdot E_{ry}(k)$$

to checking that the following equations hold, where $r \in \mathbb{F}^{\log N}$ chosen at random by the verifier over the course of the sum-check protocol:

- $\text{val}(r_z) \stackrel{?}{=} v_{\text{val}}$; and
 - $E_{rx}(r_z) \stackrel{?}{=} v_{E_{rx}}$ and $E_{ry}(r_z) \stackrel{?}{=} v_{E_{ry}}$. Here, v_{val} , $v_{E_{rx}}$ and $v_{E_{ry}}$ are values provided by the prover at the end of the sum-check protocol.
3. \mathcal{V} : check if the three obligations hold with an oracle query each to val , E_{rx} , E_{ry} .
 4. // The following steps check if E_{rx} is well-formed
 5. $\mathcal{V} \rightarrow \mathcal{P}$: $\tau, \gamma \in_R \mathbb{F}$.
 6. $\mathcal{V} \leftrightarrow \mathcal{P}$: run the layered sum-check reduction for “grand products” [Tha13, §5.3.1] to reduce the check that $\mathcal{H}_{\tau, \gamma}(WS) = \mathcal{H}_{\tau, \gamma}(RS) \cdot \mathcal{H}_{\tau, \gamma}(S)$, where RS, WS, S are as defined in Claim 3 and \mathcal{H} is defined in Claim 4 to checking if the following hold, where $r_M \in \mathbb{F}^{\log M}$, $r_N \in \mathbb{F}^{\log N}$ chosen at random by the verifier over the course of the sum-check protocol:
 - $\tilde{e}q(r_M, r_x) \stackrel{?}{=} v_{eq}$
 - $E_{rx}(r_N) \stackrel{?}{=} v_{E_{rx}}$

- $\text{row}(r_N) \stackrel{?}{=} v_{\text{row}}$; $\text{write}_{\text{row}}(r_N) \stackrel{?}{=} v_{\text{write}_{\text{row}}}$; $\text{read}_{\text{row}}(r_N) \stackrel{?}{=} v_{\text{read}_{\text{row}}}$; and $\text{final}_{\text{row}}(r_M) \stackrel{?}{=} v_{\text{final}_{\text{row}}}$
7. \mathcal{V} : directly check if the first equality holds, which can be done with $O(\log M)$ field operations; check the remaining equations hold with an oracle query to each of E_{rx} , row , $\text{write}_{\text{row}}$, read_{row} , $\text{final}_{\text{row}}$.
 8. // The following steps check if E_{ry} is well-formed
 9. $\mathcal{V} \rightarrow \mathcal{P}$: $\tau', \gamma' \in_R \mathbb{F}$.
 10. $\mathcal{V} \leftrightarrow \mathcal{P}$: run the sum-check reduction for “grand products” [Tha13, SL20] to reduce the check that $\mathcal{H}_{\tau', \gamma'}(WS') = \mathcal{H}_{\tau', \gamma'}(RS') \cdot \mathcal{H}_{\tau', \gamma'}(S')$, where RS', WS', S' are as defined in Claim 3 and \mathcal{H} is defined in Claim 4 to checking if the following hold, where $r'_M \in \mathbb{F}^{\log M}, r'_N \in \mathbb{F}^{\log N}$ chosen at random by the verifier over the course of the sum-check protocol:
 - $\tilde{e}q(r'_M, r_y) \stackrel{?}{=} v'_{eq}$
 - $E_{ry}(r'_N) \stackrel{?}{=} v_{E_{ry}}$
 - $\text{col}(r'_N) \stackrel{?}{=} v_{\text{col}}$; $\text{write}_{\text{col}}(r'_N) \stackrel{?}{=} v_{\text{write}_{\text{col}}}$; $\text{read}_{\text{col}}(r'_N) \stackrel{?}{=} v_{\text{read}_{\text{col}}}$; and $\text{final}_{\text{col}}(r'_M) \stackrel{?}{=} v_{\text{final}_{\text{col}}}$
 11. \mathcal{V} : directly check if the first equality holds, which can be done with $O(\log M)$ field operations; check the remaining equations hold with an oracle query to each of E_{ry} , col , $\text{write}_{\text{col}}$, read_{col} , $\text{final}_{\text{col}}$.

Completeness. Perfect completeness follows from perfect completeness of the sum-check protocol and the fact that the multiset equality checks using their fingerprints hold with probability 1 over the choice of τ, γ if the prover is honest.

