Counting Vampires: From Univariate Sumcheck to Updatable ZK-SNARK

Helger Lipmaa[1][0000−0001−8393−6821], Janno Siim[1][0000−0001−5824−7215], and Michał Zając[2]

[1] Simula UiB, Bergen, Norway
[2] Nethermind, London, UK

Abstract. We propose a univariate sumcheck argument $\mathfrak{Count}$ of essentially optimal communication efficiency of one group element. While the previously most efficient univariate sumcheck argument of Aurora is based on polynomial commitments, $\mathfrak{Count}$ is based on inner-product commitments. We use $\mathfrak{Count}$ to construct a new pairing-based updatable and universal zk-SNARK $\mathfrak{Vampire}$ with the shortest known argument length (four group and two finite field elements) for NP. In addition, $\mathfrak{Vampire}$ uses the aggregated polynomial commitment scheme of Boneh et al.

Keywords: Aggregatable polynomial commitment · Inner-product commitment · Sumcheck · Updatable and universal zk-SNARK

1 Introduction

Zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs, [15,21,14,25]) are zero-knowledge argument systems for NP with succinct argument length and efficient verification. In many applications, one can describe the desired NP language instance as an instance $\mathcal{R}$ of the rank-1 constraint system (R1CS) [14], and the task of the verifier is to check that $\mathcal{R}$ is satisfied on the partially-public input. Zk-SNARKs are immensely popular due to applications in, say, verifiable computation and blockchain.

Non-interactive zero-knowledge arguments, and thus also zk-SNARKs, are impossible in the plain model. To overcome this, one gives all parties access to a trusted common reference string (CRS). The most efficient zk-SNARKs have a relation-specific structured CRS (SRS). That is, they assume that there exists a trusted third party who, given the description of $\mathcal{R}$ as an input, generates an SRS $\mathsf{srs}_{\mathcal{R}}$. The most efficient zk-SNARK by Groth [16] for R1CS with a relation-specific SRS has an argument that consists of only three group elements.

A significant practical downside of such "non-universal" SNARKs is that one has to construct a new SRS for every instance of the constraint system. This observation has spurred an enormous effort to design universal zk-SNARKs, i.e., zk-SNARKs with an SRS that only depends on an upper bound on $\mathcal{R}$'s size. In addition, it is crucial to decrease the amount of trust in the SRS creator. A popular approach is to design *updatable and universal zk-SNARKs* [17,24],

where the universal SRS is updated sequentially by several parties such that the soundness holds if at least one of the updaters is honest. For brevity, by "updatable" we will sometimes mean "updatable and universal".

Plonk [13] and Marlin [10] are the first efficient universal zk-SNARKs. Marlin and many subsequent updatable and universal zk-SNARKs [26,9] work for sparse R1CS instances, where the underlying matrices contain a linear (instead of quadratic) number of non-zero elements. Chiesa et al. [10] define an information-theoretic model, algebraic holographic proof (AHP). An AHP is an interactive protocol, where at each step, the prover sends polynomial oracles, and the verifier sends to the prover random field elements. Polynomial oracles are usually implemented using polynomial commitments [19]. In the end, the verifier queries the polynomial oracles and performs low-degree tests. Then, [10] proposes a new AHP for sparse R1CS, and then compiles it to a zk-SNARK named Marlin.

Marlin relies crucially on a univariate sumcheck. A sumcheck argument aims to prove that the given polynomial sums to the given value over the given domain. The first sumcheck arguments [23] were for multivariate polynomials but small domains. Ben-Sasson et al. [5] proposed a univariate sumcheck argument for large domains and used it to construct a new zk-SNARK Aurora. Suppose the domain is a multiplicative subgroup of the given finite field. In that case, Aurora's sumcheck argument requires the prover to forward two different polynomial oracles and use a low-degree test on one of the polynomials.

Lunar [9] improves on Marlin. It defines PHPs (Polynomial Holographic IOPs), a generalization of AHPs. Lunar notes that instead of opening all the polynomial commitments, the verifier can often perform verification equations on commitments themselves, thus obtaining better efficiency. It also defines a simpler version of R1CS called R1CSLite, with one of the three characterizing matrices of $\mathcal{R}$ being the identity matrix. Moreover, it provides a more fine-grained analysis of the zero-knowledge property and several additional optimizations.

Basilisk [26] gains additional efficiency by using a different technique to obtain zero-knowledge and constructing a "free" low-degree test. In addition, [26] constructs even more efficient zk-SNARKs for somewhat more limited constraint systems. Both Lunar and Basilisk introduce new theoretical frameworks; e.g., Basilisk introduces checkable subspace sampling (CSS) arguments as a separate primitive. For simplicity (of reading), we opted not to use such frameworks in the context of the current paper.

In parallel to our work, Zhang et al. [28] proposed Vector Oracle Proofs (VOProofs), a new information-theoretic model based on vector operations. They use it to construct efficient zk-SNARKs for several well-known constraint systems such as R1CS (VOR1CS) and Plonk's constraint system (VOPlonk).

In Table 1, we overview the argument lengths of the most efficient updatable and universal zk-SNARKs. Here, $|X|$ denotes the representation length of an element from X in bits, given the BLS12-381 curve, with $|\mathbb{G}_1| = 384$, $|\mathbb{G}_2| = 768$, and $|\mathbb{F}| = 256$. Thus, even the most efficient updatable and universal zk-SNARK has an approximately two times longer argument than Groth16 [16].

Table 1. Comparison of some known updatable and universal zk-SNARKs.

Scheme	Argument length		Arithmetization
	Elements	Bits	
Updatable and universal zk-SNARKs			
Sonic [24]	$20\|\mathbb{G}_1\| + 16\|\mathbb{F}\|$	11776	[8] constraints
Marlin [10]	$13\|\mathbb{G}_1\| + 8\|\mathbb{F}\|$	7040	R1CS, sparse matrices
Basilisk [26]	$10\|\mathbb{G}_1\| + 3\|\mathbb{F}\|$	4608	R1CSLite, sparse matrices
Plonk [13]	$7\|\mathbb{G}_1\| + 7\|\mathbb{F}\|$	4480	Plonk constraints
LunarLite [9]	$10\|\mathbb{G}_1\| + 2\|\mathbb{F}\|$	4352	R1CSLite, sparse matrices
Basilisk [26]	$8\|\mathbb{G}_1\| + 4\|\mathbb{F}\|$	4096	Plonk constraints
VOR1CS* [28]	$9\|\mathbb{G}_1\| + 2\|\mathbb{F}\|$	3968	R1CS, sparse matrices
VOPlonk* [28]	$7\|\mathbb{G}_1\| + 2\|\mathbb{F}\|$	3200	Plonk constraints
Basilisk (full version, [27])	$6\|\mathbb{G}_1\| + 2\|\mathbb{F}\|$	2816	Weighted R1CS with bounded fan-out
𝔙ampire (this work)	$4\|\mathbb{G}_1\| + 2\|\mathbb{F}\|$	2048	R1CSLite, sparse matrices
Non-universal zk-SNARKs (relation-specific SRS)			
Groth16 [16]	$2\|\mathbb{G}_1\| + 1\|\mathbb{G}_2\|$	1536	R1CS

Moreover, Groth16 works for QAP [14] (i.e., full R1CS), while the most efficient variant of Basilisk works for instances of R1CS where the relation-defining matrices are limited to have a small constant number of elements per row (this corresponds to arithmetic circuits of bounded fan-out). Thus, there is still a non-trivial difference between the communication efficiency of relation-specific zk-SNARKs and updatable and universal zk-SNARKs.

Our Contributions. The current paper has three related contributions:
1. The combined use of polynomial commitments and inner-product commitments in the sumcheck and updatable and universal zk-SNARK design. The use of polynomial commitment schemes in zk-SNARKs has dramatically increased their popularity, and we hope the same will happen with inner-product commitments. In particular, ILV inner-product commitments [18] use a SRS made of non-consequent monomial powers.[3]
2. A new updatable (and universal) univariate sumcheck argument 𝔠ount that uses inner-product commitments to achieve optimal computation complexity of a single group element. Since sumchecks are used in many different zk-SNARKs (and elsewhere), we believe 𝔠ount will have wider interest.
3. A new updatable and universal zk-SNARK 𝔙ampire for sparse R1CSLite with the smallest argument length among all known updatable and universal zk-SNARKs for NP-complete languages. (See Table 1.) 𝔙ampire uses 𝔠ount and thus inner-product commitments.

---

[3] Inner-product commitments and arguments are commonly used in the zk-SNARK design. However, the way we use them is markedly different from the prior work.

1.1   Our Techniques

Non-Consequent Monomial SRSs. Groth et al. [17] proved that the SRS of an updatable zk-SNARK cannot contain non-monomial polynomials. Moreover, the SRS's correctness must be verifiable. For example, if the SRS contains[4] $[1, \sigma, \sigma^3, \sigma^4]_1 \in \mathbb{G}_1^4$ for a trapdoor σ, it must also contain $[\sigma, \sigma^2]_2 \in \mathbb{G}_2^2$, so that one can verify the consistency of the SRS elements by using pairing operations. We observe that $[\sigma^2]_1$ does not have to belong to the SRS, and thus, an updatable SRS may contain gaps. Similarly, the SRS can contain multivariate monomials. However, most of the known updatable and universal zk-SNARKs ([17,24] being exceptions) use SRSs that consist of consequent univariate monomials only, i.e., are of the shape $([(\sigma^i)_{i=0}^{m_1}]_1, [(\sigma^i)_{i=0}^{m}]_2)$ for some m_i.

One reason why efficient updatable and universal zk-SNARKs use a consequent monomial SRS is their reliance on polynomial commitment schemes like KZG [19] that have such SRSs. While many other polynomial commitment schemes are known, up to our knowledge, no efficient one uses non-consequent monomial SRSs.[5] In particular, AHP [10] and PHP [9] model polynomial commitments as polynomial oracles and allow the parties to perform operations (e.g., queries to committed oracles and low-degree tests) related to such oracles. Low-degree tests model consequent monomial SRSs: a committed polynomial is a degree-$\leq m$ polynomial iff it is in the span of X^i for $i \leq m$.

One can use non-consequent monomial SRSs to efficiently construct protocols like broadcast encryption and inner-product commitments [20,18]. We use non-consequent monomial SRSs in the context of sumchecks and updatable and universal zk-SNARKs. We will not define an information-theoretic model, but we mention two possible approaches that both have their limitations. First, the pairing-based setting can be modeled as linear interactive proofs (LIPs, [6]) or non-interactive LIPs [16]. However, either model has to be tweaked to our setting: namely, we allow the generation of updatable SRS for multi-round protocols, with the restrictions natural in such a setting (e.g., one can efficiently "span test" that a committed element is in the span of the SRS). Such a model is tailor-fit to pairings and might not be suitable in other algebraic settings. Second, one can generalize PHPs by adding an abstract model of inner-product commitment schemes and allowing for span tests. Such a model is independent of the algebraic setting but restricts one to a limited number of cryptographic tools (polynomial and inner-product commitment schemes), with a need to redefine the model when more tools are found to be helpful.

We have chosen to remain agnostic on this issue by defining new arguments without an intermediate information-theoretic model.

---

[4] We rely on the pairing-based setting and use the by now standard additive bracket notation, see Section 2 for more details.

[5] A *monomial* SRS is a SRS of the form $[(\sigma^i)_{i \in I_1}]_1, [(\sigma^i)_{i \in I_2}]_2$, where I_1, I_2 are subsets of $[1, m]$. A SRS is *consequent* if both I_1 and I_2 are intervals and non-consequent otherwise. The definition generalizes naturally to the multivariate case.

New Univariate Sumcheck Argument $\mathfrak{Count}$. Let $\mathbb{F}$ be a finite field and let $\mathbb{H} \subset \mathbb{F}$ be a fixed multiplicative subgroup. In a *univariate sumcheck argument* (for multiplicative subgroups), the prover convinces the verifier that the committed polynomial $f(X) \in \mathbb{F}[X]$ sums to the given value $v_M \in \mathbb{F}$ over $\mathbb{H}$.

Let $n_h := |\mathbb{H}|$ and $\mathcal{Z}_{\mathbb{H}}(X) := \prod_{\chi \in \mathbb{H}}(X - \chi)$. Aurora's sumcheck [5] relies on the fact that $\sum_{\chi \in \mathbb{H}} f(\chi) = n_h f(0)$, when $f \in \mathbb{F}_{\leq n_h - 1}[X]$ is a polynomial with $\deg f \leq n_h - 1$. Then, for $f \in \mathbb{F}[X]$ of arbitrarily large degree, $\sum_{\chi \in \mathbb{H}} f(\chi) = v_f$ iff there exist polynomials $R, Q \in \mathbb{F}[X]$, such that (1) $\deg R \leq n_h - 2$, and (2) $f(X) = v_f/n_h + XR(X) + Q(X)\mathcal{Z}_{\mathbb{H}}(X)$. In a cryptographic implementation of Aurora's sumcheck argument in say Marlin [10], the prover uses KZG [19] to commit to R and Q; this means the communication of two group elements. In addition, the prover uses a low-degree test to convince the verifier that (1) holds.

Based on the ILV inner-product commitment [18], we construct a new sumcheck argument $\mathfrak{Count}$ for $f \in \mathbb{F}_{\leq d}[X]$. ILV's non-consequent monomial SRS contains $([(\sigma^i)_{i=0: i \neq N+1}^{2N}]_1, [(\sigma^i)_{i=0}^{N}]_2)$, where σ is a trapdoor and N is a large integer. In ILV, the prover P commits to $\boldsymbol{\mu} \in \mathbb{Z}_p^N$ as $[\mu(\sigma)]_1 \leftarrow \sum_{j=1}^{N} \mu_j [\sigma^j]_1$. When the verifier outputs $\boldsymbol{\nu} \in \mathbb{Z}_p^N$, P returns the inner product $v \leftarrow \boldsymbol{\mu}^\top \boldsymbol{\nu}$ together with a short evaluation proof (a single group element $[\mathsf{op}]_1$) that v is correctly computed. ILV's security relies on $[\sigma^{N+1}]_1$ not belonging to the SRS.

We present an alternative extension of the equality $\sum_{\chi \in \mathbb{H}} f(\chi) = n_h f(0)$ to the case when $d = \deg f$ is arbitrarily large. Namely, we prove that if $f(X) = \sum_{i=0}^{d} f_i X^i \in \mathbb{F}_{\leq d}[X]$, then $\sum_{\chi \in \mathbb{H}} f(\chi) = n_h \cdot (\sum_{i=0}^{\lfloor d/n_h \rfloor} f_{n_h i})$. (See Lemma 1.) Alternatively, $\sum_{\chi \in \mathbb{H}} f(\chi) = v_f$ iff $\boldsymbol{f}^\top \boldsymbol{s} = v_f$, where $\boldsymbol{f} = (f_i)$ and $\boldsymbol{s}$ is a Boolean vector that has ones in positions $n_h \cdot i$ for $i \leq \lfloor d/n_h \rfloor$.

In $\mathfrak{Count}$, the prover first ILV-commits to $\boldsymbol{f}$ and then ILV-opens the commitment to $\boldsymbol{f}^\top \boldsymbol{s}$. Thus, the prover has to output one ILV commitment (one group element) instead of two polynomial commitments (two group elements) in Aurora's sumcheck. Moreover, there is no need for a low-degree test, making $\mathfrak{Count}$ even more efficient. In addition, in the application to $\mathfrak{Vampire}$, $\boldsymbol{s}$ has a small constant number of non-zero elements. Thus, differently from Aurora's sumcheck, the prover's computation is linear in both field operations and group operations. Importantly, the prover does not have to use FFT or polynomial division. An explicit cost of using ILV is that the SRS becomes larger: if the SRS, without $\mathfrak{Count}$, contains $[(\sigma^i)_{i=0}^{d}]_1$ (where d is some constant, fixed by the rest of the zk-SNARK), it now has to contain also $[(\sigma^i)_{i=d+2}^{2d}]_1$ and $[(\sigma^i)_{i=0}^{d}]_2$. (Although, in our construction, we will add significantly less elements to $\mathbb{G}_2$.)

Since sumchecks have ubiquitous applications, $\mathfrak{Count}$ is of independent interest because of both excellent communication and linear-time prover. Linear-time sumchecks are important per se. In particular, univariate sumcheck is used in both updatable and universal zk-SNARKs and transparent zk-SNARKs. As an important application, we will design a new updatable and universal zk-SNARK. We leave it an open question to apply $\mathfrak{Count}$ in transparent zk-SNARKs.

New zk-SNARK. We use $\mathfrak{Count}$ to design a new pairing-based updatable and universal zk-SNARK $\mathfrak{Vampire}$ for the sparse R1CSLite constraint system [9].

𝔙𝔞𝔪𝔭𝔦𝔯𝔢's argument length is four elements of $\mathbb{G}_1$ and two elements of $\mathbb{F}$, which is less than in any known updatable and universal zk-SNARK. While Basilisk [26] (as improved in the full version, [27]) has just 37.5% larger communication than 𝔙𝔞𝔪𝔭𝔦𝔯𝔢, it works for a version of R1CSLite with additional restrictions on the underlying matrices; the version of Basilisk for the arithmetization handled by 𝔙𝔞𝔪𝔭𝔦𝔯𝔢 is less communication-efficient than LunarLite or VOR1CS*.

Let us now give a very brief glimpse to the structure of 𝔙𝔞𝔪𝔭𝔦𝔯𝔢. (The real description, with a very long intuition behind 𝔙𝔞𝔪𝔭𝔦𝔯𝔢's construction, is given in Section 4.) Following Lunar and Basilisk, we use the R1CSLite constraint system, where an instance consists of two matrices $\mathbf{L}$ and $\mathbf{R}$ (the left and right inputs to all constraints) over $\mathbb{F}$ instead of three in the case of R1CS. Let m be the number of constraints. Following Marlin, Lunar, and Basilisk, we use the setting of sparse matrices, where $\mathbf{L}$ and $\mathbf{R}$ have together at most $|\mathbb{K}| = \Theta(m)$ non-zero entries. Here, $\mathbb{K}$ is a multiplicative subgroup of $\mathbb{F}$.

Let z be the interpolating polynomial of $(\mathbb{x}, \mathbb{w}, \boldsymbol{r}_z)$, where $\boldsymbol{r}_z$ is a short random vector needed for zero-knowledge. The prover starts by committing to $\tilde{z}$, where $\tilde{z}$ is a polynomial related to z. Using $\tilde{z}$ helps one efficiently check that the prover used the correct public input. The verifier replies with a random field element α. We reformulate the check that $(\mathbb{x}, \mathbb{w})$ (where $\mathbb{w}$ is encoded in $\tilde{z}$) satisfies the R1CSLite instance as a univariate sumcheck argument that $\sum_{y \in \mathbb{H}} \psi_\alpha(y) = 0$, for a well-chosen polynomial ψ_α. We then run 𝔠𝔬𝔲𝔫𝔱, letting the prover send an ILV-opening $[\psi_{\mathsf{ipc}}(\sigma)]_1$ of ψ_α to the verifier. The verifier replies with another random field element β. The prover's final message consists of two field elements and two group elements. These elements are needed to batch-open three polynomial commitments at different locations, two of which are related to β. It involves a complicated but by now standard step of proving the correctness of the arithmetization of a sparse matrix. This step involves using a univariate sumcheck the second time. However, since here the summed polynomial is of a small degree, we do not need to use 𝔠𝔬𝔲𝔫𝔱. We refer to Section 4 for more details.

𝔙𝔞𝔪𝔭𝔦𝔯𝔢 is based on the ideas of Marlin (e.g., we use a similar arithmetization of sparse matrices), but it uses optimizations of both Lunar [9] and Basilisk [26]. These optimizations (together with an apparently novel combination of the full witness to a single commitment) result in the argument length of 7 elements of $\mathbb{G}_1$ and 2 finite field elements, which is already comparable to prior shortest updatable and universal zk-SNARKs for any NP-complete constraint system.

