

Testudo: Linear Time Prover SNARKs with Constant Size Proofs and Square Root Size Universal Setup

Matteo Campanelli^{1(⊠)}, Nicolas Gailly², Rosario Gennaro³, Philipp Jovanovic⁴, Mara Mihali⁵, and Justin Thaler⁶

Abstract. We present Testudo, a new FFT-less SNARK with a near linear-time prover, constant-time verifier, constant-size proofs and a square-root-size universal setup. Testudo is based on a variant of Spartan [28]—and hence does not require FFTs—as well as a new, fast multivariate polynomial commitment scheme (PCS) with a square-root-sized trusted setup that is derived from PST [25] and IPPs [9]. To achieve constant-size SNARK proofs in Testudo we then combine our PCS openings proofs recursively with a Groth16 SNARK. We also evaluate Testudo and its building blocks: to compute a PCS opening proof for a polynomial of size 2²⁵, our new scheme opening procedure achieves a 110x speed-up compared to PST and 3x compared to Gemini [6], since opening computations are heavily parallelizable and operate on smaller polynomials. Furthermore, a Testudo proof for a witness of size $2^{30} (\approx 1 \text{ GB})$ requires a setup of size only 2^{15} (\approx tens of kilobytes). Finally, we show that a Testudo variant for proving data-parallel computations is almost 10x faster at verifying 2¹⁰ Poseidon-based Merkle tree opening proofs than the regular version.

1 Introduction

Succinct Non-Interactive Arguments of Knowledge (SNARKs) have been a prolific area of research in the last decade: a SNARK allows a prover to prove to a verifier that a certain (non-deterministic) computation F has been performed correctly, or more specifically that there exists a witness w such that y = F(x, w)

N. Gailly—Work done mainly while the author was affiliated with Protocol Labs. M. Mihali—Work done mainly while the author was affiliated with UCL and Protocol Labs.

[©] The Author(s), under exclusive license to Springer Nature Switzerland AG 2023 A. Aly and M. Tibouchi (Eds.): LATINCRYPT 2023, LNCS 14168, pp. 331–351, 2023. https://doi.org/10.1007/978-3-031-44469-2_17

where x is a public input. The crucial property of SNARKs is that the size of the proof and the verification time should be *short*, i.e., sublinear in the size of the computation F and of the witness w. Otherwise a simple proof would be to send w and have the verifier recompute F(x, w). Additionally, SNARKs can be zero-knowledge, i.e., they do not reveal any information about w.

SNARKs are evaluated according to various performance metrics where the three most important ones are (1) the time it takes the prover to generate a proof, (2) the size of the proof, and (3) the time it takes the verifier to validate a proof. There are various design trade-offs that can be explored to optimize those metrics and an essential distinguishing factor that impacts performance are the preprocessing phases of SNARK systems.

On the one end of the spectrum, there is, for example, the Groth16 SNARK [18] which produces proofs consisting of only 3 group elements and which has a verification that is independent of the complexity of F, only requiring the evaluation of a few pairings. This makes Groth16 essentially optimal in terms of proof size and verifier time but it comes at the expense of a superlinear overhead in prover time and a function-specific trusted setup, the latter of which is a serious drawback in practice. On the other end are (transparent) SNARKs that do not require a trusted setup [5,8,32] at all but they tend to have larger proofs and verifier times, e.g., at least logarithmic in the size of the witness. Finally, there are (universal) SNARKS which offer a compromise, as they permit smaller proofs in comparison to transparent systems at the cost of a single setup that is universal enabling them to prove any circuit up to a certain size [19]. However, these schemes are still suboptimal as they have slower provers than Groth16 [12], which becomes particularly evident for large circuits [34], and they do not have constant size proofs or verification times. Finally, all of the schemes with a trusted setup produce a linear-size common reference string (CRS) which is particularly problematic for large circuits (again), as the CRS has to be downloaded, stored, and moved into RAM at proving time. For example, Filecoin [21] uses Groth16 for circuits of size $\approx 2^{30}$ but producing a trusted setup for this size was practically infeasible. As a workaround, provers work with (sub-)circuits of size $\approx 2^{27}$ and generate ≈ 10 proofs per (large) circuit.

In summary, all of these observations led us to the following research question:

Can we design a SNARK with a small universal trusted setup, constant size proofs and verification time, and a fast prover?¹

1.1 Contributions

In this section we present our main contributions together with an informal overview of the techniques used to achieve them.

Testudo: Near-Linear Time Prover with Succinct Verification. To achieve this goal, we start our design from Spartan [28] a sumcheck-based argument requiring only field arithmetic (e.g. much faster that its point counterpart) and

¹ To maximize backward compatibility to already deployed systems, we require that our SNARK system works with R1CS-based circuits.

a single multilinear polynomial opening. The original Spartan is a transparent scheme which relies only on discrete based log curves and transparent polynomial commitment schemes, giving substantially larger proof sizes and verification times ($\approx \sqrt{N}$). To improve on those, the main idea is to have a Groth16 prover verifying the Spartan proof. Below we describe the technical challenges of this approach, but one important thing to note here is that in order to obtain a ZK-SNARK it is sufficient to run the (much simpler and more efficient) non-ZK version of Spartan since the outer Groth16 proof will hide any information possibly leaked by the non-ZK inner Spartan proof.

Embedding the verification of a Spartan-based SNARK in a Groth16 circuit presents various implementation challenges. The sumcheck component of the proof operates only on field elements and requires the use of a hash function to make it non-interactive via Fiat Shamir. Field operations can be natively encoded in R1CS constraints and we adapted our codebase to use Poseidon, a SNARK friendly hash function, having the advantage of a more efficient representation in a circuit. However, things get more complicated for writing the R1CS constraints for the verification of the polynomial commitment opening, since it requires point arithmetic which, if naively encoded in the circuit, would massively increase the number of constraints by multiple orders of magnitude. This issue would be amplified by the relative large size of Spartan's original commitment's proof size and verification $-O(\sqrt{N})$ – which can potentially be reduced if we leverage using a different PCS with a trusted setup.

WHY THE NAME TESTUDO? Testudo was a type of battle formation that ancient Rome adopted, where its soldiers operated "under the hood" of their shields. Testudo, the proof scheme, is similar: a Spartan prover working under the hood of Groth16.