Soundness. Applying a standard union bound to the soundness error introduced by probabilistic multiset equality checks with the soundness error of the sum-check protocol [LFKN90], we conclude that the soundness error for the depicted polynomial IOP as at most $O(N)/|\mathbb{F}|$.

Round and communication complexity. There are three sum-check reductions. First, it is applied on a polynomial with $\log N$ variables where the degree is at most 3 in each variable, so the round complexity is $O(\log N)$ and the communication cost is $O(\log N)$ field elements.

Second, it is applied to compute four “grand products” in parallel. Two of the grand products are over vectors of size M and the remaining two are over vectors of size N . We use the interactive proof for grand products of [Tha13], for which the round complexity is $O(\log^2 N)$ with a communication cost of $O(\log^2 N)$ field elements.

Third, the depicted IOP runs four additional “grand products”, which incurs the same costs as above.

In total, the round complexity of the depicted IOP is $O(\log^2 N)$ and the communication cost is $O(\log^2 N)$ field elements.¹⁰

Verifier time. The verifier’s runtime is dominated by its runtime in the grand product sum-check reductions, which is $O(\log^2 N)$.

Prover Time. Using linear-time sum-checks [Tha13, XZZ⁺19] in all three sum-check reductions (and exploiting the linear-time prover in the grand product interactive proof [Tha13]), the prover’s time is $O(N)$ finite field operations.

¹⁰Employing the sum-check reduction for grand products in [SL20] results in a complexity of $O(\log N)$ with a communication cost of $O(\log N)$ field elements and a verifier runtime of $O(\log N)$.

Finally, to prove Theorem 5, applying the compiler of [BFS20] to the depicted polynomial IOP with the given polynomial commitment primitive, followed by the Fiat-Shamir transformation [FS86], provides the desired non-interactive argument of knowledge for proving evaluations of committed sparse multilinear polynomials, with efficiency claimed in the theorem statement.

6 Linear-time SNARKs for R1CS from polynomial commitments

This section describes our first route to construct linear-time SNARKs for R1CS (§1). The following theorem captures our main result.

Theorem 6. *Assuming that $|\mathbb{F}| = 2^{\Theta(\lambda)}$ there exists a preprocessing SNARK for $\mathcal{R}_{\text{R1CS}}$ in the random oracle model, with the following efficiency characteristics, where M denotes the dimensions of the R1CS matrices, N denotes the number of non-zero entries, and a fixed positive integer t :*

- the preprocessing cost to the verifier is $O(N)$ \mathbb{F} -ops;
- the running time of the prover is $O(N)$ \mathbb{F} -ops;
- the running time of the verifier is $O_\lambda(N^{1/t})$ \mathbb{F} -ops; and
- the proof size is $O_\lambda(N^{1/t})$.

Proof. From applying [BFS20, Theorem 8] to the polynomial IOP in Theorem 1 using polynomial commitment schemes from Theorem 4 and Corollary 1, there exists a public-coin interactive argument for $\mathcal{R}_{\text{R1CS}}$ with witness-extended emulation. Applying the Fiat-Shamir transform [FS86] to the public-coin interactive argument results in the claimed SNARK for $\mathcal{R}_{\text{R1CS}}$.

The verifier, in a preprocessing step, commits to three $2 \log M$ -variate polynomials that evaluate to a non-zero value at at most N locations over the Boolean hypercube $\{0, 1\}^{2 \log M}$; this costs $O(N)$ \mathbb{F} -ops.

The prover: (1) commits to a $O(\log M)$ -variate polynomial (which costs $O(M)$ \mathbb{F} -ops); (2) participates in the sum-check protocol in the polynomial IOP in Theorem 1 (which costs $O(N)$ \mathbb{F} -ops); and (3) proves evaluations of one $(\log M - 1)$ -variate multilinear polynomial and three $2 \log M$ -variate multilinear polynomials from the preprocessing step (which costs $O(N)$ \mathbb{F} -ops). Together, the prover incurs $O(N)$ \mathbb{F} -ops.

The verifier to verify a proof: (1) participates in the sum-check protocol in the polynomial IOP in Theorem 1 (which costs $O(\log M)$ \mathbb{F} -ops); and (2) verifies the proofs of evaluations of one $(\log M - 1)$ -variate multilinear polynomial and three $2 \log M$ -variate multilinear polynomials from the preprocessing step (which costs $O_\lambda(N^{1/t})$ \mathbb{F} -ops). Together, the verifier incurs $O_\lambda(N^{1/t})$ \mathbb{F} -ops.

