Shorter and Faster Post-Quantum Designated-Verifier zkSNARKs from Lattices

Yuval Ishai

Technion

Haifa, Israel

yuvali@cs.technion.ac.il

Hang Su

University of Virginia
Charlottesville, VA, USA
hs2nu@virginia.edu

David J. Wu

University of Texas at Austin
Austin, TX, USA
dwu4@cs.utexas.edu

ABSTRACT

Zero-knowledge succinct arguments of knowledge (zkSNARKs) enable efficient privacy-preserving proofs of membership for general NP languages. Our focus in this work is on *post-quantum* zkSNARKs, with a focus on minimizing proof size. Currently, there is a $1000\times$ gap in the proof size between the best pre-quantum constructions and the best post-quantum ones. Here, we develop and implement new *lattice-based* zkSNARKs in the designated-verifier preprocessing model. With our construction, after an initial preprocessing step, a proof for an NP relation of size 2^{20} is just over 16 KB. Our proofs are $10.3\times$ shorter than previous post-quantum zkSNARKs for general NP languages. Compared to previous lattice-based zkSNARKs (also in the designated-verifier preprocessing model), we obtain a $42\times$ reduction in proof size and a $60\times$ reduction in the prover's running time, all while achieving a much higher level of soundness. Compared to the shortest pre-quantum zkSNARKs by Groth (Eurocrypt 2016), the proof size in our lattice-based construction is $131\times$ longer, but *both* the prover and the verifier are *faster* (by $1.2\times$ and $2.8\times$, respectively).

Our construction follows the general blueprint of Bitansky et al. (TCC 2013) and Boneh et al. (Eurocrypt 2017) of combining a linear probabilistically checkable proof (linear PCP) together with a linear-only vector encryption scheme. We develop a concretely-efficient lattice-based instantiation of this compiler by considering quadratic extension fields of moderate characteristic and using linear-only vector encryption over rank-2 module lattices.

CCS CONCEPTS

• Security and privacy → Cryptography; Privacy-preserving protocols.

KEYWORDS

zkSNARKs; succinct arguments; lattice-based SNARKs

ACM Reference Format:

Yuval Ishai, Hang Su, and David J. Wu. 2021. Shorter and Faster Post-Quantum Designated-Verifier zkSNARKs from Lattices. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '21, November 15–19, 2021, Virtual Event, Republic of Korea

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8454-4/21/11...\$15.00

<https://doi.org/10.1145/3460120.3484572>

(CCS '21), November 15–19, 2021, Virtual Event, Republic of Korea. ACM, New York, NY, USA, 23 pages. <https://doi.org/10.1145/3460120.3484572>

1 INTRODUCTION

A zero-knowledge proof of knowledge [71] for an NP relation $\mathcal{R}$ enables a prover to convince a verifier that a statement is true without revealing anything more about the statement. In a zero-knowledge succinct argument of knowledge (zkSNARK) [69, 79, 84], we additionally require that the proof consist of a single message π from the prover to the verifier, and moreover, that the length of the proof π and the verification complexity be sublinear (ideally, polylogarithmic) in the size of the circuit computing $\mathcal{R}$. Zero-knowledge SNARKs have applications to delegating and verifying computations [107] and for constructing privacy-preserving cryptocurrencies [18]. In the last few years, there have been numerous works studying constructions from different assumptions and on optimizing the asymptotic and concrete efficiency of zkSNARKs (e.g., [5, 15, 21, 30, 34, 43, 44, 46, 47, 49, 63, 65, 75, 81, 87, 99, 101, 106, 111]).

Post-quantum zkSNARKs. Many existing constructions of practical zkSNARKs for NP rely on group-based and pairing-based assumptions [30, 34, 43, 46, 64, 74, 75, 83, 87, 97, 99] and are insecure against quantum adversaries. Several recent works have introduced new concretely-efficient post-quantum zkSNARKs based on cryptographic hash functions [5, 15, 21, 28, 47, 101] or lattice-based assumptions [65]. However, compared to their pre-quantum analogs, current post-quantum constructions have substantially longer proofs. As a point of comparison, in the most succinct pre-quantum construction by Groth [75], proofs are just 128 bytes while those in the most succinct post-quantum constructions [21, 28, 47, 101] are generally $1000\times$ longer (see Table 1). The increase in parameter sizes is not entirely surprising since a similar, although smaller, gap exists between the sizes of group-based pre-quantum signatures [33, 93] and hash-based [27, 45] or lattice-based post-quantum signatures [55, 60].

This work: lattice-based designated-verifier zkSNARKs. Our focus in this work is on new approaches for constructing shorter (and faster) post-quantum zkSNARKs from *lattice-based assumptions*. Like recent works [21, 47, 64, 97–99], we focus on the NP-complete language of rank-1 constraint satisfiability (R1CS), which generalizes Boolean and arithmetic circuit satisfiability and enjoys efficient compiler support from other program representations [19, 20, 24, 42, 87, 100]. While recent works have introduced post-quantum zkSNARKs from lattice-based assumptions [31, 32, 65, 86], to our knowledge, only the construction of Gennaro et al. [65] has been implemented. Its proof sizes are significantly worse compared to

alternative post-quantum constructions based on interactive oracle proofs (IOPs) and the Fiat-Shamir heuristic (e.g., 640 KB for the lattice-based approach [65] vs. 169 KB for an IOP-based approach [21]). The prover time for current lattice-based instantiations is also over $10\times$ worse than the alternative constructions.

Similar to the previous lattice-based constructions, we design our zkSNARKs in the designated-verifier preprocessing model where there is an (expensive but practically feasible) setup algorithm that samples public parameters and a *secret* verification key (needed to verify proofs). While the designated-verifier model is a relaxation of the conventional setting of zkSNARKs, it nonetheless suffices for applications to verifiable computation and other privacy-preserving protocols.

Our results. Our main result is a new designated-verifier zkSNARK from lattice-based assumptions where the proof size (for verifying an R1CS instance of size 2^{20}) is just over 16 KB. This is a $10.3\times$ reduction in proof size compared to Aurora [21], a post-quantum IOP-based SNARK with short proofs. If we restrict our attention to post-quantum zkSNARKs with *sublinear* verification, our construction is $13.1\times$ shorter than Fractal [47]. Compared to the specialized ethSTARK [101] construction, a post-quantum STARK [15] for verifying a STARK-friendly hash chain, our proofs are $7.7\times$ shorter. Finally, compared to the lattice-based construction of Gennaro et al. [65], our zkSNARKs are $42.1\times$ shorter. However, there remains a large gap ($131\times$) compared to the shortest *pre-quantum* zkSNARK by Groth [75]. We refer to Table 1 for the full comparison and describe our experimental setup in detail in Section 4.3.

The prover and verifier complexities of our new zkSNARK compare favorably with other post-quantum schemes for verifying general NP computations. Our construction is over $4.5\times$ faster for the prover compared to Aurora and Fractal on R1CS instances of similar size. Compared to the Gennaro et al. lattice-based candidate [65], our construction is $60\times$ faster for the prover and $5.1\times$ faster for the verifier. Compared to the pre-quantum pairing-based construction of Groth [75], our construction is $1.2\times$ faster for the prover and $2.8\times$ faster for the verifier. Using an alternative instantiation of our construction with *longer* proofs (20.8 KB vs. 16.4 KB), our construction is $1.4\times$ faster than the pairing-based construction for the prover and $7.9\times$ faster for the verifier.

Another appealing feature of our lattice-based zkSNARK is the simplicity of proof verification: it only requires evaluating a matrix-vector product followed by a few simple arithmetic tests. This leads to a concretely faster verification procedure compared to previous constructions (which either required pairing computations or multiple invocations of a cryptographic hash function) and also makes our construction well-suited for verifying proofs on lightweight or energy-constrained devices that can only support a limited number of arithmetic operations. We note that for verifying small computations (e.g., an R1CS system with 2^{14} constraints) with a larger soundness error (e.g., $1/128$), the group-based designated-verifier SNARK of Barta et al. [10] can plausibly achieve even faster verification. However, this comes at the price of needing a long CRS and a high prover cost (both scale *quadratically* with the size of the R1CS system).

Further improvements to the proof size and prover complexity are possible if we relax zero knowledge. For instance, a variant of

our construction that is sound but *not* provably zero knowledge is over $2.3\times$ faster for the prover than the pairing-based construction of Groth and has a proof size of 11.1 KB (for verifying an R1CS instance with 2^{20} constraints). This construction is suitable for applications that do not require zero knowledge, or alternatively, can tolerate a small amount of leakage. Note that while we do not *prove* zero knowledge of this variant, the construction can still provide full zero knowledge assuming that the underlying information-theoretic building block we use (linear PCPs) remains zero knowledge in the presence of leakage. We provide more details in the full version of this paper [78]. We leave the question of determining whether the linear PCPs we use (see Appendix B) or variants thereof satisfy the stronger notion of zero knowledge with leakage to future work. In Section 4.3, we show trade-offs between the number of bits of *provable* zero knowledge provided by our construction and its concrete efficiency.

Compared to other post-quantum constructions based on “MPC-in-the-head” [5, 28, 77], the “GKR” approach [70, 109, 110], or constructions tailored for specific computations [101], our prover times are generally higher. For instance, compared to Ligerio [5], our prover is about $1.8\times$ more expensive, but our proof size and verification times are over $874\times$ better (see Table 1). If we compare against the ethSTARK scheme [101] for verifying a STARK-friendly hash chain [4, 25], the running time of the ethSTARK prover is $15\times$ smaller than the running time of our prover for verifying the R1CS representation of the same computation. In general, these alternative approaches typically enjoy smaller concrete prover costs, but often have longer proofs or higher verification costs when considering general, *unstructured* computations. We provide more details and comparisons with other zkSNARKs in Section 5.

The IOP-based constructions have the advantage of being *publicly-verifiable* and *transparent*. Our scheme is designated-verifier and requires an expensive trusted setup. For verifying R1CS systems with 2^{20} constraints, we need to sample a CRS of size 5.3 GB which takes 37 minutes. An alternative instantiation of our construction over a larger finite field reduces the CRS size to 1.9 GB and the setup time to 15 minutes. This leads to a modest increase in proof size from 16.4 KB to 20.8 KB (see Table 1).

Limitations of our construction. While our lattice-based zkSNARK achieve better succinctness compared to other post-quantum zkSNARK candidates, they have several limitations that give rise to natural directions for improvement. We highlight some of these here:

- **Reusable soundness and public verification.** As noted above, our lattice-based zkSNARK is in the *designated-verifier* model where a *secret* verification key is needed to verify proofs. Moreover, like existing lattice-based designated-verifier zkSNARKs [32, 65], our construction does not provide *reusable soundness* where soundness holds even against a malicious prover who can make an arbitrary polynomial number of queries to the verification oracle. Constructing a lattice-based zkSNARK with comparable concrete efficiency and reusable soundness is an interesting direction. As we discuss in greater detail in the full version of this paper [78], even without a provable notion of reusable soundness, our construction still suffices for some applications to verifiable computation. In particular, breaking soundness requires the

Scheme	Structure	PQ	TP	PV	R1CS Size	Size		Time		
						CRS	Proof	Setup	Prover	Verifier
Groth [75]	Pairings	○	○	●	2^{16}	12.4 MB	128 B	5.6 s	5.5 s	3.3 ms
					2^{20}	199 MB	128 B	72 s	79 s	3.4 ms
Gennaro et al. [65]*	Lattices	●	○	○	2^{16}	17.3 MB [†]	640 KB	167 s	235 s	3.5 ms
Ligero [5]	Random Oracle	●	●	●	2^{16}	—	4.3 MB	—	2.5 s	1.3 s
					2^{20}	—	14 MB	—	38 s	22 s
Aurora [21]	Random Oracle	●	●	●	2^{16}	—	121 KB	—	18 s	380 ms
					2^{20}	—	169 KB	—	304 s	6.3 s
Fractal [47] [‡]	Random Oracle	●	●	●	2^{16}	1.4 GB	178 KB	12 s	21 s	8.3 ms
					2^{19}	11 GB	215 KB	116 s	184 s	9.5 ms
ethSTARK [101] [§]	Random Oracle	●	●	●	2^{16}	—	77.5 KB	—	0.3 s	2.5 ms
					2^{20}	—	127 KB	—	4.5 s	4.1 ms
This work (Shorter Proofs)	Lattices	●	○	○	2^{16}	191 MB	15.2 KB	88 s	3.9 s	0.69 ms
					2^{20}	5.3 GB	16.4 KB	2240 s	68 s	1.2 ms
This work (Shorter CRS)	Lattices	●	○	○	2^{16}	104 MB	19.9 KB	53 s	3.4 s	0.37 ms
					2^{20}	1.9 GB	20.8 KB	877 s	56 s	0.43 ms

*As we discuss in Appendix E (Remark E.4), the parameter instantiation proposed in Gennaro et al. [65] only provides 15 bits of provable soundness. If we use parallel repetition to amplify to 128-bits of soundness, then all of the parameters should be scaled by a factor of 8.5X. In the table, we report the numbers as they were presented in the original paper. Their work also does not provide measurements for instances with more than 2^{16} gates.

[†]Gennaro et al. [65] do not report the CRS size for an instance of size 2^{16} . We estimate the size by doubling the size of the CRS for an instance of size 2^{15} .

[‡]The “Setup” time and “CRS” size for Fractal refers to the running time of the indexer and the size of the resulting proving state. Our system ran out of memory when running Fractal on an R1CS instance of size 2^{20} . Thus, we report the results for an instance of size 2^{19} instead.

[§]Performance numbers for ethSTARK are based on verifying a Rescue hash chain [4, 25] (specifically Rescue₁₂₂). The length of the hash chain is chosen to match the size of the corresponding R1CS system. Specifically, we use hash chains of length 270 and 4200 to represent R1CS systems with 2^{16} and 2^{20} constraints, respectively (see Section 4.3 for more detail). The ethSTARK implementation [102] does not currently support verifying general computations.

Table 1: Concrete performance comparison of our zkSNARK to the pairing-based construction of Groth [75] and several recent post-quantum zkSNARKs with polylogarithmic-size proofs. For each scheme, we report the running time and parameter sizes for an R1CS instance with 2^{16} and 2^{20} constraints. We measure the running times for an R1CS instance over each scheme’s preferred field. With the exception of the Gennaro et al. [65] construction, all measurements are taken on the *same* system (see Section 4.3 for details of our setup). For our scheme, we consider two different parameter settings. The “Shorter Proofs” instantiation works over the field $\mathbb{F}_{p^2}$ where $p = 2^{13} - 1$ and the “Shorter CRS” instantiation works over the field $\mathbb{F}_{p^2}$ where $p = 2^{19} - 1$ (see Table 2 for the lattice parameters in these instantiations). The “PQ” column specifies whether the construction is post-quantum (●) or pre-quantum (○), the “TP” column specifies whether the construction has a transparent setup (●) or relies on a trusted setup (○), and the “PV” column specifies whether the scheme is publicly-verifiable (●) or designated-verifier (○).

prover to submit a *super-constant* number of “bad” proofs to the verifier. This means that the verifier is able to detect a malicious prover trying to attack the scheme (this is reminiscent of the notion of *covert security* from [8]). Alternatively, these “selective failure” attacks can be avoided altogether if the verifier does not reveal whether a proof is valid or not to the prover.

More generally, it is a fascinating question to construct *publicly-verifiable* lattice-based zkSNARKs with comparable concrete efficiency. Existing lattice-based constructions [11, 37] that are publicly verifiable have an expensive verifier (i.e., the verifier runs in time linear in the size of the underlying NP relation). We refer to Section 5 for further comparison with related work.

- **Field characteristic.** In this work, we consider lattice-based zkSNARKs for R1CS systems over finite fields of moderate characteristic (i.e., between 12 and 20 bits). Specifically, we consider quadratic extension fields, which enable a number of concrete optimizations (see Sections 1.2 and 4.3).

For some applications, it may be helpful to consider R1CS systems over fields of higher characteristic. For example, validity of a 32-bit addition gate (on 32-bit inputs) can be expressed as a single R1CS constraint over any finite field with characteristic

$p > 2^{32}$. Thus, when verifying computations that make extensive use of 32-bit or 64-bit integer arithmetic, it can be advantageous to encode the computation in an R1CS system over a higher characteristic field. Both the pairing-based construction [75] as well as the hash-based constructions [5, 21, 47, 75] operate over base fields of high characteristic (i.e., at least 128 bits). The hash-based constructions [5, 21, 47] also efficiently support R1CS systems on high-degree extensions of the binary field, or more generally, any field that supports efficient fast Fourier transforms.

While our lattice-based instantiation can in principle support fields of higher characteristic, doing so will require using larger lattice parameters, which in turn, increases the proof size. Moreover, some of our concrete optimizations (e.g., implementing all arithmetic operations using 128-bit integer arithmetic) can no longer be applied when the field characteristic increases. We refer to Section 4.3 and Fig. 2 for more discussion on how the field characteristic affects the concrete efficiency of our scheme.

1.1 Background

The basis of our work is the compiler of Bitansky et al. [30] (also implicit in the work of Gennaro et al. [64]), and more specifically, the

vector generalization by Boneh et al. [31]. These works provided a general template for constructing SNARKs in the preprocessing model by combining a “linear PCP” with a “linear-only” encryption scheme. A linear PCP [76] for an NP language $\mathcal{L}$ is defined by a linear oracle $\pi: \mathbb{F}^\ell \rightarrow \mathbb{F}$ over a finite field $\mathbb{F}$. On input a statement x , a verifier can submit a query matrix $Q \in \mathbb{F}^{\ell \times k}$ to the oracle and obtain the responses $\mathbf{a} \leftarrow Q^T \pi \in \mathbb{F}^k$. Based on the responses, the verifier decides whether to accept or reject. We refer to k as the number of queries and ℓ as the query length of the linear PCP. The linear PCP is sound if for a false statement $x \notin \mathcal{L}$ and any proof vector $\pi \in \mathbb{F}^\ell$, the probability that the verifier accepts is negligible (where the probability is taken over the choice of Q). Concretely-efficient 4-query linear PCPs for R1CS can be constructed using the quadratic arithmetic programs (QAPs) introduced by Gennaro et al. [64]. QAPs are the basis for the most succinct pairing-based preprocessing zkSNARKs [20, 30, 64, 75, 87].

To obtain a preprocessing zkSNARK for $\mathcal{L}$ from a linear PCP for $\mathcal{L}$, the Bitansky et al. compiler encrypts the linear PCP queries (i.e., the entries of Q) using a “linear-only” encryption scheme and publishes the resulting ciphertexts as part of the common reference string (CRS). As the name suggests, a linear-only encryption scheme is an encryption scheme that only supports linear homomorphism (i.e., it is possible to add ciphertexts, but no other homomorphic operation on ciphertexts is supported). Given the encrypted queries, the prover can homomorphically compute the encrypted responses $\mathbf{a} = Q^T \pi$. Here, the linear-only property restricts the prover to linear strategies and by semantic security, the prover’s choice of linear combination is essentially independent of the linear PCP queries. This binds the prover to respect the constraints of the linear PCP model. To verify the proof, the verifier decrypts the encrypted responses and evaluates the linear PCP verification procedure. This yields a designated-verifier preprocessing SNARK. For zero knowledge, it suffices that the linear PCP be *honest-verifier* zero knowledge and the linear-only encryption scheme be “re-randomizable” (i.e., ciphertexts output by the homomorphic evaluation are computationally indistinguishable from fresh ciphertexts).

Lattice-based instantiations of Bitansky et al. Gennaro et al. [65], following Boneh et al. [31, 32], introduced candidate linear-only encryption schemes based on lattices. In these works, the underlying linear-only encryption scheme is adapted from basic Regev encryption [91]. For our purposes, a Regev-based encryption of a value $x \in \mathbb{Z}_p$ is a pair $(\mathbf{a}, \mathbf{c})$ where $\mathbf{a} \in \mathbb{Z}_q^n$ and $\mathbf{c} = \mathbf{s}^T \mathbf{a} + pe + x \in \mathbb{Z}_q$, where $\mathbf{s} \in \mathbb{Z}_q^n$ is the secret key, $e \in \mathbb{Z}_q$ is an error term, and n, q are lattice parameters. Observe that this scheme is linearly homomorphic: if $(\mathbf{a}_1, \mathbf{c}_1)$ and $(\mathbf{a}_2, \mathbf{c}_2)$ encrypt values x_1, x_2 , respectively, then $(\mathbf{a}_1 + \mathbf{a}_2, \mathbf{c}_1 + \mathbf{c}_2)$ encrypts the value $x_1 + x_2 \bmod p$, albeit with slightly larger error. As long as the error magnitude in the final ciphertext is less than $q/(2p)$, decryption succeeds.

Gennaro et al. [65] provided the first lattice-based implementation of the Bitansky et al. compiler using Regev encryption.¹ Compared to the best pairing-based constructions that followed a similar methodology [75], the lattice-based implementation is significantly less efficient. For an R1CS instance of size 2^{16} , the

proof size is 640 KB, over 5000 $\times$ larger than the pairing-based construction of Groth [75]; similarly, the prover time for a similar-sized instance is roughly 40 $\times$ slower than the pairing-based analog. In fact, as we discuss in Appendix E, because Gennaro et al. consider linear PCPs over a *small* field $\mathbb{F}$ ($\log |\mathbb{F}| = 32$), the specific parameter instantiation they consider provides at most 15 bits of provable soundness. Working over a larger field or using parallel repetition for soundness amplification would incur even more overhead.