𝔠𝔬𝔲𝔫𝔱 helps to remove one more group element from the argument of 𝔙𝔞𝔪𝔭𝔦𝔯𝔢. This step is not trivial: the sumcheck argument requires that the sumchecked polynomial f is committed to, which is not the case in 𝔙𝔞𝔪𝔭𝔦𝔯𝔢. We solve this issue using a batching technique similar to Lunar and Basilisk, asking the prover to open two polynomial commitments. The second committed polynomial is a linear combination of other polynomial commitments with coefficients known to the prover and the verifier after opening the first polynomial.

Our second innovation is the use of polynomial commitment aggregation at different points [13,7]. Intuitively, we commit to a single polynomial $\tilde{z}$ that encodes both the left and right inputs of all constraints; this allows us to save one

more group element. When combining the result with the batching technique of the previous paragraph, we need to open three polynomials at different points. In particular, we aggregate the commitment of the second sumcheck, further reducing the proof size by one group element. For batching, we use a technique of [13,7]. However, differently from [13,7], our batching is not randomized since the two opening points are different.

In Theorem 1, we prove that $\mathfrak{Vampire}$ is knowledge-sound in the Algebraic Group Model (AGM, [12]). The proof structure is standard, involving two branches depending on whether the verifier's equations hold as polynomials (we get a reduction to the well-known Power Discrete Logarithm assumption if not). However, the proof of the former case is quite complicated, partially since one has to consider several different polynomials sent by the prover, which depend on different verifier's challenge values.

We prove that $\mathfrak{Vampire}$ is perfectly zero-knowledge by constructing a simulator that uses the trapdoor to make the sumcheck argument acceptable for any, even an all-zero witness. For a simulated argument to be indistinguishable from the real one, we add random terms ($\boldsymbol{r}_z$) to polynomial $\tilde{z}(X)$ which, in the case of real argument, encodes the witness, and, in the case of a simulated argument, encodes a (mostly) zero vector. This assures that even an unbounded adversary who knows the instance and witness cannot tell apart commitments to $\tilde{z}(X)$ in real and simulated arguments. In the full version [22], we prove that $\mathfrak{Vampire}$ is also Sub-ZK (i.e., zero-knowledge even if the SRS generation is compromised, [4,1,11,2]) under the BDH-KE knowledge assumption [1].

On Efficiency. We study how much the argument length can be reduced in updatable and universal SNARKs while only allowing minimal relaxations in other efficiency parameters. We achieve the shortest argument by far. The SRS size of our zk-SNARK is a constant factor larger than in the previous work, which we believe is a reasonable compromise as the SRS needs to be transferred only once. Importantly, the verifier has only to execute $O(|\mathbb{x}|)$ field operations as opposed to $O(|\mathbb{x}|)$ group operations in Groth's zk-SNARK [16].

However, differently from the prior work, prover's computation time in $\mathfrak{Vampire}$ depends on the largest supported R1CSLite size. We discuss this issue further and give a thorough efficiency comparison in the full version [22].

Demaking $\mathfrak{Vampire}$. It is possible to "demake" $\mathfrak{Vampire}$ by removing some of the aggressive length-optimization to obtain a larger argument size but better (say) the SRS size. We leave it as an open question about which optimization should be removed first or whether this is needed at all.

2   Preliminaries

Let $\mathbb{F} = \mathbb{Z}_p$ be a finite field of prime order p, and let $\mathbb{F}_{\leq d}[X] \subset \mathbb{F}[X]$ be the set of degree $\leq d$ polynomials. Define the set of (d, d_{gap})-*punctured* univariate

polynomials over $\mathbb{F}$ as

$$\mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, X) := \{f(X) = \textstyle\sum_{i=0}^{d_{\mathsf{gap}}+d} f_i X^i \in \mathbb{F}_{\le d_{\mathsf{gap}}+d}[X] : f_{d_{\mathsf{gap}}} = 0\} \ .$$

Let $\boldsymbol{x} \circ \boldsymbol{y}$ be the elementwise product of vectors $\boldsymbol{x}$ and $\boldsymbol{y}$, $\forall i.(\boldsymbol{x} \circ \boldsymbol{y})_i = x_i y_i$. Let $\mathsf{I}_n \in \mathbb{F}^{n \times n}$ be the n-dimensional identity matrix. Denote matrix and vector elements by using square brackets as in $\mathbf{A}[i, j]$ and $\boldsymbol{a}[i]$.

Interpolation. Let ω be the n_h-th primitive root of unity in $\mathbb{F}$ and let $\mathbb{H} = \{\omega^j : 0 \le j < n_h\}$ be a multiplicative subgroup of $\mathbb{F}$. Then,
- For any $T \subset \mathbb{F}$, the *vanishing polynomial* $\mathcal{Z}_T(X) := \prod_{i \in T}(X - i)$ is the degree-$|T|$ monic polynomial, such that $\mathcal{Z}_T(i) = 0$ for all $i \in T$. $\mathcal{Z}_{\mathbb{H}}(Y) = Y^{n_h} - 1$ can be computed in $\Theta(\log n_h)$ field operations.
- For $i \in [1, n_h]$, $\ell_i^{\mathbb{H}}(Y)$ is the *ith Lagrange polynomial*, i.e., the unique degree $n_h - 1$ polynomial, such that $\ell_i^{\mathbb{H}}(\omega^{i-1}) = 1$ and $\ell_i^{\mathbb{H}}(\omega^{j-1}) = 0$ for $i \ne j$. It is well known that

$$\ell_i^{\mathbb{H}}(Y) = \frac{\mathcal{Z}_{\mathbb{H}}(Y)}{(\mathcal{Z}'_{\mathbb{H}}(\omega^{i-1}) \cdot (Y - \omega^{i-1}))} = \frac{\mathcal{Z}_{\mathbb{H}}(Y)\omega^{i-1}}{(n_h(Y - \omega^{i-1}))} \ .$$

  Here, $\mathcal{Z}'_{\mathbb{H}}(X) = d\mathcal{Z}_{\mathbb{H}}(X)/dX$.
- $L_X^{\mathbb{H}}(Y) := \mathcal{Z}_{\mathbb{H}}(Y)X/(n_h(Y - X)) \in \mathbb{F}(X, Y)$ (a lifted Lagrange rational function), with $L_{\omega^{i-1}}(Y) = \ell_i^{\mathbb{H}}(Y)$ for $i \in [1, n_h]$.

For $f \in \mathbb{F}[X]$, let $\widehat{f}^{\mathbb{H}}(X) := \sum_{i=1}^{n_h} f(\omega^{i-1})\ell_i^{\mathbb{H}}(X)$ be its low-degree extension. To simplify notation, we often omit the accent $\widehat{\cdot}$ and the superscript $\mathbb{H}$.

R1CSLite. R1CSLite [9,26] is a variant of the Rank 1 Constraint System [14,10]. An R1CSLite instance $\mathcal{I} = (\mathbb{F}, m, m_0, \mathbf{L}, \mathbf{R})$ consists of a field $\mathbb{F}$, instance size m, input size m_0, and matrices $\mathbf{L}, \mathbf{R} \in \mathbb{F}^{m \times m}$. An R1CSLite instance is *sparse* if $\mathbf{L}$ and $\mathbf{R}$ have $O(m)$ non-zero elements.

$\mathcal{I} = (\mathbb{F}, m, m_0, \mathbf{L}, \mathbf{R})$ defines the following relation $\mathcal{R} = \mathcal{R}_{\mathcal{I}}$:

$$\mathcal{R} := \begin{cases} (\mathbb{x}, \mathbb{w}) : \mathbb{x} = (z_1, \dots, z_{m_0})^\top \ \wedge \ \mathbb{w} = \left(\begin{smallmatrix} \boldsymbol{z}_{\mathsf{a}} \\ \boldsymbol{z}_{\mathsf{b}} \end{smallmatrix}\right) \wedge \boldsymbol{z}_{\mathsf{a}}, \boldsymbol{z}_{\mathsf{b}} \in \mathbb{F}^{m-m_0-1} \wedge \\ \boldsymbol{z}_{\mathsf{l}} = \left(\begin{smallmatrix} 1 \\ \mathbb{x} \\ \boldsymbol{z}_{\mathsf{a}} \end{smallmatrix}\right) \wedge \boldsymbol{z}_{\mathsf{r}} = \left(\begin{smallmatrix} \mathbf{1}_{m_0+1} \\ \boldsymbol{z}_{\mathsf{b}} \end{smallmatrix}\right) \wedge \boldsymbol{z}_{\mathsf{l}} = \mathbf{L}(\boldsymbol{z}_{\mathsf{l}} \circ \boldsymbol{z}_{\mathsf{r}}) \wedge \boldsymbol{z}_{\mathsf{r}} = \mathbf{R}(\boldsymbol{z}_{\mathsf{l}} \circ \boldsymbol{z}_{\mathsf{r}}) \end{cases} \ .$$

Equivalently, $\mathbf{W}\boldsymbol{z}^* = \mathbf{0}$, where

$$\mathbf{W} = \left(\begin{smallmatrix} \mathsf{I}_m & \mathbf{0} & -\mathbf{L} \\ \mathbf{0} & \mathsf{I}_m & -\mathbf{R} \end{smallmatrix}\right) \in \mathbb{F}^{2m \times 3m} \ , \quad \boldsymbol{z}^* = \left(\begin{smallmatrix} \boldsymbol{z}_{\mathsf{l}} \\ \boldsymbol{z}_{\mathsf{r}} \\ \boldsymbol{z} = \boldsymbol{z}_{\mathsf{l}} \circ \boldsymbol{z}_{\mathsf{r}} \end{smallmatrix}\right) \ . \tag{1}$$

Basic Cryptography. We denote the security parameter by λ. For any algorithm $\mathcal{A}$, $r \leftarrow_{\$} \mathsf{RND}_\lambda(\mathcal{A})$ samples random coins of sufficient length for $\mathcal{A}$ for fixed λ. By $y \leftarrow \mathcal{A}(x; r)$, we denote that $\mathcal{A}$ outputs y on input x and random coins r. PPT means probabilistic polynomial time.

Pairings. A bilinear group generator $\mathsf{Pgen}(1^\lambda)$ returns $\mathsf{p} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \hat{e}, [1]_1, [1]_2)$, where p is a prime, $\mathbb{G}_1, \mathbb{G}_2$, and $\mathbb{G}_T$ are three additive cyclic groups of order p, $\hat{e} : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a non-degenerate efficiently computable bilinear pairing, and $[1]_\iota$ is a generator of $\mathbb{G}_\iota$ for $\iota \in \{1, 2, T\}$ with $[1]_T = \hat{e}([1]_1, [1]_2)$. In this paper, $\mathbb{F} = \mathbb{Z}_p$ has always two large multiplicative subgroups $\mathbb{H}$ and $\mathbb{K}$. Thus, we assume implicitly that $|\mathbb{H}|, |\mathbb{K}| \mid (p - 1)$. We require the bilinear pairing to be Type-3, that is, not to have an efficient isomorphism between $\mathbb{G}_1$ and $\mathbb{G}_2$. In practice, one uses a fixed pairing-friendly curve like BLS-381; then, $|\mathbb{K}|, |\mathbb{H}| \mid 2^{32}$.

We use the by now standard additive bracket notation, by writing $[a]_\iota$ to denote $a[1]_\iota$ for $\iota \in \{1, 2, T\}$. We denote $\hat{e}([x]_1, [y]_2)$ by $[x]_1 \bullet [y]_2$. Thus, $[x]_1 \bullet [y]_2 = [xy]_T$. We freely use the bracket notation together with matrix notation; for example, if $\mathbf{A} \cdot \mathbf{B} = \mathbf{C}$ then $[\mathbf{A}]_1 \bullet [\mathbf{B}]_2 = [\mathbf{C}]_T$.

Polynomial Commitment Schemes. In a polynomial commitment scheme [19], the prover commits to a polynomial $f \in \mathbb{F}_{\leq d}[X]$ and later opens it to $f(\beta)$ for $\beta \in \mathbb{F}$ chosen by the verifier. The (non-randomized) KZG [19] polynomial commitment scheme consists of the following algorithms:

Setup: Given 1^λ, return $\mathsf{p} \leftarrow \mathsf{Pgen}(1^\lambda)$.

Commitment key generation: Given a system parameter p and an upper-bound d on the polynomial degree, compute the trapdoor $\mathsf{tk} = \sigma \leftarrow_\$ \mathbb{Z}_p^*$ and the commitment key $\mathsf{ck} \leftarrow (\mathsf{p}, [(\sigma^i)_{i=0}^d]_1, [1, \sigma]_2)$. Return $(\mathsf{ck}, \mathsf{tk})$.

Commitment: Given a commitment key ck and a polynomial $f \in \mathbb{F}_{\leq d}[X]$, return the commitment $[f(\sigma)]_1 \leftarrow \sum_{j=0}^d f_j [\sigma^j]_1$.

Opening: Given a commitment key ck, a commitment $[f(\sigma)]_1$, an evaluation point $\beta \in \mathbb{F}$, and a polynomial $f \in \mathbb{F}_{\leq d}[X]$, set $v \leftarrow f(\beta)$ and $f_{\mathsf{pc}}(X) \leftarrow (f(X) - v)/(X - \beta)$. The evaluation proof is $[f_{\mathsf{pc}}(\sigma)]_1 \leftarrow \sum_{j=0}^{d-1} (f_{\mathsf{pc}})_j [\sigma^j]_1$. Return $(v, [f_{\mathsf{pc}}(\sigma)]_1)$.

Verification: Given a commitment key ck, a commitment $[f(\sigma)]_1$, an evaluation point β, a purported evaluation $v = f(\beta)$, and an evaluation proof $[f_{\mathsf{pc}}(\sigma)]_1$, check $[f(\sigma) - v]_1 \bullet [1]_2 = [f_{\mathsf{pc}}(\sigma)]_1 \bullet [\sigma - \beta]_2$.

KZG's security is based on the fact that $(X - \beta) \mid (f(X) - v) \Leftrightarrow f(\beta) = v$.

Inner-Product Commitment Schemes. In an inner-product commitment scheme [20,18], the prover commits to a vector $\boldsymbol{\mu} \in \mathbb{F}^N$ and later opens it to the inner product $\boldsymbol{\mu}^\top \boldsymbol{\nu}$ for $\boldsymbol{\nu} \in \mathbb{F}^N$ chosen by the verifier. The (non-randomized) ILV [18] inner-product commitment scheme consists of the following algorithms:

Setup: Given 1^λ, return $\mathsf{p} \leftarrow \mathsf{Pgen}(1^\lambda)$.

Commitment key generation: Given a system parameter p and a vector dimension N, compute the trapdoor $\mathsf{tk} = \sigma \leftarrow_\$ \mathbb{Z}_p^*$ and the commitment key $\mathsf{ck} \leftarrow ([(\sigma^i)_{i=0:i\neq N+1}^{2N}]_1, [(\sigma^i)_{i=0}^N]_2)$. Return $(\mathsf{ck}, \mathsf{tk})$.

Commitment: Given a commitment key ck and a vector $\boldsymbol{\mu} \in \mathbb{F}^N$, compute the coefficients of $\mu(X) \leftarrow \sum_{j=1}^N \mu_j X^j \in \mathbb{F}_{\leq N}[X]$; $[\mu(\sigma)]_1 = \sum_{j=1}^N \mu_j [\sigma^j]_1$. Return the commitment $[\mu(\sigma)]_1$.

Opening: Given a commitment key ck, a commitment $[\mu(\sigma)]_1$, the vector $\boldsymbol{\mu}$, and a vector $\boldsymbol{\nu}$, let $v \leftarrow \boldsymbol{\mu}^\top \boldsymbol{\nu}$. Set $\nu^*(X) \leftarrow \sum_{j=1}^N \nu_j X^{N+1-j} \in \mathbb{F}_{\leq N}[X]$,

and $\mu_{\mathsf{ipc}}(X) \leftarrow \mu(X)\nu^*(X) - vX^{N+1} \in \mathsf{PolyPunc}_{\mathbb{F}}(N-1, N+1, X)$. The evaluation proof is $[\mu_{\mathsf{ipc}}(\sigma)]_1 \leftarrow \sum_{i=1, i\neq N+1}^{2N} \mu_{\mathsf{ipc}}[\sigma^i]_1$. Return $(v, [\mu_{\mathsf{ipc}}(\sigma)]_1)$.

Verification: Given a commitment key ck, a commitment $[\mu(\sigma)]_1$, a vector $\boldsymbol{\nu}$, a purported inner product $v = \boldsymbol{\mu}^\top \boldsymbol{\nu}$, and an evaluation proof $[\mu_{\mathsf{ipc}}(\sigma)]_1$, check $[\mu_{\mathsf{ipc}}(\sigma)]_1 \bullet [1]_2 = [\mu(\sigma)]_1 \bullet \sum_{j=1}^{N} \nu_j [\sigma^{N+1-j}]_2 - v[\sigma^N]_1 \bullet [\sigma]_2$.

ILV's security follows since the coefficient of X^{N+1} in $\mu_{\mathsf{ipc}}(X)$ is $\boldsymbol{\mu}^\top \boldsymbol{\nu} - v = 0$ iff v is correctly computed. In this paper, the vector $\boldsymbol{\nu}$ is public and known in advance. Then, the verifier only has to compute two pairings and no exponentiations.

Succinct Zero-Knowledge Arguments. The following definition is based on [9]. Groth et al. [17] introduced the notion of (preprocessing) zk-SNARKs with specializable universal structured reference string (SRS). This notion formalizes the idea that the key generation for $\mathcal{R} \in \mathcal{UR}$, where $\mathcal{UR}$ is a universal relation, can be seen as the sequential combination of two steps. First, a probabilistic algorithm generating an SRS for $\mathcal{UR}$ and second, a deterministic algorithm specializing this universal SRS into one for a specific $\mathcal{R}$.

We consider relation families $(\mathsf{Pgen}, \{\mathcal{UR}_{\mathsf{p},N}\}_{\mathsf{p}\in\mathrm{range}(\mathsf{Pgen}), N\in\mathbb{N}})$ parametrized by $\mathsf{p} \in \mathsf{Pgen}(1^\lambda)$ and a size bound $N \in \mathsf{poly}(\lambda)$.[6] A *succinct zero-knowledge argument $\Pi = (\mathsf{Pgen}, \mathsf{KGen}, \mathsf{Derive}, \mathsf{P}, \mathsf{V})$ with specializable universal SRS for a relation family* $(\mathsf{Pgen}, \{\mathcal{UR}_{\mathsf{p},N}\}_{\mathsf{p}\in\{0,1\}^*, N\in\mathbb{N}})$ consists of the following algorithms.

Setup: Given 1^λ, return $\mathsf{p} \leftarrow \mathsf{Pgen}(1^\lambda)$.

Universal SRS Generation: a probabilistic algorithm $\mathsf{KGen}(\mathsf{p}, N) \to (\mathsf{srs}, \mathsf{td})$ that takes as input public parameters p and an upper bound N on the relation size, and outputs $\mathsf{srs} = (\mathsf{ek}, \mathsf{vk})$ together with a trapdoor. We assume implicitly that elements like ek and vk contain p.

SRS Specialization: a deterministic algorithm $\mathsf{Derive}(\mathsf{srs}, \mathcal{R}) \to (\mathsf{ek}_\mathcal{R}, \mathsf{vk}_\mathcal{R})$ that takes as input a universal SRS srs and a relation $\mathcal{R} \in \mathcal{UR}_{\mathsf{p},N}$, and outputs a specialized SRS $\mathsf{srs}_\mathcal{R} := (\mathsf{ek}_\mathcal{R}, \mathsf{vk}_\mathcal{R})$.

Prover/Verifier: a pair of interactive algorithms $\langle \mathsf{P}(\mathsf{ek}_\mathcal{R}, \mathbb{x}, \mathbb{w}), \mathsf{V}(\mathsf{vk}_\mathcal{R}, \mathbb{x}) \rangle \to b$, where P takes a proving key $\mathsf{ek}_\mathcal{R}$ for a relation $\mathcal{R}$, a statement $\mathbb{x}$, and a witness $\mathbb{w}$, s.t. $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$, and V takes a verification key for a relation $\mathcal{R}$ and a statement $\mathbb{x}$, and either accepts ($b=1$) or rejects ($b=0$) the argument.

Π must satisfy the following four requirements.