Testudo-Comm: New Multivariate Polynomial Commitment Scheme. To reduce the size of the PCS opening, we devise a new polynomial commitment scheme based on PST [25] and inner pairing product [9] that avoids the use of FFTs. Testudo, as Spartan, considers circuits of size $N=2^n$ where the polynomial representation of the circuit has n variables, n is logarithmic and so N will be linear in the size of the circuit. The high level idea is to express the coefficients of the witness multivariate polynomial from Spartan as a square matrix of size $\sqrt{N} \times \sqrt{N}$. To commit to this polynomial, the prover commits to each row of the matrix using PST, leading to a vector \vec{A} of size \sqrt{N} . Then the prover commits to \vec{A} using the MIPP commitment in [9] (e.g., a pairing product between commitments and random base) to create the final commitment T - a single group element. To open, the prover carefully performs PST and MIPP opening on \sqrt{N} sized polynomials with many operations in parallel. Both the MIPP and PST part operate on \sqrt{N} sized polynomials.

Since the 2 opening operations can be done in parallel we obtain a considerable speedup in practice (about 2 orders of magnitude faster) than PST, even though it requires heavier operations like $\log N$ pairings to create the combined commitment. Moreover, when comparing with Gemini [6], we estimate our opening procedure to be 3x faster for large N such as 2^{25} , where [6] is likely faster

for small sizes, with the additional cost that it requires to perform FFTs for practical deployment.

Usage of 2-Chain Curves for Efficient Verification. To enable verifying group operations in a circuit, we use the standard approach of running Spartan over a 2-chain curve (such as BLS12-377) where group elements can be encoded as field elements in a companion curve over which the Groth16 prover is then implemented. For backwards compatibility reason, we also explored the possibility to run Testudo on curves without a companion curve (such as BLS12-381, which is currently used by Filecoin proofs).

Aggregation. Additionally, we show how Testudo proofs can be aggregated. Since the outer layer is a Groth16 proof, we can use standard aggregation tools such as SnarkPack [16] to aggregate several Groth16 proofs together. Another option is to aggregate proofs at the inner level and then run a single Groth16 proof on top of the aggregated inner proofs.

Note that an interesting point of this design is that the Groth16 prover can be outsourced to more powerful machines or a "prover-as-a-service" infrastructure. Indeed, if the "Spartan" prover ran the sumcheck and the commitment opening, then resulting proof already hides the witness thanks to the zero knowledge property of Spartan. As a consequence, more powerful machine can aggregate many of these "semi" proofs inside one Groth16 proof, without ever seeing the witness, similar to how Snarkpack works. This settings has practical implications in terms of deployment that we believe are worth exploring.

Analysis and Experimental Results. As explained in the body of the paper, our SNARK avoids the use of FFT altogether and obtains a nearly linear-time prover². In practice, we show that:

- Our polynomial commitment scheme has a commitment time comparable to PST while producing opening proofs at two order of magnitude faster (at the cost of larger proof sizes).
- Our experimental results show that for data parallel circuits, we can estimate Testudo to run more than $\approx 5x$ to $\approx 10x$ faster than the fastest Groth16 implementation (i.e., Bellperson [4]), depending on the size of the small subcircuit. For example, if the sub-circuit is of size 2^{15} , as an upper bound to a circuit verifying a Poseidon based Merkle Tree opening proofs with 32 layers, then Testudo can verify 2^{10} such proofs $\approx 9.7x$ faster than the Groth16 equivalent.

1.2 Related Work

The literature on SNARKs is very large and we refer the reader to Thaler's monograph [30] for a comprehensive survey. In this section we focus on a few works that are relevant to Testudo.

Our prover runs N multi-exponentiations of size N, which is roughly $O(N \frac{\lambda}{\log N})$ group operations with $\lambda > \log N$ for security reason.

We were inspired to use a 1-level recursion with Groth16 verifying a faster inner SNARK by the work of Belling et al. [3] where the verification of a GKR proof [17] for hash computations is outsourced to a Groth16 prover. Concurrently to our efforts, a similar approach was also taken in the ZKBridge paper [31], where the verification of a Virgo [32] proof is outsourced to a Groth16 prover. Because Virgo is also GKR-based, the underlying SNARK in either case is known to be efficient for large "parallel" computations. We believe we are the first to apply this approach to a general purpose SNARK like Spartan.

When it comes to universal trusted setup proofs, many systems today do not use R1CS but rather "custom gates" (sometimes also called Plonkish arithmetization), and apply SNARKs such as Plonk [15] (or alternatives such as Hyperplonk [11]) to the resulting constraint systems. The use of "custom gates" makes a comparison to pure R1CS-based schemes not immediate. We are still working on achieving meaningful comparisons but we estimate that Testudo is competitive with approaches that do use custom gates. We point out that many applications (including our main motivating one – Filecoin proofs) are already encoded as R1CS systems, and therefore it is very useful to have an efficient SNARK with universal trusted setup that can be used off-the-shelves.

Our new PCS Testudo-Comm leverages ideas from [16] and [9] to reduce the size of the trusted setup for the KZG univariate polynomial commitment [20] to square-root size from linear. We adapted them to achieve the same reduction for the PST commitment. We note that, as far as we know, we are the first to implement these techniques. We also point out that the reduction of the trusted setup size comes at the expense of larger opening proofs: however in our case that drawback is "absorbed" by the outer Groth16 proof, which compresses the final proof down to constant.

The work in [6] presents a generic transformation to turn a univariate polynomial commitment into a multilinear one. In Sect. 4 we discuss why we believe using Testudo-Comm is a better choice for us.

2 Preliminaries

We assume the reader is familiar with the definitions of R1CS, Polynomial Commitment Schemes and SNARKs.

2.1 Notation

We assume we have cyclic groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ of order q generated by g and equipped with a bilinear pairing $e: \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$. We denote by $p(x_1, \ldots, x_n)$ a multilinear polynomial with n variables. For $s_1, \ldots, s_n \in \mathbb{Z}_q$ we write $\vec{s} = (s_1, \ldots, s_n) \in \mathbb{Z}_q^n$. Let $i \in \{0, 1\}^n$, we can denote $i = (i_1 \ldots i_n)$ as $i_j \in \{0, 1\}$. We denote the value $\prod_j s_j^{i_j}$ by \vec{s}^i

2.2 Cryptographic Assumptions

The security of our constructions holds in the Generic Group Model (GGM) [29]. In Sect. 4 and Sect. 5 we rely on the security of the underlying building to claim that of our protocols. The security of these building blocks can be argued from assumptions implied by the GGM. In more detail: • for PST we require the $(\mu + 1)\delta$ -Strong Diffie-Hellman and the (δ, μ) -Extended Power Knowledge of Exponent assumption (see [33] and discussion in [10, E.1]). • for MIPP we require a variant of the (q, m)-Auxiliary Structured Single Group Pairing (see [9]).