Lattice parameter sizes. The main obstacle to the concrete efficiency of lattice-based zkSNARKs following the Bitansky et al. compiler [30] is the size of the lattice parameters. The length of a QAP for an R1CS instance with N constraints over a finite field $\mathbb{F}$ is $O(N)$, and the soundness error is $O(N/|\mathbb{F}|)$. This means we need to work over a field $\mathbb{F}$ where $|\mathbb{F}| > N$, and we need a linear-only encryption scheme over $\mathbb{F}$ that supports $O(N)$ homomorphic operations. In the Gennaro et al. construction [65], they consider a prime field $\mathbb{F}_p$ where $p > N$. For correctness then, the modulus q for a Regev-based encoding must satisfy $q > 2p^2N$. Zero knowledge adds a further multiplicative factor of 2^κ where κ is a statistical security parameter.

To achieve 128 bits of soundness, one approach is to set $p > 2^{128}N$. If we take $q \approx 2^{300}$ (and a typical error distribution), then the lattice dimension n needs to be at least 10^4 at the 128-bit security level (based on [1]). A *single* ciphertext is over 350 KB in this setting. This is a lower bound on the proof size since in the basic instantiation, the proof contains at least one ciphertext for each linear PCP response.

Alternatively, instead of working over a large field, we can work over a small field $\mathbb{F}_p$ where $p \approx N$ and amplify soundness through parallel repetition. For instance, if we take $p \approx 2^{20}$ and $q \approx 2^{100}$, then a single Regev ciphertext is roughly 45 KB. However, soundness amplification increases the proof size (and all other metrics), again leading to parameter sizes that are significantly worse than non-lattice-based zkSNARKs. The scheme of Gennaro et al. [65] considers a finite field of size 2^{32} *without* soundness amplification, and so their concrete instantiation provides very few bits of provable soundness (see Appendix E and Remark E.4). But even with this choice of parameters, the proof size in their construction is already 640 KB.

1.2 Technical Overview

The primary enablers of our concretely-efficient lattice-based zkSNARK are (1) using *vector encryption* [89] instead of vanilla Regev encryption as our linear-only encryption scheme; and (2) working over *extension fields of moderate characteristic*. We provide an overview of our techniques and construction here.

Vector encryption. Our starting point in this work is the adaptation of the Bitansky et al. compiler using linear-only *vector encryption* introduced by Boneh et al. [31]. As the name suggests, a vector encryption scheme (over a field $\mathbb{F}$) supports encrypting a vector of field elements. Instead of encrypting each entry in the linear PCP query matrix $Q \in \mathbb{F}^{\ell \times k}$ separately, the Boneh et al. compiler encrypt rows of Q . The proof then consists of a single ciphertext encrypting the vector of linear PCP responses. The advantage of this approach is that we can take advantage of amortization to

¹Gennaro et al. used square span programs [53] instead of QAPs as the underlying linear PCP, but this distinction is not important for the main discussion here.

reduce the ciphertext expansion for lattice-based vector encryption. In more detail, with vanilla Regev encryption, the overhead of encrypting a single $\mathbb{Z}_p$ value is $O(n)$, where n is the lattice dimension. Using the extension by Peikert et al. [89], we can encrypt a vector of ℓ $\mathbb{Z}_p$ -values with a ciphertext containing $(n + \ell)$ $\mathbb{Z}_q$ -elements. This approach confers several improvements for concrete efficiency:

- **Soundness amplification:** We can now amplify soundness of the linear PCP using parallel repetition (i.e., using multiple independent sets of linear PCP queries). This increases the dimensions of the vectors we encrypt, but using the Peikert et al. vector encryption scheme, the overhead is *additive* in the dimension rather than multiplicative (as with vanilla Regev encryption).
- **Number of lattice ciphertexts:** For an encryption scheme to plausibly satisfy the “linear-only” property, the ciphertext space must be sparse, and in particular, the adversary should not be able to obviously sample a valid ciphertext *without* knowledge of the corresponding plaintext value. The heuristic from earlier works [30, 31, 64] is to use “double encryption” where a valid ciphertext encrypting a message x consists of a pair of independent ciphertexts encrypting x (Gennaro et al. [65] also use a variant of this encoding method). In Section 3.3, we describe an alternative approach to sparsify the ciphertext space by first embedding the plaintext vector $\mathbf{v} \in \mathbb{Z}_p^\ell$ within a (secret) subspace $T \subseteq \mathbb{Z}_p^{\ell+\tau}$ and then encrypting the embedded vector $\mathbf{v}' \in \mathbb{Z}_p^{\ell+\tau}$. Here, τ is a “sparsification” parameter. A ciphertext is valid only if it encrypts an element of the subspace T . When T has negligible density in $\mathbb{Z}_p^{\ell+\tau}$, we conjecture that an adversary (that does not know T) cannot obviously sample a valid ciphertext without knowledge of the corresponding plaintext vector. The advantage of this approach is that it incurs a small *additive* overhead on ciphertext/proof size rather than a multiplicative one. Thus, using vector encryption, we can encrypt a vector of plaintext values using a *single* lattice ciphertext (and still plausibly prevent oblivious sampling of ciphertexts).

Reducing ciphertext size with modulus switching. Homomorphic operations on lattice ciphertexts increase the noise in the ciphertexts. To ensure decryption correctness, the ciphertext modulus q must be large enough to accommodate the accumulated noise from the homomorphic operations. The modulus switching technique developed in the context of fully homomorphic encryption [3, 40, 41, 52, 56] provides a way to reduce the size of the ciphertexts *after* performing homomorphic operations. Specifically, modulus switching takes a ciphertext with respect to a modulus q and scales it down to a new ciphertext with respect to a modulus $q' < q$ (while preserving decryption correctness). This technique applies to most Regev-based encryption schemes, including the vector encryption scheme we use. In our specific setting, after the prover homomorphically computes the encrypted vector of linear PCP responses, the prover applies modulus switching to the resulting ciphertext. For our parameter settings, this yields a $2\times$ to $3\times$ reduction in ciphertext size (and correspondingly, in proof size).

Instantiating our vector encryption scheme over $\mathbb{F}_p$ using a 23-bit characteristic p yields a zkSNARK where the proof size is 27 KB and the CRS size is 9.6 GB (for verifying R1CS instances with 2^{20} constraints). Using a larger 28-bit characteristic, the proof size increases to 29 KB and the CRS size decreases to 2.7 GB for

the same setting. Without modulus switching, the proof sizes for these two settings are 66 KB and 72 KB, respectively. While these basic instantiations already improve on previous post-quantum zkSNARKs in terms of proof size, the improvements come at the expense of needing a very large CRS. Below, we show how to use extension fields to obtain instantiations with a shorter CRS and a shorter proof.

Extension fields of moderate characteristic. The second ingredient in our construction is a way to reduce the lattice parameters themselves by considering linear PCPs over extension fields of moderate characteristic. The key observation we make is that the size of the modulus q (and other lattice parameters) scale with the plaintext modulus (i.e., the field *characteristic*) but *not* necessarily the *size* of the field. To take advantage of this, we first note that linear PCPs based on QAPs are agnostic to the choice of the field, and work equally well over extension fields $\mathbb{F}_{p^k}$. We develop two instantiations of this approach:

- **Compile linear PCPs over $\mathbb{F}_{p^k}$ to $\mathbb{F}_p$:** Our first instantiation shows how to compile a linear PCP over $\mathbb{F}_{p^k}$ to a zkSNARK using linear-only vector encryption over the base field $\mathbb{F}_p$ (i.e., the same encryption scheme from above). To do so, we first show how to transform a linear PCP over $\mathbb{F}_{p^k}$ to a linear PCP over $\mathbb{F}_p$. The transformation increases the query length and the number of queries by a factor of k , and relies on the fact that $\mathbb{F}_{p^k}$ -operations correspond to linear transformations over the vector space $\mathbb{F}_p^k$. We describe our construction in Section 3.1. For concrete efficiency reasons, we focus exclusively on quadratic extensions. Using one instantiation of this approach, we obtain a construction with shorter proofs (21 KB) and a shorter CRS (3.8 GB) compared to working over the prime field.² With a longer CRS (10.5 GB), we can bring the proof size down to just 16 KB.
- **Vector encryption over extension fields.** We next consider a *direct* compilation from linear PCPs over the extension field to a zkSNARK using a linear-only vector encryption scheme whose plaintext space coincides with the extension field. To do so, we generalize our variant of the Peikert et al. [89] encryption scheme to operate over the cyclotomic ring $R = \mathbb{Z}[x]/(x^2 + 1)$. In this case, the plaintext space is $R_p = R/pR$. When $p = 3 \bmod 4$, $R_p \cong \mathbb{F}_{p^2}$. Under the conjecture that the vector encryption scheme is linear-only over R_p , this gives a direct compilation from a linear PCP over a quadratic extension $\mathbb{F}_{p^2}$ to a zkSNARK over $\mathbb{F}_{p^2}$. By relying on linear PCPs and linear-only vector encryption over the quadratic extension, we obtain a zkSNARK with similar proof size as the above construction, but with a $2\times$ reduction in the CRS size (previously incurred by transforming the linear PCP from $\mathbb{F}_{p^2}$ to $\mathbb{F}_p$). We show the concrete performance in Table 1 and in Section 4.3. Leveraging encryption schemes over extension fields and higher-rank modules has also been useful for improving the asymptotic and concrete efficiency of other lattice-based constructions [66–68].

Parameter selection. In this work, we consider quadratic extension fields with two different characteristics: (1) $p = 2^{13} - 1$ which

²Even though the linear PCP transformation doubles the query length of the linear PCP, working over the extension field allows us to achieve the same level of soundness with fewer parallel repetitions, and *reduces* the overall size of the CRS.

yields a construction with shorter proofs but a longer CRS; and (2) $p = 2^{19} - 1$ which yields a construction with a shorter CRS and slightly longer proofs. We choose p of the form $2^t - 1$ so that $\mathbb{F}_{p^2}^*$ has a multiplicative subgroup of order 2^{t+1} (i.e., the subgroup of 2^{t+1} -th roots of unity). This enables us to take advantage of fast Fourier transforms (FFT) to implement the linear PCP prover [20]. Note that when p is sufficiently small (e.g., $p = 2^{13} - 1$), the extension field does not contain a sufficiently-large subgroup of roots of unity to directly leverage power-of-two FFTs for the linear PCP. In Section 4.1, we describe a simple approach using multiple small power-of-two FFTs on different cosets of the roots of unity that still enables an efficient implementation of the linear PCP prover.

Working over extension fields also allows us to use a smaller ciphertext modulus q in the lattice-based encryption scheme. When $q < 2^{128}$, we can use compiler intrinsic types for 128-bit integer arithmetic for our computations. This is significantly faster than using multi-precision arithmetic or even fixed-precision arithmetic over slightly larger integers. We provide more discussion and microbenchmarks in Section 4.2.

Zero knowledge and circuit privacy. As noted above, the Bitansky et al. compiler yields a zero-knowledge SNARK if the underlying linear PCP is honest-verifier zero knowledge and the linear-only encryption scheme is re-randomizable. However, the lattice-based schemes are not directly re-randomizable (due to the accumulation of noise through homomorphic operations). In this work, we show that a weaker notion of circuit privacy [66] suffices to argue zero knowledge for the SNARK (i.e., the ciphertext obtained from taking a linear combination of ciphertexts hide the coefficients of the linear combination). Using noise smudging [7, 66, 85] and the module learning with errors assumption (MLWE) [40, 80], it is straightforward to augment our linear-only vector encryption scheme to provide circuit privacy. We give the details in Section 3.3. We additionally note in the full version of this paper [78] that even without circuit privacy, a direct compilation from a linear PCP satisfying honest-verifier zero knowledge to a zkSNARK can still provide full zero knowledge if the underlying linear PCP remains zero knowledge given some additional information on the linear PCP coefficients. This variant without provable zero knowledge enables a further 30-40% reduction in prover time and a 45-50% reduction in proof size.

Implementation and evaluation. In Section 4, we describe our implementation of our lattice-based zkSNARK. We provide a comprehensive evaluation of the different trade-offs in parameter sizes and computational costs for the different settings described here. We also give fine-grained microbenchmarks of the different components of our system in Section 4.3. Finally, we conclude with additional comparisons against other zkSNARK candidates in Section 5.

2 PRELIMINARIES

Throughout this work, we write λ (oftentimes implicitly) to denote the security parameter. For a positive integer $n \in \mathbb{N}$, we write $[n]$ to denote the set $\{1, \dots, n\}$. We write $\{x_i\}_{i \in [n]}$ to denote the ordered multi-set of values $x_1, \dots, x_n$. We will typically use bold lowercase letters (e.g., $\mathbf{v}$, $\mathbf{w}$) to denote vectors and bold uppercase

letters (e.g., $\mathbf{A}$, $\mathbf{B}$) to denote matrices. For a vector $\mathbf{v} \in \mathbb{Z}_p^n$, we will use non-boldface letters to refer to its components; namely, we write $\mathbf{v} = (v_1, \dots, v_n)$. For a vector $\mathbf{v} \in \mathbb{R}^n$, we write $\|\mathbf{v}\|_\infty$ to denote the ℓ_∞ norm of $\mathbf{v}$. For a finite set S , we write $x \xleftarrow{\mathcal{R}} S$ to denote that x is sampled uniformly from S . For a distribution $\mathcal{D}$, we write $x \leftarrow \mathcal{D}$ to denote that x is sampled from $\mathcal{D}$.

We say that a function f is negligible in λ if $f(\lambda) = o(1/\lambda^c)$ for all $c \in \mathbb{N}$; we denote this $f(\lambda) = \text{negl}(\lambda)$. We write $\text{poly}(\lambda)$ to denote a function bounded by a fixed polynomial in λ . We say an event happens with negligible probability if the probability that the event occurs is negligible, and that it happens with overwhelming probability if its complement occurs with negligible probability. We say an algorithm $\mathcal{A}$ is efficient if it runs in probabilistic polynomial time in the length of its input. We say that two families of distributions $\mathcal{D}_1 = \{\mathcal{D}_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}_2 = \{\mathcal{D}_{2,\lambda}\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable if no efficient adversary can distinguish samples from $\mathcal{D}_1$ and $\mathcal{D}_2$ except with negligible probability. We say that $\mathcal{D}_1$ and $\mathcal{D}_2$ are statistically indistinguishable if the statistical distance between $\mathcal{D}_1$ and $\mathcal{D}_2$ is negligible; we denote this by writing $\mathcal{D}_1 \stackrel{s}{\approx} \mathcal{D}_2$. For an algorithm $\mathcal{A}$, we write $\mathcal{A}(x; r)$ to denote the output of running $\mathcal{A}$ on input x and randomness r . In settings where we do not need to specify the randomness explicitly, we write $\mathcal{A}(x)$ to denote the output distribution of $\mathcal{A}$ on input x where the randomness is drawn from the uniform distribution. We recall additional preliminaries, including the formal definition of linear PCPs and zkSNARKs in Appendix A.

3 LATTICE-BASED SUCCINCT ARGUMENTS

In this section, we introduce the main information-theoretic building block (linear PCPs over extension fields) and the cryptographic compiler (linear-only vector encryption) that underlie our lattice-based zkSNARK. We then show how to combine these ingredients to obtain our designated-verifier zkSNARK by invoking the Bitansky et al. [30, 31] compiler (see Section 1.2).

3.1 Linear PCPs over Extension Fields

Claim A.6 gives a linear PCP for R1CS over any (sufficiently-large) field $\mathbb{F}$. In our work, we consider linear PCPs over quadratic extensions $\mathbb{F}_{p^2}$. As discussed in Section 1.2, we consider compilers based on vector encryption over the extension $\mathbb{F}_{p^2}$ as well as over the base field $\mathbb{F}_p$. For the latter setting, we need to first transform a linear PCP over $\mathbb{F}_{p^2}$ to a linear PCP over $\mathbb{F}_p$. We describe this transformation in Appendix C.1.

3.2 Linear-Only Vector Encryption

We begin with the definition of a vector encryption scheme (adapted from [31]), and then define the linear-only [30, 31] property we rely on for our zkSNARK constructions.

Definition 3.1 (Vector Encryption). Let $\mathbb{F}$ be a finite field. A secret-key additively-homomorphic vector encryption scheme over a vector space $\mathbb{F}^\ell$ consists of a tuple of algorithms $\Pi_{\text{Enc}} = (\text{Setup}, \text{Encrypt}, \text{Decrypt}, \text{Add})$ with the following properties:

- $\text{Setup}(1^\lambda, 1^\ell) \rightarrow (\text{pp}, \text{sk})$: On input the security parameter λ and the plaintext dimension ℓ , the setup algorithm outputs public parameters pp and a secret key sk .

- **Encrypt**(sk, v) → ct: On input the secret key sk and a vector $v \in \mathbb{F}^\ell$, the encryption algorithm outputs a ciphertext ct.
- **Decrypt**(sk, ct) → v/⊥: On input the secret key sk and a ciphertext ct, the decryption algorithm either outputs a vector $v \in \mathbb{F}^\ell$ or a special symbol ⊥.
- **Add**(pp, {ct_i}_{i∈[n]}, {c_i}_{i∈[n]}) → ct*: On input the public parameters, a collection of ciphertexts ct₁, ..., ct_n and scalars c₁, ..., c_n ∈ $\mathbb{F}$, the addition algorithm outputs a new ciphertext ct*.

Moreover, Π_{Enc} should satisfy the following properties:

- **Additive homomorphism:** For all security parameters $\lambda \in \mathbb{N}$, vectors $v_1, \dots, v_k \in \mathbb{F}^\ell$, and scalars $y_1, \dots, y_k \in \mathbb{F}$, where $k = k(\lambda)$,

$$\Pr \left[\text{Decrypt}(\text{sk}, \text{ct}^*) = \sum_{i \in [k]} y_i v_i \right] = 1 - \text{negl}(\lambda), \quad (3.1)$$

where $(\text{pp}, \text{sk}) \leftarrow \text{Setup}(1^\lambda, 1^\ell)$, $\text{ct}_i \leftarrow \text{Encrypt}(\text{sk}, v_i)$ for all $i \in [k]$, and $\text{ct}^* \leftarrow \text{Add}(\text{pp}, \{\text{ct}_i\}_{i \in [k]}, \{y_i\}_{i \in [k]})$. We say that Π_{Enc} is additively homomorphic with respect to a set $S \subseteq R_p^k$ if Eq. (3.1) holds for all $(y_1, \dots, y_k) \in S$. Note that additive homomorphism implies correctness of decryption.

- **CPA security:** For all security parameters $\lambda \in \mathbb{N}$ and all efficient adversaries $\mathcal{A}$,

$$\Pr \left[\mathcal{A}^{O_b(\text{sk}, \cdot, \cdot)}(1^\lambda, \text{pp}) = b \right] = 1/2 + \text{negl}(\lambda), \quad (3.2)$$

where $(\text{pp}, \text{sk}) \leftarrow \text{Setup}(1^\lambda, 1^\ell)$, $b \xleftarrow{R} \{0, 1\}$, and oracle O_b takes inputs (sk, v_0, v_1) and outputs $\text{ct}_b \leftarrow \text{Encrypt}(\text{sk}, v_b)$. If Eq. (3.2) holds against all efficient adversaries $\mathcal{A}$ making at most Q queries to O_b , then we say Π_{Enc} is Q -query CPA secure.

Definition 3.2 (Linear-Only Vector Encryption [30, adapted]). A vector encryption scheme $\Pi_{\text{Enc}} = (\text{Setup}, \text{Encrypt}, \text{Decrypt}, \text{Add})$ over $\mathbb{F}^\ell$ is *strictly linear-only* if for all polynomial-size adversaries $\mathcal{A}$, there is a polynomial-size extractor $\mathcal{E}$ such that for all security parameters $\lambda \in \mathbb{N}$, auxiliary inputs $z \in \{0, 1\}^{\text{poly}(\lambda)}$, and any efficient plaintext generator $\mathcal{M}$,

$$\Pr[\text{ExptLinearExt}_{\Pi_{\text{Enc}}, \mathcal{A}, \mathcal{M}, \mathcal{E}, z}(1^\lambda) = 1] = \text{negl}(\lambda),$$

where the experiment $\text{ExptLinearExt}_{\Pi_{\text{Enc}}, \mathcal{A}, \mathcal{M}, \mathcal{E}, z}(1^\lambda)$ is defined as follows:

- (1) The challenger starts by sampling $(\text{pp}, \text{sk}) \leftarrow \text{Setup}(1^\lambda, 1^\ell)$ and $(v_1, \dots, v_m) \leftarrow \mathcal{M}(1^\lambda, \text{pp})$. It computes $\text{ct}_i \leftarrow \text{Encrypt}(\text{sk}, v_i)$ for each $i \in [m]$ and runs $\mathcal{A}(\text{pp}, \text{ct}_1, \dots, \text{ct}_m; z)$ to obtain a tuple $(\text{ct}'_1, \dots, \text{ct}'_k)$.
- (2) The challenger computes $\Pi \leftarrow \mathcal{E}(\text{pp}, \text{ct}_1, \dots, \text{ct}_m; z)$ and $V' \leftarrow \Pi \cdot [v_1 \mid \dots \mid v_m]^\top$, where $\Pi \in \mathbb{F}^{k \times m}$ and $V' \in \mathbb{F}^{k \times \ell}$. The experiment outputs 1 if there exists an index $i \in [k]$ such that $\text{Decrypt}(\text{sk}, \text{ct}'_i) \neq \perp$ and $\text{Decrypt}(\text{sk}, \text{ct}'_i) \neq v'_i$, where $v'_i \in \mathbb{F}^\ell$ is the i^{th} row of V' . Otherwise, the experiments outputs 0.

We provide additional discussion of these definitions and compare them to previous definitions [30, 31] in the full version of this paper [78].

Circuit privacy. In addition to the above properties, we additionally require a *circuit privacy* property [66]. Circuit privacy says that the ciphertext output by Add can be simulated given only the underlying plaintext value, *without* knowledge of the linear combination used to construct the ciphertext. This is important for arguing zero knowledge (see Section 3.4). We give the formal definition in Appendix C.2.