<u>Completeness.</u> For all $\mathsf{p} \in \mathrm{range}(\mathsf{Pgen})$, $N \in \mathbb{N}$, $\mathcal{R} \in \mathcal{UR}_{\mathsf{p},N}$, and $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$,

$$\Pr\left[\langle \mathsf{P}(\mathsf{ek}_\mathcal{R}, \mathbb{x}, \mathbb{w}), \mathsf{V}(\mathsf{vk}_\mathcal{R}, \mathbb{x}) \rangle = 1 \,\middle|\, \begin{array}{l} (\mathsf{srs}, \mathsf{td}) \leftarrow \mathsf{KGen}(\mathsf{p}, N); \\ (\mathsf{ek}_\mathcal{R}, \mathsf{vk}_\mathcal{R}) \leftarrow \mathsf{Derive}(\mathsf{srs}, \mathcal{R}) \end{array} \right] = 1 \ .$$

<u>Succinctness.</u> Π is *succinct* if the running time of V is $\mathsf{poly}(\lambda + |\mathbb{x}| + \log |\mathbb{w}|)$ and the communication size is $\mathsf{poly}(\lambda + \log |\mathbb{w}|)$.

---

[6] 𝕮𝖔𝖚𝖓𝖙 and 𝖁𝖆𝖒𝖕𝖎𝖗𝖊 have several size bounds. The definitions generalize naturally.

Knowledge-Soundness. Π is *knowledge-sound*, if for every non-uniform PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a non-uniform PPT extractor $\mathsf{Ext}_\mathcal{A}$ [7], s.t.

$$\Pr\left[\begin{array}{c} \langle \mathcal{A}_2\left(st; r\right), \mathsf{V}(\mathsf{vk}_\mathcal{R}, \mathbb{x})\rangle = 1 \\ \wedge \neg\mathcal{R}(\mathbb{x}, \mathbb{w}) \end{array} \middle| \begin{array}{c} \mathsf{p} \leftarrow \mathsf{Pgen}(1^\lambda); (\mathsf{srs}, \mathsf{td}) \leftarrow \mathsf{KGen}(\mathsf{p}, N); \\ r \leftarrow_{\$} \mathsf{RND}_\lambda(\mathcal{A}); (\mathcal{R}, \mathbb{x}, st) \leftarrow \mathcal{A}_1(\mathsf{srs}; r); \\ \mathbb{w} \leftarrow \mathsf{Ext}_\mathcal{A}(\mathsf{srs}; r); \\ (\mathsf{ek}_\mathcal{R}, \mathsf{vk}_\mathcal{R}) \leftarrow \mathsf{Derive}(\mathsf{srs}, \mathcal{R}) \end{array}\right] = \mathsf{negl}(\lambda) \ .$$

Zero-Knowledge. Π is (statistical) *zero-knowledge* if there exists a PPT simulator Sim, s.t. for all unbound $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, all $\mathsf{p} \in \mathrm{range}(\mathsf{Pgen})$, all $N \in \mathsf{poly}(\lambda)$,

$$\Pr\left[\begin{array}{c} \langle \mathsf{P}\left(\mathsf{ek}_\mathcal{R}, \mathbb{x}, \mathbb{w}\right), \mathcal{A}_2(st)\rangle = 1 \wedge \\ \mathcal{R}(\mathbb{x}, \mathbb{w}) \wedge \mathcal{R} \in \mathcal{UR}_{\mathsf{p}, N} \end{array} \middle| \begin{array}{c} (\mathsf{srs}, \mathsf{td}) \leftarrow \mathsf{KGen}(\mathsf{p}, N); \\ (\mathcal{R}, \mathbb{x}, \mathbb{w}, st) \leftarrow \mathcal{A}_1(\mathsf{srs}); \\ (\mathsf{ek}_\mathcal{R}, \mathsf{vk}_\mathcal{R}) \leftarrow \mathsf{Derive}(\mathsf{srs}, \mathcal{R}) \end{array}\right] \approx_s$$

$$\Pr\left[\begin{array}{c} \langle \mathsf{Sim}\left(\mathsf{srs}, \mathsf{td}, \mathcal{R}, \mathbb{x}\right), \mathcal{A}_2(st)\rangle = 1 \wedge \\ \mathcal{R}(\mathbb{x}, \mathbb{w}) \wedge \mathcal{R} \in \mathcal{UR}_{\mathsf{p}, N} \end{array} \middle| \begin{array}{c} (\mathsf{srs}, \mathsf{td}) \leftarrow \mathsf{KGen}(\mathsf{p}, N); \\ (\mathcal{R}, \mathbb{x}, \mathbb{w}, st) \leftarrow \mathcal{A}_1(\mathsf{srs}); \\ (\mathsf{ek}_\mathcal{R}, \mathsf{vk}_\mathcal{R}) \leftarrow \mathsf{Derive}(\mathsf{srs}, \mathcal{R}) \end{array}\right] \ .$$

Here, $\approx_s$ denotes the statistical distance as a function of λ. Π is *perfect zero-knowledge* if the above probabilities are equal.

Π is *subversion zero-knowledge* (Sub-ZK, [4]), if it is zero-knowledge even in the case the SRS is maliciously generated. For perfect zero-knowledge arguments, Sub-ZK follows from the usual zero-knowledge (with trusted SRS), SRS verifiability (there exists a PPT algorithm that checks that the SRS belongs to $\mathrm{range}(\mathsf{KGen})$), and a SNARK-specific knowledge assumption, [1,2]. We will provide the formal definition in the full version [22].

Π is *updatable* [17], if the SRS can be sequentially updated by many updaters, such that knowledge-soundness holds if either the original SRS creator or one of the updaters is honest. Groth et al. [17] showed that an updatable SRS cannot contain non-monomial polynomial evaluations. Moreover, an updatable SRS must be verifiable in the same sense as in the case of Sub-ZK.

Since $\mathfrak{Vampire}$ is public-coin and has a constant number of rounds, we can apply the Fiat-Shamir heuristic to obtain a zk-SNARK.

Sumcheck Arguments. In a sumcheck argument [23] over $\mathbb{F}$, the prover convinces the verifier that for $\mathbb{H} \subseteq \mathbb{F}$, $f \in \mathbb{F}[X_1, \ldots, X_c]$, and $v_f \in \mathbb{F}$, it holds that $\sum_{(x_1, \ldots, x_c) \in \mathbb{H}^c} f(x_1, \ldots, x_c) = v_f$. Multivariate sumcheck has many applications, with usually relatively small $|\mathbb{H}|$ but large c. In the context of efficient updatable zk-SNARKs, one is often interested in *univariate sumcheck*, where $c = 1$ but $|\mathbb{H}|$ is large. Univariate sumcheck arguments are most efficient when $\mathbb{H}$ is either an affine subspace or a multiplicative subgroup [5].

The univariate sumcheck relation for multiplicative subgroups is the set of all pairs $\mathcal{R}_{\mathsf{sum}} := \{((\mathbb{F}, d, \mathbb{H}, v_f), f)\}$, where $\mathbb{F}$ is a finite field, d is a positive integer, $\mathbb{H}$ is a multiplicative subgroup of $\mathbb{F}$, $v_f \in \mathbb{F}$, $f \in \mathbb{F}_{\leq d}[X]$, and $\sum_{\chi \in \mathbb{H}} f(\chi) = v_f$.

---

[7] Note that although the protocol is interactive, extraction is done non-interactively. This is sometimes called straight-line extractability.

Aurora's Sumcheck. As a part of the zk-SNARK Aurora, Ben-Sasson et al. [5] proposed an efficient univariate sumcheck ("Aurora's sumcheck") for multiplicative subgroups. Since the new univariate sumcheck relies on similar techniques, we next recall Aurora's sumcheck.

As before, let $\mathbb{H} = \langle \omega \rangle = \{\omega^i : i \in [0, n_h - 1]\}$ be a cyclic multiplicative subgroup of order $n_h = |\mathbb{H}|$. Fact 1 underlies Aurora's sumcheck.

Fact 1 *Let $f \in \mathbb{F}[X]$ with $\deg f \leq n_h - 1$. Then $\sum_{\chi \in \mathbb{H}} f(\chi) = n_h f(0)$.*

In the case of a large-degree f, Ben-Sasson et al. [5] used Fact 2 to construct Aurora's sumcheck argument for proving that $\sum_{\chi \in \mathbb{H}} f(\chi) = v_f$.

Fact 2 (Core Lemma of Aurora's Sumcheck) *Let $f \in \mathbb{F}[X]$ with $d = \deg f \geq n_h$. Then, $\sum_{\chi \in \mathbb{H}} f(\chi) = v_f$ iff there exist $R \in \mathbb{F}_{\leq n_h - 2}[X]$ and $Q \in \mathbb{F}_{\leq d - n_h}[X]$, such that $f(X) = v_f/n_h + R(X)X + Q(X)\mathcal{Z}_{\mathbb{H}}(X)$.*

Assume that $d = \deg f = \mathsf{poly}(\lambda)$ while $p = 2^{\Theta(\lambda)}$. In Aurora's sumcheck argument, the prover sends to the verifier polynomial commitments to f, R, and Q. The verifier accepts if (1) R has a low degree $\leq n_h - 2$ and (2) $f(X) = v_f/n_h + R(X)X + Q(X)\mathcal{Z}_{\mathbb{H}}(X)$.

On top of two polynomial commitments (two group elements), one has to implement a low-degree test to check that $\deg R \leq n_h - 2$. As the low-degree test, Aurora uses an interactive oracle proof for testing proximity to the Reed–Solomon code, resulting in additional costs. The full version of Basilisk [27] implementes a low-degree test in a partially costless way (without added argument size or verifier's computation); however, one may need to add a large number of elements to the SRS for their low-degree test to succeed.

Assumptions. Let $d_1(\lambda), d_2(\lambda) \in \mathsf{poly}(\lambda)$. Pgen is (d_1, d_2)-*PDL (Power Discrete Logarithm [21]) secure* if for any non-uniform PPT $\mathcal{A}$, $\mathsf{Adv}^{\mathrm{pdl}}_{d_1, d_2, \mathsf{Pgen}, \mathcal{A}}(\lambda) :=$

$$\Pr\left[\mathcal{A}\left(\mathsf{p}, [(x^i)_{i=0}^{d_1}]_1, [(x^i)_{i=0}^{d_2}]_2\right) = x \,\middle|\, \mathsf{p} \leftarrow \mathsf{Pgen}(1^\lambda); x \leftarrow_\$ \mathbb{F}^*\right] = \mathsf{negl}(\lambda) \ .$$

Algebraic Group Model (AGM). AGM is an idealized model [12] for security proofs. In the AGM, adversaries are restricted to be *algebraic* in the following sense: if $\mathcal{A}$ inputs some group elements and outputs a group element, it provides an algebraic representation of the latter in terms of the former. More precisely, if $\mathcal{A}$ has received group elements $[\boldsymbol{x}_1]_1, [\boldsymbol{x}_2]_2$ so far and outputs $[y_1]_1, [y_2]_2$, then there exists an extractor $\mathsf{Ext}_{\mathcal{A}}$ which on the same input and random coins outputs integer vectors $\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2$ such that $[y_1]_1 = \sum_i \gamma_{1,i}[\boldsymbol{x}_{1,i}]_1$ and $[y_2]_2 = \sum_j \gamma_{2,j}[\boldsymbol{x}_{2,j}]_2$.

3   𝕮𝖔𝖚𝖓𝖙: New Univariate Sumcheck Argument

In this section, we propose 𝕮𝖔𝖚𝖓𝖙, a new sumcheck argument with improved online efficiency (including the argument size) but a larger SRS size than Aurora's univariate sumcheck. We first prove the following generalization of Fact 1, an alternative to Fact 2 in the case f has degree larger than $n_h - 1$.

Lemma 1. *Let $f(X) = \sum_{i=0}^{d} f_i X^i$ for $d \geq 0$. Then, $\sum_{\chi \in \mathbb{H}} f(\chi) = n_h \cdot \sum_{i=0}^{\lfloor d/n_h \rfloor} f_{n_h i}$.*

Proof. Write $f(X) = R(X) + Q(X)\mathcal{Z}_{\mathbb{H}}(X)$ for $\deg R \leq n_h - 1$. Based on Fact 1, $\sum_{\chi \in \mathbb{H}} f(\chi) = \sum_{\chi \in \mathbb{H}} R(\chi) = n_h R(0)$. Since $X^{n_h} \equiv 1 \pmod{\mathcal{Z}_{\mathbb{H}}(X)}$, $f(X) = \sum_{i=0}^{d} f_i X^i \equiv \sum_{j=0}^{n_h - 1} (\sum_{i=0:n_h|(i-j)}^{d} f_i) X^j \pmod{\mathcal{Z}_{\mathbb{H}}(X)}$. Since $f(X) \equiv R(X) \pmod{\mathcal{Z}_{\mathbb{H}}(X)}$, $R(0) = \sum_{i=0}^{\lfloor d/n_h \rfloor} f_{n_h i}$. Thus, $\sum_{\chi \in \mathbb{H}} f(\chi) = n_h \cdot \sum_{i=0}^{\lfloor d/n_h \rfloor} f_{n_h i}$. $\square$

$\mathfrak{Count}$ is based on the following result.

Lemma 2 (Core Lemma of $\mathfrak{Count}$). *Let $\mathbb{H}$ be an order-$n_h > 1$ multiplicative subgroup of $\mathbb{F}^*$. Let $d_{\mathsf{gap}}, d > 0$ with $d_{\mathsf{gap}} \geq n_h \cdot \lfloor d/n_h \rfloor$, and $f \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, X)$. Define*

$$S(X) := \sum_{i=0}^{\lfloor d/n_h \rfloor} X^{d_{\mathsf{gap}} - n_h i} \in \mathbb{F}_{\leq d_{\mathsf{gap}}}[X] \ .$$

Then, $\sum_{\chi \in \mathbb{H}} f(\chi) = v_f$ and $\deg f \leq d$ iff there exists $f_{\mathsf{ipc}} \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, X)$, s.t.

$$f(X)S(X) - f_{\mathsf{ipc}}(X) = \frac{v_f}{n_h} \cdot X^{d_{\mathsf{gap}}} \ . \tag{2}$$

Here, d_{gap} is a parameter fixed by the master protocol (in our case, $\mathfrak{Vampire}$) that uses $\mathfrak{Count}$ as a subroutine.

Proof. Clearly, we need $d_{\mathsf{gap}} \geq n_h \cdot \lfloor d/n_h \rfloor$ for S to be a polynomial.

($\Rightarrow$) Define $f_{\mathsf{ipc}}(X) := f(X)S(X) - v_f/n_h \cdot X^{d_{\mathsf{gap}}}$. We must only show that $f_{\mathsf{ipc}} \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, X)$. Since $\deg f \leq d$ and $\deg S = d_{\mathsf{gap}}$, we have $\deg f_{\mathsf{ipc}} \leq d_{\mathsf{gap}} + d$. Since $f(X)S(X) = (\sum_{i=0}^{d} f_i X^i)(\sum_{i=0}^{\lfloor d/n_h \rfloor} X^{d_{\mathsf{gap}} - n_h i})$, the coefficient of $X^{d_{\mathsf{gap}}}$ in $f(X)S(X)$ is $\sum_{i=0}^{\lfloor d/n_h \rfloor} f_{n_h i}$. By Lemma 1, $\sum_{i=0}^{\lfloor d/n_h \rfloor} f_{n_h i} = v_f/n_h$. Thus, the coefficient of $X^{d_{\mathsf{gap}}}$ in f_{ipc} is 0 and $f_{\mathsf{ipc}} \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, X)$.

($\Leftarrow$) Suppose Eq. (2) holds for $f_{\mathsf{ipc}} \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, X)$. Since $\deg S = d_{\mathsf{gap}}$ and $\deg f_{\mathsf{ipc}} \leq d_{\mathsf{gap}} + d$, we have $\deg f \leq d$. As in ($\Rightarrow$), the coefficient of $X^{d_{\mathsf{gap}}}$ in $f(X)S(X)$ is $\sum_{i=0}^{\lfloor d/n_h \rfloor} f_{n_h i}$, which is equal to $(\sum_{\chi \in \mathbb{H}} f(\chi))/n_h$ due to Lemma 1. Since f_{ipc} is missing the monomial $X^{d_{\mathsf{gap}}}$, we get that $v_f = \sum_{\chi \in \mathbb{H}} f(\chi)$. $\square$

It is important that f_{ipc} has degree $\leq d_{\mathsf{gap}} + d$. Thus, one cannot add elements $[\sigma^i]_1$ for $i > d_{\mathsf{gap}} + d$ to the SRS of a master argument that uses $\mathfrak{Count}$.

Description of $\mathfrak{Count}$. Next, we describe $\mathfrak{Count}$ as a zk-SNARK for the sumcheck relation; if needed, it is straightforward to modify it to the language of polynomial oracles. In $\mathfrak{Count}$, the common input is $([f(\sigma)]_1, v_f)$. The prover sends to the verifier a polynomial commitment to $[f_{\mathsf{ipc}}(\sigma)]_1$, and the verifier accepts that $\sum_{\chi \in \mathbb{H}} f(\chi) = v_f$ iff a naturally modified version of Eq. (2) holds on committed polynomials. See Fig. 1 for the full argument. Here, Derive does only preprocessing and does not do any specialization.

Since we only use $\mathfrak{Count}$ as a sub-argument of $\mathfrak{Vampire}$, we do not formally have to prove that it is knowledge-sound or zero-knowledge. Nevertheless, for the sake of completeness, we provide proof sketches.

Pgen(p): generate p as usually. We implicitly assume $n_h \mid (p-1)$.

KGen(p, n_h, d, d_{gap}): $\mathcal{S}_1(X) \leftarrow \{(X^i)_{i=0:i\neq d_{\mathsf{gap}}}^{d_{\mathsf{gap}}+d}\}$; $\mathcal{S}_2(X) \leftarrow \{1, X, (X^{d_{\mathsf{gap}}-n_h i})_{i=0}^{\lfloor d/n_h \rfloor}\}$;

$\quad\quad \sigma \leftarrow_{\$} \mathbb{F}^*$; td $\leftarrow \sigma$; srs $\leftarrow (\mathsf{p}, n_h, d, d_{\mathsf{gap}}, [g(\sigma) : g \in \mathcal{S}_1(X)]_1, [g(\sigma) : g \in \mathcal{S}_2(X)]_2)$

Derive(srs): $S(X) \leftarrow \sum_{i=0}^{\lfloor d/n_h \rfloor} X^{d_{\mathsf{gap}}-n_h i} \in \mathbb{F}_{\leq d_{\mathsf{gap}}}[X]$; $\mathsf{ek}_{\mathcal{R}} \leftarrow$ srs;

$\quad\quad \mathsf{vk}_{\mathcal{R}} \leftarrow (\mathsf{srs}, [S(\sigma)]_2, [\sigma^{d_{\mathsf{gap}}}]_T)$; return $(\mathsf{ek}_{\mathcal{R}}, \mathsf{vk}_{\mathcal{R}})$;

P($\mathsf{ek}_{\mathcal{R}}, \mathbb{x}, \mathbb{w} = f$) $\quad\quad\quad$ /* $\mathbb{x} = ([f(\sigma)]_1, v_f)$ */ $\quad\quad\quad\quad\quad$ V($\mathsf{vk}_{\mathcal{R}}, \mathbb{x}$)

. . . . .Online phase. . . . .