2.3 PST Polynomial Commitments

We refer the reader to Sect. 2.1 for the notation we use in this section. In Fig. 1 we describe the PST polynomial commitment modified to work over the Lagrange basis [25].

Trusted setup: Let t_1,\ldots,t_n be random values in \mathbb{Z}_q . The CRS consists of the N values $g_1^{\chi_i(t)}$ where $\chi_i(X_1,\ldots,X_n)=[\Pi_{j:i_j=1}X_j][\Pi_{j:i_j=0}(1-X_j)]$ for $i\in\{0,1\}^n$ the multilinear Lagrange polynomial. The CRS also includes the values $g_2^{t_j}$ for $j=1,2,\ldots,n$.

Commitment: $Com(p) = g_1^{p(\bar{t})} = C$ which can be evaluated given the values of p on the Boolean hypercube and the CRS.

Opening: To prove that $y = p(\vec{a})$ where $\vec{a} = (a_1, \dots, a_n)$ the prover computes the polynomials $q_i(\vec{x})$ such that

$$p(\vec{x}) - y = \Sigma_i(x_i - a_i)q_i(\vec{x})$$

by repeated polynomial division (note that in reality $q_i(\vec{x}) = q_i(x_i, \dots, x_n)$, i.e., only of the last n-i variables). The proof is then the vector $\vec{w} = (w_1, \dots, w_n)$ where $w_i = g_1^{q_i(\vec{t})}$ which can be computed given the polynomials q_i and the CRS.

Verification: To verify that $y = p(\vec{w})$ given \vec{w} , the verifier checks that

$$e(Cg^{-y},g) = \Pi_i e(w_i, g_2^{t_i - a_i})$$
.

Fig. 1. The PST commitment scheme in the Lagrange basis.

Note that if $n = \log N$ where N is the size of the R1CS, then the trusted setup is linear in the size of the circuit, and that verification of the opening requires O(n) (i.e., logarithmic in the size of the circuit) work.

```
P \text{ sends the polynomial } p_1(x) = \sum_{i \in \{0,1\}^{n-1}} p(x,i).
V \text{ checks that } p_1(0) + p_1(1) = a \text{ and sends back } r_1 \in_R \mathbb{F}.
P \text{ sends the polynomial } p_2(x) = \sum_{i \in \{0,1\}^{n-2}} p(r_1,x,i).
V \text{ checks that } p_2(0) + p_2(1) = p_1(r_1) \text{ and sends back } r_2 \in_R \mathbb{F}.
\dots
At \text{ round } j \text{ } P \text{ sends the polynomial } p_j(x) = \sum_{i \in \{0,1\}^{n-j}} p(r_1,\dots,r_{j-1},x,i).
V \text{ checks that } p_j(0) + p_j(1) = p_{j-1}(r_{j-1}) \text{ and sends back } r_j \in_R \mathbb{F}.
\dots
At \text{ the last round } P \text{ sends the polynomial } p_{n-1}(x) = p(r_1,\dots,r_{n-1},x).
V \text{ checks that } p_{n-1}(0) + p_{n-1}(1) = p_{n-2}(r_{n-2}), \text{ selects } r_n \in_R \mathbb{F} \text{ and checks that } p_{n-1}(r_n) = p(r_1,\dots,r_n) \text{ via a single query to } p.
```

Fig. 2. The Sumcheck Protocol

2.4 Sumcheck

Let $p(x_1, \ldots, x_n)$ be a multilinear³ polynomial in n variables defined over a field \mathbb{F} . Consider the value $a = \sum_{i \in \{0,1\}^n} p(i)$, i.e., the sum of the value of p on all the vertices of the Boolean hypercube. This computation takes $N = 2^n$ time and the sumcheck protocol [22] described in Fig. 2, is a way for a Prover to convince a Verifier that a is correct in O(n) time, plus a *single* query to the polynomial p on a random point in \mathbb{F}^n .

2.5 Spartan Overview

In this section we review Spartan [28], a transparent SNARK for R1CS. For space reasons, ours is a very high level review and the reader is referred to [28] for details.

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

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

where z = (x, 1, w), · is the matrix-vector product, and o is the Hadamard (entry-wise) product.

The first step in Spartan is to encode the matrices A, B, C and the vector z via their multilinear polynomial extensions. Let $n = \log N$. For the matrix A consider the unique multilinear polynomial in 2n variable $\tilde{A}(t_1, \ldots, t_n, u_1, \ldots, u_n)$ such that $\tilde{A}(i_1, \ldots, i_n, j_1, \ldots, j_n) = A(i, j)$ where (i_1, \ldots, i_n) is the binary expansion of i and (j_1, \ldots, j_n) is the binary expansion of j. The polynomials \tilde{B}, \tilde{C} are defined similarly, as well as the polynomial $\tilde{Z}(u_1, \ldots, u_n)$ where $Z(i_1, \ldots, i_n) = z(i)$.

³ We only care about multilinear polynomials for Testudo but the sumcheck protocol can be run on any multivariate polynomial.

The satisfiability condition is then equivalent to the following polynomial $F(t_1, \ldots, t_n)$ being zero on all the points of the Boolean hypercube

$$F(\vec{t}) = \left(\sum_{\vec{u} \in \{0,1\}^n} \tilde{A}(\vec{t}, \vec{u}) \tilde{Z}(\vec{u})\right) \cdot \left(\sum_{\vec{u} \in \{0,1\}^n} \tilde{B}(\vec{t}, \vec{u}) \tilde{Z}(\vec{u})\right) - \sum_{\vec{u} \in \{0,1\}^n} \tilde{C}(\vec{t}, \vec{u}) \tilde{Z}(\vec{u})$$

Consider now the multilinear extension⁴ of $F(\cdot)$, that is the polynomial $Q(\vec{s}) = \sum_{\vec{t} \in \{0,1\}^n} F(\vec{t}) eq(\vec{t}, \vec{s})$ where $eq(\vec{t}, \vec{s}) = \prod_{i=1}^n s_i t_i + (1 - s_i)(1 - t_i)$ is the multilinear polynomial which is equal to 1 if and only if $\vec{t} = \vec{s}$ and otherwise is equal to 0.