3.3 Candidate Linear-Only Vector Encryption

Our constructions work over the ring $R = \mathbb{Z}[x]/(x^d + 1)$ where d is a power of 2. We specifically consider the cases where $d = 1$ ($R = \mathbb{Z}$) and $d = 2$ ($R = \mathbb{Z}[x]/(x^2 + 1)$). For a positive integer $p \in \mathbb{N}$, we write $R_p = R/pR$. We represent elements of R as a vector of coefficients (i.e., as a vector $\mathbb{Z}^d$). For an element $r \in R$, we write $\|r\|_\infty$ to denote the ℓ_∞ norm of the vector of coefficients of r . We write γ_R to denote the expansion constant where for all $r, s \in R$, we have that $\|rs\|_\infty \leq \gamma_R \|r\|_\infty \|s\|_\infty$. In particular, $\gamma_R = 1$ when $d = 1$ and $\gamma_R = 2$ when $d = 2$. Finally, for a vector $v \in R^n$, we write $\|v\|_p$ to denote the ℓ_p norm $\|v'\|_p$ of the vector $v' \in \mathbb{Z}^{dn}$ formed by concatenating the vector of coefficients of each element in v .

(Module) learning with errors. Security of our construction relies on the module learning with errors (MLWE) assumption [40, 80] (in addition to our linear-only conjecture). We state the MLWE assumption in “normal form” where the secret is sampled from the error distribution. This form of the problem is as hard as the version where the secret key is sampled uniformly at random [6].

Definition 3.3 (Module Learning With Errors (MLWE) [40, 80]). Fix a security parameter λ , integers $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$, $d = d(\lambda)$ where d is a power of two. Let $R = \mathbb{Z}[x]/(x^d + 1)$, $R_q = R/qR$, and $\chi = \chi(\lambda)$ be an error distribution over R_q . The (decisional) module learning with errors (MLWE) assumption $\text{MLWE}_{n,m,d,q,\chi}$ states that for $A \xleftarrow{R} R_q^{n \times m}$, $s \leftarrow \chi^n$, $e \leftarrow \chi^m$, and $u \xleftarrow{R} R_q^m$, the following two distributions are computationally indistinguishable:

$$(A, s^\top A + e^\top) \text{ and } (A, u^\top)$$

Remark 3.4 (Relation to LWE and RLWE). The module LWE assumption generalizes both the classic learning with errors (LWE) assumption [91] as well as the ring learning with errors (RLWE) assumption [82]. In particular, LWE is MLWE instantiated with $d = 1$ and RLWE is MLWE instantiated with $n = 1$.

Vector encryption construction. We now describe our vector encryption scheme. Our scheme is an adaptation of the Regev-based [91] scheme of Peikert et al. [89], generalized to modules and with the following additions/modifications:

- **Secret-key encryption:** Since a secret-key vector encryption suffices for our designated-verifier zkSNARK,³ we consider a secret-key version of the scheme. This reduces the concrete cost for encryption (we can substitute a random vector in each ciphertext in place of a matrix-vector product with the public key). Note that there are still public parameters in our scheme. These

³Using a public-key encryption scheme does *not* imply a publicly-verifiable zkSNARK in this setting. There is no advantage to using a public-key encryption scheme to instantiate the underlying encryption scheme.

are used for *re-randomization* of homomorphically-evaluated ciphertexts, and are *not* used for encryption.

- **Message encoding:** We encode the message in the least significant bits of the ciphertext rather than the most significant bits. When the plaintext modulus p and ciphertext modulus q are coprime, these approaches are equivalent up to scaling [2]. In our implementation, encoding a value k in the least significant bits of the ciphertext is more convenient since we avoid the need to compute the value $\lfloor k \cdot q/p \rfloor \bmod q$ (which if implemented improperly, can overflow our integer representation).
- **Ciphertext re-randomization:** For zero knowledge, we require an additional circuit privacy property. Ciphertexts in this scheme consist of pairs of vectors $\text{ct} = (\mathbf{a}, \mathbf{c})$. Homomorphic operations on ciphertexts correspond to computing component-wise linear combinations. In our construction, we include a public MLWE matrix as part of the public parameters to re-randomize the vector $\mathbf{a}$, and we use standard noise smudging techniques (see Lemma A.2) to re-randomize the vector $\mathbf{c}$. Previously, Gennaro et al. [65] suggest that the first component $\mathbf{a}$ is already random by appealing to the leftover hash lemma; unfortunately, this only applies in the setting where the coefficients of the linear combination have sufficient min-entropy (which is not necessarily the case in the zkSNARK construction). We show that in our case and under the MLWE assumption,⁴ our construction provably satisfies circuit privacy without needing any additional assumption on the choice of linear combination.
- **Ciphertext sparsification.** Our linear-only definition (Definition 3.2) essentially requires that the only way an efficient adversary can generate a *valid* ciphertext is by taking linear combinations of valid ciphertexts. This means that the set of valid ciphertexts must be *sparse* (to prevent *oblivious* sampling of a valid ciphertext). Previous works [30, 31, 64] suggest *double encryption* to realize this property. With double encryption, valid ciphertexts $\text{ct} = (\text{ct}_1, \text{ct}_2)$ are defined as pairs of ciphertexts that both encrypt identical messages. While this approach is applicable in our setting, it doubles the length of the ciphertexts.

We propose a similar, but more efficient, approach tailored for vector encryption. Namely, if our goal is to encrypt elements from a vector space $\mathbb{F}^\ell$, we enlarge the plaintext space to $\mathbb{F}^{\ell+\tau}$, where τ is a sparsification parameter. During setup, we sample a random matrix $\mathbf{T} \xleftarrow{\mathbb{R}} \mathbb{F}^{\ell \times \tau}$ which is included as part of the secret key. Then, to encrypt a vector $\mathbf{v} \in \mathbb{F}^\ell$, we instead encrypt the vector $\mathbf{u}^\top = [\mathbf{v}^\top \mid (\mathbf{T}\mathbf{v})^\top]$. During decryption, after recovering $\mathbf{u}^\top = [\mathbf{u}_1^\top \mid \mathbf{u}_2^\top]$, the decryption algorithm outputs $\perp$ if $\mathbf{u}_2 \neq \mathbf{T}\mathbf{u}_1$. Semantic security of the vector encryption scheme ensures that the secret transformation $\mathbf{T}$ is computationally hidden from the view of the adversary. By setting the sparsification parameter τ accordingly, we can ensure that for any fixed vector $\mathbf{u}^\top = [\mathbf{u}_1^\top \mid \mathbf{u}_2^\top]$, the probability that $\mathbf{u}_2 = \mathbf{T}\mathbf{u}_1$ is negligible (over the randomness of $\mathbf{T}$). We conjecture that our approach also yields an encryption scheme that satisfies the linear-only assumption. The advantage of this approach is that the ciphertext size in the underlying vector encryption scheme grows *additively* with the

plaintext dimension (i.e., the resulting ciphertext size is $n + \ell + \tau$ rather than $2(n + \ell)$ as with “encrypting twice”).

We now describe our vector encryption scheme:

Construction 3.5 (Vector Encryption). Let $d = d(\lambda)$ be a power of two and let $R = \mathbb{Z}[x]/(x^d + 1)$. Fix lattice parameters $p = p(\lambda)$, $q = q(\lambda)$, $n = n(\lambda)$ and an error distribution $\chi = \chi(\lambda)$ over R_q . We additionally define the following parameters:

- ℓ : the plaintext dimension
- τ : the sparsification parameter
- B : the noise smudging bound

Let $\ell' = \ell + \tau$. We construct a secret-key vector encryption scheme $\Pi_{\text{Enc}} = (\text{Setup}, \text{Encrypt}, \text{Decrypt}, \text{Add})$ over R_p as follows:

- **Setup**($1^\lambda, 1^\ell$): Sample matrices $\mathbf{A} \xleftarrow{\mathbb{R}} R_q^{n \times n}$, $\mathbf{S} \xleftarrow{\mathbb{R}} \chi^{n \times \ell'}$, $\mathbf{T} \xleftarrow{\mathbb{R}} R_p^{\tau \times \ell}$, and $\mathbf{E} \xleftarrow{\mathbb{R}} \chi^{n \times \ell'}$. Compute $\mathbf{D} \leftarrow \mathbf{S}^\top \mathbf{A} + p\mathbf{E}^\top \in R_q^{\ell' \times n}$. Output the secret key $\text{sk} = (\mathbf{S}, \mathbf{T})$ and the public parameters $\text{pp} = (\mathbf{A}, \mathbf{D})$.
- **Encrypt**($\text{sk}, \mathbf{v}$): On input the secret key $\text{sk} = (\mathbf{S}, \mathbf{T})$ and a vector $\mathbf{v} \in R_p^\ell$, construct the concatenated vector $\mathbf{u}^\top = [\mathbf{v}^\top \mid (\mathbf{T}\mathbf{v})^\top] \in R_p^{\ell'}$. Sample $\mathbf{a} \xleftarrow{\mathbb{R}} R_q^n$, $\mathbf{e} \xleftarrow{\mathbb{R}} \chi^{\ell'}$ and compute $\mathbf{c} \leftarrow \mathbf{S}^\top \mathbf{a} + p\mathbf{e} + \mathbf{u} \in R_q^{\ell'}$. Output the ciphertext $\text{ct} = (\mathbf{a}, \mathbf{c})$.
- **Add**($\text{pp}, \{\text{ct}_i\}_{i \in [k]}, \{y_i\}_{i \in [k]}$): On input the public parameters $\text{pp} = (\mathbf{A}, \mathbf{D})$, ciphertexts $\text{ct}_i = (\mathbf{a}_i, \mathbf{c}_i)$ for $i \in [k]$, and scalars $y_i \in R_p$, sample $\mathbf{r} \xleftarrow{\mathbb{R}} \chi^n$, $\mathbf{e}_a \xleftarrow{\mathbb{R}} \chi^n$, $\mathbf{e}_c \xleftarrow{\mathbb{R}} [-B, B]^{d\ell'}$ and output the ciphertext

$$\text{ct}^* = \left(\sum_{i \in [k]} y_i \mathbf{a}_i + \mathbf{A}\mathbf{r} + p\mathbf{e}_a, \sum_{i \in [k]} y_i \mathbf{c}_i + \mathbf{D}\mathbf{r} + p\mathbf{e}_c \right). \quad (3.3)$$

- **Decrypt**(sk, ct): On input the secret key $\text{sk} = (\mathbf{S}, \mathbf{T})$ and a ciphertext $\text{ct} = (\mathbf{a}, \mathbf{c})$, compute $\mathbf{z} \leftarrow \mathbf{c} - \mathbf{S}^\top \mathbf{a} \in R_q^{\ell'}$. Compute $\mathbf{u} = \mathbf{z} \bmod p$, and parse $\mathbf{u}^\top = [\mathbf{v}_1^\top \mid \mathbf{v}_2^\top]$ where $\mathbf{v}_1 \in R_p^\ell$ and $\mathbf{v}_2 \in R_p^\tau$. Output $\mathbf{v}_1$ if $\mathbf{v}_2 = \mathbf{T}\mathbf{v}_1 \in R_p^\tau$ and $\perp$ otherwise.

Correctness and security analysis. Below, we state our main theorems on the correctness and security of Construction 3.5. We defer the formal analysis to the full version of this paper [78].

Theorem 3.6 (Additive Homomorphism). *Let λ be a security parameter and p, q, n, ℓ', χ, B be as defined in Construction 3.5. Suppose χ is subgaussian⁵ with parameter s . If $n, \ell', s, d, \gamma_R = \text{poly}(\lambda)$, then for all $k = k(\lambda)$, there exists $q = (pB + kp^2) \cdot \text{poly}(\lambda)$ such that Construction 3.5 is additively homomorphic with respect to R_p^k . Concretely, let C be a correctness parameter and let B_1, B_2 be bounds. Define the set $S = \{\mathbf{y} \in R_p^k : \|\mathbf{y}\|_1 \leq B_1 \text{ and } \|\mathbf{y}\|_2 \leq B_2\}$. If $\ell', n > 8$, and*

$$q > 2p(B + \gamma_R B_2 C s + \gamma_R B_1 / 2 + 2\gamma_R n C^2 s^2) + p, \quad (3.4)$$

then Eq. (3.1) holds with probability $1 - (4n + 2)d\ell' \exp(-\pi C^2)$ for all $(y_1, \dots, y_k) \in S$.

Theorem 3.7 (CPA Security). *Fix a security parameter λ and let p, q, n, ℓ', χ be as defined in Construction 3.5. Take any $Q = \text{poly}(\lambda)$ and suppose that p, q are coprime. Under the $\text{MLWE}_{n,m,d,q,\chi}$ assumption with $m = n + Q$, Construction 3.5 is Q -query CPA secure.*

⁴We could make this step statistical by relying on the leftover hash lemma, but this requires much larger parameters. Instead, we rely on MLWE and settle for computational circuit privacy (which translates to computational zero knowledge).

⁵When $d > 1$, we assume that χ is the concatenation of d independent subgaussian distributions over $\mathbb{Z}$, each with parameter at most s . This is true for discrete Gaussian distributions over a power-of-two cyclotomic ring.

Theorem 3.8 (Circuit Privacy). *Let λ be a security parameter and p, q, n, ℓ', χ be as defined in Construction 3.5. Suppose that χ is subgaussian with parameter s . If $n, \ell', s, d, \gamma_R = \text{poly}(\lambda)$, and $B = 2^{\omega(\log \lambda)}$. $k p^2$, then under the $\text{MLWE}_{n,m,d,q,\chi}$ assumption with $m = n$, Construction 3.5 is circuit private with respect to the set $S = R_p^k$. Concretely, let C be a correctness parameter and let B_1, B_2 be bounds. Let $S = \{y \in R_p^k : \|y\|_1 \leq B_1 \text{ and } \|y\|_2 \leq B_2\}$. Then under the $\text{MLWE}_{n,m,d,q,\chi}$ assumption with $m = n$, for every efficient adversary $\mathcal{A}$ restricted to strategies in S , there exists an efficient simulator S where*

$$\Pr[\text{ExptCircuitPriv}_{\Pi_{\text{Enc}}, \mathcal{A}, S}(1^\lambda) = 1] \leq 1/2 + \varepsilon + \text{negl}(\lambda),$$

and

$$\varepsilon = (4n + 2)d\ell' \exp(-\pi C^2) + \frac{d\ell'(\gamma_R B_2 C s + \gamma_R B_1 / 2 + 2\gamma_R n C^2 s^2)}{B}. \quad (3.5)$$

Conjecture 3.9 (Linear-Only). *Fix a security parameter λ and let p, d, τ be defined as in Construction 3.5. If $|R_p|^\tau = p^{\tau d} = \lambda^{\omega(1)}$, then Construction 3.5 is strictly linear-only (Definition 3.2).*

Extensions and variants. In the full version of this paper [78], we discuss the plausibility of Conjecture 3.9 and describe an extension of Construction 3.5 to higher-degree extensions and an alternative approach based on bit decompositions to reduce noise growth.

Modulus switching. The size of the ciphertext in Construction 3.5 is determined by three main parameters: the ring dimension d , the module dimension n , and the ciphertext modulus q . According to Theorem 3.6, the modulus q must be sufficiently large to support the required number of homomorphic operations. However, the modulus switching technique developed in the context of fully homomorphic encryption [3, 40, 41, 52, 56] provides a way to reduce the size of the ciphertexts *after* performing homomorphic operations. Specifically, modulus switching allows one to take a ciphertext over R_q and convert it to one over $R_{q'}$ where $q' < q$ while preserving the correctness of decryption. This technique applies to most Regev-based encryption schemes, including Construction 3.5. Reducing the size of the ciphertexts after homomorphic evaluation translates to a reduction in the proof size of the resulting zkSNARK. We begin by defining the ciphertext rescaling operation Scale from Brakerski et al. [40]:

- $\text{Scale}(\mathbf{x}, q, q', p) \rightarrow \mathbf{x}'$: On input integers $q > q' > p$ and a vector $\mathbf{x} \in R_q^n$, the scale operation outputs the vector $\mathbf{x}' \in R_{q'}^n$ that is closest to $(q'/q) \cdot \mathbf{x}$ such that $\mathbf{x}' = \mathbf{x} \pmod{p}$.

We now state the main theorem, adapted from [40]. We provide the proof in the full version of this paper [78].

Theorem 3.10 (Modulus Switching [40, adapted]). *Let λ be a security parameter and p, q, n, d, ℓ', χ be as defined in Construction 3.5. Let C be a correctness parameter. Suppose that χ is subgaussian with parameter s . Let $q' < q$ be a positive integer where $q' = q \pmod{p}$. Take any vector $\mathbf{a} \in R_q^n$, $\mathbf{c} \in R_{q'}^d$, and let $\mathbf{a}' \leftarrow \text{Scale}(\mathbf{a}, q, q', p)$, $\mathbf{c}' \leftarrow \text{Scale}(\mathbf{c}, q, q', p)$. Sample $\mathbf{S} \leftarrow \chi^{n \times \ell'}$. Let $\mathbf{z} = \mathbf{c} - \mathbf{S}^T \mathbf{a} \in R_q^d$ and suppose that*

$$\|\mathbf{z}\|_\infty < q/2 - (1 + n\gamma_R C s) \cdot (p/2) \cdot (q/q'). \quad (3.6)$$

Then, with probability $1 - 2dn\ell' \exp(-\pi C^2)$, $\mathbf{z} = \mathbf{z}' \pmod{p}$, where $\mathbf{z}' = \mathbf{c}' - \mathbf{S}^T \mathbf{a}' \in R_{q'}^d$.

In our construction, after performing homomorphic operations (via Add), the evaluator takes the final ciphertext $\text{ct} = (\mathbf{a}, \mathbf{c})$, computes $\mathbf{a}' \leftarrow \text{Scale}(\mathbf{a}, q, q', p)$ and $\mathbf{c}' \leftarrow \text{Scale}(\mathbf{c}, q, q', p)$, and outputs the rescaled ciphertext $\text{ct}' = (\mathbf{a}', \mathbf{c}')$. In particular, the components of ct' are elements of $R_{q'}$ rather than R_q . We additionally conjecture that Conjecture 3.9 holds even if we introduce this additional rescaling operation.

3.4 zkSNARKs from Linear-Only Encryption

In this section, we state the result of Bitansky et al. [30] for constructing zkSNARKs from linear PCPs and linear-only vector encryption. We specifically describe the variant by Boneh et al. [31] based on linear-only vector encryption. We refer to Section 1.1 for an overview of the transformation, and to the full version of this paper [78] for a formal description of the construction.

Theorem 3.11 (SNARKs from Linear-Only Vector Encryption [30, 31]). *Let $\mathbb{F}$ be a finite field, CS be an R1CS system over $\mathbb{F}$, and $\mathcal{R}_{CS}$ be the associated relation. Let Π_{LPCP} be a k -query input-oblivious linear PCP for $\mathcal{R}_{CS}$ and Π_{Enc} be a secret-key vector encryption scheme for $\mathbb{F}^k$. If Π_{LPCP} is statistically sound against linear provers and Π_{Enc} is CPA-secure and strictly linear-only, then there is a designated-verifier succinct argument of knowledge Π_{SNARK} for $\mathcal{R}_{CS}$ in the preprocessing model.*

Zero knowledge. Bitansky et al. [30] showed that combining a linear PCP satisfying HVZK with *re-randomizable* linear-only encryption yields a zkSNARK. An encryption scheme is *re-randomizable* if there is a public procedure that transforms *any* valid encryption of m into a *fresh* encryption of m . Our lattice-based vector encryption does not satisfy this property (due to the variability in the amount of noise in ciphertexts). Instead, we show that the *weaker* property of circuit privacy suffices to argue zero knowledge.

At a high-level, the argument goes as follows. First, by HVZK of the linear PCP, the linear PCP responses can be simulated given only the statement. Circuit privacy then says that the encrypted linear PCP responses can be simulated given only the simulated LPCP responses. We state the theorem below, but defer the analysis (and additional discussion) to the full version of this paper [78].

Theorem 3.12 (Zero Knowledge from Circuit Privacy). *Let Π_{LPCP} , and Π_{Enc} be as defined in Theorem 3.11. If Π_{LPCP} satisfies perfect honest-verifier zero knowledge and Π_{Enc} is CPA-secure and computationally (resp., statistically) circuit private, then Π_{SNARK} from Theorem 3.11 is computationally (resp., statistically) zero knowledge.*

4 IMPLEMENTATION AND EVALUATION

In this section, we provide an overview of our lattice-based zkSNARK implementation (by combining Claim A.6 with Construction 3.5) and then describe our experimental evaluation.

4.1 Linear PCP Implementation

The prover's computation in the zkSNARK from Theorem 3.11 consists of two main components: computing the linear PCP proof

(Claim A.6) and homomorphically computing the encrypted linear PCP responses. Computing the linear PCP responses (over a finite field $\mathbb{F}$) requires the prover to compute the coefficients of a polynomial $H(z) := (A(z)B(z) - C(z))/Z(z)$, where A, B, C are polynomials of degree $N_g - 1$ (over $\mathbb{F}$) determined by the R1CS system (which has N_g constraints), the statement, and the witness, and Z is a fixed polynomial. We refer to the full version of this paper [78] for the full details of this construction.