$S(X) \leftarrow \sum_{i=0}^{\lfloor d/n_h \rfloor} X^{d_{\mathsf{gap}}-n_h i} \in \mathbb{F}_{\leq d_{\mathsf{gap}}}[X]$; $f_{\mathsf{ipc}}(X) \leftarrow f(X)S(X) - v_f/n_h \cdot X^{d_{\mathsf{gap}}}$

$\quad\quad\quad\quad\quad [f_{\mathsf{ipc}}(\sigma)]_1$

$\quad\quad\quad\quad\quad \xrightarrow{\hspace{3cm}}$

$\quad\quad\quad\quad$ Check $[f(\sigma)]_1 \bullet [S(\sigma)]_2 - [f_{\mathsf{ipc}}(\sigma)]_1 \bullet [1]_2 = v_f/n_h \cdot [\sigma^{d_{\mathsf{gap}}}]_T$
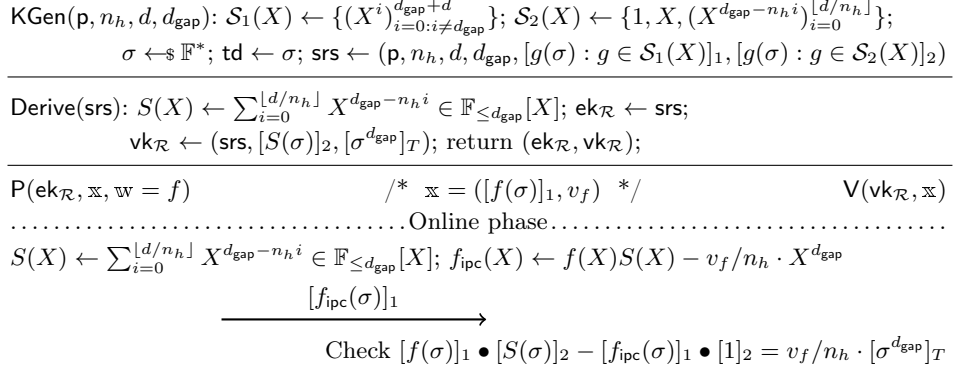
Fig. 1. The new univariate sumcheck zk-SNARK $\mathfrak{Count}$ for $\sum_{\chi\in\mathbb{H}} f(\chi) = v_f$.

Lemma 3. *The sumcheck zk-SNARK $\mathfrak{Count}$ in Eq. (2) is complete and perfectly zero-knowledge. Additionally, the probability that any algebraic $\mathcal{A}$ can break knowledge-soundness is bounded by $\mathsf{Adv}^{\mathrm{pdl}}_{d_1,d_2,\mathsf{Pgen},\mathcal{B}}(\lambda)$, where $\mathcal{B}$ is some PPT adversary, $d_1 = d_{\mathsf{gap}} + d$, and $d_2 = d_{\mathsf{gap}}$.*

Proof. Completeness follows from Lemma 2.

We sketch a knowledge-soundness proof in the AGM [12]. Since $\mathcal{A}$ is algebraic, $f(X), f_{\mathsf{ipc}}(X)$ are in the span of X^i for $i \in \mathcal{S}_1(X)$, i.e., $f, f_{\mathsf{ipc}} \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, X)$. If Eq. (2) holds, then by Lemma 2, the prover is honest. Otherwise, we have a non-zero polynomial $\mathcal{V}(X) := f(X)S(X) - f_{\mathsf{ipc}}(X) - v_f/n_h \cdot X^{d_{\mathsf{gap}}}$ (its coefficients are known since the adversary is algebraic), such that (since the verifier accepts) σ is a root of $\mathcal{V}$. We construct a (d_1, d_2)-PDL adversary $\mathcal{B}$ that gets $(\mathsf{p}, [(\sigma^i)_{i=0}^{d_1}]_1, [(\sigma^i)_{i=0}^{d_2}]_2)$ as an input. $\mathcal{B}$ constructs srs from the challenge input, and runs $\mathcal{A}$ and its extractor $\mathsf{Ext}_{\mathcal{A}}$ to obtain $\mathcal{V}(X)$. Whenever $\mathcal{V}(X) \neq 0$, $\mathcal{B}$ can find the root σ and break the PDL assumption.

We construct a simulator that on input $(\mathsf{srs}, \mathsf{td} = \sigma, ([f(\sigma)]_1, v_f))$ outputs an argument indistinguishable from the real argument. The simulator just computes $[f_{\mathsf{ipc}}(\sigma)]_1$, such that the verification equation holds. That is, $[f_{\mathsf{ipc}}(\sigma)]_1 \leftarrow S(\sigma)[f(\sigma)]_1 - v_f/n_h \cdot \sigma^{d_{\mathsf{gap}}}[1]_1$. Zero-knowledge follows since in the real argument, $[f_{\mathsf{ipc}}(\sigma)]_1$ is computed the same way. $\quad\square$

SRS Verifiability. As noted in Section 2, for both Sub-ZK and updatability, it is required that the SRS is verifiable, i.e., that there exists a PPT algorithm that checks that the SRS belongs to the span of KGen. One can verify $\mathfrak{Count}$'s SRS by checking that $[\sigma]_1 \bullet [1]_2 = [1]_1 \bullet [\sigma]_2$, $[\sigma^i]_1 \bullet [1]_2 = [\sigma^{i-1}]_1 \bullet [\sigma]_2$ for $i \in [1, d_{\mathsf{gap}} + d] \setminus \{d_{\mathsf{gap}}, d_{\mathsf{gap}}+1\}$, $[\sigma^{d_{\mathsf{gap}}+1}]_1 \bullet [1]_2 = [\sigma]_1 \bullet [\sigma^{d_{\mathsf{gap}}}]_2$, $[\sigma^{d_{\mathsf{gap}}-1}]_1 \bullet [\sigma]_2 = [1]_1 \bullet [\sigma^{d_{\mathsf{gap}}}]_2$, and $[\sigma^{n_h i}]_1 \bullet [\sigma^{d_{\mathsf{gap}}-n_h i}]_2 = [1]_1 \bullet [\sigma^{d_{\mathsf{gap}}}]_2$ for $i \in [1, \lfloor d/n_h \rfloor]$. Since, in addition, $\mathfrak{Count}$'s SRS consists of monomial evaluations only, $\mathfrak{Count}$ is updatable.

Efficiency. In $\mathfrak{Count}$, the prover outputs a single group element instead of two in Aurora's univariate sumcheck argument. The latter also requires one to implement a low-degree test, while there is no need for a low-degree test in $\mathfrak{Count}$.

Another important aspect of $\mathfrak{Count}$ is the prover's computation. In Aurora's univariate sumcheck, the prover computes polynomials R and Q, such that $f(X) = v_f/n_h + X R(X) + Q(X) \mathcal{Z}_{\mathbb{H}}(X)$; this can be done in quasilinear number of field operations. On the other hand, since in $\mathfrak{Vampire}$, S only has a small number of non-zero coefficients, the prover of $\mathfrak{Count}$ only executes a linear number of field operations. Both univariate sumchecks however require the prover to use a linear number of $\mathbb{G}_1$ operations. Linear-time *multivariate* sumchecks are well-known, and important in applications.

We emphasize that d_{gap} needs to satisfy $d_{\mathsf{gap}} \geq n_h \cdot \lfloor d/n_h \rfloor$, but it can be bigger. In $\mathfrak{Vampire}$, $d_{\mathsf{gap}} = d$.

As a drawback, $\mathfrak{Count}$'s SRS contains more elements than in Aurora's sumcheck. This is a consequence of using the ILV inner-product commitment scheme.

4 $\mathfrak{Vampire}$: New Updatable And Universal zk-SNARK

In this section, we will use $\mathfrak{Count}$ to construct an efficient updatable and universal zk-SNARK $\mathfrak{Vampire}$ for the sparse R1CSLite constraint system. At a very high level, we use the general approach of Marlin [10], taking into account optimizations of Lunar [9] and Basilisk [26]. On top of already aggressive optimization, we use three novel techniques.

First, Marlin uses Aurora's univariate sumcheck twice. We replace it with $\mathfrak{Count}$ in one of the instantiations. (In another one, the sumcheck is for a low-degree polynomial; thus, we just use Fact 1.) Second, we use a variant of the aggregated polynomial commitment scheme of Boneh et al. [7] to batch the openings of three different polynomials at different points. While Boneh et al. [7] proposed only a randomized batch-opening protocol, we observe that in our case, it can be deterministic. Third, we use a single commitment to commit to left and right inputs of each constraint. All the techniques together remove four group elements from the communication. In the end, $\mathfrak{Vampire}$ is the most communication-efficient updatable and universal zk-SNARK for any NP-complete constraint system. (See Table 1 and the full version [22] for an efficiency comparison.)

4.1   Formulating R1CSLite as Sumcheck

Let $\mathbb{F} = \mathbb{Z}_p$. As in [10,26,9], let $\mathbb{H} = \langle \omega \rangle$ and $\mathbb{K}$ be two multiplicative subgroups of $\mathbb{F}$. We use $\mathbb{H}$ to index the rows (and columns) and $\mathbb{K}$ to index the non-zero elements of specific matrices. From now on, we assume that the R1CSLite instance $\mathcal{I} = (\mathbb{F}, \mathbb{H}, \mathbb{K}, m, m_0, \mathsf{L}, \mathsf{R})$ includes descriptions of $\mathbb{H}$ and $\mathbb{K}$.

We want to demonstrate the satisfiability of $\mathcal{I}$. Recall from Eq. (1) that for this we need to show that $\mathsf{W} \cdot z^* = \mathbf{0}$, where $\mathsf{W} = (\mathsf{I}_{2m} \| - \mathsf{M})$, $\mathsf{M} = \binom{\mathsf{L}}{\mathsf{R}}$, and $z^* = (z_{\mathsf{l}}^\top \| z_{\mathsf{r}}^\top \| (z_{\mathsf{l}} \circ z_{\mathsf{r}})^\top)^\top$, where z_{l} and z_{r} are the vectors of all left and right inputs of all R1CSLite constraints.

Zero-Knowledge. To obtain zero-knowledge, we use a technique motivated by [27]. Let $|\mathbb{H}| = n_h := 2m + b$, for a randomizing parameter $b \in \mathbb{N}$ (to be fixed to $b = 4$ in Theorem 2) that helps us to achieve zero knowledge. We add new random elements to $\boldsymbol{z}^*$ and zero elements to $\mathbf{W}$; the latter are needed not to disturb the knowledge-soundness proof. More precisely, for $\boldsymbol{r}_z \leftarrow_\$ \mathbb{F}^b$, let

$$\boldsymbol{z}_\mathsf{l} := \left(\begin{smallmatrix} 1 \\ \mathsf{x} \\ \boldsymbol{z}_\mathsf{a} \end{smallmatrix}\right) \in \mathbb{F}^m \quad, \quad \boldsymbol{z}_\mathsf{r} := \left(\begin{smallmatrix} \mathbf{1}_{m_0+1} \\ \boldsymbol{z}_\mathsf{b} \end{smallmatrix}\right) \in \mathbb{F}^m \quad, \text{ and } \boldsymbol{z} := \left(\begin{smallmatrix} \boldsymbol{z}_\mathsf{l} \\ \boldsymbol{z}_\mathsf{r} \\ \boldsymbol{r}_z \end{smallmatrix}\right) \quad.$$

Let $\mathbf{I}^b := \left(\begin{smallmatrix} \mathbf{I}_m & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_m & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{smallmatrix}\right)$ and $\mathbf{M}^b := \left(\begin{smallmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{smallmatrix}\right)$ be $n_h \times n_h$ matrices. Let $\boldsymbol{z}' := \left(\begin{smallmatrix} \boldsymbol{z}_\mathsf{r} \\ \mathbf{0}_{n_h-m} \end{smallmatrix}\right)$. Our goal is to show

$$\mathbf{W}^b \cdot \left(\begin{smallmatrix} \boldsymbol{z} \\ \boldsymbol{z} \circ \boldsymbol{z}' \end{smallmatrix}\right) = \mathbf{0} \quad, \tag{3}$$

where $\mathbf{W}^b := \left(\mathbf{I}^b \| - \mathbf{M}^b\right)$. Clearly, Eq. (3) is equivalent to $\mathbf{W} \cdot \boldsymbol{z}^* = \mathbf{0}$.

Next, Eq. (3) holds iff $\mathbf{I}^b \boldsymbol{z} - \mathbf{M}^b (\boldsymbol{z} \circ \boldsymbol{z}') = \mathbf{0}$, i.e.,

$$\forall x \in \mathbb{H}. P[x] := \sum_{y \in \mathbb{H}} \left(\mathbf{I}^b[x,y] - \mathbf{M}^b[x,y] z'[y]\right) z[y] = 0 \quad.$$

Language of Polynomials. Next, we replace vectors with their low-degree encodings, with say $z(Y) := \sum_{\chi \in \mathbb{H}} z[\chi] L_\chi^\mathbb{H}(Y) \in \mathbb{F}_{\leq n_h - 1}[Y]$. Let $\Lambda_\mathbb{H}^b(X, Y)$ and M^b be polynomials, fixed later, that interpolate the matrices $\mathbf{I}^b$ and $\mathbf{M}^b$. That is, $\Lambda_\mathbb{H}^b(x, y) = \mathbf{I}^b[x, y]$ and $M^b(x, y) = \mathbf{M}^b[x, y]$ for $x, y \in \mathbb{H}$. Thus, $\mathbf{I}^b[x,y] z[y] = \Lambda_\mathbb{H}^b(x, y) z(y)$ for any $x, y \in \mathbb{H}$. Moreover, since $z(y\omega^m) = z[y\omega^m] = z'[y]$ for $y \in \{\omega^0, \ldots, \omega^{m-1}\}$, we get $\mathbf{M}^b[x, y] z'[y] z[y] = M^b(x, y) z(y\omega^m) z(y)$. On the other hand, for $x \in \mathbb{H}$ and $y \in \{\omega^m, \ldots, \omega^{n_h-1}\}$, the value of $z[y\omega^m]$ does not matter since we multiply it by $M^b(x, y) = 0$.

Thus, Eq. (3) is equivalent to $\forall x \in \mathbb{H}. P(x) = 0$, where

$$P(X) := \sum_{y \in \mathbb{H}} \psi(X, y), \tag{4}$$

$$\psi(X, Y) := \left(\Lambda_\mathbb{H}^b(X, Y) - M^b(X, Y) z(Y\omega^m)\right) z(Y) \quad. \tag{5}$$

To simplify it further, $\Lambda_\mathbb{H}^b(X, Y)$ and $M^b(X, Y)$ have to satisfy additional conditions that we define in the rest of this subsection.

Interpolating $\mathbf{I}^b$. Following Lunar [9], we interpolate $\mathbf{I}$ with the function

$$\Lambda_\mathbb{H}(X, Y) := \frac{\mathcal{Z}_\mathbb{H}(X) Y - \mathcal{Z}_\mathbb{H}(Y) X}{n_h (X - Y)} \quad. \tag{6}$$

$\Lambda_\mathbb{H}$ satisfies the following properties: (1) $\Lambda_\mathbb{H}(x, y)$ is PPT computable, (2) $\Lambda_\mathbb{H}$ is a polynomial (this follows since $\mathcal{Z}_\mathbb{H}(X) Y - \mathcal{Z}_\mathbb{H}(Y) X = X - Y + XY(X^{n_h-1} - Y^{n_h-1}) = (X - Y)(1 + XY(\sum_{i=0}^{n_h-2} X^{n_h-2-i} Y^i))$ divides by $X - Y$), (3) $\Lambda_\mathbb{H}$ is symmetric, $\Lambda_\mathbb{H}(X, Y) = \Lambda_\mathbb{H}(Y, X)$, (4) $\Lambda_\mathbb{H}(x, y)$ interpolates $\mathbf{I}$ over $\mathbb{H}^2$, i.e., $\forall x, y \in \mathbb{H}. \Lambda_\mathbb{H}(x, y) = \mathbf{I}[x, y]$ (this follows since $\mathcal{Z}_\mathbb{H}(x) y - \mathcal{Z}_\mathbb{H}(y) x = 0$ for all $x \neq y \in \mathbb{H}$ and $1 + XY(\sum_{i=0}^{n_h-2} X^{n_h-2-i} Y^i) = 1 + (n_h - 1) x^{n_h} = 1 + n_h - 1 = n_h$ when $X = Y = x \in \mathbb{H}$), (5) $\Lambda_\mathbb{H}(x, y) = L_x^\mathbb{H}(y)$ for any $x \in \mathbb{H}, y \in \mathbb{F}$. Thus, $\{\Lambda_\mathbb{H}(x, Y)\}_{x \in \mathbb{H}}$ is a basis of $\mathbb{F}_{\leq |\mathbb{H}|-1}[Y]$.

It is natural to define the interpolating polynomial of I^b as

$$\Lambda_{\mathbb{H}}^b(X,Y) := \Lambda_{\mathbb{H}}(X,Y) - \sum_{i=1}^b \ell_{n_h-b+i}^{\mathbb{H}}(X)\ell_{n_h-b+i}^{\mathbb{H}}(Y) \ .$$

Clearly, if b is small, then $\Lambda_{\mathbb{H}}^b(X,Y)$ is efficiently computable. Moreover, $\Lambda_{\mathbb{H}}^b(X,Y)$ is symmetric since $\Lambda_{\mathbb{H}}(X,Y)$ is symmetric.

Interpolating M^b. We use the sparse matrix encoding of M^b from Marlin [10] that keeps track of the matrix's non-zero entries. Let $\mathcal{NZ} := \{(i,j) \in \mathbb{H} \times \mathbb{H} : \mathsf{M}^b[i,j] \neq 0\}$ be the set of indices where M^b is non-zero. Let $\mathbb{K}$ be the minimum-size multiplicative subgroup of $\mathbb{F}$ such that $n_k := |\mathbb{K}| \geq |\mathcal{NZ}|$.[8] We encode M^b by using polynomials row and col to keep track of the indices of its non-zero entries while using a polynomial val for the values of these entries. That is, $\forall \kappa \in \mathbb{K}$, $\mathsf{row}(\kappa) \in \mathbb{H}$ is the row index of the κth element (by using the natural ordering of $\mathbb{H}^2$) of $\mathcal{NZ}$, $\mathsf{col}(\kappa) \in \mathbb{H}$ is the column index of the κth element of $\mathcal{NZ}$, and $\mathsf{val}(\kappa) = \mathsf{M}^b[\mathsf{row}(\kappa),\mathsf{col}(\kappa)] \in \mathbb{F}$ is the corresponding matrix entry. Let

$$\mathsf{row}(Z) := \sum_{\kappa \in \mathbb{K}} \mathsf{row}(\kappa)L_\kappa^{\mathbb{K}}(Z) \in \mathbb{F}_{\leq n_k-1}[Z]$$

be the low-degree extension of the vector $(\mathsf{row}(\kappa))_{\kappa \in \mathbb{K}}$. Let $\mathsf{col}(Z)$ and $\mathsf{val}(Z)$ be the low-degree extensions of $(\mathsf{col}(\kappa))_{\kappa \in \mathbb{K}}$ and $(\mathsf{val}(\kappa))_{\kappa \in \mathbb{K}}$. Let $\mathsf{zcv}(Z)$, $\mathsf{rcv}(Z)$, $\mathsf{zrow}(Z)$, $\mathsf{zcol}(Z)$, $\mathsf{rc}(Z)$, and $\mathsf{zrc}(Z)$ be the low-degree encodings of $Z\mathsf{col}(Z)\mathsf{val}(Z)$, $\mathsf{row}(Z)\mathsf{col}(Z)\mathsf{val}(Z)$, $Z\mathsf{row}(Z)$, $Z\mathsf{col}(Z)$, $\mathsf{row}(Z)\mathsf{col}(Z)$, and $Z\mathsf{row}(Z)\mathsf{col}(Z)$. For example,

$$\mathsf{rcv}(Z) := \sum_{\kappa \in \mathbb{K}} \mathsf{row}(\kappa)\mathsf{col}(\kappa)\mathsf{val}(\kappa)L_\kappa^{\mathbb{K}}(Z) \in \mathbb{F}_{\leq n_k-1}[Z] \ .$$

We define $M^b \in \mathbb{F}[X,Y]$ that interpolates M^b, as the low-degree extension of

$$\forall x,y \in \mathbb{H}.M^b(x,y) := \mathsf{M}^b[x,y] = \sum_{\kappa \in \mathbb{K}} \mathsf{val}(\kappa)\Lambda_{\mathbb{H}}(\mathsf{row}(\kappa),x)\Lambda_{\mathbb{H}}(\mathsf{col}(\kappa),y) \ .$$

Next, $\Lambda_{\mathbb{H}}(\mathsf{row}(\kappa),x) = (\mathcal{Z}_{\mathbb{H}}(\mathsf{row}(\kappa))x - \mathcal{Z}_{\mathbb{H}}(x)\mathsf{row}(\kappa))/(n_h(\mathsf{row}(\kappa) - x))$. Since $\mathcal{Z}_{\mathbb{H}}(\mathsf{row}(\kappa)) = 0$, $\Lambda_{\mathbb{H}}(\mathsf{row}(\kappa),x) = \mathcal{Z}_{\mathbb{H}}(x)\mathsf{row}(\kappa)/(n_h(x - \mathsf{row}(\kappa)))$. Similarly, $\Lambda_{\mathbb{H}}(\mathsf{col}(\kappa),y) = \mathcal{Z}_{\mathbb{H}}(y)\mathsf{col}(\kappa)/(n_h(y - \mathsf{col}(\kappa)))$. Thus,

$$\forall x,y \in \mathbb{H}.M^b(x,y) = \sum_{\kappa \in \mathbb{K}} \mathsf{val}(\kappa) \cdot \frac{\mathcal{Z}_{\mathbb{H}}(x)\mathsf{row}(\kappa)}{n_h(x-\mathsf{row}(\kappa))} \cdot \frac{\mathcal{Z}_{\mathbb{H}}(y)\mathsf{col}(\kappa)}{n_h(y-\mathsf{col}(\kappa))}$$

$$= \frac{\mathcal{Z}_{\mathbb{H}}(x)\mathcal{Z}_{\mathbb{H}}(y)}{n_h^2} \sum_{\kappa \in \mathbb{K}} \frac{\mathsf{rcv}(\kappa)}{(x-\mathsf{row}(\kappa))(y-\mathsf{col}(\kappa))}$$

$$\stackrel{(*)}{=} \frac{\mathcal{Z}_{\mathbb{H}}(x)\mathcal{Z}_{\mathbb{H}}(y)}{n_h^2} \sum_{\kappa \in \mathbb{K}} \frac{\mathsf{rcv}(\kappa)}{xy-x\mathsf{col}(\kappa)-y\mathsf{row}(\kappa)+\mathsf{rc}(\kappa)} \ ,$$

where (*) follows from $\forall \kappa \in \mathbb{K}.\mathsf{rc}(\kappa) = \mathsf{col}(\kappa)\mathsf{row}(\kappa)$. Thus, we define

$$M^b(X,Y) := \frac{\mathcal{Z}_{\mathbb{H}}(X)\mathcal{Z}_{\mathbb{H}}(Y)}{n_h^2} \sum_{\kappa \in \mathbb{K}} \frac{\mathsf{rcv}(\kappa)}{XY-X\mathsf{col}(\kappa)-Y\mathsf{row}(\kappa)+\mathsf{rc}(\kappa)} \ . \tag{7}$$

Since $\deg_X M^b(X,Y) \leq |\mathbb{H}| - 1$, $\forall y \in \mathbb{H}.M^b(X,y) = \sum_{\chi \in \mathbb{H}} M^b(\chi,y)\Lambda_{\mathbb{H}}(\chi,X)$. Clearly, M^b interpolates M^b.