Since $F(\vec{t})$ is zero on the Boolean hypercube, $Q(\vec{s})$ is then identical to the zero polynomial by Schwartz-Zippel lemma. This condition can be verified by testing $Q(\vec{s})$ on a random point. Spartan is a way to check this evaluation in an efficient way. More precisely, to verify the satisfiability of the original R1CS Spartan performs the following steps:

- 1. Proves that $Q(\vec{r}) = 0$ for a random point $\vec{r} \in \mathbb{F}^n$. Note that due to the definition of $Q(\cdot)$ this can be done via a sumcheck protocol.
- 2. The above sumcheck protocol reduces to proving that $\sigma = F(\vec{\rho})$ for a random $\vec{\rho} \in \mathbb{F}^n$. Due to the definition of F this reduces to proving the value of three summations $\sum_{\vec{u} \in \{0,1\}^n} \tilde{A}(\vec{\rho}, \vec{u}) \tilde{Z}(\vec{u}), \sum_{\vec{u} \in \{0,1\}^n} \tilde{B}(\vec{t}, \vec{u}) \tilde{Z}(\vec{u})$, and $\sum_{\vec{u} \in \{0,1\}^n} \tilde{C}(\vec{t}, \vec{u}) \tilde{Z}(\vec{u})$. Each one of them can also be proven via a sumcheck, and in Spartan these 3 sumchecks are aggregated into a single one.
- 3. Finally the above sumchecks reduce to proving the values of the multilinear extensions on random points, i.e., the values of $\tilde{A}(\vec{r_x}, \vec{r_y})$, $\tilde{B}(\vec{r_x}, \vec{r_y})$, $\tilde{C}(\vec{r_x}, \vec{r_y})$, and $\tilde{Z}(\vec{r_y})$.

The final point is achieved via the use of polynomial commitments. The prover commits to the polynomials $\tilde{A}, \tilde{B}, \tilde{C}$ (these are called computation commitments since they encode the computation), and \tilde{Z} (witness commitment, since it encodes the witness).

A major contribution of Spartan is to show how to efficiently commit to $\tilde{A}, \tilde{B}, \tilde{C}$ to leverage their sparseness (recall that in R1CS matrices have N^2 entries but only m are non-zero). This requires a non-trivial use of memory checking techniques, and introduces a substantial overhead which can be avoided in practice for uniform circuits where the Verifier can evaluate $\tilde{A}, \tilde{B}, \tilde{C}$ on their own.

Spartan's focus was to obtain a transparent SNARK, and therefore it uses a multidimensional Pedersen's commitment together with an inner product proof to implement the polynomial commitment. Because we are already using a trusted setup for the Groth16 layer, we changed the polynomial commitment to a different one which also has a trusted setup.

⁴ Such a polynomial of degree at most 1 in each variable always exists for any function f mapping $\{0,1\} \to \mathbb{F}$ [30].

3 A Generalized MIPP Protocol

In order to obtain a multilinear PCS with $O(\sqrt{N})$ trusted setup, in this section we show how to adapt ideas from Section 6 of [9] which were applied to the KZG univariate PCS. We generalize it to work with multivariate polynomials and the PST commitment.

Changes from Original MIPP. The protocol in this section has two changes when compared to the one in [9] (we are referring specifically to $MIPP_k$). First, we generalize MIPP to work on multivariate rather than univariate polynomials. Second, we show that the techniques also work when the polynomial is represented in the Lagrangian basis.

The Generalized MIPP Protocol: Given a vector $\vec{A} = [A_1, \dots, A_M]$ of group elements in \mathbb{G}_1 , the IPP commitment to \vec{A} with a CRS $\vec{h} = [h_1, \dots, h_M]$ of group elements in \mathbb{G}_2 is

$$T = CM(\vec{A}, \vec{h}) = \prod_{i=1}^{M} e(A_i, h_i).$$

Let $m = \log M$. In our case, $h_i = h^{\chi_i(\vec{t})}$ for $i \in \{0, 1\}^m$, where $\vec{t} = [t_1, \dots, t_m]$ is a random secret vector of field elements and h is a generator of \mathbb{G}_2 .

Our generalized MIPP protocol allows a prover to prove that given a public vector of field elements $\vec{b} = [b_1, \dots, b_m]$, we have that

$$U = \langle \vec{A}, \vec{y} \rangle = \vec{A}^{\vec{y}} = \prod_{i=1}^{M} A_i^{y_i},$$

where the vector \vec{y} is defined as $\vec{y} = [y_1, \dots, y_M]$ with $y_i = \chi_i(\vec{b})$ for $i \in \{0, 1\}^m$, where $\chi_i(X)$ is the i^{th} Lagrange polynomial defined as

$$\chi_i(X_1,\ldots,X_m) = \prod_{j:i_j=1} X_j \cdot \prod_{j:i_j=0} (1-X_j).$$

This proof has size and verification time O(m), which means that the verifier needs only to read the vector \vec{b} and not construct the entire vector \vec{y} , which is only implicitly defined.

The protocol is described in Fig. 3.

Note that there will be m levels of recursion. Also note that the Verifier cannot compute the vectors $\vec{A'}, \vec{y'}, \vec{h'}$ since they are too big. Only the prover will compute those and provide the final value at the end of the recursion to the Verifier. We show later how the Verifier can check that they are correct. The Verifier can compute T', U'.

Properties of the Construction. We make the following claims about the construction above which are easily proven by induction.

Claim:
$$T' = CM(\vec{A'}, \vec{h'})$$

```
- If M=1, the Prover sends A_1,y_1, and the verifier checks that T=e(A_1,h_1) and U=A_1^{y_1}.

- Assume now that M\geq 2 and is a power of 2. Let M'=M/2.

• The prover sets the following:

- \vec{A_L}=[A_1\dots A_{M'}] and \vec{A_R}=[A_{M'+1}\dots A_M]

- \vec{y_L}=[y_1\dots y_{M'}] and \vec{y_R}=[y_{M'+1}\dots y_M]

- \vec{h_L}=[h_1\dots h_{M'}] and \vec{h_R}=[h_{M'+1}\dots h_M]

• The Prover computes and sends to the Verifier the following values U_L=\langle \vec{A_L},\vec{y_R}\rangle; U_R=\langle \vec{A_R},\vec{y_L}\rangle; T_L=CM(\vec{A_L},\vec{h_R}); T_R=CM(\vec{A_R},\vec{h_L})

• The Verifier chooses a random field element x and sends it to the prover

• They recurse on the following values (note that the vectors have now size M'=M/2):

\vec{A'}=\vec{A_L}*\vec{A_R}^x; \vec{y'}=\vec{y_L}*\vec{y_R}^{x^{-1}}; \vec{h'}=\vec{h_L}*\vec{h_R}^{x^{-1}}; T'=T*T_L^{x^{-1}}*T_R^x; U'=U*U_L^{x^{-1}}*U_R^x
```

Fig. 3. Generalized MIPP Protocol.

```
Claim: U' = \langle \vec{A'}, \vec{y'} \rangle
Claim: The Verifier will work only in O(m) = O(\log M) time
```

How can the verifier compute the vectors $\vec{y'}$, $\vec{h'}$ without reading them? The trick is that they are "structured". It is easy to see by induction that at the end of the recursion the value \hat{y} (the collapsed version of $\vec{y'}$ at the end of the recursion) is equal to $(1 - b_1 + x_1^{-1}b_1), \ldots, (1 - b_m + x_m^{-1}b_m)$ which the verifier can compute in O(m) time on their own.