Ben-Sasson et al. [20] described an efficient approach to compute the coefficients of H using fast Fourier transforms (FFTs) over $\mathbb{F}$. To use standard Cooley-Tukey FFTs for powers of two [50] (which we refer to as “radix-2 FFTs”), we require that $\mathbb{F}$ contains a multiplicative subgroup of order 2^d where $2^d > N_g$. Indeed, the construction of Ben-Sasson et al. uses a specially-chosen elliptic curve group whose order is divisible by a large power of 2. In our setting, we consider linear PCPs over a quadratic extension $\mathbb{F}_{p^2}$, whose order is $p^2 - 1 = (p + 1)(p - 1)$. In the best case, if $p = 2^d \pm 1$, then $\mathbb{F}_{p^2}$ has a subgroup of order 2^{d+1} . However, if $N_g > 2p$, the field $\mathbb{F}_{p^2}$ never has a sufficiently large subgroup to directly compute radix-2 FFTs.⁶

Our approach. When the field $\mathbb{F}$ contains a multiplicative subgroup whose order is a moderately large power of two (e.g., 2^d), we can still leverage (multiple) radix-2 FFTs to efficiently implement multipoint polynomial evaluation and interpolation over a domain $D \subset \mathbb{F}$ of size $|D| = k \cdot 2^d$ for a (small) $k > 1$. We give a brief overview of our approach here and defer the full details to Appendix D. Let $\omega \in \mathbb{F}$ be a primitive 2^d -th root of unity and let $H = H_1 = \langle \omega \rangle \subset \mathbb{F}$ be the subgroup of order 2^d generated by ω (corresponding to the 2^d -th roots of unity). We define our domain D (for multipoint evaluation and interpolation) to be $\bigcup_{i \in [k]} H_i$, where $H_2, \dots, H_k$ are pairwise disjoint cosets of H . Polynomial evaluation over D can be implemented using k degree- 2^d FFTs (over H_1), along with 2^d multipoint evaluations of polynomials of degree- k (over a fixed basis determined by the cosets). An analogous result holds for interpolation. As long as $k < 2^d$, the smaller evaluation/interpolations can be implemented in $k \log k$ time using standard FFTs. In this case, the running time of our algorithm for evaluating a polynomial on a domain of size $2^d k$ is $O(2^d k (d + \log k))$, which matches the asymptotic complexity of a standard FFT over a domain of the same size. We give more details in Appendix D. We use this approach to implement the linear PCP prover when working over fields with insufficient roots of unity to support a standard radix-2 FFT.

4.2 Lattice-Based zkSNARK Implementation

In this section, we describe our overall zkSNARK implementation. We begin by describing our methodology for setting the lattice parameters n, q, χ for our lattice-based vector encryption scheme (Construction 3.5). We then describe a few optimizations to improve the concrete efficiency of the resulting construction.

⁶While more general algorithms for FFT can be used for multipoint evaluation and interpolation over a domain whose size is a prime power [90] or a product of coprime values [72, 104], these algorithms are more complex to implement and worse in terms of concrete efficiency compared to basic radix-2 FFTs. We show how to implement our approach using a small number of radix-2 FFTs.

Fields*	λ_q	λ_c	p	(n, d)	$\log q$	s	ℓ	τ
$\mathbb{F}_p, \mathbb{F}_p$	128	138	$5 \cdot 2^{25} + 1$	(4700, 1)	123	80	82	5
$\mathbb{F}_{p^2}, \mathbb{F}_p$	128	138	$2^{19} - 1$	(4050, 1)	107	36	71	7
$\mathbb{F}_{p^2}, \mathbb{F}_{p^2}$	128	138	$2^{13} - 1$	(1815, 2)	98	64	109	5
	128	138	$2^{19} - 1$	(2045, 2)	108	40	36	4

* The first field listed is the base field for the linear PCP Π_{LPCP} and the second is the plaintext field for the linear-only vector encryption scheme Π_{Enc} .

Table 2: Lattice parameters for zkSNARK instantiations obtained by combining the linear PCP Π_{LPCP} from Claim A.6 and Remark A.7 with the linear-only vector encryption scheme Π_{Enc} from Construction 3.5. Here, λ_q is the estimated number of bits of quantum security, λ_c is the estimated number bits of classical security, p is the plaintext modulus, n is the ring dimension, d is the module rank, q is the ciphertext modulus, s is the width parameter for the discrete Gaussian noise distribution, ℓ is the dimension of the plaintext space, and τ is the sparsification parameter. Parameters shown are based on supporting an R1CS system with 2^{20} constraints. The final two rows correspond to the “Shorter Proofs” and the “Shorter CRS” instantiations in Table 1, respectively.

Lattice parameter selection. Due to space limitations, we defer our parameter selection methodology to the full version of this paper [78]. We describe several example parameter sets and their estimated bits of classical and post-quantum security in Table 2. For our main experiments, we use $\kappa = 40$ for the statistical zero-knowledge parameter; the remaining parameters are chosen to provide 128-bits of post-quantum security (based on the analysis from Section 3.3 and hardness estimates from the LWE Estimator tool et al. [1]).

Reducing the CRS size. Like most Regev-based encryption schemes, the ciphertexts in Construction 3.5 have the form $\text{ct} = (\mathbf{a}, \mathbf{c})$ where $\mathbf{a} \in R_q^n$ is uniformly random and $\mathbf{c} \in R_q^{\ell'}$ encodes the message. A heuristic approach to reduce the ciphertext size is to derive the random vector $\mathbf{a}$ as the output of a pseudorandom function (PRF) and include the PRF key in place of the vector $\mathbf{a}$ (or alternatively, take them to be the outputs of a public hash function). Security of these heuristics can be justified in the random oracle model [62]. We adopt this approach in our implementation. In our implementation, we use AES (in counter mode) as the underlying PRF. Similar approaches for reducing the size of the public components of lattice-based cryptosystems has been used for both lattice-based key-exchange [38] as well as previous lattice-based zkSNARKs [65].

Noise distribution. We take our noise distribution χ to be a discrete Gaussian distribution $\chi = \chi_s$ with mean 0 and width s (Eq. (A.1)). We use inversion sampling to sample from χ_s given a uniformly-random 64-bit value. We refer to the full version of this paper [78] for more details. This is similar to the approach used in lattice-based key-exchange [38].

Big integer support. In our implementation, the ciphertext modulus q is around 100 bits. We implement all of the homomorphic operations (over the ring R_q) using 128-bit arithmetic. Since we choose q to be a power-of-two, we can compute over $\mathbb{Z}_{2^{128}}$ and defer the modular reduction to the end of the computation.

We use the compiler intrinsic type `__uint128_t` for 128-bit arithmetic on a 64-bit architecture. Internally, each 128-bit value is represented by two 64-bit words. The intrinsic representation is $16\times$ faster than using a general multi-precision arithmetic and $8\times$ – $9\times$ than fixed-point arithmetic for slightly larger bitlengths (i.e., 192-bit or 256-bit fixed-precision arithmetic from Boost). We provide more detailed microbenchmarks in the full version of this paper [78].

4.3 Experimental Evaluation

We now describe our implementation and experimental evaluation of our lattice-based zkSNARK from Section 3.4.

System implementation. Our implementation is written in C++.⁷ We use `libsark` [96] and `libfqfft` [94] to implement the linear PCP for R1CS satisfiability (Claim A.6). In particular, we use the linear PCP implementation from `libsark` (with the minor changes from Appendix B), and the implementation of the standard radix-2 FFT [50] (over a finite field) as well as the Bostan-Schost algorithms for multipoint evaluation and interpolation on points from a geometric sequence [39] from `libfqfft`. These building blocks suffice to implement our approach described in Section 4.1 for evaluating the linear PCP.

Metrics and evaluation methodology. Following previous works [21, 47, 99], we measure the performance of our system on R1CS systems with different number of constraints m (ranging from $m = 2^{10}$ to $m = 2^{20}$). Like previous works, we keep the number of variables n in each R1CS system to be roughly m (i.e., $n \approx m$), and we consider statements of a fixed length $k = 100$. The statement length only has a mild effect on the verification complexity (which is already very fast) and we do not focus on it in our evaluation.

We run all of our experiments on an Amazon EC2 c5.4xlarge instance running Ubuntu 20.04. The machine has 16 vCPUs (Intel Xeon Platinum 8124M at 3.0 GHz) and 32 GB of RAM. The processor supports the AES-NI instruction set. We compile our code using `gcc 9.3.0` for a 64-bit x86 architecture. All of our measurements are taken with a single-threaded execution.

General benchmarks. In Fig. 1, we compare the performance of different instantiations of our zkSNARK on R1CS instances of varying sizes. We consider two instantiations using linear PCPs and vector encryption over the extension field $\mathbb{F}_{p^2}$ (for $p = 2^{13} - 1$ and $p = 2^{19} - 1$), as well as two alternative instantiations where we use a vector encryption over the base field $\mathbb{F}_p$. For the latter instantiations, we consider both the instantiation where we first compile a linear PCP over the extension field to a linear PCP over the base field (Construction C.1) and a second instantiation where we directly construct a linear PCP over the base field. Across the board, the verifier time is small so we focus our discussion on the other metrics.

For our main instantiations (working over the extension field), the field size provides a trade-off in CRS size vs. proof size. Using a larger field decreases the CRS size (fewer repetitions needed for knowledge amplification at the linear PCP level), but leads to longer proofs (due to larger parameters). Concretely, for R1CS systems with 2^{20} constraints, increasing the characteristic from $p = 2^{13} - 1$

to $p = 2^{19} - 1$ decreases the CRS size by $2.8\times$ (with a corresponding decrease in setup time), but increases the proof size by $1.2\times$. The prover complexity is essentially the same in the two cases.

In the case where we take a linear PCP over $\mathbb{F}_{p^2}$ and first apply Construction C.1 to obtain a linear PCP over $\mathbb{F}_p$, the proof size still remains comparable to the case where we work exclusively over $\mathbb{F}_{p^2}$. However, the CRS size is doubled since Construction C.1 increases the query length by the degree of the field extension, as is the prover complexity. The advantage of this construction is that it is based on *standard* lattices as opposed to *module* lattices, and thus, plausibly has better security.

Finally, if we consider the direct compilation of a linear PCP over the base field $\mathbb{F}_p$, the proof size is $1.4\times$ to $1.8\times$ longer than the constructions that use the extension field.

Extension field vs. base field. To quantify the concrete performance trade-off enabled by extension fields, we also compare our zkSNARKs over $\mathbb{F}_{p^2}$ with an instantiation over $\mathbb{F}_p$ (i.e., compile the linear PCP from Claim A.6 over $\mathbb{F}_p$ using the linear-only vector encryption from Construction 3.5 over $\mathbb{F}_p$). The results are summarized in Fig. 2. We first note that most of the instantiations over $\mathbb{F}_p$ require working over a ring R_q with $q > 2^{128}$. As discussed in Section 4.2, this will incur considerable computational overhead for the big-integer arithmetic. Working over the extension field allows us to consider instantiations over much larger fields without incurring the cost of big-integer arithmetic.

Fig. 2 shows that working over a quadratic extension field introduces a modest increase in the CRS size (by a factor of $1.4\times$ to $1.6\times$) compared to working over a prime-order base field of similar size. In return, working over the extension field *reduces* the proof size by $1.7\times$ to $2.4\times$ (specifically, from nearly 70 KB to under 30 KB when considering a 56-bit field). As discussed in Section 1.2 (see also the formal analysis in Section 3.3), all of the lattice parameters grow with the field *characteristic*. Thus, for fields of comparable size, all of the lattice parameters will be larger if we work over a base field than if we work over an extension field (of smaller characteristic). This leads to longer proofs. The size of the CRS is smaller because of the CRS compression technique from Section 4.2. In particular, each lattice ciphertext in the CRS only consists of the *message-embedding* component. In this case, an element in R_q is represented by a *pair* of integers when $R = \mathbb{Z}[X]/(x^2 + 1)$ and by a *single* integer when $R = \mathbb{Z}$. Moreover, the dimension of the message space depends only on the field size and thus, is the same regardless of whether we work over a base field or an extension field.⁸ As a result, when comparing instantiations over a base field vs. an extension field of similar size, the CRS in the extension field instantiation is longer.

If we consider an alternative instantiation over $\mathbb{F}_p$ where the prover decomposes each $\mathbb{F}_p$ coefficient in the linear PCP proof into two separate coefficients, each of magnitude $\sqrt{p}$, then we can support a larger field size (i.e., up to 38 bits) without requiring a modulus q that exceeds 128 bits. The reduced parameter sizes translate to slightly shorter proofs ($1.1\times$ – $1.2\times$) compared to the

⁷Available here: <https://github.com/lattice-based-zkSNARKS/lattice-zkSNARK>.

⁸The dimension (and size) of the *full* ciphertext is smaller when working over the extension field because the lattice parameters are smaller. If we only consider the message-embedding component of the ciphertext (which is a small fraction of the full ciphertext), then the size is smaller when working over a base field compared to working over the extension field.

Figure 1: Performance comparison for different instantiations of our scheme for supporting R1CS instances of different sizes. The solid lines correspond to our primary instantiations using a linear PCP over $\mathbb{F}_{p^2}$ with a vector encryption scheme over $\mathbb{F}_{p^2}$. The dashed lines represent alternative instantiations using a vector encryption over the base field $\mathbb{F}_p$. In the case where the linear PCP is over the extension field and the vector encryption is over the base field, we first apply Construction C.1 to obtain a linear PCP over the base field. We also consider a direct compilation from a linear PCP over $\mathbb{F}_p$ using a vector encryption scheme over $\mathbb{F}_p$.

Figure 2: CRS size and proof size as a function of the field size $|\mathbb{F}|$, where $\mathbb{F}$ is either a quadratic extension $\mathbb{F}_{p^2}$ or a base field $\mathbb{F}_p$. The characteristic p is chosen so $\mathbb{F}$ has the prescribed size. Parameters based on a SNARK over $\mathbb{F}$ for an R1CS system with 2^{20} constraints. For the $\mathbb{F} = \mathbb{F}_p$ setting, we also consider the case where each coefficient in the linear PCP is represented by two digits, each of size $\sqrt{p}$. Elements with a non-filled marker (and a dotted line) denote parameter settings where the modulus q exceeds 128 bits.

setting without the digit decomposition. However, this comes at the drawback of needing a longer CRS that is $1.7\times$ – $1.8\times$ longer (since each component of the CRS is now decomposed into two components). Indeed, in this setting, the CRS size is comparable to the CRS size for the extension field instantiation; it is slightly worse due to the larger lattice parameters (some of which still scale based on the field characteristic). Despite the improvements in proof size obtained via the digit decomposition, the overall proof size is still $1.5\times$ – $2\times$ longer than the proof size obtained from working over extension fields.

We provide additional benchmarks in Appendix F.

Comparison with other schemes. Finally, we compare the performance of our scheme with the most succinct pairing-based zk-SNARK of Groth [75] as well as several recent post-quantum zk-SNARKs: Ligero [5], Aurora [21], Fractal [47], ethSTARK [15, 101], and Gennaro et al. [65]. With the exception of the lattice-based scheme of Gennaro et al. [65], we measure the performance of each scheme on the same system and with a single-threaded execution. We use libsnark [96] for the implementation of Groth’s pairing-based construction [75] and libiop [95] for the implementations

of Ligero [5], Aurora [21], and Fractal [47]. We use the ethSTARK library [102] for the STARK implementation [101]. For each scheme, we consider the default implementation provided by the library. We note that these schemes export different base fields for the R1CS which makes a direct comparison challenging. With the exception of ethSTARK, we measure the performance of each scheme over their preferred field for an R1CS system with a fixed number of constraints. In the case of ethSTARK, the current implementation only supports verifying a hash chain computation (with the Rescue₁₂₂ hash function [4, 25]). In our benchmarks, we choose the length of the hash chain so that the size of the corresponding R1CS system has the prescribed size. Specifically, the Rescue₁₂₂ hash function consists of $r = 10$ rounds and operates over a state with $m = 12$ field elements. The computation over each round can be encoded as an R1CS system with $2m = 24$ constraints. Thus, each hash computation can be encoded as an R1CS system with 240 constraints. We summarize our benchmarks in Table 1 and refer to Section 1 for further discussion.

5 RELATED WORK

There has been a flurry of recent works studying the asymptotic and concrete efficiency of succinct arguments. We survey several families of constructions here and also include a comparison with several representative schemes in Table 3. In the following, we use N to denote the size of the NP relation being verified.

Linear PCPs and QAP-based constructions. Gennaro et al. [64] and Bitansky et al. [30] described general frameworks for constructing *constant-size* zkSNARKs from linear PCPs (specifically, from QAPs). Several works have extended these frameworks [10, 31, 32, 53, 65, 75]. These constructions are the basis of numerous systems and implementations [9, 18, 20, 23, 24, 42, 48, 54, 58, 59, 87, 105]. These constructions offer the best *succinctness*, but this comes at the expense of needing an expensive, trusted, and language-dependent setup, as well as a quasilinear-time prover.

Interactive oracle proofs. Following the seminal works of Kilian [79] and Micali [84], a recent line of works [13–17, 21, 35, 36, 46, 47, 81] have shown how to construct zkSNARKs from short PCPs [26], and their generalization, interactive oracle proofs (IOPs) [22,

	PQ	TP	PV	Proof Size		Runtime		Cryptographic Structure
				Asymptotic	Concrete	Prover	Verifier	
Groth [75]	○	○	●	1	128 B	$N \log N$	$ x $	Pairings
Marlin [46]	○	●	●	1	704 B	$N \log N$	$ x + \log N$	Pairings
Sonic [83]	○	●	●	1	1.1 KB	$N \log N$	$ x + \log N$	Pairings
Xiphos [99]	○	●	●	$\log N$	61 KB	N	$ x + \log N$	Pairings
Spartan [97]	○	●	●	$\sqrt{N}$	142 KB	N	$ x + \sqrt{N}$	Groups
Fractal [47]	●	●	●	$\log^2 N$	215 KB [†]	$N \log N$	$ x + \log^2 N$	Random Oracle
Gennaro et al. [65] [*]	●	○	○	1	640 KB [‡]	$N \log N$	$ x $	Lattices
STARK [15]	●	●	●	$\log^2 N$	127 KB [§]	$N \text{polylog}(N)$	$ x + \log^2 N$	Random Oracle
This work[*]	●	○	○	1	16 KB	$N \log N$	$ x $	Lattices

^{*}For the *asymptotic* estimates for the lattice-based constructions, we consider an instantiation over a field of size $2^{\Omega(\lambda)}$ (i.e., similar to the field sizes in the group-based and pairing-based constructions).

[†]Proof sizes for Fractal measured using the implementation from libiop [95] with the default configuration over a 181-bit prime field. The largest RICS instance we could measure has 2^{19} constraints, so this is the proof size we report here.

[‡]This number is for a circuit with 2^{16} gates since the paper does not provide measurements for larger circuit sizes.

[§]This is the proof size for verifying a Rescue₁₂₂ hash chain [4, 25] of length 4200 using the ethSTARK implementation [101, 102]. This computation can be expressed as an RICS instance with roughly 2^{20} constraints (see Section 4.3). Since the ethSTARK implementation does not currently support verifying general computations, we do not report performance metrics for the general setting.

Table 3: Comparison with recent zkSNARKs for verifying an NP relations of size N and statements of length $|x|$. For brevity, we focus on schemes that have *sublinear* proof size and *sublinear* verification for general NP relations. Asymptotic running times and parameter sizes are given up to multiplicative $\text{poly}(\lambda)$ factors (where λ is the security parameter). For the “Concrete Proof Size” column, we report the approximate size of a proof for verifying an NP relation of size $N \approx 2^{20}$ at the 128-bit security level (as reported in the respective works unless noted otherwise). The “PQ” column specifies whether the construction is post-quantum secure (●) or only classically secure (○). The “TP” column denotes whether the scheme is transparent (●), relies on a trusted setup for a *universal* CRS (●), or relies on a *trusted* sampling of a *language-dependent* CRS (○). The “PV” column specifies whether the argument is publicly-verifiable (●) or designated-verifier (○). The “Cryptographic Structure” column describes the primary (algebraic) structure underlying the construction. We distinguish between pairing groups and pairing-free groups by using “Groups” to denote the latter.

92]. These constructions rely on the Fiat-Shamir heuristic [57] to obtain a non-interactive argument in the random oracle model. Many IOP constructions have a *transparent* (i.e., non-trusted) setup, and moreover, are plausibly post-quantum. Proof sizes for IOP-based constructions typically range in the hundreds of kilobytes.

Bünz et al. [44] introduced *polynomial IOPs*, a generalization of linear PCPs to the IOP setting, where on each round, the verifier has oracle access to a bounded-degree polynomial. Polynomial IOPs can be compiled into succinct arguments [46, 61, 83, 99] via polynomial commitments. These schemes have excellent concrete succinctness (a few hundred bytes to a few kilobytes), a universal or transparent setup, but generally rely on pre-quantum assumptions.

MPC-in-the-head. Ishai et al. [77] introduced the “MPC-in-the-head” paradigm for building zero-knowledge proofs from general multiparty computation. The Ligero system [5] was the first argument with $\sqrt{N}$ size proofs in this framework. Bhadauria et al. [28] combined Ligero with IOPs to reduce the proof size to $\text{polylog}(N)$. Both constructions support sublinear verification for *structured* circuits, but verification is linear for general circuits.

GKR-constructions. Another line of work starts from the succinct interactive argument for verifying arithmetic circuits by Goldwasser, Kalai, and Rothblum (GKR) [70]. A sequence of works [51, 97, 103, 106, 108, 110, 111] have built on GKR to obtain efficient *non-interactive* arguments for (layered) circuits (and often tailoring to special structures for better concrete efficiency). In these constructions, the size of the proof (and the verifier complexity) typically scale with the depth of the circuit. An appealing feature of these constructions is their low prover complexities: namely, the cost of the prover scale *linearly* in the size of the NP relation (over

large fields). Zhang et al. [109] recently showed how to leverage GKR to verify *general* arithmetic circuits while retaining a linear-time prover and sublinear verification (for structured circuits). The proof size in their construction scales with the depth of the circuit.

Inner product arguments. Building on works by Bayer and Groth [12, 73], Bootle et al. [34] introduced zero-knowledge arguments for arithmetic circuit satisfiability based on *inner product arguments*. Bünz et al. [43] improved the construction to achieve shorter proofs and verification times. While the proofs are short, the verification time scales linearly with the circuit size and these constructions rely on pre-quantum assumptions.