Finally, the proof size is the sum of the proof sizes from the sum-check protocol in the polynomial IOP from Theorem 1 and the proof sizes from polynomial commitment schemes. In total, the proof size is $O(\log M) + O_\lambda(N^{1/t}) = O_\lambda(N^{1/t})$. \square

We obtain zkSNARKs with the cost profiles and cryptographic assumptions asserted in the final three rows of Figure 1 by composing the SNARK of Theorem 6 with known zkSNARKs [Set20, SL20, Gro16]. Specifically, the prover in the composed SNARKs proves that it knows a proof π that would convince the SNARK verifier in Theorem 6 to accept. Perfect zero-knowledge of the resulting composed SNARK is immediate from the zero-knowledge property of the SNARKs from these prior works [Set20, SL20, Gro16]. Perfect completeness follows from the perfect completeness properties of these prior works and of Theorem 6. Knowledge soundness follows from a standard argument [Val08, BCCT13]: one composes the knowledge extractors of the two constituent SNARKs to get a knowledge extractor for the composed SNARK.

7 Linear-time SNARKs for R1CS from compiling tensor queries

This section describes our second route to construct linear-time SNARKs for R1CS (§1). Below, we rely on the definition of tensor IOPs and how to represent polynomial evaluations with tensor queries from Section 4.

Polynomial IOPs are Tensor IOPs. For the moment, as we did in Theorem 1, let us assume that the verifier has query access to the polynomials \widetilde{A} , \widetilde{B} , and \widetilde{C} capturing an R1CS instance $(\mathbb{F}, A, B, C, M, N, \text{io})$. Under this assumption, Theorem 1 gave a linear-time polynomial IOP in which the prover begins by sending a multilinear, $(\log M - 1)$ -variate polynomial \widetilde{W} , and then the prover and verifier apply the sum-check protocol interactive proof to $O(1)$ many polynomials. Across all invocations of the sum-check protocol, the verifier runs in $O(\log M)$ time and queries \widetilde{W} at a single point.

Equation (3) implies that this polynomial IOP can be transformed, in a straightforward manner, into a (\mathbb{F}, k, t) -tensor IOP for any constant $t > 0$. Specifically, the honest prover's first message Π_1 in the tensor IOP is simply the coefficient vector of \widetilde{W} over the Lagrange basis, which in turn is simply the witness vector $W \in \mathbb{F}^{M/2}$ (the prover sets m_1 to the empty string). When the verifier needs to learn the evaluation of \widetilde{W} at a point $r \in \mathbb{F}^{\log M - 1}$, Equation (3) shows that the evaluation can be obtained with a single tensor query (q_1, \dots, q_t) to Π_1 . Moreover, in terms of time complexity, the verifier can compute the set of vectors (q_1, \dots, q_t) , in $O(N^{1/t})$ operations over \mathbb{F} .

The following theorem captures the tensor IOP obtained from a straightforward transformation of the polynomial IOP in Theorem 1.

Theorem 7. *For any finite field \mathbb{F} and positive integers k, t , there exists a (\mathbb{F}, k, t) -tensor IOP for \mathcal{R}_{R1CS} with the following parameters, where $N = 2^\ell = k^t$.*

- completeness error is 0;
- soundness error is $O(\log(M))/|\mathbb{F}|$;
- round complexity is $O(\log M)$;
- query complexity is $O(1)$;
- the proof length is $O(M)$ elements of \mathbb{F} ;
- the prover performs $O(N)$ operations over \mathbb{F} ; and
- the verifier performs $O(N^{1/t})$ operations over \mathbb{F} , given a tensor-query access to the R1CS matrices.

Section 5 provides a polynomial IOP for proving evaluations of sparse multilinear polynomials, where the verifier is given an oracle access to a dense representation of the sparse multilinear polynomial. From the aforementioned discussion, this polynomial IOP can be viewed as a holographic tensor IOP for proving evaluations of sparse multilinear polynomials. Using this tensor IOP to implement the tensor-query access to the R1CS matrices results in the following holographic tensor IOP:

Theorem 8. *For any finite field \mathbb{F} and positive integers k, t , there exists a (\mathbb{F}, k, t) -holographic-tensor IOP for \mathcal{R}_{R1CS} with the following parameters, where $N = 2^\ell = k^t$.*

- completeness error is 0;
- soundness error is $O(N)/|\mathbb{F}|$;
- round complexity is $O(\log^2 N)$;
- query complexity is $O(1)$;
- the index length is $O(N)$ elements of \mathbb{F} ;
- the indexer performs $O(N)$ operations over \mathbb{F} ;
- the proof length is $O(N)$ elements of \mathbb{F} ;
- the prover performs $O(N)$ operations over \mathbb{F} ; and
- the verifier performs $O(N^{1/t})$ operations over \mathbb{F} .