Similarly the value \hat{h} (the collapsed version of $\vec{h'}$ at the end of the recursion) can be seen to be equal to $\hat{h} = h^{\prod_i (1 - t_i + x_i^{-1} t_i)}$.

Note that \hat{h} is a PST commitment of a multilinear polynomial in m variables. The Verifier does not compute it itself (it would be too expensive) but receives it at the end of the recursion from the Prover. To check that it is correct, the verifier computes the polynomial in a random point and it asks the prover to open this PST commitment. The verification time of this construction is O(m).

4 Testudo-Comm: Our PCS with Square Root Trusted Setup

Now we show how to reduce the size of the PST trusted setup to $O(\sqrt{N})$ using the generalized MIPP in Sect. 3. See Fig. 4.

Theorem 1. Testudo-Comm (Fig. 4) is secure in the GGM.

Efficiency. Our commitment scheme improves in proving time trading against proof size and verification time. The key observation for proving efficiency is that, even though prover has to do more expensive operations (pairings, \mathbb{G}_t multiplications etc.), it does them on a \sqrt{N} sized polynomial, which makes a large difference in practice for large N. For example, in Gemini [6], for 2^{25} , it takes at least 36 s to create an opening proof [23] while we evaluate it takes only 11 s using Testudo's commitment (see Sect. 7.1). However, on smaller N, the Gemini transformation is

Trusted Setup We perform the PST trusted setup for m=n/2 variables. Let t_1,\ldots,t_m be random values in \mathbb{Z}_q . The CRS consists of the $M=\sqrt{N}$ values $g_1^{\chi_i(\vec{t})}$.

There is also the trusted setup to generate the vector \vec{h} for the MIPP proof also of size $M = \sqrt{N}$.

Commitment Let $p(x_1, ..., x_n)$ be a multilinear polynomial with n variables. We split the variables in two sets X and Y each of size m = n/2, and denote p(X, Y) accordingly. Let $p_i(X) = \left(\sum_{j \in \{0,1\}^m} p(j,i) \cdot \chi_j(X)\right)$, so that

$$p(X,Y) = \sum_{i \in \{0,1\}^m} \left(\sum_{j \in \{0,1\}^m} p(j,i) \cdot \chi_j(X) \right) \chi_i(Y)$$

The prover "in its own head" PST-commits to each $p_i(X)$ (it can do this because each p_i is multilinear and in m = n/2 variables). This yields a vector of M group elements in G_1 , say $A = (A_1, ..., A_M)$, one commitment A_i for each p_i .

The actual commitment the prover sends to the verifier is a MIPP commitment T to the vector A, defined as above i.e. $T = CM(\vec{A}, \vec{h}) = \Pi_i e(A_i, h_i)$

Opening Now suppose the verifier asks for $p(\vec{a}, \vec{b})$ where $\vec{a}, \vec{b} \in \mathbb{F}^m$. This can be written as:

$$p(\vec{a}, \vec{b}) = \sum_{i \in \{0,1\}^m} p_i(\vec{a}) * \chi_i(\vec{b})$$

Let us define the m-variate multilinear polynomial

$$q(X) := \sum_{i \in \{0,1\}^m} p_i(X) \chi_i(\vec{b})$$

Note that q is multilinear, and $p(\vec{a}, \vec{b}) = q(\vec{a})$.

Let \vec{y} denote the length \sqrt{N} vector consisting of $y_i = \chi_i(\vec{a}) \in \mathbb{F}_r$ as i ranges over $\{0,1\}^m$. To prove that $v = p(\vec{a}, \vec{b})$ the prover proceeds in three steps as in Section 6 of the IPP paper [9]:

- 1. P sends U a PST-commitment to q using the same SRS as that used to commit to each p_i . Note that U is the inner product of \vec{A} and \vec{y} i.e., $U = \Pi_i A_i^{y_i}$.
- 2. Second, P proves using the generalized MIPP protocol that U is indeed the inner product of the vector \vec{A} and the vector \vec{y} , where \vec{A} is the opening vector to T.
- 3. Third, P uses the PST-evaluation protocol to prove that given the commitment U the opening polynomial q evaluated at \vec{a} yields v (i.e. $q(\vec{a}) = v$).

Verification The verifier receives v claimed to be $v = p(\vec{a}, \vec{b})$, the value U, the generalized MIPP proof π_1 and the PST proof π_2 and it performs the MIPP and PST verifications. Note that all verification steps are $O(m) = O(\log N)$.

Fig. 4. Testudo-Comm

likely to outperform Testudo's commitment because of the time required to perform the pairings and \mathbb{G}_t commutations in our case.

We summarize the efficiency properties of the prover in Table 1 assuming a circuit of size $N = 2^n$ and security parameter λ . We then compare the efficiency for the verifier in Table 2.

- **PST:** To open, the prover computes for each of the n rounds, a polynomial division of size 2^{n-1} leading to a $O(2^{n-1})$ -sized polynomial division complexity. While this operates on field elements, we found out that this division, because it doesn't use FFTs, is actually a bottleneck on large sizes (such as 2^{25}).
- **Gemini:** To open, the prover computes for each of the n rounds, 1 KZG openings of size 2^{n-i} and 2 of size 2^{n-i-1} , leading to a complexity of

Scheme	Setup Size	Committing	Opening
Testudo	$O(\sqrt{N}), \mathbb{G}_1,$	$O(\sqrt{N})\mathbb{G}_1,$	$6\sqrt{N}\mathbb{G}_1,$
	$O(\sqrt{N})\mathbb{G}_2$	$O(\sqrt{N})$ pairings,	$O(\sqrt{N})\mathbb{G}_2,$
		$O(\sqrt{N})\mathbb{G}_t$	$4\sqrt{N}$ pairings,
			$4\sqrt{N}\mathbb{G}_t$
PST	$O(N)\mathbb{G}_1$	$O(N)\mathbb{G}_1$	$2N\mathbb{G}_1$
			O(N) poly division
Gemini	$O(N)\mathbb{G}_1$	$O(N)\mathbb{G}_1$	$4N\mathbb{G}_1$