Lattice-based constructions. In the lattice-based setting, there have been several instantiations in the designated-verifier model based on linear PCPs [31, 32, 65]. Baum et al. [11] gave the first publicly-verifiable argument from standard lattice assumptions with $\tilde{O}(\sqrt{N})$ -size proofs. Bootle et al. [37] reduced the proof size further to $\text{polylog}(N)$. In both these cases, the verifier is not succinct and runs in *linear* time.

ACKNOWLEDGMENTS

We thank Brennan Shacklett and Samir Menon for their help with an early prototype implementation of lattice-based zkSNARKs. We thank Eli Ben-Sasson for helpful comments and pointers. Y. Ishai is supported by ERC Project NTSC (742754), BSF grant 2018393, and ISF grant 2774/20. D. J. Wu is supported by NSF CNS-1917414, NSF CNS-2045180, and a Microsoft Research Faculty Fellowship.

REFERENCES

- [1] Martin R. Albrecht, Rachel Player, and Sam Scott. 2015. On the concrete hardness of Learning with Errors. *J. Math. Cryptol.* 9, 3 (2015), 169–203.
- [2] Jacob Alperin-Sheriff and Chris Peikert. 2013. Practical Bootstrapping in Quasi-linear Time. In *CRYPTO*. 1–20.
- [3] Jacob Alperin-Sheriff and Chris Peikert. 2014. Faster Bootstrapping with Polynomial Error. In *CRYPTO*. 297–314.
- [4] Abdelrahman Aly, Tomer Ashur, Eli Ben-Sasson, Siemen Dhooghe, and Alan Szepieniec. 2020. Design of Symmetric-Key Primitives for Advanced Cryptographic Protocols. *IACR Trans. Symmetric Cryptol.* 2020, 3 (2020), 1–45.
- [5] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. 2017. Liger: Lightweight Sublinear Arguments Without a Trusted Setup. In *ACM CCS*. 2087–2104.
- [6] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. 2009. Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems. In *CRYPTO*. 595–618.
- [7] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. 2012. Multiparty Computation with Low Communication, Computation and Interaction via Threshold FHE. In *EUROCRYPT*. 483–501.
- [8] Yonatan Aumann and Yehuda Lindell. 2007. Security Against Covert Adversaries: Efficient Protocols for Realistic Adversaries. In *TCC*. 137–156.
- [9] Michael Backes, Manuel Barbosa, Dario Fiore, and Raphael M. Reischuk. 2015. ADSNARK: Nearly Practical and Privacy-Preserving Proofs on Authenticated Data. In *IEEE Symposium on Security and Privacy*. 271–286.
- [10] Ohad Barta, Yuval Ishai, Rafail Ostrovsky, and David J. Wu. 2020. On Succinct Arguments and Witness Encryption from Groups. In *CRYPTO*. 776–806.
- [11] Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafael del Pino, Jens Groth, and Vadim Lyubashevsky. 2018. Sub-linear Lattice-Based Zero-Knowledge Arguments for Arithmetic Circuits. In *CRYPTO*. 669–699.
- [12] Stephanie Bayer and Jens Groth. 2012. Efficient Zero-Knowledge Argument for Correctness of a Shuffle. In *EUROCRYPT*. 263–280.
- [13] Eli Ben-Sasson, Iddo Bentov, Alessandro Chiesa, Ariel Gabizon, Daniel Genkin, Matan Hamilis, Evgenya Pergament, Michael Riabzev, Mark Silberstein, Eran Tromer, and Madars Virza. 2017. Computational Integrity with a Public Random String from Quasi-Linear PCPs. In *EUROCRYPT*. 551–579.
- [14] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. 2018. Fast Reed-Solomon Interactive Oracle Proofs of Proximity. In *ICALP*. 14:1–14:17.
- [15] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. 2018. Scalable, transparent, and post-quantum secure computational integrity. *IACR Cryptol. ePrint Arch.* 2018 (2018), 46.
- [16] Eli Ben-Sasson, Alessandro Chiesa, Michael A. Forbes, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. 2017. Zero Knowledge Protocols from Succinct Constraint Detection. In *TCC*. 172–206.
- [17] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, and Madars Virza. 2016. Quasi-Linear Size Zero Knowledge from Linear-Algebraic PCPs. In *TCC*. 33–64.
- [18] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. 2014. Zerocash: Decentralized Anonymous Payments from Bitcoin. In *IEEE Symposium on Security and Privacy*. 459–474.
- [19] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. 2013. Fast reductions from RAMs to delegatable succinct constraint satisfaction problems: extended abstract. In *ITCS*. 401–414.
- [20] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. 2013. SNARKs for C: Verifying Program Executions Succinctly and in Zero Knowledge. In *CRYPTO*. 90–108.
- [21] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. 2019. Aurora: Transparent Succinct Arguments for R1CS. In *EUROCRYPT*. 103–128.
- [22] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. 2016. Interactive Oracle Proofs. In *TCC*. 31–60.
- [23] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. 2014. Scalable Zero Knowledge via Cycles of Elliptic Curves. In *CRYPTO*. 276–294.
- [24] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. 2014. Succinct Non-Interactive Zero Knowledge for a von Neumann Architecture. In *USENIX Security Symposium*. 781–796.
- [25] Eli Ben-Sasson, Lior Goldberg, and David Levit. 2020. STARK Friendly Hash - Survey and Recommendation. *IACR Cryptol. ePrint Arch.* 2020 (2020), 948.
- [26] Eli Ben-Sasson and Madhu Sudan. 2008. Short PCPs with Polylog Query Complexity. *SIAM J. Comput.* 38, 2 (2008), 551–607.
- [27] Daniel J. Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Michael Schneider, Peter Schwabe, and Zooko Wilcox-O’Hearn. 2015. SPHINCS: Practical Stateless Hash-Based Signatures. In *EUROCRYPT*. 368–397.
- [28] Rishabh Bhaduria, Zhiyong Fang, Carmit Hazay, Muthuramakrishnan Venkatasubramanian, Tiancheng Xie, and Yupeng Zhang. 2020. Liger++: A New Optimized Sublinear IOP. In *ACM CCS*. 2025–2038.
- [29] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. 2013. Recursive composition and bootstrapping for SNARKS and proof-carrying data. In *STOC*. 111–120.
- [30] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. 2013. Succinct Non-interactive Arguments via Linear Interactive Proofs. In *TCC*. 315–333.
- [31] Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. 2017. Lattice-Based SNARKs and Their Application to More Efficient Obfuscation. In *EUROCRYPT*. 247–277.
- [32] Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. 2018. Quasi-Optimal SNARKs via Linear Multi-Prover Interactive Proofs. In *EUROCRYPT*. 222–255.
- [33] Dan Boneh, Ben Lynn, and Hovav Shacham. 2001. Short Signatures from the Weil Pairing. In *ASIACRYPT*. 514–532.
- [34] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. 2016. Efficient Zero-Knowledge Arguments for Arithmetic Circuits in the Discrete Log Setting. In *EUROCRYPT*. 327–357.
- [35] Jonathan Bootle, Alessandro Chiesa, and Jens Groth. 2020. Linear-Time Arguments with Sublinear Verification from Tensor Codes. In *TCC*. 19–46.
- [36] Jonathan Bootle, Alessandro Chiesa, and Siqi Liu. 2020. Zero-Knowledge Succinct Arguments with a Linear-Time Prover. *IACR Cryptol. ePrint Arch.* 2020 (2020), 1527.
- [37] Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. 2020. A Non-PCP Approach to Succinct Quantum-Safe Zero-Knowledge. In *CRYPTO*. 441–469.
- [38] Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. 2016. Frodo: Take off the Ring! Practical, Quantum-Secure Key Exchange from LWE. In *ACM CCS*. 1006–1018.
- [39] Alin Bostan and Éric Schost. 2005. Polynomial evaluation and interpolation on special sets of points. *J. Complex.* 21, 4 (2005), 420–446.
- [40] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. 2012. (Leveled) fully homomorphic encryption without bootstrapping. In *ITCS*. 309–325.
- [41] Zvika Brakerski and Vinod Vaikuntanathan. 2011. Efficient Fully Homomorphic Encryption from (Standard) LWE. In *FOCS*. 97–106.
- [42] Benjamin Braun, Ariel J. Feldman, Zuoqiang Ren, Sriniath T. V. Setty, Andrew J. Blumberg, and Michael Walfish. 2013. Verifying computations with state. In *SOSP*. 341–357.
- [43] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. 2018. Bulletproofs: Short Proofs for Confidential Transactions and More. In *IEEE Symposium on Security and Privacy*. 315–334.
- [44] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. 2020. Transparent SNARKs from DARK Compilers. In *EUROCRYPT*. 677–706.
- [45] Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. 2017. Post-Quantum Zero-Knowledge and Signatures from Symmetric-Key Primitives. In *ACM CCS*. 1825–1842.
- [46] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas P. Ward. 2020. Marlin: Preprocessing zkSNARKs with Universal and Updatable SRS. In *EUROCRYPT*. 738–768.
- [47] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. 2020. Fractal: Post-quantum and Transparent Recursive Proofs from Holography. In *EUROCRYPT*. 769–793.
- [48] Alessandro Chiesa, Eran Tromer, and Madars Virza. 2015. Cluster Computing in Zero Knowledge. In *EUROCRYPT*. 371–403.
- [49] Alessandro Chiesa and Eylon Yogev. 2021. Subquadratic SNARKs in the Random Oracle Model. (2021).
- [50] James W Cooley and John W Tukey. 1965. An algorithm for the machine calculation of complex Fourier series. *Mathematics of computation* 19, 90 (1965), 297–301.
- [51] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. 2012. Practical verified computation with streaming interactive proofs. In *ITCS*. 90–112.
- [52] Jean-Sébastien Coron, David Naccache, and Mehdi Tibouchi. 2012. Public Key Compression and Modulus Switching for Fully Homomorphic Encryption over the Integers. In *EUROCRYPT*. 446–464.
- [53] George Danezis, Cédric Fournet, Jens Groth, and Markulf Kohlweiss. 2014. Square Span Programs with Applications to Succinct NIZK Arguments. In *ASIACRYPT*. 532–550.
- [54] Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, and Bryan Parno. 2016. Cinderella: Turning Shabby X.509 Certificates into Elegant Anonymous Credentials with the Magic of Verifiable Computation. In *IEEE Symposium on Security and Privacy*. 235–254.
- [55] Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. 2018. CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2018, 1 (2018), 238–268.
- [56] Léo Ducas and Daniele Micciancio. 2015. FHEW: Bootstrapping Homomorphic Encryption in Less Than a Second. In *EUROCRYPT*. 617–640.
- [57] Amos Fiat and Adi Shamir. 1986. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *CRYPTO*. 186–194.

- [58] Dario Fiore, Cédric Fournet, Esha Ghosh, Markulf Kohlweiss, Olga Ohrimenko, and Bryan Parno. 2016. Hash First, Argue Later: Adaptive Verifiable Computations on Outsourced Data. In *ACM CCS*. 1304–1316.
- [59] Dario Fiore, Rosario Gennaro, and Valerio Pastro. 2014. Efficiently Verifiable Computation on Encrypted Data. In *ACM CCS*. 844–855.
- [60] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. 2020. Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU (Specification v1.2). (2020).
- [61] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. 2019. PLONK: Permutations over Lagrange-bases for Oecumenical Noninteractive arguments of Knowledge. *IACR Cryptol. ePrint Arch.* 2019 (2019), 953.
- [62] Steven D Galbraith. 2013. Space-efficient variants of cryptosystems based on learning with errors. (2013).
- [63] Chaya Ganesh, Anca Nitulescu, and Eduardo Soria-Vazquez. 2021. Rinocchio: SNARKs for Ring Arithmetic. *IACR Cryptol. ePrint Arch.* 2021 (2021), 322.
- [64] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. 2013. Quadratic Span Programs and Succinct NIZKs without PCPs. In *EUROCRYPT*. 626–645.
- [65] Rosario Gennaro, Michele Minelli, Anca Nitulescu, and Michele Orrù. 2018. Lattice-Based zk-SNARKs from Square Span Programs. In *ACM CCS*. 556–573.
- [66] Craig Gentry. 2009. *A fully homomorphic encryption scheme*. Ph.D. Dissertation. Stanford University. crypto.stanford.edu/craig.
- [67] Craig Gentry, Shai Halevi, and Nigel P. Smart. 2012. Fully Homomorphic Encryption with Polylog Overhead. In *EUROCRYPT*. 465–482.
- [68] Craig Gentry, Shai Halevi, and Nigel P. Smart. 2012. Homomorphic Evaluation of the AES Circuit. In *CRYPTO*. 850–867.
- [69] Craig Gentry and Daniel Wichs. 2011. Separating succinct non-interactive arguments from all falsifiable assumptions. In *STOC*. 99–108.
- [70] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. 2008. Delegating computation: interactive proofs for muggles. In *STOC*. 113–122.
- [71] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. 1985. The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract). In *STOC*. 291–304.
- [72] Irving John Good. 1958. The interaction algorithm and practical Fourier analysis. *Journal of the Royal Statistical Society: Series B (Methodological)* 20, 2 (1958), 361–372.
- [73] Jens Groth. 2009. Linear Algebra with Sub-linear Zero-Knowledge Arguments. In *CRYPTO*. 192–208.
- [74] Jens Groth. 2010. Short Pairing-Based Non-interactive Zero-Knowledge Arguments. In *ASIACRYPT*. 321–340.
- [75] Jens Groth. 2016. On the Size of Pairing-Based Non-interactive Arguments. In *EUROCRYPT*. 305–326.
- [76] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. 2007. Efficient Arguments without Short PCPs. In *CCC*. 278–291.
- [77] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. 2007. Zero-knowledge from secure multiparty computation. In *STOC*. 21–30.
- [78] Yuval Ishai, Hang Su, and David J. Wu. 2021. Shorter and Faster Post-Quantum Designated-Verifier zkSNARKs from Lattices. *IACR Cryptol. ePrint Arch.* 2021 (2021), 1.
- [79] Joe Kilian. 1992. A Note on Efficient Zero-Knowledge Proofs and Arguments (Extended Abstract). In *STOC*. 723–732.
- [80] Adeline Langlois and Damien Stehlé. 2015. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.* 75, 3 (2015), 565–599.
- [81] Jonathan Lee, Srinath Setty, Justin Thaler, and Riad Wahby. 2021. Linear-time zero-knowledge SNARKs for R1CS. *IACR Cryptol. ePrint Arch.* 2021 (2021), 30.
- [82] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. 2010. On Ideal Lattices and Learning with Errors over Rings. In *EUROCRYPT*. 1–23.
- [83] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. 2019. Sonic: Zero-Knowledge SNARKs from Linear-Size Universal and Updateable Structured Reference Strings. *IACR Cryptol. ePrint Arch.* 2019 (2019), 99.
- [84] Silvio Micali. 2000. Computationally Sound Proofs. *SIAM J. Comput.* 30, 4 (2000), 1253–1298.
- [85] Pratyay Mukherjee and Daniel Wichs. 2016. Two Round Multiparty Computation via Multi-key FHE. In *EUROCRYPT*. 735–763.
- [86] Anca Nitulescu. 2019. Lattice-Based Zero-Knowledge SNARKs for Arithmetic Circuits. In *LATINCRYPT*. 217–236.
- [87] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. 2013. Pinocchio: Nearly Practical Verifiable Computation. In *IEEE Symposium on Security and Privacy*. 238–252.
- [88] Chris Peikert. 2016. A Decade of Lattice Cryptography. *Found. Trends Theor. Comput. Sci.* 10, 4 (2016), 283–424.
- [89] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. 2008. A Framework for Efficient and Composable Oblivious Transfer. In *CRYPTO*. 554–571.
- [90] Charles M Rader. 1968. Discrete Fourier transforms when the number of data samples is prime. *Proc. IEEE* 56, 6 (1968), 1107–1108.
- [91] Oded Regev. 2005. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*. 84–93.
- [92] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. 2016. Constant-round interactive proofs for delegating computation. In *STOC*. 49–62.
- [93] Jacob T. Schwartz. 1980. Fast Probabilistic Algorithms for Verification of Polynomial Identities. *J. ACM* 27, 4 (1980).
- [94] SCIPR Lab. 2021. libfqfft: C++ library for FFTs in finite fields. <https://github.com/scipr-lab/libfqfft/>.
- [95] SCIPR Lab. 2021. libiop: a C++ library for IOP-based zkSNARKs. <https://github.com/scipr-lab/libiop>.
- [96] SCIPR Lab. 2021. libsnark: a C++ library for zkSNARK proofs. <https://github.com/scipr-lab/libsnark/>.
- [97] Srinath Setty. 2020. Spartan: Efficient and General-Purpose zkSNARKs Without Trusted Setup. In *CRYPTO*. 704–737.
- [98] Srinath T. V. Setty, Benjamin Braun, Victor Vu, Andrew J. Blumberg, Bryan Parno, and Michael Walfish. 2013. Resolving the conflict between generality and plausibility in verified computation. In *EuroSys*. 71–84.
- [99] Srinath T. V. Setty and Jonathan Lee. 2020. Quarks: Quadruple-efficient transparent zkSNARKs. *IACR Cryptol. ePrint Arch.* 2020 (2020), 1275.
- [100] Srinath T. V. Setty, Victor Vu, Nikhil Panpalia, Benjamin Braun, Andrew J. Blumberg, and Michael Walfish. 2012. Taking Proof-Based Verified Computation a Few Steps Closer to Practicality. In *USENIX*. 253–268.
- [101] StarkWare Team. 2021. ethSTARK Documentation. *IACR Cryptol. ePrint Arch.* 2021 (2021), 582.
- [102] StarkWare Team. 2021. ethSTARK. <https://github.com/starkware-libs/ethSTARK>.
- [103] Justin Thaler. 2013. Time-Optimal Interactive Proofs for Circuit Evaluation. In *CRYPTO*. 71–89.
- [104] Llewellyn H Thomas. 1963. Using a computer to solve problems in physics. *Applications of digital computers* (1963), 44–45.
- [105] Riad S. Wahby, Srinath T. V. Setty, Zuoqiang Ren, Andrew J. Blumberg, and Michael Walfish. 2015. Efficient RAM and control flow in verifiable outsourced computation. In *NDSS*.
- [106] Riad S. Wahby, Ioanna Tzialla, Abhi Shelat, Justin Thaler, and Michael Walfish. 2018. Doubly-Efficient zkSNARKs Without Trusted Setup. In *IEEE Symposium on Security and Privacy*. 926–943.
- [107] Michael Walfish and Andrew J. Blumberg. 2015. Verifying computations without reexecuting them. *Commun. ACM* 58, 2 (2015), 74–84.
- [108] Tiancheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. 2019. Libra: Succinct Zero-Knowledge Proofs with Optimal Prover Computation. In *CRYPTO*. 733–764.
- [109] Jiaheng Zhang, Weijie Wang, Yinyu Zhang, and Yupeng Zhang. 2020. Doubly Efficient Interactive Proofs for General Arithmetic Circuits with Linear Prover Time. *IACR Cryptol. ePrint Arch.* 2020 (2020), 1247.
- [110] Jiaheng Zhang, Tiancheng Xie, Yupeng Zhang, and Dawn Song. 2020. Transparent Polynomial Delegation and Its Applications to Zero Knowledge Proof. In *IEEE Symposium on Security and Privacy*. 859–876.
- [111] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. 2017. A Zero-Knowledge Version of vSQL. *IACR Cryptol. ePrint Arch.* 2017 (2017), 1146.
- [112] Richard Zippel. 1979. Probabilistic algorithms for sparse polynomials. In *EUROSAM*.

A ADDITIONAL PRELIMINARIES

In this section, we recall additional preliminaries.

Lemma A.1 (Schwartz-Zippel [93, 112]). *Let $f \in \mathbb{F}[x_1, \dots, x_n]$ be a multivariate polynomial of total degree at most d over $\mathbb{F}$, not identically zero. Then for any set $S \subseteq \mathbb{F}$,*

$$\Pr[f(\alpha_1, \dots, \alpha_n) = 0 \mid \alpha_1, \dots, \alpha_n \xleftarrow{R} S] \leq d/|S|.$$

Lemma A.2 (Smudging Lemma). *Let B, B' be integers. Fix any value $|e_1| \leq B'$ and sample $e_2 \xleftarrow{R} [-B, B]$. The statistical distance between the distributions of $e_1 + e_2$ and e_2 is at most B'/B .*

Discrete Gaussians and tail bounds. We also recall some preliminaries on the discrete Gaussian distribution. We refer to Peikert’s survey [88] for additional details and references. For a real value $s > 0$, the Gaussian function $\rho_s : \mathbb{R} \rightarrow \mathbb{R}^+$ with width s is the function $\rho_s(x) := \exp(-\pi x^2/s^2)$. The discrete Gaussian distribution $D_{\mathbb{Z},s}$ over $\mathbb{Z}$ with mean 0 and width s is the distribution where

$$\Pr[X = x : X \leftarrow D_{\mathbb{Z},s}] = \frac{\rho_s(x)}{\sum_{y \in \mathbb{Z}} \rho_s(y)}. \quad (\text{A.1})$$

A real random variable X is *subgaussian* with parameter s if for every $t \geq 0$, $\Pr[|X| > t] \leq 2 \exp(-\pi t^2/s^2)$. The following two facts will be useful in our analysis.

- If X is subgaussian with parameter s and $a \in \mathbb{R}$, then aX is subgaussian with parameter $|a|s$.
- If $X_1, \dots, X_m$ are independent subgaussian random variables with parameters $s_1, \dots, s_m$, respectively, then $\sum_{i \in [m]} X_i$ is subgaussian with parameter $\|\mathbf{s}\|_2$ where $\mathbf{s} = (s_1, \dots, s_m)$.