Bootle, Chiesa, and Groth [BCG20a] show how to transform any tensor IOP such as the above into a point-query IOP while preserving holography, which is in turn known to be transformable into preprocessing

SNARKs. Both transformations preserve a linear-time prover. Specifically, applying these transformations to the tensor IOP of Theorem 8 yields the following two results.

Theorem 9. *For security parameter λ , given an (\mathbb{F}, k, t) -tensor IOP for \mathcal{R}_{R1CS} with $N = k^t$ and fixed value of t , and a linear code over \mathbb{F} with rate $\rho = k/n$, relative distance $\delta = d/n$, and encoding time $O(k)$, there exists a point-query IOP for \mathcal{R}_{R1CS} with the following parameters:*

- soundness error is $O(N/|\mathbb{F}| + d^t/|\mathbb{F}| + (1 - \frac{\delta^t}{2})^\lambda)$;
- round complexity is $O(\log^2 N)$;
- query complexity is $O(N^{1/t})$;
- the index length is $O(N)$ elements of \mathbb{F} ;
- the indexer performs $O(N)$ operations over \mathbb{F} ;
- the proof length is $O(N)$ elements of \mathbb{F} ;
- the prover performs $O(N)$ operations over \mathbb{F} ;
- the verifier performs $O(N^{1/t})$ operations over \mathbb{F} .

Proof. Applying [BCG20a, Theorem 3] to the tensor IOP from Theorem 8 provides the desired result. \square

Theorem 10. *For security parameter λ and a positive integer t , given a hash function that can compute a Merkle-hash of N elements of \mathbb{F} with the same time complexity as $O(N)$ \mathbb{F} -ops, there exists a preprocessing SNARK in the random oracle model with the following efficiency properties.*

- the running time of the prover is $O(N)$ operations over \mathbb{F} ;
- the preprocessing time for the verifier is $O(N)$ operations over \mathbb{F} ;
- after pre-processing, the running time of the verifier is $O_\lambda(N^{1/t})$ operations over \mathbb{F} ; and
- the proof size is $O_\lambda(N^{1/t})$.

Proof. It is known that, over any finite field, there are linear-time encodable linear codes with constant rate and relative distance [Spi96, DI14]. The theorem then follows by combining the resulting point-query IOP of Theorem 9 with the BCS transform [BCS16] from point-query IOPs to SNARKs in the random oracle model (with a linear-time hash function [AHI⁺17, BCG⁺17]). \square

Acknowledgments

We thank Jonathan Bootle, Alessandro Chiesa, and Jens Groth for answering our questions on [BCG20a], and Eran Tromer for helpful discussions on recursive composition. Justin Thaler was supported in part by NSF CAREER award CCF-1845125. Both Justin Thaler and Riad Wahby were supported in part by DARPA under Agreement No. HR00112020022. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the United States Government or DARPA.

References

- [AHI⁺17] Benny Applebaum, Naama Haramaty, Yuval Ishai, Eyal Kushilevitz, and Vinod Vaikuntanathan. Low-complexity cryptographic hash functions. In *Proceedings of the Innovations in Theoretical Computer Science (ITCS)*, 2017.
- [AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthu Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2018.
- [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 *Proceedings of the Innovations in Theoretical Computer Science (ITCS)*, 2012.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKs and proof-carrying data. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2013.
- [BCG⁺17] Jonathan Bootle, Andrea Cerulli, Essam Ghadafi, Jens Groth, Mohammad Hajiabadi, and Sune K. Jakobsen. Linear-time zero-knowledge proofs for arithmetic circuit satisfiability. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*, 2017.
- [BCG20a] Jonathan Bootle, Alessandro Chiesa, and Jens Groth. Linear-time arguments with sub-linear verification from tensor codes. In *Theory of Cryptography Conference (TCC)*, 2020.
- [BCG⁺20b] Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. Zexe: Enabling decentralized private computation. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2020.
- [BCGT13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. Fast reductions from RAMs to delegatable succinct constraint satisfaction problems: Extended abstract. In *Proceedings of the Innovations in Theoretical Computer Science (ITCS)*, 2013.
- [BCL20] Jonathan Bootle, Alessandro Chiesa, and Siqi Liu. Zero-knowledge succinct arguments with a linear-time prover. ePrint Report 2020/1527, 2020.
- [BCMS20] Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data from accumulation schemes. In *Theory of Cryptography Conference (TCC)*, 2020.
- [BCR⁺19] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for R1CS. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2019.
- [BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive Oracle Proofs. In *Theory of Cryptography Conference (TCC)*, 2016.
- [BDFG20] Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Halo Infinite: Recursive zk-SNARKs from any Additive Polynomial Commitment Scheme. Cryptology ePrint Archive, Report 2020/1536, 2020.
- [BEG⁺91] Manuel Blum, Will Evans, Peter Gemmell, Sampath Kannan, and Moni Naor. Checking the correctness of memories. In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*, 1991.