Table 1. Comparison of Prover Efficiency

Table 2. Comparison of Verifier Efficiency

Scheme	Proof Size	Verification Time
Testudo	$\approx \log(N)/2(\mathbb{G}_t + \mathbb{G}_1 + \mathbb{G}_2)$	$\approx \log(N)/2(\mathbb{G}_t + \mathbb{G}_1 + \mathbb{G}_2)$
PST	$O(\log N)\mathbb{G}_2$	$O(\log N)$
Gemini	$3n\mathbb{G}_2$	$8n$ pairings, $3n\mathbb{G}_2$, $3n\mathbb{G}_1$

 $O(2N + N + N) = O(4N)\mathbb{G}_1$ scalar multiplications. For verification, it therefore requires performing O(8n) pairings (or 4n pairing checks).

- **Testudo:** To open, the prover must compute:
 - A PST commitment to the q(X) polynomial, so $O(\sqrt{N})\mathbb{G}_1$
 - A MIPP opening proof, consisting of n rounds where prover computes (a) $O(2\sqrt{N}/2^i)\mathbb{G}_1$ scalar multiplication to compute the reduced vectors, and (b) $O(2\sqrt{N}/2^i)\mathbb{G}_t$ and pairings operations to compute the commitment to each reduced vectors. This leads to $O(4\sqrt{N})\mathbb{G}_1$ and $O(4\sqrt{N})\mathbb{G}_t$ + pairings.
 - Two PST opening proofs, each of size $O(\sqrt{N})$, one on \mathbb{G}_1 and one on \mathbb{G}_2

Remark 1 (Distributed trusted setup). Our construction requires a trusted setup for the polynomial commitment of a specific form. It needs to encode in particular a secret tuple of points and their (multivariate) monomial evaluation. We can obtain an MPC for such a setup by straightforwardly adapting the techniques from [7]. We will detail these techniques in the full version of the paper.

Remark 2 (Proof Size). The proof size for our commitment scheme are 8x bigger at 2^{25} than PST. To reduce the size, we can compress the G_t elements on the torus as in [24]. This could potentially reduce by half the proof size, bringing it to the same order of magnitude as a PST opening proof. Note however, that proof sizes do not matter much in the Testudo SNARK as they are verified by another Groth16 proof on top.

5 Testudo: Our Construction

At this point we recap the general structure of Testudo. Let A, B, C be the input R1CS of size N.

Trusted Setup. We assume that a trusted party (or a distributed multiparty computation protocol, aka ceremony) generates the trusted setup for Testudo-Comm (which is of size \sqrt{N}) and the Groth16 trusted setup for an R1CS corresponding to the verification algorithm of the Spartan sumchecks and the Testudo-Comm opening proofs (this R1CS has size $O(\log N)$). This trusted setup is independent of A, B, C and therefore universal.

Computation Commitments. As in Spartan, in a preprocessing stage, the prover encodes A, B, C as sparse polynomials $\tilde{A}, \tilde{B}, \tilde{C}$ and commits to them via polynomial commitments (computational commitments). We note that for uniform circuits (e.g., data-parallel, with many sub-circuits repeating in regular patterns), this step is not necessary or much reduced in complexity, since the verifier can efficiently compute $\tilde{A}, \tilde{B}, \tilde{C}$ on their own or is only required to compute the computational commitment for the subcircuit.

Witness Commitments. In the online phase, the prover computes \tilde{w} , a multilinear extension of the witness w and commits to it using Testudo-Comm. Note that the polynomials are of size O(N) here, corresponding to the number of R1CS constraints.

Prover. The Testudo prover:

- Executes the Spartan prover to prove the satisfiability of A, B, C (see Sect. 2.5), with the only difference that it uses Testudo-Comm as the underlying polynomial commitment.
- Produces the appropriate openings of the Testudo-Comm PCS.
- Produces and outputs a Groth16 proof that it knows the above modified Spartan proof.

Verifier. The verifier checks the output Groth16 proof and accept/rejects accordingly.

Theorem 2. Assuming that Testudo-Comm is an extractable PCS, and Groth16 is a SNARK, then Testudo is a SNARK.

Informally the proof follows from the fact that if Groth16 is a SNARK we can extract a "modified Spartan" proof – modified to use Testudo-Comm ass the underlying PCS. But if Testudo-Comm is extractable, then we know that we can extract the witness (Spartan is a SNARK as long as the underlying PCS is extractable).

As with all recursive SNARKs we have to heuristically assume that we can instantiate the random oracle in Spartan to a very specific hash function (in our case Poseidon) and not lose security. This is because the code of the hash function has to be embedded in the outer Groth16 proof.

6 Practical Considerations

Choice of Curves. The original version of Spartan (the starting point for the Testudo) uses a custom version of curve25519-dalek, which provides an efficient implementation of a prime-order Ristretto group [26], an abstraction that facilitates implementations of prime-order groups with strong security guarantees. However, as this elliptic curve does not support pairings, composing the original Spartan with Groth16 is not possible and, thus, we had to find a pairing-friendly alternative. We opted for BLS12-377 combined with BW6-761 because they represent the most efficient pair that further supports 2-chaining of pairing-equipped elliptic curves [1,13,14] which is required for our design. See appendix in the full version for more details.

Testudo for Data-Parallel Computation. We can make Testudo particularly efficient for data-parallel computation. Consider a relation R^* composed of several repetitions of the same relation $R(\vec{x}^{(1)}, \vec{w}^{(1)}) \wedge \cdots \wedge R(\vec{x}^{(K)}, \vec{w}^{(K)})$ $(\vec{x}^{(1)}, \ldots, \vec{x}^{(K)})$ We are able to amortize the proving costs related to the wiring of the circuit whenever the circuit is of this form.

In the Spartan lingo the building block for proving the wiring of the circuit refers to a computation commitment. A computation commitment is a polynomial commitment opening to polynomials encoding the structure of the circuit. If we apply Testudo naively to such a relation we would need to open a computation commitment of size roughly K|w|. Instead, we modify our building blocks appropriately to leverage the structure of the circuit and we require computation commitments whose opening grows only linearly in the size of the small subrelation R. We expand on this construction in the full version.