Rank-1 constraint satisfiability. We recall the definition of the R1CS language introduced implicitly by Gennaro et al. [64] and formalized explicitly in [20, 21, 98]:

Definition A.3 (Rank-1 Constraint Satisfiability [21, 64, 98]). A rank-1 constraint satisfiability (R1CS) system over a finite field $\mathbb{F}$ is specified by a tuple $CS = (n, N_g, N_w, \{\mathbf{a}_i, \mathbf{b}_i, \mathbf{c}_i\}_{i \in [N_g]})$ where $n, N_g, N_w \in \mathbb{N}$, $n \leq N_w$, and $\mathbf{a}_i, \mathbf{b}_i, \mathbf{c}_i \in \mathbb{F}^{N_w+1}$. The system CS is *satisfiable* for a *statement* $\mathbf{x} \in \mathbb{F}^n$ if there exists a *witness* $\mathbf{w} \in \mathbb{F}^{N_w}$ such that

- $\mathbf{x} = (w_1, \dots, w_n)$ and
- $[1 \mid \mathbf{w}^T] \mathbf{a}_i \cdot [1 \mid \mathbf{w}^T] \mathbf{b}_i = [1 \mid \mathbf{w}^T] \mathbf{c}_i$ for all $i \in [N_g]$.

We denote this by writing $CS(\mathbf{x}, \mathbf{w}) = 1$, and refer to n as the statement size, N_w as the number of variables, and N_g as the number of constraints. Given an R1CS system CS , we define the corresponding relation $\mathcal{R}_{CS} = \{(\mathbf{x}, \mathbf{w}) \in \mathbb{F}^n \times \mathbb{F}^{N_w} : CS(\mathbf{x}, \mathbf{w}) = 1\}$.

Remark A.4 (Boolean and Arithmetic Circuit Satisfiability). As shown in [20, 64], the language of R1CS capture Boolean and arithmetic circuit satisfiability as special cases. Namely, a Boolean circuit satisfiability instance for a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ with α wires and β bilinear gates yields an R1CS instance with $N_w = \alpha$ variables and $N_g = \beta + h + 1$ constraints. Similarly, an arithmetic circuit $C: \mathbb{F}^n \times \mathbb{F}^h \rightarrow \mathbb{F}^\ell$ with α wires and β bilinear gates corresponds to an R1CS instance with $N_w = \alpha$ variables and $N_g = \beta + \ell$ constraints. In this work, we focus exclusively on linear PCPs and SNARKs for R1CS.

A.1 Linear PCPs

We now recall the notion of a linear PCP (LPCP) from [30, 76]. In this work, we only consider linear PCPs for R1CS systems, so we specialize all of our definitions to this setting:

Definition A.5 (Linear PCP [30, 76, adapted]). Let p be a polynomial and let $CS = \{CS_\kappa\}_{\kappa \in \mathbb{N}}$ be a family of R1CS systems over a finite field $\mathbb{F}$ where $CS_\kappa = (n_\kappa, N_{g,\kappa}, N_{w,\kappa}, \{\mathbf{a}_{i,\kappa}, \mathbf{b}_{i,\kappa}, \mathbf{c}_{i,\kappa}\}_{i \in [N_{g,\kappa}]})$ has size at most $|CS_\kappa| \leq p(\kappa)$. In the following, we write $n = n(\kappa)$ to denote a polynomially-bounded function where $n(\kappa) = n_\kappa$ for all $\kappa \in \mathbb{N}$. We define $N_g = N_g(\kappa)$ and $N_w = N_w(\kappa)$ similarly. A k -query input-independent linear PCP for CS with query length $\ell = \ell(\kappa)$ and knowledge error $\varepsilon = \varepsilon(\kappa)$ is a tuple of algorithms $\Pi_{LPCP} = (Q_{LPCP}, P_{LPCP}, V_{LPCP})$ with the following properties:

- $Q_{LPCP}(1^\kappa) \rightarrow (\text{st}, Q)$: The query-generation algorithm takes as input the system index $\kappa \in \mathbb{N}$ and outputs a query matrix $Q \in \mathbb{F}^{\ell \times k}$ and a verification state st .
- $P_{LPCP}(1^\kappa, \mathbf{x}, \mathbf{w}) \rightarrow \pi$: On input the system index $\kappa \in \mathbb{N}$, a statement $\mathbf{x} \in \mathbb{F}^n$, and a witness $\mathbf{w} \in \mathbb{F}^{N_w}$, the prove algorithm outputs a proof $\pi \in \mathbb{F}^\ell$.

- $V_{LPCP}(\text{st}, \mathbf{x}, \mathbf{a})$: On input the verification state st , the statement $\mathbf{x} \in \mathbb{F}^n$, and a vector of responses $\mathbf{a} \in \mathbb{F}^k$, the verification algorithm outputs a bit $b \in \{0, 1\}$.

In addition, Π_{LPCP} should satisfy the following properties:

- **Completeness:** For all $\kappa \in \mathbb{N}$, $\mathbf{x} \in \mathbb{F}^n$, and $\mathbf{w} \in \mathbb{F}^{N_w}$ where $CS_\kappa(\mathbf{x}, \mathbf{w}) = 1$,

$$\Pr \left[V_{LPCP}(\text{st}, \mathbf{x}, Q^T \pi) = 1 \mid \begin{array}{l} (\text{st}, Q) \leftarrow Q_{LPCP}(1^\kappa), \\ \pi \leftarrow P_{LPCP}(1^\kappa, \mathbf{x}, \mathbf{w}) \end{array} \right] = 1.$$

- **Knowledge:** There exists an efficient extractor $\mathcal{E}_{LPCP}$ such that for all $\kappa \in \mathbb{N}$, $\mathbf{x} \in \mathbb{F}^n$, and $\pi^* \in \mathbb{F}^\ell$, if

$$\Pr[V_{LPCP}(\text{st}, \mathbf{x}, Q^T \pi^*) = 1 \mid (\text{st}, Q) \leftarrow Q_{LPCP}(1^\kappa)] > \varepsilon,$$

then

$$\Pr[CS_\kappa(\mathbf{x}, \mathbf{w}) = 1 \mid \mathbf{w} \leftarrow \mathcal{E}_{LPCP}^{(\pi^*, \cdot)}(1^\kappa, \mathbf{x})] = 1.$$

We refer to ε as the *knowledge error* of the linear PCP.

- **Perfect honest-verifier zero knowledge (HVZK):** There exists an efficient simulator $S_{LPCP} = (S_1, S_2)$ such that for all $\kappa \in \mathbb{N}$ and all instances $(\mathbf{x}, \mathbf{w})$ where $CS_\kappa(\mathbf{x}, \mathbf{w}) = 1$,

$$\{(\text{st}, Q, Q^T \pi)\} \equiv \{(\tilde{\text{st}}, \tilde{Q}, \tilde{\pi})\},$$

where $(\text{st}, Q) \leftarrow Q_{LPCP}(1^\kappa)$, $\pi \leftarrow P_{LPCP}(1^\kappa, \mathbf{x}, \mathbf{w})$, $(\tilde{\text{st}}, \tilde{Q}, \tilde{\pi}) \leftarrow S_1(1^\kappa)$, and $\tilde{\pi} \leftarrow S_2(\text{st}, \mathbf{x})$.

Linear PCPs for R1CS. The quadratic arithmetic programs (QAPs) introduced by Gennaro et al. [64] immediately imply a 4-query linear PCP for R1CS [20]. Note that Ben-Sasson et al. [20] described the construction as a 5-query linear PCP with statistical HVZK (over large fields); however, it is straightforward to adapt the construction to obtain a 4-query LPCP with perfect HVZK (over any field). These changes incur a slight increase in the verification complexity and the knowledge error. We state the main result below and describe the construction from [20] and our modifications in Appendix B.

Claim A.6 (Linear PCPs for R1CS [20, 64, adapted]). Let $CS = \{CS_\kappa\}_{\kappa \in \mathbb{N}}$ be a family of R1CS instances over a finite field $\mathbb{F}$, where $CS_\kappa = (n_\kappa, N_{g,\kappa}, N_{w,\kappa}, \{\mathbf{a}_{i,\kappa}, \mathbf{b}_{i,\kappa}, \mathbf{c}_{i,\kappa}\}_{i \in [N_{g,\kappa}]})$. We write $n = n(\kappa)$ to denote a function where $n(\kappa) = n_\kappa$ for all $\kappa \in \mathbb{N}$; we define $N_g = N_g(\kappa)$ and $N_w = N_w(\kappa)$ correspondingly. Then, there exists a 4-query linear PCP for CS with knowledge error $2N_g/(|\mathbb{F}| - N_g)$, query length $4 + N_w + N_g - n$, and satisfying perfect HVZK.

Remark A.7 (Knowledge Amplification for Linear PCPs). Claim A.6 gives a 4-query linear PCP for any R1CS system with N_g constraints that has knowledge error $\varepsilon = 2N_g/(|\mathbb{F}| - N_g)$. To achieve negligible knowledge error, this necessitates working over a field of super-polynomial size. In our lattice-based instantiation, it is more efficient to work over smaller fields. To amplify knowledge, we use standard parallel repetition. Namely, for a k -query LPCP with query length m and knowledge error ε , we can obtain a $(k\rho)$ -query LPCP with the same query length and knowledge error ε^ρ . In more detail, the setup algorithm samples ρ independent sets of queries $Q_1, \dots, Q_\rho \in \mathbb{F}^{m \times k}$ and constructs its query matrix Q as $Q = [Q_1 \mid \dots \mid Q_\rho] \in \mathbb{F}^{m \times k\rho}$. The verifier accepts a response $\mathbf{a} = [\mathbf{a}_1 \mid \dots \mid \mathbf{a}_\rho]$ only if all ρ sets of responses are valid.

A.2 Succinct Non-Interactive Arguments

We recall the definitions of a succinct non-interactive argument of knowledge (SNARK) for R1CS:

Definition A.8 (Succinct Non-Interactive Argument of Knowledge). Let $CS = \{CS_\kappa\}_{\kappa \in \mathbb{N}}$ be a family of R1CS systems over a finite field $\mathbb{F}$, where $|CS_\kappa| \leq s(\kappa)$ for some fixed polynomial $s(\cdot)$. A succinct non-interactive argument (SNARK) in the preprocessing model⁹ for CS is a tuple $\Pi_{\text{SNARK}} = (\text{Setup}, \text{Prove}, \text{Verify})$ with the following properties:

- $\text{Setup}(1^\lambda, 1^\kappa) \rightarrow (\text{crs}, \text{st})$: On input the security parameter λ and the system index κ , the setup algorithm outputs a common reference string crs and verification state st .
- $\text{Prove}(\text{crs}, \mathbf{x}, \mathbf{w}) \rightarrow \pi$: On input a common reference string crs , a statement $\mathbf{x}$, and a witness $\mathbf{w}$, the prove algorithm outputs a proof π .
- $\text{Verify}(\text{st}, \mathbf{x}, \pi) \rightarrow \{0, 1\}$: On input the verification state st , a statement $\mathbf{x}$ and a proof π , the verification algorithm outputs a bit $b \in \{0, 1\}$.

Moreover, Π_{SNARK} should satisfy the following properties:

- **Completeness**: For all security parameters $\lambda \in \mathbb{N}$, system indices $\kappa \in \mathbb{N}$, and instances $(\mathbf{x}, \mathbf{w})$ where $CS_\kappa(\mathbf{x}, \mathbf{w}) = 1$,

$$\Pr[\text{Verify}(\text{st}, \mathbf{x}, \pi) = 1] = 1,$$

where $(\text{crs}, \text{st}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa)$, $\pi \leftarrow \text{Prove}(\text{crs}, \mathbf{x}, \mathbf{w})$.

- **Knowledge**: For all polynomial-size provers $\mathcal{P}^*$, there exists a polynomial-size extractor $\mathcal{E}$ such that for all security parameters $\lambda \in \mathbb{N}$, system indices $\kappa \in \mathbb{N}$, and auxiliary inputs $z \in \{0, 1\}^{\text{poly}(\lambda)}$,

$$\Pr[\text{Verify}(\text{st}, \mathbf{x}, \pi) = 1 \wedge CS_\kappa(\mathbf{x}, \mathbf{w}) \neq 1] = \text{negl}(\lambda),$$

where $(\text{crs}, \text{st}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa)$, $(\mathbf{x}, \pi) \leftarrow \mathcal{P}^*(1^\lambda, 1^\kappa, \text{crs}; z)$, and $\mathbf{w} \leftarrow \mathcal{E}(1^\lambda, 1^\kappa, \text{crs}, \text{st}, \mathbf{x}; z)$.

- **Efficiency**: There exist a universal polynomial p (independent of CS) such that Setup and Prove run in time $p(\lambda + |CS_\kappa|)$, Verify runs in time $p(\lambda + |\mathbf{x}| + \log |CS_\kappa|)$, and the proof size is $p(\lambda + \log |CS_\kappa|)$.

Remark A.9 (Public Verification vs. Designated Verifier). We say a SNARK is *publicly-verifiable* if st can be efficiently computed from crs (i.e., verification only depends on the public common reference string). Otherwise, the SNARK is *designated-verifier* (i.e., only the holder of the *secret* verification state st can check proofs). In this work, we focus on designated-verifier SNARKs.

Definition A.10 (Zero Knowledge). A SNARK $\Pi_{\text{SNARK}} = (\text{Setup}, \text{Prove}, \text{Verify})$ for an R1CS system $CS = \{CS_\kappa\}_{\kappa \in \mathbb{N}}$ is computational zero knowledge (i.e., a zkSNARK) if there exists an efficient simulator $\mathcal{S}_{\text{SNARK}} = (\mathcal{S}_1, \mathcal{S}_2)$ such that for all $\kappa \in \mathbb{N}$ and all efficient and stateful adversaries $\mathcal{A}$, we have that

$$\Pr[\text{ExptZK}_{\Pi_{\text{SNARK}}, \mathcal{A}, \mathcal{S}_{\text{SNARK}}}(1^\lambda, 1^\kappa) = 1] \leq 1/2 + \text{negl}(\lambda), \quad (\text{A.2})$$

where the experiment $\text{ExptZK}_{\Pi_{\text{SNARK}}, \mathcal{A}, \mathcal{S}_{\text{SNARK}}}(1^\lambda, 1^\kappa)$ is defined as follows:

⁹In the preprocessing model, we allow for a *statement-independent* setup algorithm that runs in time polynomial in the size of the instance CS_κ . In contrast, a “fully-succinct” SNARK also requires that the setup run in time sublinear (or polylogarithmic) in the size of CS_κ . Using recursive composition [29], it is possible to obtain fully succinct SNARKs from preprocessing SNARKs.

- (1) The challenger samples $b \xleftarrow{\mathcal{R}} \{0, 1\}$. If $b = 0$, the challenger computes $(\text{crs}, \text{st}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa)$ and gives (crs, st) to $\mathcal{A}$. If $b = 1$, the challenger computes $(\widetilde{\text{crs}}, \widetilde{\text{st}}, \text{st}_S) \leftarrow \mathcal{S}_1(1^\lambda, 1^\kappa)$ and gives $(\widetilde{\text{crs}}, \widetilde{\text{st}})$ to $\mathcal{A}$.
- (2) The adversary $\mathcal{A}$ outputs a statement $\mathbf{x}$ and a witness $\mathbf{w}$.
- (3) If $CS_\kappa(\mathbf{x}, \mathbf{w}) \neq 1$, then the experiment halts with output 0. Otherwise, the challenger proceeds as follows:
 - If $b = 0$, the challenger replies with $\pi \leftarrow \text{Prove}(\text{crs}, \mathbf{x}, \mathbf{w})$.
 - If $b = 1$, the challenger replies with $\tilde{\pi} \leftarrow \mathcal{S}_2(\text{st}_S, \mathbf{x})$.
 At the end of the experiment, $\mathcal{A}$ outputs a bit $b' \in \{0, 1\}$. The output of the experiment is 1 if $b' = b$ and is 0 otherwise.

When the probability in Eq. (A.2) is bounded by $1/2 + \epsilon$, we say that the scheme satisfies ϵ -computational zero knowledge.

B LINEAR PCP FOR R1CS

In this section, we describe the linear PCP we use for R1CS. The construction is based on the quadratic arithmetic programs of Genaro et al. [64], and is adapted from the 5-query linear PCP construction by Ben-Sasson et al. [20]. There are two minor differences in our construction:

- We remove the statement-dependent query and have the verifier introduce the statement-dependent components during verification. This yields a *4-query* linear PCP with shorter query length at the expense of a slightly more expensive verification step. A similar approach is used implicitly in [24, 64].
- The LPCP query-generation samples the random point from a smaller subset of the field. This introduces some knowledge error, but enables *perfect* HVZK. The construction of Ben-Sasson et al. provided statistical HVZK where the statistical distance was inversely proportional to the field size. The difference between statistical HVZK and perfect HVZK is negligible for super-polynomial size fields, but not for the moderate-size fields we use in this work.

For completeness, we provide the full description and analysis below. Our presentation and analysis is adapted from [20, Appendix E].

Construction B.1 (Linear PCP for R1CS [20, 64, adapted]). Let $CS = \{CS_\kappa\}_{\kappa \in \mathbb{N}}$ be a family of R1CS instances over a finite field $\mathbb{F}$, where $CS_\kappa = (n_\kappa, N_g, N_w, \kappa, \{\mathbf{a}_{i,\kappa}, \mathbf{b}_{i,\kappa}, \mathbf{c}_{i,\kappa}\}_{i \in [N_g]})$, $\mathbf{a}_{i,\kappa}, \mathbf{b}_{i,\kappa}, \mathbf{c}_{i,\kappa} \in \mathbb{F}^{N_{w,\kappa}+1}$ (and entries indexed from 0 to $N_{w,\kappa}$). For notational convenience, we write $n = n(\kappa)$ to denote a function where $n(\kappa) = n_\kappa$ for all $\kappa \in \mathbb{N}$. We define $N_g = N_g(\kappa)$, $N_w = N_w(\kappa)$, $\mathbf{a}_i = \mathbf{a}_i(\kappa)$, $\mathbf{b}_i = \mathbf{b}_i(\kappa)$ and $\mathbf{c}_i = \mathbf{c}_i(\kappa)$ similarly. We additionally define the following components:

- Let $S = \{\alpha_1, \dots, \alpha_{N_g}\} \subset \mathbb{F}$ be an arbitrary subset of $\mathbb{F}$.
- For each $i \in \{0, \dots, N_w\}$, let $A_i, B_i, C_i: \mathbb{F} \rightarrow \mathbb{F}$ be the unique polynomial of degree $N_g - 1$ where for all $j \in [N_g]$,

$$A_i(\alpha_j) = \mathbf{a}_{j,i}, \quad B_i(\alpha_j) = \mathbf{b}_{j,i}, \quad C_i(\alpha_j) = \mathbf{c}_{j,i}.$$

- Let $Z_S: \mathbb{F} \rightarrow \mathbb{F}$ be the polynomial $Z_S(z) := \prod_{j \in [N_g]} (z - \alpha_j)$. Namely, Z_S is the polynomial whose roots are the elements of S .

The 4-query linear PCP $\Pi_{\text{LPCP}} = (\mathcal{Q}_{\text{LPCP}}, \mathcal{P}_{\text{LPCP}}, \mathcal{V}_{\text{LPCP}})$ for CS is defined as follows:

- $\mathcal{Q}_{\text{LPCP}}(1^\kappa)$: On input $\kappa \in \mathbb{N}$, sample $\tau \xleftarrow{\mathcal{R}} \mathbb{F} \setminus S$. Define vectors $\mathbf{a} = (A_1(\tau), \dots, A_n(\tau))$, $\mathbf{b} = (B_1(\tau), \dots, B_n(\tau))$, and $\mathbf{c} =$

$$\mathbf{Q} = \begin{bmatrix} Z_S(\tau) & 0 & 0 & A_{n+1}(\tau) & \cdots & A_{N_w}(\tau) & 0 & 0 & \cdots & 0 \\ 0 & Z_S(\tau) & 0 & B_{n+1}(\tau) & \cdots & B_{N_w}(\tau) & 0 & 0 & \cdots & 0 \\ 0 & 0 & Z_S(\tau) & C_{n+1}(\tau) & \cdots & C_{N_w}(\tau) & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 & \tau & \cdots & \tau^{N_g} \end{bmatrix}^T \in \mathbb{F}^{(4+N_w+N_g-n) \times 4}.$$

Figure 3: The query matrix $\mathbf{Q}$ output by $\mathcal{Q}_{\text{LPCP}}$ in Construction B.1. Here, $\tau \xleftarrow{\mathbf{R}} \mathbb{F} \setminus S$.

$(C_1(\tau), \dots, C_n(\tau))$. Output $\text{st} = (A_0(\tau), B_0(\tau), C_0(\tau), \mathbf{a}, \mathbf{b}, \mathbf{c}, Z_S(\tau))$ and the query matrix $\mathbf{Q}$ as defined in Fig. 3.

- $\mathcal{P}_{\text{LPCP}}(1^\kappa, \mathbf{x}, \mathbf{w})$: On input $\kappa \in \mathbb{N}$ and an instance $(\mathbf{x}, \mathbf{w})$ where $C\mathcal{S}_\kappa(\mathbf{x}, \mathbf{w}) = 1$, sample $\delta_1, \delta_2, \delta_3 \xleftarrow{\mathbf{R}} \mathbb{F}$. Construct polynomials $A, B, C: \mathbb{F} \rightarrow \mathbb{F}$, each of degree N_g , where

$$\begin{aligned} A(z) &:= \delta_1 Z_S(z) + A_0(z) + \sum_{i \in [N_w]} w_i A_i(z) \\ B(z) &:= \delta_2 Z_S(z) + B_0(z) + \sum_{i \in [N_w]} w_i B_i(z) \\ C(z) &:= \delta_3 Z_S(z) + C_0(z) + \sum_{i \in [N_w]} w_i C_i(z). \end{aligned} \quad (\text{B.1})$$

Let $H(z) := (A(z)B(z) - C(z))/Z_S(z)$, and let $\mathbf{h} = (h_0, \dots, h_{N_g}) \in \mathbb{F}^{N_g+1}$ be the coefficients of H . Parse $\mathbf{w}^T = [\mathbf{x}^T \mid \tilde{\mathbf{w}}^T]$. Output the proof vector $\boldsymbol{\pi} = (\delta_1, \delta_2, \delta_3, \tilde{\mathbf{w}}, \mathbf{h}) \in \mathbb{F}^{4+N_w+N_g-n}$.