---

[8] $\mathbb{H}$ and $\mathbb{K}$ can be arbitrary subsets of $\mathbb{F}$, but the most efficient algorithms are known when they are multiplicative subgroups. One can assume $\mathbb{K} = \mathbb{H}$ by adding all-zero rows and columns to the matrix, but we generally do not need that $\mathbb{K} = \mathbb{H}$. Keeping $|\mathbb{K}|$ and $|\mathbb{H}|$ flexible allows us to achieve different trade-offs.

Getting to Sumcheck. Next, we show that, under mild conditions on interpolating matrices that the above encodings satisfy, $\forall x \in \mathbb{H}.P(x) = 0$ (and thus also Eq. (3)) is equivalent to $\sum_{y \in \mathbb{H}} \psi(X, y) = 0$.

Lemma 4. *Assume* $\deg_X \Lambda_{\mathbb{H}}(X, Y), \deg_X M^b(X, Y) \leq |\mathbb{H}| - 1$. *Then,* $\forall x \in \mathbb{H}.P(x) = 0$ *iff* $\sum_{y \in \mathbb{H}} \psi(X, y) = 0$.

Proof. ($\Rightarrow$) Assume $\forall x \in \mathbb{H}.P(x) = 0$. Recall from Eq. (5) that $\psi(X, y) = (\Lambda_{\mathbb{H}}^b(X, y) - M^b(X, y)z(y\omega^m))z(y)$. Since $\deg_X \Lambda_{\mathbb{H}}(X, Y), \deg_X M^b(X, Y) \leq |\mathbb{H}| - 1$, then also $\deg_X \psi(X, y) \leq |\mathbb{H}| - 1$. Thus,

$$\sum_{y \in \mathbb{H}} \psi(X, y) = \sum_{y \in \mathbb{H}} \sum_{x \in \mathbb{H}} \psi(x, y)L_x(X) \overset{4}{=} \sum_{x \in \mathbb{H}} P(x)L_x(X) \overset{(*)}{=} 0 \ ,$$

where $(*)$ follows from $\forall x \in \mathbb{H}.P(x) = 0$.
($\Leftarrow$) Let $\sum_{y \in \mathbb{H}} \psi(X, y) = 0$. By Eq. (4), $\forall x \in \mathbb{H}.P(x) = \sum_{y \in \mathbb{H}} \psi(x, y) = 0$. $\square$

To enable efficient verification that the public input was correctly computed, the prover transmits $[\tilde{z}(\sigma)]_1$, for the polynomial $\tilde{z}(Y)$ defined as follows. Let

$$\begin{aligned}
\mathcal{Z}_{\mathsf{inp}}(Y) &:= \prod_{i=1}^{m_0+1}(Y - \omega^{i-1})(Y - \omega^{m+i-1}) \in \mathbb{F}_{\leq 2(m_0+1)}[Y] \ , \\
\mathsf{inp}(Y) &:= \ell_1^{\mathbb{H}}(Y) + \sum_{i=1}^{m_0} \mathbb{x}_i \ell_{i+1}^{\mathbb{H}}(Y) + \sum_{i=1}^{m_0+1} \ell_{m+i}^{\mathbb{H}}(Y) \in \mathbb{F}_{\leq n_h-1}[Y] \ , \\
\tilde{z}(Y) &:= \sum_{i=1}^{m-m_0-1} \boldsymbol{z}_{\mathsf{a}}[i]\frac{\ell_{m_0+1+i}^{\mathbb{H}}(Y)}{\mathcal{Z}_{\mathsf{inp}}(Y)} + \sum_{i=1}^{m-m_0-1} \boldsymbol{z}_{\mathsf{b}}[i]\frac{\ell_{m+m_0+1+i}^{\mathbb{H}}(Y)}{\mathcal{Z}_{\mathsf{inp}}(Y)} + \\
&\quad \sum_{i=1}^{b} \boldsymbol{r}_z[i]\frac{\ell_{2m+i}^{\mathbb{H}}(Y)}{\mathcal{Z}_{\mathsf{inp}}(Y)} \ .
\end{aligned} \quad (8)$$

Since $\ell_i^{\mathbb{H}}(Y) = \prod_{j \neq i}(Y - \omega^{j-1})/(\omega^{i-1} - \omega^{j-1})$, $\tilde{z}(Y) \in \mathbb{F}_{\leq n_h-2m_0-3}[Y]$. Thus, $\mathcal{Z}_{\mathsf{inp}}(Y)\tilde{z}(Y) = \sum_{i=1}^{m-m_0-1} \boldsymbol{z}_{\mathsf{a}}[i]\ell_{m_0+1+i}^{\mathbb{H}}(Y) + \sum_{i=1}^{m-m_0-1} \boldsymbol{z}_{\mathsf{b}}[i]\ell_{m+m_0+1+i}^{\mathbb{H}}(Y) + \sum_{i=1}^{b} \boldsymbol{r}_z[i]\ell_{2m+i}^{\mathbb{H}}(Y)$ interpolates $(\mathbf{0}_{m_0+1}^\top \| \boldsymbol{z}_{\mathsf{a}}^\top \| \mathbf{0}_{m_0+1}^\top \| \boldsymbol{z}_{\mathsf{b}}^\top \| \boldsymbol{r}_z^\top)^\top$. Moreover,

$$z(Y) = \mathcal{Z}_{\mathsf{inp}}(Y)\tilde{z}(Y) + \mathsf{inp}(Y) \in \mathbb{F}_{\leq n_h-1}[Y] \ . \quad (9)$$

Thus, the existence of a polynomial $\tilde{z}(Y)$, such that Eq. (9) holds, guarantees that $z(Y)$ interpolates $(1\|\mathbb{x}^\top\|\boldsymbol{z}_{\mathsf{a}}^\top\|\mathbf{1}_{m_0+1}^\top\|\boldsymbol{z}_{\mathsf{b}}^\top\|\boldsymbol{r}_z^\top)^\top$ for some $\boldsymbol{z}_{\mathsf{a}}$, $\boldsymbol{z}_{\mathsf{b}}$, and $\boldsymbol{r}_z$.

4.2   From Sumcheck to $\mathfrak{Vampire}$

According to the preceding discussion, one can handle R1CSLite by proving that $\sum_{y \in \mathbb{H}} \psi(X, y) = 0$. In the current subsection, we construct an argument for the latter. We replace X with a random α chosen by the verifier, obtaining the polynomial $\psi_\alpha(Y) := \psi(\alpha, Y)$. We use $\mathfrak{Count}$ to show that $\sum_{y \in \mathbb{H}} \psi_\alpha(y) = 0$. For this, as in Section 3, the prover computes the polynomial ψ_{ipc} and the verifier checks $\varphi(Y) := \psi_\alpha(Y)S(Y) - \psi_{\mathsf{ipc}}(Y)$ is a zero polynomial. The latter can be done by KZG-opening all involved polynomials (e.g., $\tilde{z}(Y)$; see Eq. (5)), but this is inefficient. Instead, the prover KZG-opens $\tilde{z}(Y)$ at $Y = \beta\omega^m$ and $\Phi(Y)$, $M^b(\alpha, Y)$ at $Y = \beta$, where (1) Φ is a polynomial defined so that $\Phi(\beta) = \varphi(\beta) = 0$, and (2) one can verify efficiently the correctness, given $v_z \leftarrow \tilde{z}(\beta\omega^m)$ and $v_M \leftarrow$

$M^b(\alpha, \beta)$. This requires us to open a polynomial related to the ILV-opening of $\psi_\alpha(Y)$. We aggregate two KZG-openings by using the technique of Boneh et al. [7]. Finally, we use a univariate sumcheck to check the correctness of v_M; this step is complicated, but it follows closely [10,9]. Importantly, we also show that one of the two commitments from the second sumcheck can be considered an aggregated KZG-opening and thus batched with other KZG-openings.

To simplify some formulas, we assume always $n_h > 3$. This is w.l.o.g., since $n_h = 2m + b$, $m \geq 1$, and $b \geq 2$.

Details. Let $\alpha \leftarrow\!\!\$ \, \mathbb{F} \backslash \mathbb{H}$ be sampled by the verifier. (We explain later why $\alpha \notin \mathbb{H}$.) To test that $\sum_{y \in \mathbb{H}} \psi(X, y) = 0$, define

$$\psi_\alpha(Y) := \psi(\alpha, Y) \in \mathbb{F}_{\leq d}[Y] \ .$$

From Eqs. (5) and (9), we get $\psi_\alpha(Y) = \left(\varLambda_{\mathbb{H}}^b(\alpha, Y) - M^b(\alpha, Y)z(Y\omega^m)\right) \cdot (\mathbb{Z}_{\mathsf{inp}}(Y)\tilde{z}(Y) + \mathsf{inp}(Y))$. Clearly, one can set

$$d := \deg \psi_\alpha = 3(n_h - 1) \ . \tag{10}$$

We use $\mathfrak{Count}$ to prove that $\sum_{y \in \mathbb{H}} \psi_\alpha(y) = 0$. As in Lemma 2, we define

$$S(Y) := \sum_{i=0}^{\lfloor d/n_h \rfloor} Y^{d_{\mathsf{gap}} - n_h i} \in \mathbb{F}_{\leq d_{\mathsf{gap}}}[Y] \ ,$$
$$\psi_{\mathsf{ipc}}(Y) := \psi_\alpha(Y)S(Y) \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, Y) \ . \tag{11}$$

Here, $d_{\mathsf{gap}} \in \mathbb{N}$ is some integer, such that $S(Y)$ and $\psi_{\mathsf{ipc}}(Y)$ are polynomials, i.e., $d_{\mathsf{gap}} \geq n_h \cdot \lfloor d/n_h \rfloor = n_h \cdot \lfloor 3(n_h - 1)/n_h \rfloor = 2n_h$. (This holds for $n_h \geq 3$.) Taking into account later considerations, we set

$$d_{\mathsf{gap}} := 3(n_h - 1) \ . \tag{12}$$

Thus, $S(Y) = Y^{d_{\mathsf{gap}}} + Y^{d_{\mathsf{gap}} - n_h} + Y^{d_{\mathsf{gap}} - 2n_h} = Y^{3n_h - 3} + Y^{2n_h - 3} + Y^{n_h - 3}$.

According to Lemma 2, we need to check that the coefficient of $Y^{d_{\mathsf{gap}}}$ in $\psi_\alpha(Y)S(Y)$ is 0. We do it by checking that
  (i) $\psi_{\mathsf{ipc}}(Y) \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, Y)$, and
  (ii) $\psi_{\mathsf{ipc}}(Y)$ is the correct ILV-opening polynomial, i.e.,

$$\begin{aligned}\varphi(Y) :=& \psi_\alpha(Y)S(Y) - \psi_{\mathsf{ipc}}(Y) \\ =& \left(\varLambda_{\mathbb{H}}^b(\alpha, Y) - M^b(\alpha, Y)z(Y\omega^m)\right)\left(\mathbb{Z}_{\mathsf{inp}}(Y)\tilde{z}(Y) + \mathsf{inp}(Y)\right) \cdot S(Y) - \psi_{\mathsf{ipc}}(Y)\end{aligned}$$

is a zero polynomial.
The prover sends to the verifier KZG-commitments to $\tilde{z}(Y)$ and $\psi_{\mathsf{ipc}}(Y)$. Checking i is free in the pairing-based setting. To check ii, we verify that $\varphi(\beta) = 0$, where $\beta \in C_\beta \subset \mathbb{F} \backslash \mathbb{H}$ is sampled by the verifier. (We will define and motivate C_β later.) More precisely, we verify that $\varphi(\beta) = 0$, where $M^b(\alpha, \beta)$ is substituted by a value v_M computed by the prover. (The latter means that the verifier does not have to compute $M^b(\alpha, \beta)$ itself.) We first describe how to check that $\varphi(\beta) = 0$, assuming v_M is correct. After that, we use another sumcheck instantiation to prove that v_M is correctly computed.

First: checking $\varphi(\beta) = 0$. A straightforward check that $\varphi(\beta) = 0$ requires, on top of sending v_M, the prover to KZG-open $\tilde{z}(Y)$ both at $Y = \beta$ and $Y = \beta\omega^m$ and $\psi_{\mathsf{ipc}}(Y)$ at $Y = \beta$. (The verifier can efficiently evaluate other polynomials like $\Lambda^b_{\mathbb{H}}(X, Y)$, $\mathcal{Z}_{\mathsf{inp}}(Y)$, and $S(Y)$ at $(X, Y) = (\alpha, \beta)$ itself.)

To improve on efficiency, we *implicitly* KZG-commit to Φ, where

$$\Phi(Y) := (\Psi(Y)S(Y) - \psi_{\mathsf{ipc}}(Y))/S(Y) = \Psi(Y) - \psi_\alpha(Y) \in \mathbb{F}_{\leq d}[Y] \ , \text{ and}$$

$$\Psi(Y) := \left(\Lambda^b_{\mathbb{H}}(\alpha, \beta) - v_M \cdot z(\beta\omega^m)\right)\left(\mathcal{Z}_{\mathsf{inp}}(\beta)\tilde{z}(Y) + \mathsf{inp}(\beta)\right) \in \mathbb{F}_{\leq n_h - 2m_0 - 3}[Y] \ . \tag{13}$$

Ψ is obtained from ψ_α by replacing $M^b(\alpha, \beta)$ with v_M and all but one occurrences of Y with β. Φ is a low-degree polynomial satisfying $\Phi(\beta) = \varphi(\beta) = 0$.

We open KZG-commitments to $\tilde{z}(Y)$ at $Y = \beta\omega^m$ (in order to compute $z(\beta\omega^m)$) and $\Phi(Y)$ at $Y = \beta$. For this, the prover sends

$$v_z \leftarrow \tilde{z}(\beta\omega^m) \in \mathbb{F} \ .$$

Since $\Phi(\beta) = 0$, $\Phi(\beta)$ is not transferred. We can open and verify the KZG-commitment to Φ (see Eq. (13)) since we have KZG-commitments to $\tilde{z}$ and ψ_{ipc} (the need for the latter becomes apparent soon), KZG is homomorphic, and the verifier knows all other information present in Φ like $\mathsf{inp}(\beta)$ and v_M. More precisely, the prover batch-opens the two KZG-commitments by computing the KZG-opening polynomials

$$\tilde{z}_{\mathsf{pc}}(Y) := \frac{\tilde{z}(Y) - \tilde{z}(\beta\omega^m)}{Y - \beta\omega^m} \in \mathbb{F}_{\leq n_h - 2m_0 - 4}[Y] \ ,$$

$$\Phi_{\mathsf{pc}}(Y) := \frac{\Phi(Y) - \Phi(\beta)}{Y - \beta} = \frac{\Psi(Y) - \psi_\alpha(Y)}{Y - \beta} \in \mathbb{F}_{\leq d - 1}[Y] \ .$$

Since the prover batches these openings together with one more opening, we will explain the batching process later.

Second (correctness of v_M). We modify a technique from [10,9] by using batching. Recall that M^b satisfies Eq. (7). Moreover, $\deg_X M^b(X, Y)$, $\deg_Y M^b(X, Y) \leq n_h - 1$. Thus, $M^b(\alpha, \beta) = \sum_{\kappa \in \mathbb{K}} T(\kappa) \in \mathbb{F}$, where

$$\mathsf{num}(Z) := \mathcal{Z}_{\mathbb{H}}(\alpha)\mathcal{Z}_{\mathbb{H}}(\beta)/n_h^2 \cdot \mathsf{rcv}(Z) \in \mathbb{F}_{\leq n_k - 1}[Z] \ ,$$

$$\mathsf{den}(Z) := \alpha\beta - \alpha \cdot \mathsf{col}(Z) - \beta \cdot \mathsf{row}(Z) + \mathsf{rc}(Z) \in \mathbb{F}_{\leq n_k - 1}[Z] \ , \tag{14}$$

$$T(Z) := \frac{\mathsf{num}(Z)}{\mathsf{den}(Z)} \in \mathbb{F}(Z) \ .$$

Here, we need $\mathsf{den}(\kappa) = (\alpha - \mathsf{row}(\kappa))(\beta - \mathsf{col}(\kappa)) \neq 0$ for any $\kappa \in \mathbb{K}$. This explains why we chose $\alpha, \beta \notin \mathbb{H}$.