6.1 Parallelization and Aggregation of Testudo Proofs

We observe that this framework enables aggregation of proofs at different levels, each with their pros and cons, but all being compatible with each other, resulting in a system that can scale to large instances in practice because it enables parallelization of the proof generation.

Aggregation at Spartan Level: Assume a prover is running different sumchecks + PCS openings in parallel using different witnesses, on different machines (otherwise, the prover should use the data parallel version that requires the whole witness to be present). In this setting, aggregating the verification of the sumcheck can be done either via (a) aggregating the different Groth16 sumcheck-verifier proofs together using Snarkpack-like constructions, or (b) having one Groth16 proof that verifies multiple instances of the sumcheck. Aggregation for the polynomial commitment scheme could be done by the prover (a) at the beginning, by committing to a random linear combination of the different multilinear extensions of the witnesses and (b) by opening at a random point this combined polynomial. This would require communication between the machines

to aggregate the polynomials together; given it's a single round of communication, we believe it can still be a useful for a practical deployement inside a cluster of machines.

Aggregation at the Groth16 Level: Instead of verifying a single sumcheck instance Groth16 proof and a single PCS opening proof, the outer proof can verify multiple of those. We need further work to estimate the complexity of the final outer circuit but our current estimation (4M constraints for the outer circuit) seem to indicate that it is possible to verify in the order of 5–10 proofs together in a reasonable timeframe, depending on the application.

Note that this aggregation does not require knowledge of the witness and therefore can be done by more expensive prover machines.

Aggregation on Top: Because Testudo's final proof is a Groth16 proof, one can use Snarkpack to efficiently pack thousands of such proofs together. Similar to previous category, this aggregation level *does not require knowledge of the witness* and therefore can be done by more expensive prover machines.

7 Implementation and Evaluation

7.1 Implementation

We have a working Testudo implementation⁵ that features the sumcheck verifier proof and our Testudo commitment scheme. We based our work on the Spartan [28] codebase, which has been adapted to use the Arkworks [2] framework to enable support for any pairing based curves. We also have started an effort on parallelizing the Spartan codebase, although there are still many low hanging fruits to optimize for. On top of this, we implemented a wrapper around a BLS12-381 library that supports GPU operations and released it open source on Github (we do not provide a link in order to keep anonymity of submission).

7.2 Testudo Commitment

We first evaluate our new multivariate commitment scheme compared to the standard PST algorithm. This has been run over a c5a.12xlarge AWS instance, i.e., 48 cores with 96 threads. We note that the structure of the commitment allows for heavy use of parallelism, which we exploited in our implementation.

Figure 5 shows that the Testudo commitment maintains the performance of the PST commitment and, for large circuit instances, **it is 2 orders of magnitude faster** for opening by significantly reducing the size of the MSM required. Indeed, it operates on \sqrt{N} size MSM. However, verification is slower, due to the logarithmic number of pairings required to verify the inner pairing product argument proof. There are still many low-hanging fruits in the codebase to speed up verification such as batching the pairing operations in MIPP and PST.

⁵ The current version of the repository is available at https://github.com/cryptonetlab/testudo.

Moreover, we continue to avoid the use of FFT due to the usage of multilinear polynomials

Verification Speed: As mentioned above, there are many low hanging fruits for optimizing the codebase. For example, to speed up verification, one can bundle the MIPP and the PST part together (e.g. run the pairings check all at once). Currently both MIPP and PST codebase are quite separate.

Proof Size: The Testudo Commitment brings an increase to the proof size in comparison to the simple PST by a factor of 3 but we ensure this is not an issue for the communication cost through recursion. Using the BLS12-377, we are able to efficiently verify commitment openings inside a Groth16 circuit as outlined in Sect. 6.1.

7.3 Testudo Groth16 Constraints

In this section, we estimate the number of R1CS constraints necessary to verify the core of Testudo (the sumchecks, the PCS opening and the computation commitment). Figure 6 shows the number of constraints for each parts according to the user circuit size (e.g. the circuit that the user writes on Testudo). In this estimation, we:

- Use Testudo Commitment as the multilinear PCS scheme for the computation commitment part of Spartan. In the original design, it uses Hyrax.
- Thanks to the previous point, we can now do a random linear combination of all the polynomials the prover must perform and having the prover only compute a single Testudo Commitment opening proof
- We verify the core Spartan sumchecks (steps 2 and 3 in Sect. 2.5) and the grand product sumchecks ([27, bottom of pg. 27]) from the computation commitment inside the same Groth16 verifier that checks the sumcheck in the satisfiability proof. This is possible since both are operating on the same fields.
- Note that we need to verify a PST opening both on \mathbb{G}_1 and \mathbb{G}_2 during the Testudo commitment (for the PST part and for the MIPP part respectively).

The biggest contributor of the R1CS constraints number is the MIPP part, because it requires to compute log(N) \mathbb{G}_t operations (exponentiation is almost 40k in the library we used). We expect this number to drastically go down by roughly 30–50% with optimizations on \mathbb{G}_t computations, such as using the torus arithmetic version and endomorphisms optimizations [24].

7.4 Testudo on Data-Parallel Circuits

We have the necessary building blocks to estimate accurately the proving time of data-parallel circuits (even though the implementation does not yet offer that feature). Specifically, to estimate the time of proving for uniform circuits, we need to add the time for

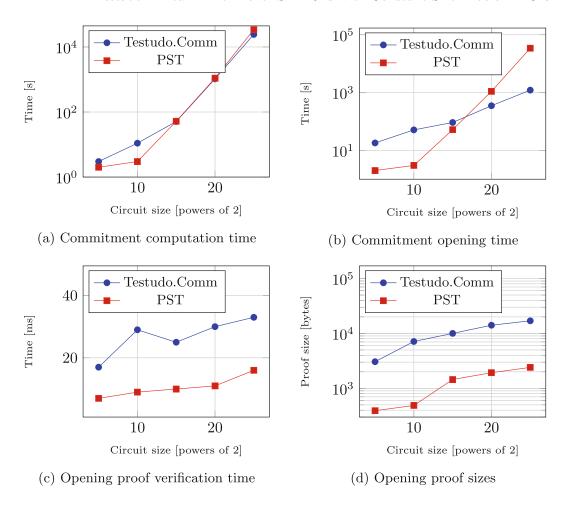


Fig. 5. Comparison between the Testudo and PST commitment schemes. While the time to compute a commitment is similar for Testudo and PST, see (a), we can see in (b) that Testudo outperforms PST when opening a commitment for large circuit sizes by almost two orders of magnitude. On the other hand, graph (c) shows that PST is faster than Testudo by a factor of 2x which, however, can likely be addressed as there are various straightforward ways to further optimize the Testudo code. Finally, graph (d) shows that Testudo opening proofs are about one order of magnitude larger than PST proofs, which, however, will not have any practical impact on the overall sizes of Testudo SNARK proofs, as the opening proof will be ultimately verified by a constant-size Groth16 proof.