- $\mathcal{V}_{\text{LPCP}}(\text{st}, \mathbf{x}, \mathbf{a})$: On input $\text{st} = (a_0, b_0, c_0, \mathbf{a}, \mathbf{b}, \mathbf{c}, z)$, $\mathbf{x} \in \mathbb{F}^n$ and $\mathbf{a} \in \mathbb{F}^4$, the verifier computes $a'_1 = a_1 + a_0 + \mathbf{x}^T \mathbf{a}$, $a'_2 = a_2 + b_0 + \mathbf{x}^T \mathbf{b}$, and $a'_3 = a_3 + c_0 + \mathbf{x}^T \mathbf{c}$. It accepts if

$$a'_1 a'_2 - a'_3 - a_4 z = 0. \quad (\text{B.2})$$

Theorem B.2 (Linear PCP for QAPs). *Construction B.1 is complete, has knowledge error $2N_g/(|\mathbb{F}| - N_g)$, and is perfect HVZK.*

PROOF. Let $C\mathcal{S} = \{C\mathcal{S}_\kappa\}_{\kappa \in \mathbb{N}}$ be an R1CS system over $\mathbb{F}$. We consider each property separately:

- **Completeness:** Take any $\kappa \in \mathbb{N}$ and $(\mathbf{x}, \mathbf{w})$ where $C\mathcal{S}_\kappa(\mathbf{x}, \mathbf{w}) = 1$. Let $(\text{st}, \mathbf{Q}) \leftarrow \mathcal{Q}_{\text{LPCP}}(1^\kappa)$, $\boldsymbol{\pi} \leftarrow \mathcal{P}_{\text{LPCP}}(1^\kappa, \mathbf{x}, \mathbf{w})$, $\mathbf{a} \leftarrow \mathbf{Q}^T \boldsymbol{\pi}$. Consider the value of $\mathcal{V}_{\text{LPCP}}(\text{st}, \mathbf{x}, \mathbf{a})$. Let a'_1, a'_2, a'_3 be the values computed by $\mathcal{V}_{\text{LPCP}}$. By definition,

$$\begin{aligned} a'_1 &= a_1 + a_0 + \mathbf{x}^T \mathbf{a} \\ &= \delta_1 Z_S(\tau) + A_0(\tau) + \sum_{i \in [n]} x_i A_i(\tau) + \sum_{i \in [N_w-n]} w_{n+i} A_{n+i}(\tau) \\ &= \delta_1 Z_S(\tau) + A_0(\tau) + \sum_{i \in [N_w]} w_i A_i(\tau) \\ &= A(\tau). \end{aligned}$$

since $w_i = x_i$ for $i \in [n]$, A is the polynomial in Eq. (B.1), and $\tau \in \mathbb{F} \setminus S$ is the element sampled by $\mathcal{Q}_{\text{LPCP}}$. Similarly, we have that $a'_2 = B(\tau)$ and $a'_3 = C(\tau)$. Finally $a_4 = h_0 + \sum_{i \in [N_g]} h_i \tau^i = H(\tau)$, where $H(z) = (A(z)B(z) - C(z))/Z_S(z)$ is the polynomial constructed by the prover. The verification procedure now computes

$$a'_1 a'_2 - a'_3 - a_4 z = A(\tau)B(\tau) - C(\tau) - H(\tau)Z_S(\tau) = 0,$$

by definition of the polynomial H . Completeness follows.

- **Knowledge:** Define $\mathcal{E}_{\text{LPCP}}^{(\boldsymbol{\pi}^*, \cdot)}$ to be the algorithm that on input a statement $\mathbf{x}$ and given linear access to a proof vector $\boldsymbol{\pi}^* = (\delta_1^*, \delta_2^*, \delta_3^*, \tilde{\mathbf{w}}^*, \mathbf{h}^*)$, outputs $\mathbf{w}^T = [\mathbf{x}^T \mid (\tilde{\mathbf{w}}^*)^T] \in \mathbb{F}^{N_w}$. To show

that this extractor works, take any $\boldsymbol{\pi}^* = (\delta_1^*, \delta_2^*, \delta_3^*, \tilde{\mathbf{w}}^*, \mathbf{h}^*)$ where

$$\Pr[\mathcal{V}_{\text{LPCP}}(\text{st}, \mathbf{x}, \mathbf{Q}^T \boldsymbol{\pi}^*) = 1 : (\text{st}, \mathbf{Q}) \leftarrow \mathcal{Q}_{\text{LPCP}}(1^\kappa)] > \frac{2N_g}{|\mathbb{F}| - N_g}.$$

We use $\boldsymbol{\pi}^*$ and $C\mathcal{S}$ to define polynomials $A, B, C, H: \mathbb{F} \rightarrow \mathbb{F}$:

$$\begin{aligned} A(z) &= \delta_1^* Z_S(z) + A_0(z) + \sum_{i \in [n]} x_i A_i(z) + \sum_{i \in [N_w-n]} \tilde{w}_i^* A_{n+i}(z) \\ B(z) &= \delta_2^* Z_S(z) + B_0(z) + \sum_{i \in [n]} x_i B_i(z) + \sum_{i \in [N_w-n]} \tilde{w}_i^* B_{n+i}(z) \\ C(z) &= \delta_3^* Z_S(z) + C_0(z) + \sum_{i \in [n]} x_i C_i(z) + \sum_{i \in [N_w-n]} \tilde{w}_i^* C_{n+i}(z) \\ H(z) &= h_0^* + \sum_{i \in [N_g]} h_i^* z^i \end{aligned}$$

Let $\mathbf{Q}$ be the query matrix output by $\mathcal{Q}_{\text{LPCP}}$, $\mathbf{a} \leftarrow \mathbf{Q}^T \boldsymbol{\pi}^*$ and a'_1, a'_2, a'_3 be the components computed by $\mathcal{V}_{\text{LPCP}}$. By construction, $a'_1 = A(\tau)$, $a'_2 = B(\tau)$, $a'_3 = C(\tau)$ and $a_4 = H(\tau)$. Define the polynomial $P: \mathbb{F} \rightarrow \mathbb{F}$ where $P(z) = A(z)B(z) - C(z) - H(z)Z_S(z)$. By construction, $\deg(P) \leq 2N_g$. Next, $\mathcal{V}_{\text{LPCP}}$ accepts if $a'_1 a'_2 - a'_3 - a_4 z = 0$, where $z = Z(\tau)$, or equivalently, if

$$0 = A(\tau)B(\tau) - C(\tau) - H(\tau)Z_S(\tau) = P(\tau). \quad (\text{B.3})$$

Suppose Eq. (B.3) holds with probability $\varepsilon > 2N_g/(|\mathbb{F}| - N_g)$; that is, the verifier accepts with probability greater than ε . Since $\mathcal{Q}_{\text{LPCP}}$ samples τ uniformly from $\mathbb{F} \setminus S$ and $\deg(P) \leq 2N_g$, we conclude by the Schwartz-Zippel lemma (Lemma A.1) that $P \equiv 0$. In particular, this means that for all $j \in [N_g]$,

$$P(\alpha_j) = A(\alpha_j)B(\alpha_j) - C(\alpha_j) = 0,$$

since $Z_S(\alpha_j) = 0$ for all $j \in [N_g]$. Equivalently, this means that $A(\alpha_j)B(\alpha_j) = C(\alpha_j)$ for all $j \in [N_g]$. By construction of A, B, C , this means that

$$[1 \mid \tilde{\mathbf{u}}^T] \mathbf{a}_j \cdot [1 \mid \tilde{\mathbf{u}}^T] \mathbf{b}_j = [1 \mid \tilde{\mathbf{u}}^T] \mathbf{c}_j,$$

where $\tilde{\mathbf{u}}^T = [\mathbf{x}^T \mid (\tilde{\mathbf{w}}^*)^T]$. Since this holds for all $j \in [N_g]$, we have that $C\mathcal{S}_\kappa(\mathbf{x}, \tilde{\mathbf{w}}^*) = 1$, as required.

- **HVZK:** We first construct a simulator $\mathcal{S}_{\text{LPCP}} = (\mathcal{S}_{\text{LPCP},1}, \mathcal{S}_{\text{LPCP},2})$:
 - $\mathcal{S}_{\text{LPCP},1}(1^\kappa)$: The statement-independent algorithm samples $(\tilde{\text{st}}, \tilde{\mathbf{Q}}) \leftarrow \mathcal{Q}_{\text{LPCP}}(1^\kappa)$. It outputs $\tilde{\text{st}}$, $\tilde{\mathbf{Q}}$, and $\text{st}_S = \tilde{\text{st}}$.
 - $\mathcal{S}_{\text{LPCP},2}(\text{st}_S, \mathbf{x})$: On input the state $\text{st}_S = (\tilde{a}_0, \tilde{b}_0, \tilde{c}_0, \tilde{\mathbf{a}}, \tilde{\mathbf{b}}, \tilde{\mathbf{c}}, \tilde{z})$ and the statement $\mathbf{x}$, the statement-dependent algorithm samples $\tilde{a}_1, \tilde{a}_2, \tilde{a}_3 \xleftarrow{\mathbf{R}} \mathbb{F}$. It computes $\tilde{a}'_1 = \tilde{a}_1 + \tilde{a}_0 + \mathbf{x}^T \tilde{\mathbf{a}}$, $\tilde{a}'_2 = \tilde{a}_2 + \tilde{b}_0 + \mathbf{x}^T \tilde{\mathbf{b}}$, and $\tilde{a}'_3 = \tilde{a}_3 + \tilde{c}_0 + \mathbf{x}^T \tilde{\mathbf{c}}$. Compute $\tilde{a}_4 = \tilde{z}^{-1}(\tilde{a}'_1 \tilde{a}'_2 - \tilde{a}'_3)$. It outputs $\tilde{\mathbf{a}} = (\tilde{a}_1, \tilde{a}_2, \tilde{a}_3, \tilde{a}_4)$.

To complete the proof, it suffices to show that the simulated distribution is identical to the real distribution for any $(\mathbf{x}, \mathbf{w})$ where $C\mathcal{S}_\kappa(\mathbf{x}, \mathbf{w}) = 1$. By construction, the verification state and query matrix (output by $\mathcal{S}_{\text{LPCP},1}$) are identically distributed in

the two cases, so it suffices to analyze the distribution of the responses. Let $(st, Q) \leftarrow Q_{LPCP}(1^K)$, $\pi \leftarrow \mathcal{P}_{LPCP}(1^K, x, w)$, and $a \leftarrow Q^T \pi$. Write $st = (a_0, b_0, c_0, a, b, c, z)$. First, $z = Z_S(\tau)$ for some $\tau \in \mathbb{F} \setminus S$. Since $Z_S(x) = \prod_{\alpha \in S} (x - \alpha)$ and $\tau \notin S$, we have that $z = Z_S(\tau) \neq 0$. Then the following holds:

- In the real distribution, Eq. (B.2) holds (by completeness). Since $z \neq 0$, the value of a_4 is uniquely defined given a_1, a_2, a_3 and st . The value of a_4 that satisfies Eq. (B.2) precisely coincides with the value $\tilde{a}_4$ sampled by $S_{LPCP,2}$ (for the choice of $\tilde{a}_1, \tilde{a}_2, \tilde{a}_3$ chosen by the simulator).
- In the real distribution, $a_1 = \delta_1 Z_S(\tau) + \sum_{i \in [N_w - n]} w_{n+i} A_{n+i}(\tau)$, where δ_1 is uniform over $\mathbb{F}$ and independent of all other components. Since $Z_S(\tau) \neq 0$, this means a_1 is uniform over $\mathbb{F}$. A similar argument holds for a_2 and a_3 (by appealing to the randomness of δ_2 and δ_3 , respectively). This is precisely the distribution of $\tilde{a}_1, \tilde{a}_2, \tilde{a}_3$ in the simulation.

Thus, the simulated response is identically distributed as the real response, and perfect HVZK holds. $\square$

Knowledge against affine strategies. While our compiler only requires a linear PCP with knowledge against linear strategies we can easily modify the linear PCP from Construction C.1 to provide knowledge against affine prover strategies *without* increasing the query complexity. We describe this modified variant in the full version of this paper [78]. This means that we can base security on the *weaker* conjecture that Construction 3.5 is “affine-only.” Using the modified linear PCP comes at a very slight increase in the concrete cost of the verifier, and has no effect on the prover complexity. Since we believe that Conjecture 3.9 holds, we do not use this modified linear PCP in our concrete implementation.

C LINEAR PCP AND ZKSNARK ANALYSIS

In this section, we provide the formal analysis of our linear-only vector encryption scheme and resulting zkSNARKs.

C.1 Linear PCPs over Extension Fields

First, we describe how to transform a linear PCP over $\mathbb{F}_{p^2}$ to a linear PCP over $\mathbb{F}_p$.

Field extensions. Recall that a degree- d field extension $\mathbb{F}_{p^d}$ of $\mathbb{F}_p$ is a d -dimensional vector space over $\mathbb{F}_p$. For a field element $s \in \mathbb{F}_{p^d}$, we write $\mathbf{v}_s \in \mathbb{F}_p^d$ to denote its representation in $\mathbb{F}_p^d$. There is an efficiently-computable isomorphism between $s \in \mathbb{F}_{p^d}$ and $\mathbf{v}_s \in \mathbb{F}_p^d$. In particular, this means that for all $s, t \in \mathbb{F}_{p^d}$, $\mathbf{v}_s + \mathbf{v}_t = \mathbf{v}_{s+t} \in \mathbb{F}_p^d$. We write $\mathbf{M}_s \in \mathbb{F}_p^{d \times d}$ to denote the linear transformation over $\mathbb{F}_p^d$ corresponding to scalar multiplication by s over $\mathbb{F}_{p^d}$. Namely, for all $s, t \in \mathbb{F}_{p^d}$, we have that $\mathbf{M}_s \mathbf{v}_t = \mathbf{v}_{st}$.

Construction C.1 ($\mathbb{F}_{p^d}$ -Linear PCP to $\mathbb{F}_p$ -Linear PCP). Let $\Pi'_{LPCP} = (Q'_{LPCP}, \mathcal{P}'_{LPCP}, \mathcal{V}'_{LPCP})$ be a k -query linear PCP for a family of R1CS systems $CS = \{CS_\kappa\}_{\kappa \in \mathbb{N}}$ over an extension field $\mathbb{F}_{p^d}$ with query length ℓ . We construct a (dk) -query linear PCP $\Pi_{LPCP} = (Q_{LPCP}, \mathcal{P}_{LPCP}, \mathcal{V}_{LPCP})$ for CS with query length $d\ell$ over the base field $\mathbb{F}_p$:

- $Q_{LPCP}(1^K)$: Run $(st, Q') \leftarrow Q'_{LPCP}(1^K)$, where $Q' \in \mathbb{F}_{p^d}^{\ell \times k}$. Let $Q \in \mathbb{F}_p^{d\ell \times dk}$ be the matrix formed by taking each component

$q'_{i,j} \in \mathbb{F}_{p^d}$ in Q' and replacing it with the transpose of the “multiplication-by- $q'_{i,j}$ ” matrix $\mathbf{M}_{q'_{i,j}}^T \in \mathbb{F}_p^{d \times d}$. Output st and Q .

- $\mathcal{P}_{LPCP}(1^K, x, w)$: Compute $\pi' \leftarrow \mathcal{P}'_{LPCP}(1^K, x, w) \in \mathbb{F}_{p^d}^\ell$. Let $\pi \in \mathbb{F}_p^{d\ell}$ be the vector formed by taking each component $\pi'_i \in \mathbb{F}_{p^d}$ in π' and replacing it with the vector $\mathbf{v}_{\pi'_i} \in \mathbb{F}_p^d$ representing π_i . Output the proof vector π .
- $\mathcal{V}_{LPCP}(st, x, a)$: First, parse $a \in \mathbb{F}_{p^d}^{dk}$ as $[a'_1 | \dots | a'_k]$ for some $a' = (a'_1, \dots, a'_k) \in \mathbb{F}_{p^d}^k$. Output $\mathcal{V}'_{LPCP}(st, x, a')$.

Theorem C.2 ($\mathbb{F}_{p^d}$ -Linear PCP to $\mathbb{F}_p$ -Linear PCP). *If Π'_{LPCP} is complete, perfect HVZK, and has knowledge error ϵ , then the same holds for Π_{LPCP} from Construction C.1.*

PROOF. We analyze each property individually:

- **Completeness:** Take any x, w where $\mathcal{R}(x, w) = 1$, and let $(st, Q') \leftarrow Q'_{LPCP}(1^K)$, $\pi' \leftarrow \mathcal{P}'_{LPCP}(1^K, x, w)$. Let $Q \in \mathbb{F}_p^{d\ell \times dk}$ and $\pi \in \mathbb{F}_p^{d\ell}$ be as specified in Q_{LPCP} and $\mathcal{P}_{LPCP}$ and let $a \leftarrow Q^T \pi$. Write $a = [a_1, \dots, a_k]$. By construction, for all $i \in [k]$,

$$a_i = \sum_{j \in [\ell]} \mathbf{M}_{q'_{j,i}} \mathbf{v}_{\pi'_j} = \sum_{j \in [\ell]} \mathbf{v}_{q'_{j,i} \pi'_j} = \mathbf{v}_{(q'_i)^T \pi'} \in \mathbb{F}_p^d,$$

where $q'_i \in \mathbb{F}_{p^d}^\ell$ denotes the i th column of Q' . This means that the vector a' computed by $\mathcal{V}_{LPCP}$ satisfies $a' = (Q')^T \pi'$. Completeness now follows by completeness of Π'_{LPCP} .

- **Knowledge:** Let $(st, Q') \leftarrow Q'_{LPCP}(1^K)$ and let Q be the matrix Q_{LPCP} constructs from Q' . Take any proof $\pi \in \mathbb{F}_p^{d\ell}$, and let $\pi' \in \mathbb{F}_{p^d}^\ell$ be the vector obtained by viewing each contiguous block of d elements of π as an element of $\mathbb{F}_{p^d}$. By construction, $\mathcal{V}_{LPCP}(st, x, Q^T \pi) = 1$ if and only if $\mathcal{V}'_{LPCP}(st, x, (Q')^T \pi') = 1$. The claim now follows by knowledge soundness of Π'_{LPCP} . Namely, the extractor $\mathcal{E}_{LPCP}$ for Π_{LPCP} simply invokes the extractor $\mathcal{E}'_{LPCP}$ for Π'_{LPCP} . Any linear query $q' \in \mathbb{F}_{p^d}^\ell$ that $\mathcal{E}'_{LPCP}$ makes to $\langle \pi', \cdot \rangle$ can be simulated via d linear queries to $\langle \pi, \cdot \rangle$ by expanding each component in q' into a matrix over $\mathbb{F}_p^{d \times d}$.

- **Perfect HVZK:** Let $\mathcal{S}'_{LPCP} = (\mathcal{S}'_{LPCP,1}, \mathcal{S}'_{LPCP,2})$ be the linear PCP simulator for Π'_{LPCP} . We define $\mathcal{S}_{LPCP} = (\mathcal{S}_{LPCP,1}, \mathcal{S}_{LPCP,2})$ for Π_{LPCP} as follows:

- $\mathcal{S}_{LPCP,1}(1^K)$: On input $\kappa \in \mathbb{N}$, run the simulator $\mathcal{S}'_{LPCP,1}(1^K)$ to obtain a pair (st, Q') where $Q' \in \mathbb{F}_{p^d}^{\ell \times k}$. The simulator constructs $Q \in \mathbb{F}_p^{d\ell \times dk}$ by expanding replacing each component $q'_{i,j}$ of Q' with $\mathbf{M}_{q'_{i,j}}^T$ (as in Q_{LPCP}). It outputs (st, Q) .
- $\mathcal{S}_{LPCP,2}(st, x)$: On input the simulation state st and a statement x , run $\mathcal{S}'_{LPCP,2}(st, x)$ to obtain $a' \in \mathbb{F}_{p^d}^k$. Then, compute and output $a \in \mathbb{F}_p^{dk}$ by expanding each component $\pi'_i \in \mathbb{F}_{p^d}$ as a vector $\mathbf{v}_{\pi'_i} \in \mathbb{F}_p^d$ (as in $\mathcal{P}_{LPCP}$).

Perfect HVZK now follows by perfect HVZK of Π'_{LPCP} . $\square$

C.2 Circuit Privacy

We give the formal definition of circuit privacy below:

Definition C.3 (Circuit Privacy). Let $\Pi_{\text{Enc}} = (\text{Setup}, \text{Encrypt}, \text{Decrypt}, \text{Add})$ be a secret-key vector encryption scheme over $\mathbb{F}^\ell$.