We use a sumcheck to check that $v_M = M^b(\alpha, \beta)$. Since this sumcheck is over a low-degree polynomial, we do not need to use $\mathfrak{Count}$'s full power. Let

$$\widehat{T}(Z) := \sum_{\kappa \in \mathbb{K}} T(\kappa) L^{\mathbb{K}}_\kappa(Z) \in \mathbb{F}_{\leq n_k - 1}[Z] \ .$$

Clearly, $\mathsf{num}(Z) - \widehat{T}(Z)\mathsf{den}(Z) \equiv 0 \pmod{\mathcal{Z}_{\mathbb{K}}(Z)}$. Since $\sum_{\kappa \in \mathbb{K}} \widehat{T}(\kappa) = v_M$, by Fact 1, $\widehat{T}(Z) = ZR(Z) + v_M/n_k$ for

$$R(X) \leftarrow (\widehat{T}(Z) - v_M/n_k)/Z \in \mathbb{F}_{\leq n_k - 2}[Z] \ . \ .$$

Thus, $\mathsf{num}(Z) - (ZR(Z) + v_M/n_k)\mathsf{den}(Z) \equiv 0 \pmod{\mathcal{Z}_{\mathbb{K}}(Z)}$. Since this equality has to hold only when $Z \in \mathbb{K}$, we modify it as follows. Let

$$Q(Z) \leftarrow (\mathsf{num}(Z) - R(Z) \cdot \mathsf{zden}(Z) - v_M/n_k \cdot \mathsf{den}(Z))/\mathcal{Z}_{\mathbb{K}}(Z) \in \mathbb{F}_{\leq n_k - 3}[Z]$$

be such that

$$\mathsf{num}(Z) - R(Z) \cdot \mathsf{zden}(Z) - v_M/n_k \cdot \mathsf{den}(Z) = Q(Z)\mathcal{Z}_{\mathbb{K}}(Z) \ , \ \text{where} \quad (15)$$
$$\mathsf{zden}(Z) := \alpha\beta Z - \alpha\mathsf{zcol}(Z) - \beta\mathsf{zrow}(Z) + \mathsf{zrc}(Z) \in \mathbb{F}_{\leq n_k - 1}[Z] \ . \quad\quad (16)$$

Thus, $\mathsf{zden}(\kappa) = \kappa\mathsf{den}(\kappa)$ for $\kappa \in \mathbb{K}$. This rewriting minimizes the degree of polynomials (e.g., $\mathsf{zcol}(Z) \in \mathbb{F}_{\leq n_k - 1}[Z]$ while $Z\mathsf{col}(Z) \in \mathbb{F}_{\leq n_k}[Z]$).

Marlin and Lunar [10,9] now transferred polynomial commitment to $Q(Z)$. We improve on it, by interpreting Eq. (15) as saying that the polynomial $\mathsf{num}(Z) - R(Z) \cdot \mathsf{zden}(Z) - v_M/n_k \cdot \mathsf{den}(Z)$ opens to 0 at all points $Z \in \mathbb{K}$. Thus, $Q(Z)$ is an aggregated polynomial *opening* of the left-hand side of Eq. (15) at all points of $\mathbb{K}$. Importantly, we can aggregate this opening with the openings $\tilde{z}_{\mathsf{pc}}(Z)$ and $\Phi_{\mathsf{pc}}(Z)$ from before. Hence, we can save an additional one group element. (We will explain batching in a few paragraphs.)

Thus, the prover only commits to R. When we add to $\mathsf{srs}_{\mathcal{R}}$ elements like $[\mathsf{rcv}(\sigma), \mathsf{col}(\sigma)]_2$, the verifier can compute the $\mathbb{G}_2$ elements in the last equation since he knows α and β. Thus, polynomials like $[\mathsf{rcv}(\sigma)]_2$ need to be in $\mathsf{srs}_{\mathcal{R}}$, while monomials, needed for the V to be able to compute $\mathsf{srs}_{\mathcal{R}}$, need to be in srs. This explains the definition of $\mathsf{srs}_{\mathcal{R}}$ in Fig. 2.

One needs to check that $\deg R \leq n_k - 2$. To perform this test without increasing the argument size, we use a second trapdoor $\tau \leftarrow_\$ \mathbb{F}^*$. We add $[(\sigma^i\tau)_{i=0}^{n_k - 2}]_2$ to the SRS and use $[R(\sigma)\tau, Q(\sigma)\tau]_1$ instead of $[R(\sigma), Q(\sigma)]_1$. This modifies the verification equations. The idea is that if the SRS contains $[(\sigma^i)_{i \in \mathcal{S}}, (\sigma^i\tau)_{i \in \mathcal{S}'}]_1$, then a verification $[a]_1 \bullet [1]_2 = [b]_1 \bullet [\tau]_2$ guarantees in the AGM that $a \in \mathsf{span}(\sigma^i\tau)_{i \in \mathcal{S}'}$.

Batching. The prover batches the openings of $\tilde{z}(Y)$ at $Y = \beta\omega^m$, $\Phi(Y)$ at $Y = \beta$, and the left-hand side of Eq. (15) at all $Y \in \mathbb{K}$ as $[B_{\mathsf{pc}}(\sigma, \tau)]_1 \leftarrow [\tilde{z}_{\mathsf{pc}}(\sigma) + \Phi_{\mathsf{pc}}(\sigma) + Q(\sigma)\tau]_1$. Notably, since the polynomial openings are at different locations (β, $\beta\omega^m$, and all points of $\mathbb{K}$, correspondingly), one does not have to randomize this check. (See Section 5.1 for formal proof.) The latter is a general fact, not mentioned in [13,7] and is thus an independent contribution.

Following [13,7], the verifier must check that $[\tilde{z}(\sigma) - v_z]_1 \bullet [(\sigma - \beta)\mathcal{Z}_{\mathbb{K}}(\sigma)]_2 + [\Phi(\sigma)]_1 \bullet [(\sigma - \beta\omega^m)\mathcal{Z}_{\mathbb{K}}(\sigma)]_2 + [\mathsf{num}(\sigma) - R(\sigma)\mathsf{zden}(\sigma) - v_M/n_k \cdot \mathsf{den}(\sigma)]_1 \bullet [(\sigma - \beta)(\sigma - \beta\omega^m)]_2 = [B_{\mathsf{pc}}(\sigma)]_1 \bullet [(\sigma - \beta)(\sigma - \beta\omega^m)\mathcal{Z}_{\mathbb{K}}(\sigma)]_2$, where $[\Phi(\sigma)]_1 = [\Psi(\sigma)]_1 - [\psi_\alpha(\sigma)]_1$. Since the verifier does not know $[\psi_\alpha(\sigma)]_1$ but knows $[\psi_{\mathsf{ipc}}(\sigma)]_1 = [\psi_\alpha(\sigma)S(\sigma)]_1$, we multiply each term of the verification equation by $S(\sigma)$. We also modify the last addend on the left-hand side to allow the prover and the verifier to compute it given the terms given in the SRS. Finally, we use the trapdoor τ because we need to do a low-degree test.

As part of $[B_{\mathsf{pc}}(\sigma)]_1$, the prover has to compute $[\Phi_{\mathsf{pc}}(\sigma)]_1 \leftarrow [(\Phi(\sigma) - \psi_\alpha(\sigma))/(\sigma - \beta)]_1$, where $\Phi_{\mathsf{pc}} \in \mathbb{F}_{\leq d-1}[Y]$ and σ is a trapdoor. For $\mathfrak{Count}$ to

$\mathsf{Pgen}(1^\lambda)$: generate p as usually, assuming that $n_h, n_k \mid (p-1)$ and $3 \nmid n_k$;

$\mathsf{KGen}(\mathsf{p}, n_h, n_k)$: $\mathcal{S}_1(X, X_\tau) = \{(X^i)_{i=0:i\neq d_{\mathsf{gap}}}^{d_{\mathsf{gap}}+d}, (X^i X_\tau)_{i=0}^{n_k-2}\}$;

$\quad \mathcal{S}_2(X) = \{1, X, X^2, X^{n_k}, X^{n_k+1}, (X^{d_{\mathsf{gap}}-jn_h+i})_{j\in\{0,1,2\},i\in[0,n_k+2]}\}$;

$\quad \sigma, \tau \leftarrow_\$ \mathbb{F}^*; \mathsf{td} \leftarrow (\sigma, \tau)$;

$\quad \mathsf{srs} \leftarrow (\mathsf{p}, n_h, n_k, [g(\sigma, \tau) : g \in \mathcal{S}_1(X, X_\tau)]_1, [g(\sigma) : g \in \mathcal{S}_2(X)]_2)$;

$\mathsf{Derive}(\mathsf{srs}, \mathcal{I})$: $\mathsf{ek}_{\mathcal{R}} \leftarrow (\mathsf{p}, \mathcal{I}, [g(\sigma, \tau) : g \in \mathcal{S}_1(X, X_\tau)]_1)$;

$\quad \mathsf{srs}_{\mathcal{R}} \leftarrow \begin{pmatrix} \left[(\sigma^i S(\sigma))_{i=0}^3, \mathsf{rcv}(\sigma)S(\sigma), \mathsf{col}(\sigma)S(\sigma), \mathsf{row}(\sigma)S(\sigma), \mathsf{rc}(\sigma)S(\sigma), \mathcal{Z}_{\mathbb{K}}(\sigma)\right]_2, \\ \left[\sigma\mathcal{Z}_{\mathbb{K}}(\sigma), (\sigma^i \mathsf{zcol}(\sigma)S(\sigma), \sigma^i \mathsf{zrow}(\sigma)S(\sigma), \sigma^i \mathsf{zrc}(\sigma)S(\sigma), \sigma^i \mathcal{Z}_{\mathbb{K}}(\sigma)S(\sigma))_{i=0}^2\right]_2 \end{pmatrix}$;

$\quad \mathsf{vk}_{\mathcal{R}} \leftarrow (\mathsf{p}, \mathcal{I}, [1, \tau, \sigma\tau, \sigma^2\tau]_1, [1]_2, \mathsf{srs}_{\mathcal{R}})$;

Fig. 2. $\mathfrak{Vampire}$'s parameter and SRS generation, where $\mathcal{I} = (\mathbb{F}, \mathbb{H}, \mathbb{K}, m, m_0, \mathbf{L}, \mathbf{R})$.

be secure, the SRS cannot contain $[\sigma^{d_{\mathsf{gap}}}]_1$. Hence, we need to assume $d \leq d_{\mathsf{gap}}$. This motivates the choice of $d_{\mathsf{gap}} = 3(n_h - 1)$ in Eq. (10). The batch opening reduces the communication by two group elements.

4.3   Description of $\mathfrak{Vampire}$

In Figs. 2 and 3, we describe interactive $\mathfrak{Vampire}$, the new succinct interactive zero-knowledge argument with a specializable universal SRS. For the sake of completeness, Figs. 2 and 3 define all used polynomials. Since this argument is public-coin and has a constant number of rounds, we can apply the Fiat-Shamir heuristic (we omit the details) to obtain the zk-SNARK $\mathfrak{Vampire}$.

We sample the challenge β from the set

$$C_\beta = \left\{ \beta \in \mathbb{F} \middle| \begin{array}{c} \beta \notin (\mathbb{H} \cup \mathbb{K} \cup \{0, \sigma, \sigma/\omega^m\}) \wedge \\ S(\beta) \neq 0 \wedge S(\beta\omega^m) \neq 0 \wedge \beta\omega^m \notin \mathbb{K} \end{array} \right\} \ .$$

We need $\beta \notin \{\sigma, \sigma/\omega^m\}$ to get perfect zero-knowledge (see the proof of Theorem 2). One can efficiently verify that $\beta \notin \{\sigma, \sigma/\omega^m\}$, given $[\sigma]_1$ from the SRS. In addition, in the knowledge-soundness proof we need that $S(Y), \mathcal{Z}_{\mathbb{K}}(Y)$, $Y - \beta$, and $Y - \beta\omega^m$ are coprime. Hence, we need that (1) $S(\beta) \neq 0$, $S(\beta\omega^m) \neq 0$, $\mathcal{Z}_{\mathbb{K}}(\beta) \neq 0$, $\mathcal{Z}_{\mathbb{K}}(\beta\omega^m) = 0$ (the latter two conditions hold iff $\beta \notin \mathbb{K}$ and $\beta\omega^m \notin \mathbb{K}$) for coprimeness with $Y - \beta$ and $Y - \beta\omega^m$, and (2) $\beta \neq 0$ (otherwise $\beta = \beta\omega^m$, and thus $Y - \beta$ and $Y - \beta\omega^m$ cannot be coprime). As mentioned previously, $\alpha, \beta \notin \mathbb{H}$ since otherwise $\mathsf{den}(\kappa) = 0$ for any $\kappa \in \mathbb{K}$. Note that if n_h and d_{gap} are much smaller than $|\mathbb{F}|$ (which is typically the case), then $\beta \leftarrow_\$ \mathbb{F}$ is contained in C_β with an overwhelming probability. Thus, in practice, β can be sampled from $\mathbb{F}$, resulting in only a negligible security risk.

Since $\mathcal{S}_1(X, X_\tau)$ and $\mathcal{S}_2(X)$ consist of monomials and one can verify the correctness of its SRS efficiently, $\mathfrak{Vampire}$ is updatable. We will prove the latter in the full version [22]. See the full version [22] for a thorough efficiency analysis.

$\mathsf{P}(\mathsf{ek}_{\mathcal{R}}, \mathbb{x}, \mathbb{w})$ .... Init .... $\mathsf{V}(\mathsf{vk}_{\mathcal{R}}, \mathbb{x})$

$\mathcal{Z}_{\mathsf{inp}}(Y) \leftarrow \prod_{i=1}^{m_0+1}(Y - \omega^{i-1})(Y - \omega^{m+i-1}) \in \mathbb{F}_{\leq 2m_0+2}[Y]$;

$\mathsf{inp}(Y) \leftarrow \ell_1^{\mathbb{H}}(Y) + \sum_{i=1}^{m_0} \mathbb{x}_i \ell_{i+1}^{\mathbb{H}}(Y) + \sum_{i=1}^{m_0+1} \ell_{m+i}^{\mathbb{H}}(Y) \in \mathbb{F}_{\leq n_h-1}[Y]$; $\boldsymbol{r}_z \leftarrow_\$ \mathbb{F}^b$;

$\tilde{z}(Y) \leftarrow \sum_{i=1}^{m-m_0-1} \boldsymbol{z}_{\mathsf{a}}[i] \frac{\ell_{m_0+1+i}^{\mathbb{H}}(Y)}{\mathcal{Z}_{\mathsf{inp}}(Y)} + \sum_{i=1}^{m-m_0-1} \boldsymbol{z}_{\mathsf{b}}[i] \frac{\ell_{m+m_0+1+i}^{\mathbb{H}}(Y)}{\mathcal{Z}_{\mathsf{inp}}(Y)} + \sum_{i=1}^{b} \boldsymbol{r}_z[i] \frac{\ell_{2m+i}^{\mathbb{H}}(Y)}{\mathcal{Z}_{\mathsf{inp}}(Y)}$;

$z(Y) \leftarrow \mathcal{Z}_{\mathsf{inp}}(Y)\tilde{z}(Y) + \mathsf{inp}(Y)$; $\quad/\!/\quad \tilde{z}(Y) \in \mathbb{F}_{\leq n_h-2m_0-3}[Y]$; $z(Y) \in \mathbb{F}_{\leq n_h-1}[Y]$

$$\xrightarrow{\quad [\tilde{z}(\sigma)]_1 \quad}$$
$$\xleftarrow{\qquad \alpha \qquad} \qquad\qquad \alpha \leftarrow_\$ \mathbb{F} \setminus \mathbb{H}$$

Abort if $\alpha \notin \mathbb{F} \setminus \mathbb{H}$

.... $\mathfrak{Count}$: $\sum_{y \in \mathbb{H}} \psi_\alpha(y) = 0$ for $\psi_\alpha(Y) = \left(\Lambda_{\mathbb{H}}^b(\alpha, Y) - M^b(\alpha, Y)z(Y\omega^m)\right) z(Y)$ ....

$S(Y) \leftarrow \sum_{i=0}^{\lfloor d/n_h \rfloor} Y^{d_{\mathsf{gap}}-n_h i} \in \mathbb{F}_{\leq d_{\mathsf{gap}}}[Y]$; $\psi_{\mathsf{ipc}}(Y) \leftarrow \psi_\alpha(Y)S(Y) \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, Y)$;

$$\xrightarrow{\quad [\psi_{\mathsf{ipc}}(\sigma)]_1 \quad}$$
$$\xleftarrow{\qquad \beta \qquad} \qquad\qquad \beta \leftarrow_\$ C_\beta$$

.... Low-degree sumcheck for $\sum_{\kappa \in \mathbb{K}}(\mathsf{num}(\kappa)/\mathsf{den}(\kappa)) = v_M = M^b(\alpha, \beta)$ ....

Abort if $\beta \notin C_\beta$; $v_M \leftarrow M^b(\alpha, \beta) \in \mathbb{F}$; $v_z \leftarrow \tilde{z}(\beta\omega^m) \in \mathbb{F}$;

$\mathsf{num}(Z) \leftarrow \mathcal{Z}_{\mathbb{H}}(\alpha)\mathcal{Z}_{\mathbb{H}}(\beta)/n_h^2 \cdot \mathsf{rcv}(Z) \in \mathbb{F}_{\leq n_k-1}[Z]$;

$\mathsf{den}(Z) \leftarrow \alpha\beta - \alpha \cdot \mathsf{col}(Z) - \beta \cdot \mathsf{row}(Z) + \mathsf{rc}(Z) \in \mathbb{F}_{\leq n_k-1}[Z]$;

$T(Z) \leftarrow \mathsf{num}(Z)/\mathsf{den}(Z) \in \mathbb{F}(Z)$; $\widehat{T}(Z) \leftarrow \sum_{\kappa \in \mathbb{K}} T(\kappa)L_\kappa^{\mathbb{K}}(Z) \in \mathbb{F}_{\leq n_k-1}[Z]$;

$R(Z) \leftarrow (\widehat{T}(Z) - v_M/n_k)/Z \in \mathbb{F}_{\leq n_k-2}[Z]$;

$Q(Z) \leftarrow (\mathsf{num}(Z) - R(Z) \cdot \mathsf{zden}(Z) - v_M/n_k \cdot \mathsf{den}(Z))/\mathcal{Z}_{\mathbb{K}}(Z) \in \mathbb{F}_{\leq n_k-3}[Z]$;

$z(\beta\omega^m) \leftarrow \mathcal{Z}_{\mathsf{inp}}(\beta\omega^m)v_z + \mathsf{inp}(\beta\omega^m)$;

$\Psi(Y) \leftarrow \left(\Lambda_{\mathbb{H}}^b(\alpha, \beta) - v_M \cdot z(\beta\omega^m)\right)\left(\mathcal{Z}_{\mathsf{inp}}(\beta)\tilde{z}(Y) + \mathsf{inp}(\beta)\right) \in \mathbb{F}_{\leq n_h-m_0-3}[Y]$;

$\tilde{z}_{\mathsf{pc}}(Y) \leftarrow (\tilde{z}(Y) - v_z)/(Y - \beta\omega^m) \in \mathbb{F}_{\leq n_h-m_0-4}[Y]$;

$\Phi_{\mathsf{pc}}(Y) \leftarrow (\Psi(Y) - \psi_\alpha(Y))/(Y - \beta) \in \mathbb{F}_{\leq d-1}[Y]$;

$B_{\mathsf{pc}}(Y, X_\tau) \leftarrow \tilde{z}_{\mathsf{pc}}(Y) + \Phi_{\mathsf{pc}}(Y) + Q(Y)X_\tau \in \mathbb{F}_{\leq d-1}[Y] \cup (\mathbb{F}_{\leq n_k-3}[Y])[X_\tau]$;

$$\xrightarrow{\quad v_z, v_M, [R(\sigma)\tau, B_{\mathsf{pc}}(\sigma, \tau)]_1 \quad}$$

$[B_1(\sigma)]_2 \leftarrow [(\sigma - \beta)\mathcal{Z}_{\mathbb{K}}(\sigma)S(\sigma)]_2$; $[B_2(\sigma)]_2 \leftarrow [(\sigma - \beta\omega^m)\mathcal{Z}_{\mathbb{K}}(\sigma)S(\sigma)]_2$;

$[B_3(\sigma)]_2 \leftarrow [(\sigma - \beta\omega^m)\mathcal{Z}_{\mathbb{K}}(\sigma)]_2$; $[B_4(\sigma)\tau]_1 \leftarrow [(\sigma - \beta)(\sigma - \beta\omega^m)\tau]_1$;

$[\zeta_1(\sigma)]_2 \leftarrow \mathcal{Z}_{\mathbb{H}}(\alpha)\mathcal{Z}_{\mathbb{H}}(\beta)/n_h^2 \cdot [\mathsf{rcv}(\sigma)S(\sigma)]_2$; $\quad/\!/\quad = [\mathsf{num}(\sigma)S(\sigma)]_2$

$[\zeta_2(\sigma)]_2 \leftarrow \alpha\beta[S(\sigma)]_2 - \alpha[\mathsf{col}(\sigma)S(\sigma)]_2 - \beta[\mathsf{row}(\sigma)S(\sigma)]_2 + [\mathsf{rc}(\sigma)S(\sigma)]_2$; $\quad/\!/\quad = [\mathsf{den}(\sigma)S(\sigma)]_2$