- [BFR⁺13] Benjamin Braun, Ariel J. Feldman, Zuocheng Ren, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. Verifying computations with state. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, 2013.
- [BFS20] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2020.
- [BGH19] Sean Bowe, Jack Grigg, and Daira Hopwood. Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Report 2019/1021, 2019.
- [BMRS20] Carsten Baum, Alex J. Malozemoff, Marc Rosen, and Peter Scholl. Mac’n’cheese: Zero-knowledge proofs for arithmetic circuits with nested disjunctions. Cryptology ePrint Archive, Report 2020/1410, 2020.
- [CDG⁺17] Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [CFH⁺15] Craig Costello, Cedric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, May 2015.
- [CHM⁺20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2020.
- [CMT12] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In *Proceedings of the Innovations in Theoretical Computer Science (ITCS)*, 2012.
- [COS20] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2020.
- [DI14] Erez Druk and Yuval Ishai. Linear-time encodable codes meeting the Gilbert-Varshamov bound and their cryptographic applications. In *Proceedings of the Innovations in Theoretical Computer Science (ITCS)*, pages 169–182, 2014.
- [DIO20] Samuel Dittmer, Yuval Ishai, and Rafail Ostrovsky. Line-point zero knowledge and its applications. Cryptology ePrint Archive, Report 2020/1446, 2020.
- [FS86] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Proceedings of the International Cryptology Conference (CRYPTO)*, pages 186–194, 1986.
- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2013.
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2008.
- [GMO16] Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster zero-knowledge for Boolean circuits. In *Proceedings of the USENIX Security Symposium*, 2016.

- [Gro10] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*, 2010.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2016.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 99–108, 2011.
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. ePrint Report 2019/953, 2019.
- [IKO07] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Efficient arguments without short PCPs. In *IEEE Conference on Computational Complexity*, 2007.
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 1992.
- [KMP20] Abhiram Kothapalli, Elisaweta Masserova, and Bryan Parno. A direct construction for asymptotically optimal zkSNARKs. Cryptology ePrint Archive, Report 2020/1318, 2020.
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*, pages 177–194, 2010.
- [Lee20] Jonathan Lee. Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. Cryptology ePrint Archive, Report 2020/1274, 2020.
- [LFKN90] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*, October 1990.
- [Lip12] Helger Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In *Theory of Cryptography Conference (TCC)*, 2012.
- [LNS20] Jonathan Lee, Kirill Nikitin, and Srinath Setty. Replicated state machines without replicated execution. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2020.
- [Mic94] Silvio Micali. CS proofs. In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*, 1994.
- [PGHR13] Bryan Parno, Craig Gentry, Jon Howell, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, May 2013.
- [RR20] Noga Ron-Zewi and Ron Rothblum. Local proofs approaching the witness length. In *Foundations of Computer Science (FOCS)*, 2020.
- [RRR16] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 49–62, 2016.
- [SAGL18a] Srinath Setty, Sebastian Angel, Trinabh Gupta, and Jonathan Lee. Proving the correct execution of concurrent services in zero-knowledge. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, October 2018.