We say that Π_{Enc} satisfies circuit privacy if for all efficient and stateful adversaries $\mathcal{A}$, there exists an efficient simulator $\mathcal{S}$ such that for all security parameters $\lambda \in \mathbb{N}$,

$$\Pr[\text{ExptCircuitPriv}_{\Pi_{\text{Enc}}, \mathcal{A}, \mathcal{S}}(1^\lambda) = 1] = 1/2 + \text{negl}(\lambda), \quad (\text{C.1})$$

where the experiment $\text{ExptCircuitPriv}_{\Pi_{\text{Enc}}, \mathcal{A}, \mathcal{S}}(1^\lambda)$ is defined as follows:

- (1) The challenger samples $(\text{pp}, \text{sk}) \leftarrow \text{Setup}(1^\lambda, 1^\ell)$ and gives (pp, sk) to the adversary. The adversary replies with a collection of vectors $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{F}^\ell$.
- (2) The challenger constructs ciphertexts $\text{ct}_i \leftarrow \text{Encrypt}(\text{sk}, \mathbf{v}_i)$ for all $i \in [k]$ and gives $(\text{ct}_1, \dots, \text{ct}_k)$ to $\mathcal{A}$. The adversary replies with a collection of coefficients $y_1, \dots, y_k \in \mathbb{F}$.
- (3) The challenger computes $\text{ct}_0^* \leftarrow \text{Add}(\text{pp}, \{\text{ct}_i\}_{i \in [k]}, \{y_i\}_{i \in [k]})$ and $\text{ct}_1^* \leftarrow \mathcal{S}(1^\lambda, \text{pp}, \text{sk}, \sum_{i \in [k]} y_i \mathbf{v}_i)$. It also samples a random bit $b \xleftarrow{\mathbb{R}} \{0, 1\}$ and replies to the adversary with ct_b^* .
- (4) The adversary outputs a bit $b' \in \{0, 1\}$. The output of the experiment is 1 if $b' = b$ and 0 otherwise.

In this work, we also consider a weaker notion of circuit privacy where we additionally constrain the adversary to choosing the coefficients from an *a priori* specified set $S \subseteq \mathbb{F}$. In this case, we say that Π_{Enc} satisfies circuit privacy with respect to S . In addition, when the probability in Eq. (C.1) is bounded by $1/2 + \varepsilon$, we say that Π_{Enc} is ε -circuit private.

Remark C.4 (Multi-Query Circuit Privacy). We can define a multi-query variant of Definition C.3 where the adversary can adaptively choose *multiple* collections of coefficients $y_1, \dots, y_k \in R_p$ and on each query, the adversary learns either the homomorphically-evaluated ciphertext (from Add) or the simulated ciphertext (from $\mathcal{S}$). This multi-query notion is useful to argue *multi-theorem* zero knowledge when compiling a linear PCP into a preprocessing SNARG [30]. Definition C.3 implies this multi-query variant by a standard hybrid argument.

D LINEAR PCP IMPLEMENTATION DETAILS

In this section, we provide a more detailed description of our multi-point evaluation and interpolation approach outlined in Section 4.1. Following [20], for a polynomial $A(z)$ of degree less than $|D|$, we write $\text{FFT}_D(A(z))$ to denote the vector of evaluations $(A(\alpha))_{\alpha \in D}$. Similarly, we write $\text{FFT}_D^{-1}((A'(\alpha))_{\alpha \in D})$ to denote the coefficients of the polynomial A (of degree less than $|D|$) where $A(\alpha) = A'(\alpha)$ for all $\alpha \in D$.

Let $\omega \in \mathbb{F}$ be a primitive 2^d -th root of unity and $H = H_1 = \langle \omega \rangle \subset \mathbb{F}$ be the subgroup of order 2^d generated by ω (consisting of the 2^d -th roots of unity). Let $\xi_1 = 1$ and take $\xi_2, \dots, \xi_i \in \mathbb{F}^* \setminus H_1$ such that the cosets $H_i = \xi_i H_1$ are all pairwise disjoint. We define the domain to be $D = \bigcup_{i \in [k]} H_i$. For a set $S \subset \mathbb{F}$, let $\mathbf{V}_S \in \mathbb{F}^{|D| \times |D|}$ be the Vandermonde matrix associated with evaluating a polynomial of degree up to $|D| - 1$ on the points in D . Let $\hat{\mathbf{V}}_H \in \mathbb{F}^{2^d \times 2^d}$ be the Vandermonde matrix associated with evaluating a polynomial of

degree up to 2^d on H (i.e., the roots of unity). Then, we have that

$$\mathbf{V}_S = \begin{bmatrix} \hat{\mathbf{V}}_H & \hat{\mathbf{V}}_H & \cdots & \hat{\mathbf{V}}_H \\ \hat{\mathbf{V}}_H \cdot \Xi_2 & \hat{\mathbf{V}}_H \cdot \xi_2^{2^d} \Xi_2 & \cdots & \hat{\mathbf{V}}_H \cdot \xi_2^{(k-1)2^d} \Xi_2 \\ \vdots & \vdots & \ddots & \vdots \\ \hat{\mathbf{V}}_H \cdot \Xi_k & \hat{\mathbf{V}}_H \cdot \xi_k^{2^d} \Xi_k & \cdots & \hat{\mathbf{V}}_H \cdot \xi_k^{(k-1)2^d} \Xi_k \end{bmatrix},$$

where $\Xi_i = \text{diag}(1, \xi_i, \xi_i^2, \dots, \xi_i^{2^d-1})$. Take any input $\mathbf{a} \in \mathbb{F}^{k \cdot 2^d}$, and for $i \in [k]$, let $\hat{\mathbf{a}}_i = (a_{(i-1)2^d+1}, \dots, a_{i \cdot 2^d}) \in \mathbb{F}^{2^d}$. We describe an algorithm to compute $\mathbf{a}' = \mathbf{V}_S \mathbf{a}$:

- Let $\hat{\mathbf{a}}'_i = (a'_{(i-1)2^d+1}, \dots, a'_{i \cdot 2^d})$. By construction,

$$\hat{\mathbf{a}}'_i = \hat{\mathbf{V}}_H \cdot \left(\sum_{j \in [k]} \xi_i^{(j-1)2^d} \Xi_i \hat{\mathbf{a}}_j \right).$$

Let $\hat{\mathbf{b}}_i = \sum_{j \in [k]} \xi_i^{(j-1)2^d} \Xi_i \hat{\mathbf{a}}_j \in \mathbb{F}^{2^d}$. Given $\hat{\mathbf{b}}_i$, computing $\hat{\mathbf{a}}_i = \hat{\mathbf{V}}_H^{-1} \hat{\mathbf{b}}_i$ can be done using a standard radix-2 FFT in $O(d \cdot 2^d)$ time.

- Naïvely, we can compute $\hat{\mathbf{b}}_i$ in $O(k \cdot 2^d)$ time, so computing all of the entries in $\mathbf{b} = [\hat{\mathbf{b}}_1^\top \mid \cdots \mid \hat{\mathbf{b}}_k^\top]^\top \in \mathbb{F}^{k \cdot 2^d}$ requires $O(2^d k^2)$ time. However we can do so more efficiently as follows. By definition,

$$\hat{b}_{i,j} = \sum_{\ell \in [k]} \xi_i^{(\ell-1)2^d} \xi_i^{j-1} \hat{a}_{\ell,j}.$$

Now, define $\tilde{\mathbf{b}}_j = (\hat{b}_{1,j}, \dots, \hat{b}_{k,j}) \in \mathbb{F}^k$, and similarly, let $\tilde{\mathbf{a}}_j = (\hat{a}_{1,j}, \dots, \hat{a}_{k,j}) \in \mathbb{F}^k$. Then,

$$\tilde{\mathbf{b}}_j = \text{diag}(\xi_1^{j-1}, \dots, \xi_k^{j-1}) \underbrace{\begin{bmatrix} 1 & \xi_1^{2^d} & \cdots & \xi_1^{2^d(k-1)} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \xi_k^{2^d} & \cdots & \xi_k^{2^d(k-1)} \end{bmatrix}}_{\Xi'} \tilde{\mathbf{a}}_j.$$

Observe now that $\Xi' \in \mathbb{F}^{k \times k}$ is itself a Vandermonde matrix corresponding to evaluating a degree $(k-1)$ polynomial on the points $\xi_1^{2^d}, \dots, \xi_k^{2^d}$. While $\xi_1^{2^d}, \dots, \xi_k^{2^d}$ are *not* roots of unity (so standard FFTs cannot be used here), we can still solve this problem efficiently if $\xi_1, \dots, \xi_k$ form a geometric sequence (i.e., $\xi_i = \alpha \xi_{i-1}$ for some fixed $\alpha \in \mathbb{F}$) [39]. In particular, using the Bostan-Schost algorithms, multipoint evaluation on k values in a geometric sequence requires computing 2 degree- k polynomial multiplications and $O(k)$ additional work. In the case where $k < 2^{d-1}$, we can use standard radix-2 FFTs to implement the degree- k polynomial multiplications in $O(k \log k)$ time. Thus, computing each $\tilde{\mathbf{b}}_j$ can be done in just $O(k \log k)$ time. Repeating this for all $j \in [2^d]$ yields an algorithm to compute $\mathbf{b}$ in $O(2^d k \log k)$ time.

The overall running time of this algorithm is $O(2^d k(d + \log k))$, which matches the running time of a standard FFT over a domain of size $k \cdot 2^d$. While the concrete efficiency of the algorithm is worse than a standard radix-2 FFT, in fields where there are insufficient roots of unity (such as the ones we consider), this provides an efficient algorithm to implement the linear PCP prover. In all of our experiments, $k \leq 64$.

In our implementation, we set $\xi_i = g^{2^{(i-1)}\omega}$ for $i \in [k]$, where g is a multiplicative generator of $\mathbb{F}^*$. This enables efficient implementation of multipoint evaluation over the set D as well as the set $gD = \{gh \mid h \in D\}$ (needed for efficient implementation of the linear PCP prover algorithm; see [20] for further details).

Lagrange interpolation and inverse FFTs. In addition to computing FFT_D , the linear PCP prover needs to compute the inverse operation FFT_D^{-1} . This follows immediately from our algorithm above by inverting each of the steps (i.e., replace both sets of FFTs with their corresponding inverse FFTs).

The query-generation algorithm Q_{LPCP} in Claim A.6 (Construction B.1) essentially reduces to multiple Lagrange polynomial evaluations (with basis D) at a random field element. Ben-Sasson et al. [20] described an efficient implementation of this when the domain D is the roots of unity. In our setting (of working over a field with insufficient roots of unity), we augment D with cosets of the roots of unity. The Ben-Sasson et al. algorithm directly generalizes to this setting and we refer to [20, Appendix E] for the details.

E THE POWER DIFFIE-HELLMAN ASSUMPTION OVER SMALL FIELDS

In this section, we briefly recall the q -power Diffie-Hellman assumption introduced by Groth [74] and subsequently used as the basis for both pairing-based SNARKs [64, 87] as well as lattice-based SNARKs [65]. Following [65], we formulate the assumption with respect to a linear encoding scheme, which captures both the pairing-based instantiation as well as the lattice-based instantiation.

Definition E.1 (Linear Encoding Scheme). A (secret-key) linear encoding scheme Π_{Enc} over a finite field $\mathbb{F}$ is a tuple of algorithms $\Pi_{\text{Enc}} = (\text{Setup}, \text{Encode}, \text{Add})$ with the following properties:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$: On input the security parameter λ , the setup algorithm outputs a public evaluation key pk and a secret encoding key sk .
- $\text{Encode}(\text{sk}, x) \rightarrow \text{enc}_x$: On input the secret key sk and an element $x \in \mathbb{F}$, the encoding algorithm outputs an encoding enc_x of x .
- $\text{Add}(\text{pk}, (\text{enc}_1, \dots, \text{enc}_d), (\alpha_1, \dots, \alpha_d)) \rightarrow \text{enc}'$: On input the public key pk , encodings $\text{enc}_1, \dots, \text{enc}_d$ and coefficients $\alpha_1, \dots, \alpha_d \in \mathbb{F}$, the add algorithm outputs a new encoding enc' .

The encoding scheme is d -linear if for all values $k \leq d$, values $x_1, \dots, x_k \in \mathbb{F}$, scalars $\alpha_1, \dots, \alpha_k \in \mathbb{F}^d$, and sampling $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$, $\text{enc}_i \leftarrow \text{Encode}(\text{sk}, x_i)$ for all $i \in [k]$, we have that

$$\Pr[\text{Add}(\text{pk}, (\text{enc}_1, \dots, \text{enc}_k), (\alpha_1, \dots, \alpha_k)) \in S] = 1 - \text{negl}(\lambda),$$

where S denotes the support of $\text{Encode}(\text{sk}, \sum_{i \in [k]} \alpha_i x_i)$.

Definition E.2 (q -Power Diffie-Hellman Assumption [65, 74]). Fix a parameter $q \in \mathbb{N}$. A linear encoding scheme $\Pi_{\text{Enc}} = (\text{Setup}, \text{Encode}, \text{Add})$ over a field $\mathbb{F}$ satisfies the q -power Diffie-Hellman assumption (q -PDH) if for all efficient adversaries $\mathcal{A}$, and sampling $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$, $s \xleftarrow{\mathbb{R}} \mathbb{F}$, $\text{enc}_i \leftarrow \text{Encode}(\text{sk}, s^i)$ for all $i \in \{0, \dots, 2q\}$, $\sigma \leftarrow (\text{pk}, \text{enc}_0, \dots, \text{enc}_q, \text{enc}_{q+2}, \dots, \text{enc}_{2q})$, we have that

$$\Pr[\mathcal{A}(1^\lambda, \sigma) \in S] = \text{negl}(\lambda),$$

where S is the set of encodings in the support of $\text{Encode}(\text{sk}, s^{q+1})$.

Lemma E.3 (q -PDH Assumption over Small $\mathbb{F}$). Let $\Pi_{\text{Enc}} = (\text{Setup}, \text{Encode}, \text{Add})$ be a d -linear encoding scheme over a finite field $\mathbb{F}$. If $d \geq 2q$, there exists an adversary that runs in time $\text{poly}(q, \log |\mathbb{F}|)$ and wins the q -PDH security game for Π_{Enc} with advantage $2q/|\mathbb{F}|$.

PROOF. The adversary $\mathcal{A}$ starts by choosing $2q$ distinct points $z_1, \dots, z_{2q} \in \mathbb{F}$, and forms the polynomial $f(x) = \prod_{i \in [2q]} (x - z_i)$. Write this as $f(x) = \sum_{i=0}^{2q} \alpha_i x^i$. Then, for all $i \in [2q]$, $z_i^{q+1} = -\alpha_{q+1}^{-1} \sum_{j \neq q+1} \alpha_j z_i^j$. Let $(\text{pk}, \text{enc}_0, \dots, \text{enc}_q, \text{enc}_{q+2}, \dots, \text{enc}_{2q})$ be the q -PDH challenge. Here, enc_i is an encoding of s^i , where $s \in \mathbb{F}$ is sampled by the q -PDH challenger at the beginning of the experiment. Since $d \geq 2q$, the adversary can homomorphically compute an encoding of $-\alpha_{q+1}^{-1} \sum_{i \neq q+1} \alpha_i s^i$. By the above analysis, if $s \in \{z_1, \dots, z_{2q}\}$, then this quantity is exactly s^{q+1} . Since s is uniform and independent of $z_1, \dots, z_{2q}$, the probability that $s \in \{z_1, \dots, z_{2q}\}$ is exactly $2q/|\mathbb{F}|$, which proves the claim. $\square$

Remark E.4 (q -Power Diffie-Hellman Assumption over Small $\mathbb{F}$). When the q -PDH assumption is used for constructing pairing-based zkSNARKs [64, 74, 87], the size of the underlying field $\mathbb{F}$ is super-polynomial (i.e., $|\mathbb{F}| = 2^{\Omega(\lambda)}$). In this case, the attack in Lemma E.3 has negligible advantage. Indeed, the q -PDH assumption plausibly holds over standard pairing-based groups, and holds unconditionally in the generic (bilinear) group model [74].

In the lattice-based zkSNARK of Gennaro et al. [65], they consider fields of polynomial size. Unfortunately, Lemma E.3 shows that the q -PDH assumption does not hold for encoding schemes over fields of polynomial size. For the specific instantiation proposed by Gennaro et al., $q \approx 2^{16}$ and $|\mathbb{F}| \approx 2^{32}$, so Lemma E.3 gives an attack on q -PDH with advantage $2q/|\mathbb{F}| = 2^{-15}$. Since their zkSNARK relies on hardness of the q -PDH assumption for soundness, this means that their suggested parameters provide at best 15 bits of provable soundness. To obtain 128-bits of soundness, it would be necessary to either apply soundness amplification (which increases all parameters by a factor of $128/15 \approx 8.5$) or instantiate the Regev-based encoding scheme over a super-polynomial size field (which would also incur additional overhead).

In this work, we work over small (polynomial-size) fields and use parallel repetition (at the linear PCP level) for soundness amplification (see Remark A.7). This increases the number of linear PCP queries, but since we encrypt *vectors* of queries, the overhead for parallel amplification is *additive* rather than multiplicative in the number of repetitions. This yields a much more efficient construction over *small* fields compared to the Gennaro et al. construction (see Table 1).

F ADDITIONAL BENCHMARKS

In this section, we provide additional benchmarks for our lattice-based SNARK.

Microbenchmarks. For the setup and prover algorithms, we measure the concrete cost of each subcomponent. We show the breakdown for the construction over $\mathbb{F}_{p^2}$ where $p = 2^{13} - 1$ in Fig. 4 (the breakdown for other parameter settings are similar). For CRS generation, the cost is dominated by the time needed to encrypt

Figure 4: Cost breakdowns for CRS setup and prover for different R1CS instances. Measurements are based on an instantiation with a linear PCP and a vector encryption scheme over $\mathbb{F}_{p^2}$ where $p = 2^{13} - 1$.

the linear PCP queries. Namely, for an R1CS system with 2^{20} constraints, linear PCP query encryption constitutes 99% of the CRS generation time.

For the prover computation, we consider the cost of the linear PCP prover (Claim A.6 and Appendix B), the time spent on CRS expansion (i.e., deriving the random ciphertext components $\mathbf{a} \in \mathbb{R}_q^n$ from the PRF key), and the cost of the homomorphic operations for computing the encrypted linear PCP response. The microbenchmarks show that about 40% of the time is spent on CRS expansion. For an R1CS instance of size 2^{20} , the expanded CRS is over 80 GB, and CRS expansion takes just under 30 s. Note that the vectors are generated on the fly and we do *not* need to store the full CRS in memory. For the larger instances, the remaining prover computation is evenly split between the homomorphic operations and computing the coefficients of the linear PCP; specifically, each of these components constitutes roughly 30% of the overall prover computation. In the case of the linear PCP prover, the computation is dominated by computing FFTs (see Appendix B). There is a jump in the cost of the FFTs when we switch to our modified FFT procedure (Section 4.1) for implementing the prover computation (for settings where $\mathbb{F}_{p^2}$ does not have enough primitive roots of unity to use standard power-of-two FFTs). By extrapolating the performance, our approach is about 7× slower than the basic radix-2 FFT.¹⁰ When considering an R1CS system over $\mathbb{F}_{p^2}$ where $p = 2^{19} - 1$ (where there are sufficient roots of unity to invoke standard FFTs in the linear PCP prover algorithm), the linear PCP prover, homomorphic operations, and CRS expansion account for 6% (3.1 s), 38% (21.4 s), and 56% (31.8 s) of the total prover cost, respectively.

Zero knowledge. We also measure the concrete performance of our zkSNARKs for different choices of the zero-knowledge parameter κ . We provide the results in Fig. 5. In particular, when we work over $\mathbb{F}_{p^2}$ with $p = 2^{19} - 1$, and consider the setting *without* provable zero knowledge (i.e., setting $\kappa = 0$), the prover time (for an R1CS instance of size 2^{20}) is just 34 s. This represents an additional 1.6×

¹⁰When $p = 2^{13} - 1$, the field $\mathbb{F}_{p^2}$ contains a 2^{14} -th root of unity, so we can use standard radix-2 FFTs for R1CS instances with up to 2^{14} constraints. For instances of size 2^{15} , we use the approach from Section 4.1 and Appendix D, but directly inline the multipoint evaluation and interpolation on two points (this coincides with an existing implementation from libfqfft [94]). For instances larger than 2^{15} , we use the general Bostan-Schost algorithms [39] for the multipoint evaluation and interpolation. This introduces the 7× overhead in the cost of the FFT.

p	Setting	Size		Time	
		CRS	Proof	Setup	Prover
$2^{13} - 1$	PQ	5.3 GB	16.4 KB	2240 s	68 s
	Classical	5.3 GB	15.2 KB	2225 s	69 s
$2^{19} - 1$	PQ	1.9 GB	20.8 KB	877 s	56 s
	Classical	1.9 GB	19.2 KB	865 s	56 s

Table 4: Performance comparison of zkSNARKs instantiated using parameters for 128-bits of classical vs. 128-bits of post-quantum security (denoted “PQ”). For all measurements, we consider R1CS instances over $\mathbb{F}_{p^2}$ with 2^{20} constraints and compile them to zkSNARKs using linear-only vector encryption over $\mathbb{F}_{p^2}$.

speed-up over our construction with $\kappa = 40$ bits of zero knowledge. We see a 1.9× reduction in proof size (from 20.8 KB to 11.1 KB) in this setting. Working over a smaller base field, we can bring the proof size down to just 8 KB. This is around 20× shorter than previous post-quantum candidates (see Table 1). This reduction in proof size comes at the expense of a longer CRS (2.7 GB).

Classical vs. post-quantum security. If we instead instantiate our scheme to provide 128-bits of *classical* security (instead of post-quantum security), we obtain about a 5% reduction in proof size, setup time, and prover time. Realizing post-quantum security requires using a larger ring dimension n , but does not affect the modulus q . As such, the size of the CRS is unaffected (since we are deriving the random component of each ciphertext from a PRF). We provide more details in Table 4.

Figure 5: Cost breakdowns as a function of the zero-knowledge parameter κ (i.e., the zero-knowledge distinguishing advantage of any $\text{poly}(\lambda)$ adversary is bounded by $2^{-\kappa} + \text{negl}(\lambda)$). All measurements taken for an R1CS instance over $\mathbb{F}_{p^2}$ with 2^{20} constraints (and compiled using vector encryption over $\mathbb{F}_{p^2}$).