$[\zeta_3(\sigma)]_2 \leftarrow \alpha\beta[(\sigma - \beta)(\sigma - \beta\omega^m)\sigma S(\sigma)]_2 - \alpha[(\sigma - \beta)(\sigma - \beta\omega^m)\mathsf{zcol}(\sigma)S(\sigma)]_2 -$
$\beta[(\sigma - \beta)(\sigma - \beta\omega^m)\mathsf{zrow}(\sigma)S(\sigma)]_2 + [(\sigma - \beta)(\sigma - \beta\omega^m)\mathsf{zrc}(\sigma)S(\sigma)]_2$;
$\quad/\!/\quad = [(\sigma - \beta)(\sigma - \beta\omega^m)\mathsf{zden}(\sigma)S(\sigma)]_2$

$\mathsf{inp}(\beta\omega^m) \leftarrow \ell_1^{\mathbb{H}}(\beta\omega^m) + \sum_{i=1}^{m_0} \mathbb{x}_i \ell_{i+1}^{\mathbb{H}}(\beta\omega^m) + \sum_{i=1}^{m_0+1} \ell_{m+i}^{\mathbb{H}}(\beta\omega^m)$;

$\mathcal{Z}_{\mathsf{inp}}(\beta\omega^m) \leftarrow \prod_{i=1}^{m_0+1}(\beta\omega^m - \omega^{i-1})(\beta\omega^m - \omega^{m+i-1})$;

$z(\beta\omega^m) \leftarrow \mathcal{Z}_{\mathsf{inp}}(\beta\omega^m)v_z + \mathsf{inp}(\beta\omega^m)$;

$[\Psi(\sigma)]_1 \leftarrow \left(\Lambda_{\mathbb{H}}^b(\alpha, \beta) - v_M \cdot z(\beta\omega^m)\right)\left(\mathcal{Z}_{\mathsf{inp}}(\beta)[\tilde{z}(\sigma)]_1 + \mathsf{inp}(\beta)[1]_1\right)$;

($\sharp\!\sharp$) Check $([\tilde{z}(\sigma) - v_z]_1 \bullet [B_1(\sigma)]_2) + ([\Psi(\sigma)]_1 \bullet [B_2(\sigma)]_2 - [\psi_{\mathsf{ipc}}(\sigma)]_1 \bullet [B_3(\sigma)]_2) +$
$([B_4(\sigma)\tau]_1 \bullet [\zeta_1(\sigma) - v_M/n_k \cdot \zeta_2(\sigma)]_2 - [R(\sigma)\tau]_1 \bullet [\zeta_3(\sigma)]_2)$
$=^? [B_{\mathsf{pc}}(\sigma, \tau)]_1 \bullet [(\sigma - \beta)(\sigma - \beta\omega^m)\mathcal{Z}_{\mathbb{K}}(\sigma)S(\sigma)]_2$;

Fig. 3. $\mathfrak{Vampire}$'s online phase: $\mathcal{I} = (\mathbb{F}, \mathbb{H}, \mathbb{K}, m, m_0, \mathbf{L}, \mathbf{R})$ and $\mathbb{w} = \left(\begin{smallmatrix} \boldsymbol{z}_{\mathsf{a}} \\ \boldsymbol{z}_{\mathsf{b}} \end{smallmatrix}\right) \in \mathbb{F}^{2(m-m_0-1)}$.

5   Security Proofs

We first provide additional preliminaries, needed to prove $\mathfrak{Vampire}$'s security.

Fact 3 (Schwartz-Zippel Lemma) *Let $f(X_1, \ldots, X_c) \neq 0$ be a total degree-d polynomial over a field $\mathbb{F}$ and let $\mathcal{S} \subseteq \mathbb{F}^c$. Then, $\Pr[\boldsymbol{x} \leftarrow_\$ \mathcal{S} : f(\boldsymbol{x}) = 0] \leq d/|\mathcal{S}|$.*

Fact 4 (Bauer et al. [3]) *Let $\mathcal{V}(X_1, \ldots, X_c) \in \mathbb{F}[X_1, \ldots, X_c]$ be a non-zero polynomial of total degree d. Define $\mathcal{P}(Z) \in (\mathbb{F}[S_1, \ldots, S_c, R_1, \ldots, R_c])[Z]$ as $\mathcal{P}(Z) := \mathcal{V}(S_1 Z + R_1, \ldots, S_c Z + R_c)$. Then the coefficient of the leading term in $\mathcal{P}(Z)$ is a polynomial in $\mathbb{F}[S_1, \ldots, S_c]$ of degree d.*

The following lemma (based on [19,13,7]) allows to batch-open several polynomials in distinct points. The prior work [13,7] had a more general version where the points do not have to be distinct; the cost of it is a randomized verification that involves a value $\gamma \leftarrow_\$ \mathbb{F}$ sampled by the verifier. On the other hand, [13,7] did not involve the polynomial $S(Y)$ and worked only with univariate polynomials.

Lemma 5 (Aggregation lemma). *Let $f_i \in \mathbb{F}[Y, X_\tau]$, $T_i \subset \mathbb{F}$ be mutually disjoint sets, and let $T := \cup_i T_i$. Let $S(Y) \in \mathbb{F}[Y]$ be such that $\forall s \in T.S(s) \neq 0$. Fix $v_s \in \mathbb{F}$ for all $s \in T$. Let $\hat{v}_i \in \mathbb{F}[Y]$ be a polynomial, such that $\hat{v}_i(s) = v_s$ for all $s \in T_i$. Let $b_i \in \{0, 1\}$. If there exists a polynomial $B_{\mathsf{pc}} \in \mathbb{F}[Y, X_\tau]$, such that*

$$\sum_i (f_i(Y, X_\tau) - \hat{v}_i(Y)) \mathcal{Z}_{T \setminus T_i}(Y) S(Y)^{b_i} = B_{\mathsf{pc}}(Y, X_\tau) \mathcal{Z}_T(Y) S(Y) \ , \qquad (17)$$

then $\forall i. \forall s \in T_i. f_i(s, X_\tau) = v_s$.

Proof. Since T_i are disjoint and the roots of $S(Y)$ are not in T, Eq. (17) implies that $\forall i. (\mathcal{Z}_{T_i}(Y) \mid (f_i(Y, X_\tau) - \hat{v}_i(Y)))$. The lemma follows.   □

In our use, $\hat{v}_i(Y)$ is either constant or the unique monic polynomial (e.g, Lagrange's polynomial) of degree $|T_i| - 1$, such that $\hat{v}_i(s) = v_s$ for all $s \in T_i$.

Remark 1. When $T_i \cap T_j \neq \emptyset$, $\forall i. (\mathcal{Z}_{T_i}(Y) \mid (f_i(Y, X_\tau) - \hat{v}_i(Y)))$ does not follow. However, if we introduce another variable Z to Eq. (17), changing Eq. (17) to $\sum_i (f_i(Y, X_\tau) - \hat{v}_i(Y)) \mathcal{Z}_{T \setminus T_i}(Y) S(Y)^{b_i} Z^{i-1} = B_{\mathsf{pc}}(Z, Y, X_\tau) \mathcal{Z}_T(Y) S(Y)$, then the claim will again follow. This is essentially how the randomized batching in [7] works (Z is substituted by a random β).

5.1   Knowledge-Soundness Proof

We start by proving two lemmas about coprimeness of some of the polynomials used in $\mathfrak{Vampire}$. We need them later in the knowledge soundness proof.

Lemma 6. *Recall that $\mathcal{Z}_{\mathbb{K}}(Y) = \prod_{\kappa \in \mathbb{K}} (Y - \kappa)$ and $S(Y) = Y^{d_{\mathsf{gap}}} + Y^{d_{\mathsf{gap}} - n_h} + Y^{d_{\mathsf{gap}} - 2n_h}$. If $\mathrm{char}(\mathbb{F}) \neq 3$, $n_h \geq 3$, and $3 \nmid n_k$, then $\gcd(S(Y), \mathcal{Z}_{\mathbb{K}}(Y)) = 1$.*

Proof. Clearly, $\gcd(S(Y), \mathbb{Z}_{\mathbb{K}}(Y)) = 1$ iff $\mathbb{K}$ does not contain roots of $S(Y)$. Since $S(Y) = Y^{3n_h-3} + Y^{2n_h-3} + Y^{n_h-3} = Y^{n_h-3}(Y^{2n_h} + Y^{n_h} + 1)$, roots of $S(Y)$ are 0 (when $n_h > 3$), which is not in $\mathbb{K}$, and roots of $S^*(Y) := Y^{2n_h} + Y^{n_h} + 1$.

Consider the polynomial $P(X) = X^2 + X + 1$ where Y^{n_h} from $S^*(Y)$ is substituted by X. Let a be a root of $P(X)$. Since $a^2 = -a - 1$, we have $a^3 = -a^2 - a = a + 1 - a = 1$. Thus, the order of a divides three. The order cannot be one since then $a^1 = a = 1$, but $P(1) = 1 + 1 + 1 \neq 0$ when $\mathrm{char}(\mathbb{F}) \neq 3$. Thus, the order of a is three. If $a \in \mathbb{K}$, then by Lagrange's theorem, $3 \mid n_k$, violating the assumption $3 \nmid n_k$. Thus, $a \notin \mathbb{K}$. Finally, suppose that b is a root of $Y^{2n_h} + Y^{n_h} + 1$. If $b \in \mathbb{K}$, then $b^{n_h} \in \mathbb{K}$ and $P(b^{n_h}) = 0$. We already showed that $P(X)$ does not have roots in $\mathbb{K}$ and thus, $S(Y)$ does not have roots in $\mathbb{K}$.   □

Lemma 7. *If* $\mathrm{char}(\mathbb{F}) \neq 3$, $n_h \geq 3$, $3 \nmid n_k$, *and* $\beta \in C_\beta$, *then* $S(Y)$, $\mathbb{Z}_{\mathbb{K}}(Y)$, $Y - \beta$, *and* $Y - \beta\omega^m$ *are pair-wise coprime.*

Proof. Let us look at all the pairs one-by-one.
1. We proved in Lemma 6 that $S(Y)$ and $\mathbb{Z}_{\mathbb{K}}(Y)$ are coprime assuming $\mathrm{char}(\mathbb{F}) \neq 3$, $n_h \geq 3$, and $3 \nmid n_k$.
2. Suppose that $\beta = \beta\omega^m$. Then, $\beta(\omega^m - 1) = 0$ and thus either $\beta = 0$ or $\omega^m = 1$. However, $0 \notin C_\beta$. Moreover, $\omega^m \neq 1$ since $m < n_h$. Thus, $Y - \beta$ and $Y - \beta\omega^m$ are coprime.
3. $\mathbb{Z}_{\mathbb{K}}(Y)$ is coprime with $Y - \beta$ and $Y - \beta\omega^m$ since β and $\beta\omega^m$ are not roots of $\mathbb{Z}_{\mathbb{K}}(Y)$ by the definition of C_β.
4. For the same reason, $S(Y)$ is coprime with $Y - \beta$ and $Y - \beta\omega^m$.   □

Theorem 1. *Assume that* $\mathrm{char}(\mathbb{F}) \neq 3$, $n_h \geq 3$ *and* $3 \nmid n_k$. *Then,* $\mathfrak{Vampire}$ *is knowledge-sound in the AGM under the PDL assumption. More precisely, an algebraic* $\mathcal{A}$ *breaks the knowledge-soundness of* $\mathfrak{Vampire}$ *with probability at most*

$$\mathsf{Adv}^{\mathrm{pdl}}_{d_1,d_2,\mathsf{Pgen},\mathcal{B}}(\lambda) \cdot \frac{|\mathbb{F}|^2}{|\mathbb{F}|^2-q} + \frac{16n_h+4m_0-12}{|C_\beta|} + \frac{n_h-1}{|\mathbb{F}|-n_h} \ , \tag{18}$$

where $\mathcal{B}$ *is some PPT adversary,* $d_1 = \max(d_{\mathsf{gap}}+d, n_k-1)$, $d_2 = n_k+d_{\mathsf{gap}}+2$, *and* $q \leq 2 + n_k + d_{\mathsf{gap}} + d_{\max}$ *such that* $d_{\max} = \max(d_{\mathsf{gap}}+d, n_k-1)$.

Proof. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an arbitrary algebraic adversary in the knowledge soundness game and $\mathsf{Ext}_{\mathcal{A}}$ its extractor. In each round, $\mathcal{A}$ sends some elements of either $\mathbb{G}_1$ or $\mathbb{F}$. For the elements of $\mathbb{G}_1$, $\mathsf{Ext}_{\mathcal{A}}$ outputs coefficients of a polynomial where its monomials belong to $\mathcal{S}_1(X, X_\tau)$. We denote polynomials that the adversary sends as $\tilde{z}(Y, X_\tau)$, $\psi_{\mathsf{ipc}}(Y, X_\tau)$, $R(Y, X_\tau)$, and $B_{\mathsf{pc}}(Y, X_\tau)$, where each of the polynomials is in the span of $\mathcal{S}_1(Y, X_\tau)$. We denote the field elements $v_z, v_M \in \mathbb{F}$, sent by the prover, as in the honest protocol description.

In Fig. 4, we depict the knowledge extractor Ext. Ext runs $\mathsf{Ext}_{\mathcal{A}}$[9] to obtain coefficients of $\tilde{z}(Y, X_\tau)$. Ext then evaluates $\tilde{z}(Y, 0) \cdot \mathbb{Z}_{\mathsf{inp}}(Y)$ at points of $Y \in \mathbb{H}$, corresponding to $\boldsymbol{z}_\mathsf{a}$ and $\boldsymbol{z}_\mathsf{b}$ in the honest argument. Ext then returns those vectors. In the rest of this proof, we show that the value outputted by Ext is a valid witness for $\mathbb{x}$ with an overwhelming probability.

---

[9] Even though $\mathcal{A}$ is interactive, since we extract only from the first round message of $\mathcal{A}$, the knowledge soundness extractor is still non-interactive.

$$\boxed{\begin{array}{l} \mathsf{Ext}(\mathsf{srs}, \mathsf{aux}; r) \\ \hline \tilde{z}(Y, X_\tau) \leftarrow \mathsf{Ext}_{\mathcal{A}}(\mathsf{srs}, \mathsf{aux}; r); \\ \boldsymbol{z}_{\mathsf{a}} \leftarrow \big(\tilde{z}(\omega^{m_0+1}, 0) \cdot \mathcal{Z}_{\mathsf{inp}}(\omega^{m_0+1}), \ldots, \tilde{z}(\omega^{m-1}, 0) \cdot \mathcal{Z}_{\mathsf{inp}}(\omega^{m-1}) \big)^\top; \\ \boldsymbol{z}_{\mathsf{b}} \leftarrow \big(\tilde{z}(\omega^{m+m_0+1}, 0) \cdot \mathcal{Z}_{\mathsf{inp}}(\omega^{m+m_0+1}), \ldots, \tilde{z}(\omega^{2m-1}, 0) \cdot \mathcal{Z}_{\mathsf{inp}}(\omega^{2m-1}) \big)^\top; \\ \textbf{return } \mathbb{w} = \left(\begin{smallmatrix} \boldsymbol{z}_{\mathsf{a}} \\ \boldsymbol{z}_{\mathsf{b}} \end{smallmatrix} \right); \end{array}}$$

Fig. 4. The knowledge-soundness extractor Ext for $\mathfrak{Vampire}$ zk-SNARK where $\mathcal{A}$ is an algebraic adversary and $\mathsf{Ext}_{\mathcal{A}}$ its extractor.

We have one verification check that guarantees $\mathcal{V}(\sigma, \tau) = 0$, where

$$\begin{aligned} \mathcal{V}(Y, X_\tau) := &(\tilde{z}(Y, X_\tau) - v_z) \cdot (Y - \beta)\mathcal{Z}_{\mathbb{K}}(Y)S(Y) + \\ &(\Psi(Y, X_\tau)S(Y) - \psi_{\mathsf{ipc}}(Y, X_\tau)) \cdot (Y - \beta\omega^m)\mathcal{Z}_{\mathbb{K}}(Y) + \\ &\Big(\big(\mathsf{num}(Y) - \tfrac{v_M}{n_k}\mathsf{den}(Y) \big) X_\tau - R(Y, X_\tau)\mathsf{zden}(Y) \Big) \cdot (Y - \beta)(Y - \beta\omega^m)S(Y) - \\ &B_{\mathsf{pc}}(Y, X_\tau) \cdot (Y - \beta)(Y - \beta\omega^m)\mathcal{Z}_{\mathbb{K}}(Y)S(Y) \ , \end{aligned}$$

where $\Psi(Y, X_\tau) = \big(\Lambda_{\mathbb{H}}^b(\alpha, \beta) - v_M \cdot z(\beta\omega^m) \big) \big(\mathcal{Z}_{\mathsf{inp}}(\beta)\tilde{z}(Y, X_\tau) + \mathsf{inp}(\beta) \big)$ for $z(\beta\omega^m) = \mathcal{Z}_{\mathsf{inp}}(\beta\omega^m)v_z + \mathsf{inp}(\beta\omega^m)$. ($\mathcal{V}(\sigma, \tau) = 0$ follows from ($\sharp\sharp$) in Fig. 3 when one allows polynomials like $\tilde{z}(Y, X_\tau)$ to be maliciously chosen.)

Clearly, $\Pr[\mathcal{A} \text{ wins}] \leq \Pr[\mathcal{A} \text{ wins} \mid \mathcal{V}(Y, X_\tau) = 0] + \Pr[\mathcal{A} \text{ wins} \mid \mathcal{V}(Y, X_\tau) \neq 0]$. Below, we will analyze both conditional probabilities.

Lemma 8. *Assume* $\mathrm{char}(\mathbb{F}) \neq 3$, $n_h \geq 3$, *and* $3 \nmid n_k$. *For an algebraic* $\mathcal{A}$, $\Pr[\mathcal{A} \text{ wins} \mid \mathcal{V}(Y, X_\tau) = 0] \leq (16n_h + 4m_0 - 12)/|C_\beta| + (n_h - 1)/(|\mathbb{F}| - n_h)$.

Proof. Assume $\mathcal{V}(Y, X_\tau) = 0$. Recall that by Lemma 7, $S(Y), \mathcal{Z}_{\mathbb{K}}(Y), Y - \beta$, and $Y - \beta\omega^m$ are pair-wise coprime. Hence, we can use Lemma 5 with $f_1(Y) = \tilde{z}(Y, X_\tau)S(Y)$, $f_2(Y) = \Psi(Y, X_\tau)S(Y) - \psi_{\mathsf{ipc}}(Y, X_\tau)$, $f_3(Y) = (\mathsf{num}(Y) - v_M/n_k \cdot \mathsf{den}(Y))X_\tau - R(Y, X_\tau)\mathsf{zden}(Y)$, $T_1 = \{\beta\omega^m\}$, $T_2 = \{\beta\}$, $T_3 = \mathbb{K}$, $v_{\beta\omega^m} = v_z$, $v_\beta = 0$, and $v_y = 0$ for $y \in \mathbb{K}$. It follows from $\mathcal{V} = 0$ and Lemma 5 that

$$\tilde{z}(\beta\omega^m, X_\tau) = v_z \ , \tag{19}$$

$$\Psi(\beta, X_\tau)S(\beta) - \psi_{\mathsf{ipc}}(\beta, X_\tau) = 0 \ , \tag{20}$$

$$\forall y \in \mathbb{K}. \big(\mathsf{num}(y) - \tfrac{v_M}{n_k}\mathsf{den}(y) \big) X_\tau - R(y, X_\tau)\mathsf{zden}(y) = 0 \ . \tag{21}$$

We analyze each of the three equations separately.

Equation (19). Denote $\tilde{z}(Y, X_\tau) = \tilde{z}'(Y)X_\tau + \tilde{z}''(Y)$. It follows from Eq. (19) that $\tilde{z}'(\beta\omega^m)X_\tau + \tilde{z}''(\beta\omega^m) = v_z$. Thus, $\tilde{z}''(\beta\omega^m) = v_z$ and $\tilde{z}'(\beta\omega^m) = 0$.