- [SAGL18b] Srinath Setty, Sebastian Angel, Trinabh Gupta, and Jonathan Lee. Proving the correct execution of concurrent services in zero-knowledge (extended version). ePrint Report 2018/907, September 2018.
- [SBV⁺13] Srinath Setty, Benjamin Braun, Victor Vu, Andrew J. Blumberg, Bryan Parno, and Michael Walfish. Resolving the conflict between generality and plausibility in verified computation. In *Proceedings of the ACM European Conference on Computer Systems (EuroSys)*, April 2013.
- [Set20] Srinath Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In *Proceedings of the International Cryptology Conference (CRYPTO)*, 2020.
- [SL20] Srinath Setty and Jonathan Lee. Quarks: Quadruple-efficient transparent zkSNARKs. Cryptology ePrint Archive, Report 2020/1275, 2020.
- [SMBW12] Srinath Setty, Richard McPherson, Andrew J. Blumberg, and Michael Walfish. Making argument systems for outsourced computation practical (sometimes). In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, February 2012.
- [Spi96] Daniel A Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Transactions on Information Theory*, 42(6):1723–1731, 1996.
- [SVP⁺12] Srinath Setty, Victor Vu, Nikhil Panpalia, Benjamin Braun, Andrew J. Blumberg, and Michael Walfish. Taking proof-based verified computation a few steps closer to practicality. In *Proceedings of the USENIX Security Symposium*, August 2012.
- [Tha13] Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In *Proceedings of the International Cryptology Conference (CRYPTO)*, 2013.
- [Tha20] Justin Thaler. Proofs, arguments, and zero-knowledge. <http://people.cs.georgetown.edu/jthaler/ProofsArgsAndZK.html>, 2020.
- [Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *Theory of Cryptography Conference (TCC)*, pages 1–18, 2008.
- [WJB⁺17] Riad S. Wahby, Ye Ji, Andrew J. Blumberg, Abhi Shelat, Justin Thaler, Michael Walfish, and Thomas Wies. Full accounting for verifiable outsourcing. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [WSR⁺15] Riad S. Wahby, Srinath Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2015.
- [WTS⁺18] Riad S. Wahby, Ioanna Tzialla, Abhi Shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2018.
- [WYKW20] Chenkai Weng, Kang Yang, Jonathan Katz, and Xiao Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for Boolean and arithmetic circuits. Cryptology ePrint Archive, Report 2020/925, 2020.
- [XZZ⁺19] Tiancheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. Libra: Succinct zero-knowledge proofs with optimal prover computation. In *Proceedings of the International Cryptology Conference (CRYPTO)*, 2019.
- [ZGK⁺17] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. vSQL: Verifying arbitrary SQL queries over dynamic outsourced databases. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2017.
- [ZXZS20] Jiaheng Zhang, Tiancheng Xie, Yupeng Zhang, and Dawn Song. Transparent polynomial delegation and its applications to zero knowledge proof. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2020.

A Description of the polynomial commitment scheme for $t = 2$

Notation. g is a multilinear polynomial with n coefficients. We assume for simplicity that $n = m^2$ for some integer m . Let u denote the coefficient vector of g in the Lagrange basis (equivalently, u is the vector of all evaluations of g over inputs in $\{0, 1\}^{\log n}$). Recalling that $[m] = \{1, \dots, m\}$, we can naturally index entries of u by elements of the set $[m]^2$. As per Section 4, for any input r to g there exist vectors $q_1, q_2 \in \mathbb{F}^m$ such that

$$g(r) = \langle (q_1 \otimes q_2), u \rangle.$$

For each $i \in [m]$, let us view u as an $m \times m$ matrix, and let u_i denote the i th row of this matrix, i.e., $u_i = \{u_{i,j}\}_{j \in [m]}$.

Let $N = \rho^{-1} \cdot m$, and let $\text{Enc}: \mathbb{F}^m \rightarrow \mathbb{F}^N$ denote the encoding function of a linear code with constant rate $\rho > 0$ and constant relative distance $\gamma > 0$. We assume that Enc runs in time proportional to that required to perform $O(N)$ operations over \mathbb{F} . We assume for simplicity that Enc is systematic, since explicit systematic codes with the properties we require are known [Spi96].

Commitment phase. Let $\hat{u} = \{\text{Enc}(u_i)\}_{i \in [m]} \in (\mathbb{F}^N)^m$ denote the vector obtained by encoding each row of u . In the IOP setting, the commitment to u is just the vector \hat{u} , i.e., the prover sends \hat{u} to the verifier, and the verifier is given point query access to \hat{u} . In the derived polynomial commitment scheme in the plain or random oracle model, the commitment to u will be the Merkle-hash of the vector \hat{u} . As with u , we may view \hat{u} as a matrix, with $\hat{u}_i \in \mathbb{F}^N$ denoting the i th row of \hat{u} for $i \in [m]$.

Testing phase. Upon receiving the commitment message, the IOP verifier will interactively test it to confirm that each “row” of u is indeed (close to) a codeword of Enc . We describe this process as occurring in a separate “testing phase” so as to keep the commitment size constant in the plain or random oracle models. In practice, the testing phase can occur during the commit phase, during the evaluation phase, or sometime in between the two.

The verifier sends the prover a random vector $r \in \mathbb{F}^m$, and the prover sends a vector $u' \in \mathbb{F}^m$ claimed to equal the random linear combination of the m rows of u , in which the coefficients of the linear combination are given by r . The verifier reads u' in its entirety.