Equation (21). Write $R(Y, X_\tau) = R'(Y)X_\tau + R''(Y)$ and $Q(Y, X_\tau) = Q'(Y)X_\tau + Q''(Y)$. In particular, $\deg R'(Y) \leq n_k - 2$ since the only X_τ-dependent monomials in $\mathcal{S}_1(Y)$ are $(Y^i X_\tau)_{i=0}^{n_k-2}$. Thus, from Eq. (21),

$$\forall y \in \mathbb{K}. \big(\mathsf{num}(y) - \tfrac{v_M}{n_k}\mathsf{den}(y) - R'(Y)\mathsf{zden}(y) \big) X_\tau - R''(y)\mathsf{zden}(y) = 0.$$

Hence, $\forall y \in \mathbb{K}.\mathsf{num}(y) - v_M/n_k \cdot \mathsf{den}(y) - yR'(y)\mathsf{den}(y) = 0$, that is, $\forall y \in \mathbb{K}.T(y) := \mathsf{num}(y)/\mathsf{den}(y) = v_M/n_k + yR'(y)$. Since $\widehat{T}(Z) := \sum_{y \in \mathbb{K}} T(y)L_y^{\mathbb{K}}(Z)$ has degree $\leq n_k - 1$, we get that $\widehat{T}(Z) = ZR'(Z) + v_M/n_k$. By Fact 1,

$$M^b(\alpha, \beta) = \sum_{y \in \mathbb{K}} T(y) = \sum_{y \in \mathbb{K}} \widehat{T}(y) = v_M \; . \tag{22}$$

Equation (20). Denote $\psi_{\mathsf{ipc}}(Y, X_\tau) = \psi'_{\mathsf{ipc}}(Y)X_\tau + \psi''_{\mathsf{ipc}}(Y)$. Observe that $\psi''_{\mathsf{ipc}}(Y) \in \mathsf{PolyPunc}_{\mathbb{F}}(d, d_{\mathsf{gap}}, Y)$. We express $\Psi(Y, X_\tau)$ as

$$\begin{aligned}
\Psi(Y, X_\tau) &= \left(\Lambda_{\mathbb{H}}^b(\alpha, \beta) - v_M \cdot z(\beta\omega^m)\right) \cdot \left(\mathcal{Z}_{\mathsf{inp}}(\beta)\tilde{z}(Y, X_\tau) + \mathsf{inp}(\beta)\right) \\
&= \left(\Lambda_{\mathbb{H}}^b(\alpha, \beta) - v_M \cdot z(\beta\omega^m)\right) \cdot \left(\mathcal{Z}_{\mathsf{inp}}(\beta)(\tilde{z}'(Y)X_\tau + \tilde{z}''(Y)) + \mathsf{inp}(\beta)\right) \\
&= \Psi'(Y)X_\tau + \Psi''(Y) \; ,
\end{aligned}$$

where $\Psi'(Y) := \left(\Lambda_{\mathbb{H}}^b(\alpha, \beta) - v_M \cdot z(\beta\omega^m)\right) \mathcal{Z}_{\mathsf{inp}}(\beta)\tilde{z}'(Y)$ and $\Psi''(Y) = \left(\Lambda_{\mathbb{H}}^b(\alpha, \beta) - v_M \cdot z(\beta\omega^m)\right) \cdot \left(\mathcal{Z}_{\mathsf{inp}}(\beta)\tilde{z}''(Y) + \mathsf{inp}(\beta)\right)$.

Thus, Eq. (20) implies $(\Psi'(\beta)S(\beta) - \psi'_{\mathsf{ipc}}(\beta))X_\tau + \Psi''(\beta)S(\beta) - \psi''_{\mathsf{ipc}}(\beta) = 0$. Hence, $\Psi''(\beta)S(\beta) = \psi''_{\mathsf{ipc}}(\beta)$.

Denote $\psi(Y) := \left(\Lambda_{\mathbb{H}}^b(\alpha, Y) - M^b(\alpha, Y) \cdot z(Y\omega^m)\right) \cdot \left(\mathcal{Z}_{\mathsf{inp}}(Y)\tilde{z}''(Y) + \mathsf{inp}(Y)\right)$. Let $\mathcal{V}_3(Y) := \psi(Y)S(Y) - \psi''_{\mathsf{ipc}}(Y)$. By Eq. (22), $\psi(\beta) = \Psi''(\beta)$ and thus $\mathcal{V}_3(\beta) = 0$. Since $\psi(Y)$ and $\psi_{\mathsf{ipc}}(Y)$ were fixed before the adversary received β, we can apply the Schwartz-Zippel lemma to $\mathcal{V}_3$. Recall that (1) $\deg \tilde{z}'' \leq d_{\mathsf{gap}} + d$, (2) $\deg \mathsf{inp} \leq n_h - 1$, (3) $\deg \mathcal{Z}_{\mathsf{inp}} \leq 2(m_0 + 1)$, (4) $\deg z \leq d_{\mathsf{gap}} + d + 2(m_0 + 1)$, (5) $\deg_Y \Lambda_{\mathbb{H}}(\alpha, Y) \leq n_h - 1$, $\deg_Y M^b(\alpha, Y) \leq n_h - 1$, (6) $\deg \psi''_{\mathsf{ipc}} \leq d_{\mathsf{gap}} + d$, (7) $\deg \psi \leq (n_h - 1) + 2(d_{\mathsf{gap}} + d + 2(m_0 + 1)) = 13n_h + 4m_0 - 9$. Thus, $\deg \mathcal{V}_3 \leq \max(\deg \psi + d_{\mathsf{gap}}, \deg \psi''_{\mathsf{ipc}}) \leq \max(16n_h + 4m_0 - 12, 6(n_h - 1)) = 16n_h + 4m_0 - 12$. If $\mathcal{V}_3(Y) \neq 0$, then the verifier's acceptance implies that $\mathcal{V}_3(\beta) = 0$, which according to Schwartz-Zippel lemma can only happen with probability $(16n_h + 4m_0 - 12)/|C_\beta|$.

Let us consider the case $\mathcal{V}_3(Y) = 0$. Since $\psi(Y)S(Y) = \psi''_{\mathsf{ipc}}(Y)$, $\deg \psi''_{\mathsf{ipc}} \leq d_{\mathsf{gap}} + d$ and $\deg S = d_{\mathsf{gap}}$, then $\deg \psi(Y) \leq d$. Since $Y^{d_{\mathsf{gap}}} \notin \mathcal{S}_1(Y, X_\tau)$, the coefficient of $Y^{d_{\mathsf{gap}}}$ in $\psi(Y)S(Y) = \psi''_{\mathsf{ipc}}(Y) = \sum_{i=0}^{d_{\mathsf{gap}}+d}(\psi''_{\mathsf{ipc}})_i Y^i$ is 0. But this coefficient is $\psi_0 + \psi_{n_h} + \psi_{2n_h} = 0$. Thus, from Lemma 1, it follows that $\sum_{y \in \mathbb{H}} \psi(y) = 0$.

Let us express $\psi(Y)$ as $\psi(X, Y)$, where X corresponds to α. We established that $\sum_{y \in \mathbb{H}} \psi(\alpha, y) = 0$. For any $y \in \mathbb{H}$, $\deg \psi(X, y) = n_h - 1$. If $\sum_{y \in \mathbb{H}} \psi(X, y) \neq 0$, then by the Schwartz-Zippel lemma, $\sum_{y \in \mathbb{H}} \psi(\alpha, y) = 0$ with probability at most $(n_h - 1)/(|\mathbb{F}| - n_h)$. Assume that $\sum_{y \in \mathbb{H}} \psi(X, y) = 0$. By Lemma 4, $\forall x \in \mathbb{H}.P(x) = 0$, where $P(x)$ is as in Eq. (4). In the beginning of Section 4, we established that this equation is equivalent to R1CSLite. Since $z(Y) = \mathcal{Z}_{\mathsf{inp}}(Y)\tilde{z}''(Y) + \mathsf{inp}(Y) = \tilde{z}''(Y)\prod_{i=1}^{m_0+1}(Y - \omega^{i-1})(Y - \omega^{m+i-1}) + \ell_1^{\mathbb{H}}(Y) + \sum_{i=1}^{m_0} \mathbb{x}_i \ell_{i+1}^{\mathbb{H}}(Y) + \sum_{i=1}^{m_0+1} \ell_{m+i}^{\mathbb{H}}(Y)$, then $z(\omega^{i-1})$ for $i \in \{1, \ldots, m_0 + 1\}$ correctly encodes $(1, \mathbb{x}_1, \ldots, \mathbb{x}_{m_0})$. The extractor extracts $z(\omega^{i-1})$ for $i \in \{m_0 + 2, \ldots, m\} \cup \{m + m_0 + 2, \ldots, 2m\}$ which indeed corresponds to the R1CSLite witness. $\square$

Lemma 9. *Let $d_1 := \max(d_{\mathsf{gap}} + d, n_k - 1)$, $d_2 := n_k + d_{\mathsf{gap}} + 2$, and $q \leq 2 + n_k + d_{\mathsf{gap}} + d_{\max}$ for $d_{\max} := \max(d_{\mathsf{gap}} + d, n_k - 1)$. For an algebraic $\mathcal{A}$ and*

$\mathcal{V}(Y, X_\tau)$ *as above, there exists a PPT* $\mathcal{B}$, *such that* $\Pr[\mathcal{A}$ *wins* $\mid \mathcal{V}(Y, X_\tau) \neq 0] \leq$ $\mathsf{Adv}^{\mathrm{pdl}}_{d_1, d_2, \mathsf{Pgen}, \mathcal{B}}(\lambda) \cdot |\mathbb{F}|^2 / (|\mathbb{F}|^2 - q)$.

Proof. The proof is standard and similar to [12]'s proof for Groth16 SNARK. We sketch the main idea. We construct an adversary $\mathcal{B}$ that breaks the (d_1, d_2)-PDL assumption if $\mathcal{A}$ wins in the knowledge soundness game and $\mathcal{V}(Y) \neq 0$.

$\mathcal{B}$ gets as an input $(\mathsf{p}; [(x^i)_{i=0}^{d_1}]_1, [(x^i)_{i=0}^{d_2}]_2)$. $\mathcal{B}$ samples s_1, s_2, r_1, r_2 and defines $\sigma = s_1 x + r_1$ and $\tau = s_2 x + r_2$. Although $\mathcal{B}$ does not know σ or τ (they depend on the challenge x), $\mathcal{B}$ is able to homomorphically compute elements of the form $[\sigma^i]_\iota$ and $[\sigma^i \tau]_\iota$ (e.g., $[\sigma]_1 = s_1[x]_1 + r_1[1]_1$). The degrees d_1 and d_2 are sufficiently high so that $\mathcal{B}$ can compute srs where σ and τ are the trapdoors. Next, $\mathcal{B}$ runs $\mathcal{A}$ and $\mathsf{Ext}_\mathcal{A}$ on this srs to obtain the argument and related argument polynomials. $\mathcal{B}$ now knows coefficients of verification polynomial $\mathcal{V}(Y, X_\tau)$.

When $\mathcal{A}$ wins, $\mathcal{V}(\sigma, \tau) = 0$. Let $\mathcal{P}(X) := \mathcal{V}(S_1 X + R_1, S_2 X + R_2) \in (\mathbb{F}[S_1, S_2, R_1, R_2])[X]$. From Fact 4, if $\mathcal{V}(Y, X_\tau) \neq 0$ has degree q, the coefficient of the maximal degree of $\mathcal{P}(X)$ is some $C(S_1, S_2) \in \mathbb{F}[S_1, S_2]$ of degree q. Thus, the coefficient of the leading term of $\mathcal{P}'(X) := \mathcal{V}(s_1 X + r_1, s_2 X + r_2) \in \mathbb{F}[X]$ is $C(s_1, s_2)$. Since s_1 and s_2 are information-theoretically hidden from $\mathcal{A}$ (they are masked by r_1 and r_2), by the Schwartz-Zippel lemma, $C(s_1, s_2) = 0$ at most with probability $q/|\mathbb{F}|^2$. Thus, with an overwhelming probability, $C(s_1, s_2) \neq 0$ and $\mathcal{P}'(X) \neq 0$. Thus, $\mathcal{B}$ can find the roots of $\mathcal{P}'(X)$. One of the roots must be σ since $\mathcal{P}'(\sigma) = \mathcal{V}(s_1 \sigma + r_1, s_2 \sigma + r_2) = \mathcal{V}(\sigma, \tau) = 0$. Finally, $\mathcal{B}$ outputs σ.

The total degree q of $\mathcal{V}$ is $\leq 2 + d_{\mathsf{gap}} + n_k + d_{\max}$, where $d_{\max} := \max(d_{\mathsf{gap}} + d, n_k - 2)$. Thus, $\Pr[\mathcal{A}$ wins $\mid \mathcal{V}(Y) \neq 0](1 - q/|\mathbb{F}|^2) \leq \mathsf{Adv}^{\mathrm{pdl}}_{d_1, d_2, \mathsf{Pgen}, \mathcal{B}}(\lambda)$. Hence, $\Pr[\mathcal{A}$ wins $\mid \mathcal{V}(Y) \neq 0] \leq \mathsf{Adv}^{\mathrm{pdl}}_{d_1, d_2, \mathsf{Pgen}, \mathcal{B}}(\lambda) \cdot |\mathbb{F}|^2 / (|\mathbb{F}|^2 - q)$. □

It follows from these lemmas that Eq. (18) holds. This proves the claim.   □

5.2   Zero-Knowledge Proof

Theorem 2. *Let* $b = 4$. *Then,* 𝔙𝔞𝔪𝔭𝔦𝔯𝔢 *is perfectly zero-knowledge.*

We prove zero-knowledge (and subversion zero-knowledge) in the full version.

References

1. Abdolmaleki, B., Baghery, K., Lipmaa, H., Zajac, M.: A subversion-resistant SNARK. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part III. LNCS, vol. 10626, pp. 3–33. Springer, Heidelberg (Dec 2017). `https://doi.org/10.1007/978-3-319-70700-6_1`
2. Abdolmaleki, B., Lipmaa, H., Siim, J., Zajac, M.: On subversion-resistant SNARKs. Journal of Cryptology **34**(3), 17 (Jul 2021). `https://doi.org/10.1007/s00145-021-09379-y`

3. Bauer, B., Fuchsbauer, G., Loss, J.: A classification of computational assumptions in the algebraic group model. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 121–151. Springer, Heidelberg (Aug 2020). https://doi.org/10.1007/978-3-030-56880-1_5

4. Bellare, M., Fuchsbauer, G., Scafuro, A.: NIZKs with an untrusted CRS: Security in the face of parameter subversion. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part II. LNCS, vol. 10032, pp. 777–804. Springer, Heidelberg (Dec 2016). https://doi.org/10.1007/978-3-662-53890-6_26

5. Ben-Sasson, E., Chiesa, A., Riabzev, M., Spooner, N., Virza, M., Ward, N.P.: Aurora: Transparent succinct arguments for R1CS. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 103–128. Springer, Heidelberg (May 2019). https://doi.org/10.1007/978-3-030-17653-2_4

6. Bitansky, N., Chiesa, A., Ishai, Y., Ostrovsky, R., Paneth, O.: Succinct non-interactive arguments via linear interactive proofs. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 315–333. Springer, Heidelberg (Mar 2013). https://doi.org/10.1007/978-3-642-36594-2_18

7. Boneh, D., Drake, J., Fisch, B., Gabizon, A.: Efficient polynomial commitment schemes for multiple points and polynomials. Cryptology ePrint Archive, Report 2020/081 (2020), https://eprint.iacr.org/2020/081

8. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 327–357. Springer, Heidelberg (May 2016). https://doi.org/10.1007/978-3-662-49896-5_12

9. Campanelli, M., Faonio, A., Fiore, D., Querol, A., Rodríguez, H.: Lunar: A toolbox for more efficient universal and updatable zkSNARKs and commit-and-prove extensions. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part III. LNCS, vol. 13092, pp. 3–33. Springer, Heidelberg (Dec 2021). https://doi.org/10.1007/978-3-030-92078-4_1

10. Chiesa, A., Hu, Y., Maller, M., Mishra, P., Vesely, N., Ward, N.P.: Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part I. LNCS, vol. 12105, pp. 738–768. Springer, Heidelberg (May 2020). https://doi.org/10.1007/978-3-030-45721-1_26

11. Fuchsbauer, G.: Subversion-zero-knowledge SNARKs. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part I. LNCS, vol. 10769, pp. 315–347. Springer, Heidelberg (Mar 2018). https://doi.org/10.1007/978-3-319-76578-5_11

12. Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 33–62. Springer, Heidelberg (Aug 2018). https://doi.org/10.1007/978-3-319-96881-0_2

13. Gabizon, A., Williamson, Z.J., Ciobotaru, O.: PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953 (2019), https://eprint.iacr.org/2019/953

14. Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 626–645. Springer, Heidelberg (May 2013). https://doi.org/10.1007/978-3-642-38348-9_37

15. Groth, J.: Short pairing-based non-interactive zero-knowledge arguments. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 321–340. Springer, Heidelberg (Dec 2010). https://doi.org/10.1007/978-3-642-17373-8_19

16. Groth, J.: On the size of pairing-based non-interactive arguments. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 305–326. Springer, Heidelberg (May 2016). `https://doi.org/10.1007/978-3-662-49896-5_11`

17. Groth, J., Kohlweiss, M., Maller, M., Meiklejohn, S., Miers, I.: Updatable and universal common reference strings with applications to zk-SNARKs. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part III. LNCS, vol. 10993, pp. 698–728. Springer, Heidelberg (Aug 2018). `https://doi.org/10.1007/978-3-319-96878-0_24`

18. Izabachène, M., Libert, B., Vergnaud, D.: Block-wise P-signatures and non-interactive anonymous credentials with efficient attributes. In: Chen, L. (ed.) 13th IMA International Conference on Cryptography and Coding. LNCS, vol. 7089, pp. 431–450. Springer, Heidelberg (Dec 2011)

19. Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 177–194. Springer, Heidelberg (Dec 2010). `https://doi.org/10.1007/978-3-642-17373-8_11`

20. Libert, B., Yung, M.: Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 499–517. Springer, Heidelberg (Feb 2010). `https://doi.org/10.1007/978-3-642-11799-2_30`

21. Lipmaa, H.: Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 169–189. Springer, Heidelberg (Mar 2012). `https://doi.org/10.1007/978-3-642-28914-9_10`

22. Lipmaa, H., Siim, J., Zajac, M.: Counting vampires: From univariate sumcheck to updatable ZK-SNARK. Cryptology ePrint Archive, Report 2022/406 (2022), `https://eprint.iacr.org/2022/406`

23. Lund, C., Fortnow, L., Karloff, H.J., Nisan, N.: Algebraic methods for interactive proof systems. In: 31st FOCS. pp. 2–10. IEEE Computer Society Press (Oct 1990). `https://doi.org/10.1109/FSCS.1990.89518`

24. Maller, M., Bowe, S., Kohlweiss, M., Meiklejohn, S.: Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019. pp. 2111–2128. ACM Press (Nov 2019). `https://doi.org/10.1145/3319535.3339817`

25. Parno, B., Howell, J., Gentry, C., Raykova, M.: Pinocchio: Nearly practical verifiable computation. In: 2013 IEEE Symposium on Security and Privacy. pp. 238–252. IEEE Computer Society Press (May 2013). `https://doi.org/10.1109/SP.2013.47`

26. Ràfols, C., Zapico, A.: An algebraic framework for universal and updatable SNARKs. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part I. LNCS, vol. 12825, pp. 774–804. Springer, Heidelberg, Virtual Event (Aug 2021). `https://doi.org/10.1007/978-3-030-84242-0_27`

27. Ràfols, C., Zapico, A.: An Algebraic Framework for Universal and Updatable SNARKs. Tech. Rep. 2021/590, IACR (May 5, 2021), `https://ia.cr/2021/590`, last checked modification from August 19, 2021

28. Zhang, Y., Szepeniec, A., Zhang, R., Sun, S.F., Wang, G., Gu, D.: VOProof: Efficient zkSNARKs Generation for Algebra Dummies. In: ACM CCS 2022. pp. ?–? ACM, Los Angeles, USA (Nov 11–15, 2022), accepted