Next, the verifier tests u' for consistency with \hat{u} . That is, the verifier will pick $\Theta(\lambda)$ random entries of the codeword $\text{Enc}(u') \in \mathbb{F}^N$ and confirm that $\text{Enc}(u')$ is consistent with $v \in \mathbb{F}^N$ at those entries, where v is:

$$\sum_{i=1}^m r_i \hat{u}_i \in \mathbb{F}^N. \tag{6}$$

Observe that, by definition of v (Equation (6)), any individual entry v_j of v can be learned by querying m entries of \hat{u} (we refer to these m entries as the “ j ’th column” of \hat{u}). Meanwhile, since the verifier reads u' in its entirety, \mathcal{V} can compute $\text{Enc}(u')_j$ for all desired $j \in [N]$ in $O(m)$ time.

Evaluation phase. As per Section 7, let $q_1, q_2 \in \mathbb{F}^m$ be such that

$$g(r) = \langle (q_1 \otimes q_2), u \rangle.$$

The evaluation phase is identical to the testing phase, except that r is replaced with q_1 (and the verifier uses fresh randomness to choose the sets of coordinates used for consistency testing). Let $u'' \in \mathbb{F}^m$ denote the vector that the prover sends in this phase, which is claimed to equal $\sum_{i=1}^m q_{1,i} \cdot u_i$. If the prover is honest, then u'' satisfies $\langle u'', q_2 \rangle = \langle (q_1 \otimes q_2), u \rangle$. Hence, if the verifier’s consistency tests all pass in the testing and evaluation phases, the verifier outputs $\langle u'', q_2 \rangle$ as $g(r)$.

Description of polynomial commitment in the language of IOPs. Following standard transformations [Kil92, Mic94, Val08, BCS16], in the actual polynomial commitment scheme, each vector sent by the prover in the IOP is replaced with a Merkle-commitment to that vector, and each query the verifier makes to a vector is answered by the prover along with Merkle-tree authentication path for the answer. Each phase of the scheme can be rendered non-interactive using the Fiat-Shamir transformation [FS86].

Commit phase.

- $\mathcal{P} \rightarrow \mathcal{V}$: a vector $\hat{u} = (\hat{u}_1, \dots, \hat{u}_m) \in (\mathbb{F}^N)^m$. If \mathcal{P} is honest, each “row” \hat{u}_i of \hat{u} contains a codeword in Enc .

Testing phase.

- $\mathcal{V} \rightarrow \mathcal{P}$: a random vector $r \in \mathbb{F}^m$.
- $\mathcal{P} \rightarrow \mathcal{V}$ sends a vector $u' \in \mathbb{F}^m$ claimed to equal $v = \sum_{i=1}^m r_i \cdot u_i \in \mathbb{F}^m$.
- //Now \mathcal{V} probabilistically checks consistency between \hat{u} and u' (\mathcal{V} reads u' in entirety).
- \mathcal{V} : chooses Q to be a random set of size $\Theta(\lambda)$ with $Q \subseteq [N]$. For each $j \in Q$:
 - \mathcal{V} queries all m entries of the corresponding “column” of \hat{u} , namely $\hat{u}_{1,j}, \dots, \hat{u}_{m,j}$.
 - \mathcal{V} confirms that $\text{Enc}(u')_j = \sum_{i=1}^m r_i \cdot \hat{u}_{i,j}$, halting and rejecting if not.

Evaluation phase.

- Let $q_1, q_2 \in \mathbb{F}^m$ be such that $g(r) = \langle (q_1 \otimes q_2), z \rangle$.
- The evaluation phase is identical to the testing phase, except that r is replaced with q_1 (and fresh randomness is used to choose a set Q' of columns for use in consistency checking).
- If all consistency tests pass, then \mathcal{V} outputs $\langle u', q_2 \rangle$ as $g(r)$.

Concrete optimizations to the commitment scheme. Here are optimizations that can reduce the “authentication paths” provided in the testing and evaluation phases by constant factors without affecting the correctness guarantees of the commitment scheme. First, in settings where the evaluation phase will only be run once, the testing phase and evaluation phase can be run in parallel and the same query set Q can be used for both testing and evaluation. This saves roughly a factor of 2 in the proof size. Second, in settings where the commitment is trusted (e.g., applying the polynomial commitment to achieve holography as per Section 5), the testing phase can be omitted.

Soundness analysis for the testing phase. The following claim roughly states that if $\hat{u} = (\hat{u}_1, \dots, \hat{u}_m) \in (\mathbb{F}^N)^m$, then if even a single \hat{u}_i is far from all codewords in Enc , then a random linear combination of the \hat{u}_i ’s is also far from all codewords with high probability.

Claim 5. *Let $\hat{u} = (\hat{u}_1, \dots, \hat{u}_m) \in (\mathbb{F}^N)^m$ and for each $i \in [m]$ let c_i be the closest codeword in Enc to \hat{u}_i . Let E with $|E| \leq (\gamma/4)N$ be a subset of the columns $j \in [N]$ of \hat{u} on which there is even one row $i \in [m]$ such that $\hat{u}_{i,j} \neq c_{i,j}$. With probability at least*

$$1 - (|E| + 1)/|\mathbb{F}| > 1 - N/|\mathbb{F}|$$

over the choice of $r \in \mathbb{F}^m$, $\sum_{i=1}^m r_i \cdot \hat{u}_i$ has distance at least $|E|$ from any codeword in Enc .

Lemma 1. *If the prover passes all of the checks in the testing phase with probability at least*

$$N/|\mathbb{F}| + (1 - \gamma/4)^\ell,$$

then there is a sequence of m codewords c_1, \dots, c_m in Enc such that

$$E := |\{j \in [N]: \exists i \in [m] \text{ such that } c_{i,j} \neq \hat{u}_{i,j}\}| \leq (\gamma/4)N. \tag{7}$$

Proof. Let $d(b, c)$ denote the relative Hamming distance between two vectors $b, c \in \mathbb{F}^N$. Assume by way of contradiction that Equation (7) does not hold. We explain that the prover must fail the verifier’s consistency tests during the testing phase with probability at least $N/|\mathbb{F}| + (1 - \gamma/4)^\ell$.

Recall that v denotes $\sum_{i=1}^m r_i \hat{u}_i$. By Claim 5, the probability over the verifier's choice of r that there exists a codeword a satisfying $d(a, v) \leq \gamma/4$ is less than $N/|\mathbb{F}|$. If no such a exists, then $d(\text{Enc}(u'), v) \geq \gamma/4$. In this event, at least one of the verifier's consistency tests will fail with probability at least $(1 - \gamma/4)^\ell$. □

Completeness and binding of the polynomial commitment scheme. Completeness holds by design.

To argue binding, recall from the analysis of the testing phase that c_i denotes the codeword in Enc that is closest to row i of \hat{u} , and let $w := \sum_{i=1}^m q_{1,i} \cdot c_i$. We show that, if the prover passes the verifier's checks in the testing phase with probability more than $N/|\mathbb{F}| + (1 - \gamma/4)^\ell$ and passes the verifier's checks in the evaluation phase with probability more than $(1 - (3/4)\gamma)^\ell$, then $w = \text{Enc}(u'')$.

If $w \neq \text{Enc}(u'')$, then w and $\text{Enc}(u'')$ are two distinct codewords in Enc and hence they can agree on at most $(1 - \gamma) \cdot N$ coordinates. Denote this agreement set by A . The verifier rejects in the evaluation phase if there is any $j \in Q'$ such that $j \notin A \cup E$, where E is as in Equation (7). $|A \cup E| \leq |A| + |E| \leq (1 - \gamma) \cdot N + (\gamma/4)N = (1 - (3/4)\gamma)N$, and hence a randomly chosen column $j \in [N]$ is in $A \cup E$ with probability at most $1 - (3/4)\gamma$. It follows that u'' will pass the verifier's consistency checks in the evaluation phase with probability at most $(1 - (3/4)\gamma)^\ell$.

In summary, we have shown that if the prover passes the verifier's checks in the commitment phase with probability at least $N/|\mathbb{F}| + (1 - \gamma/4)^\ell$, then, in the following sense, the prover is *bound* to the polynomial g^* whose coefficients in the Lagrange basis are given by $c_{1,1}, \dots, c_{m,m}$, where $c_i \in \mathbb{F}^N$ denotes the closest codeword to row i of the vector \hat{u} sent in the commitment phase: on evaluation query r , the verifier either outputs $g^*(r)$, or else rejects in the evaluation phase with probability at least $1 - (1 - (3/4)\gamma)^\ell$. Furthermore, the polynomial commitment scheme provides standard extractability properties. This is because with the transformation of [Kil92, Mic94, Val08, BCS16], informally speaking, given a prover that convinces a verifier, one can extract the IOP proof strings from the prover; row-by-row decoding of the prover's proof string in the commitment phase provides the coefficients of the multilinear polynomial that the prover is bound to.