- The first sumcheck on the full R1CS matrix (SC1)
- The second sumcheck on the small subcircuit (SC2)
- The Testudo commitment (TC) times on the full witness size commitment and opening combined
- The computation commitment time on the small subcircuit (CC)

We have benchmarked these data for two different subcircuit of size (a) 2^{15} and (b) 2^{20} . The subcircuit corresponds roughly to the size of a Merkle Tree proof verification circuit for 32 layers using either (a) Poseidon for 2^{15} or (b) SHA256

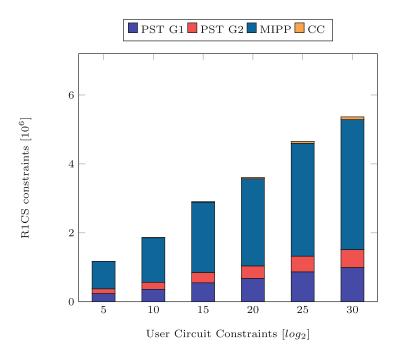


Fig. 6. Constraints for Groth16 verifier of Testudo Proof

Constraints	2^{15}	2^{20}	2^{25}
Commit (s)	0.072	0.723	11
Sumcheck 1 (s)	0.04	0.903	28
Sumcheck 2 (s)	0.036	0.941	29
TC Opening (s)	0.134	0.471	1.5
CC Opening (s)	2	32	939

Fig. 7. Individual cost of data parallel version of Testudo

Subcircuit Size	Proving Time	Speed up vs Groth16
2^{15} constraints (1024 reps)	33s	9.7x
2^{20} constraints (32 reps)	73s	4.3x
Groth16 Proving (2^{25})	322s	_

Fig. 8. Cost of the data parallel version of Testudo vs Groth16

for 2^{20} as the hash function. We're estimating the proving time to repeat these circuits to achieve overall a circuit of 2^{25} constraints: subcircuit (a) is repeated 1024 times (i.e. verify 1024 Merkle Tree proofs), and subcircuit (b) is repeated 32 times. For these parameters, we show that our data parallel Testudo version can be **9.7x faster** for (a) and **4.3x faster** for (b) than their Groth16 equivalent. The improvement is expected to be even larger as the gap of constraints between the subcircuit and the bigger circuit grows. As further work, we will compare these improvement to SNARKs based on plonkish arithmetization, although we are not aware of such speedups for other proof systems at this time (Figs. 7 and 8).

References

- 1. Aranha, D.F., El Housni, Y., Guillevic, A.: A survey of elliptic curves for proof systems. Cryptology ePrint Archive, Report 2022/586 (2022). https://eprint.iacr.org/2022/586
- 2. Arkworks contributors (2023). arkworks zksnark ecosystem
- 3. Belling, A., Soleimanian, A., Bégassat, O.: Recursion over public-coin interactive proof systems; faster hash verification. Cryptology ePrint Archive, Report 2022/1072 (2022). https://eprint.iacr.org/2022/1072
- 4. Bellperson contributors (2023). The bellperson zk-SNARK library
- 5. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Report 2018/046 (2018). https://eprint.iacr.org/2018/046
- Bootle, J., Chiesa, A., Hu, Y., Orrù, M.: Gemini: elastic SNARKs for diverse environments. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022. LNCS, vol. 13276, pp. 427–457. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07085-3_15
- 7. Bowe, S., Gabizon, A., Miers, I.: Scalable multi-party computation for zk-SNARK parameters in the random beacon model. Cryptology ePrint Archive, Report 2017/1050 (2017). https://eprint.iacr.org/2017/1050
- 8. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy. IEEE Computer Society Press (2018)
- 9. Bünz, B., Maller, M., Mishra, P., Tyagi, N., Vesely, P.: Proofs for inner pairing products and applications. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021. LNCS, vol. 13092, pp. 65–97. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-92078-4_3
- 10. Campanelli, M., Fiore, D., Querol, A.: LegoSNARK: modular design and composition of succinct zero-knowledge proofs. In: ACM CCS 2019. ACM Press (2019)
- 11. Chen, B., Bünz, B., Boneh, D., Zhang, Z.: HyperPlonk: plonk with linear-time prover and high-degree custom gates. Cryptology ePrint Archive, Report 2022/1355 (2022). https://eprint.iacr.org/2022/1355
- Chiesa, A., Hu, Y., Maller, M., Mishra, P., Vesely, N., Ward, N.: Marlin: preprocessing zkSNARKs with universal and updatable SRS. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12105, pp. 738–768. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_26
- El Housni, Y., Guillevic, A.: Optimized and secure pairing-friendly elliptic curves suitable for one layer proof composition. In: Krenn, S., Shulman, H., Vaudenay, S. (eds.) CANS 2020. LNCS, vol. 12579, pp. 259–279. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-65411-5_13
- El Housni, Y., Guillevic, A.: Families of SNARK-friendly 2-chains of elliptic curves.
 In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022. LNCS, vol. 13276,
 pp. 367–396. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07085-3_13
- 15. 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

- Gailly, N., Maller, M., Nitulescu, A.: SnarkPack: practical SNARK aggregation.
 In: Eyal, I., Garay, J. (eds.) FC 2022. LNCS, vol. 13411, pp. 203–229. Springer,
 Cham (2022). https://doi.org/10.1007/978-3-031-18283-9_10
- 17. Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: Delegating computation: interactive proofs for muggles. In: 40th ACM STOC. ACM Press (2008)
- 18. Groth, J.: On the size of pairing-based non-interactive arguments. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 305–326. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_11
- 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. LNCS, vol. 10993, pp. 698–728. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96878-0_24
- 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 (2010). https://doi.org/10.1007/978-3-642-17373-8_11
- 21. P Labs (2023). Filecoin: A Decentralized Storage Network
- 22. Lund, C., Fortnow, L., Karloff, H.J., Nisan, N.: Algebraic methods for interactive proof systems. In: 31st FOCS. IEEE Computer Society Press (1990)
- 23. Michele Orrù, G.K.: (2023). zka.lc
- Naehrig, M., Barreto, P.S.L.M., Schwabe, P.: On compressible pairings and their computation. In: Vaudenay, S. (ed.) AFRICACRYPT 2008. LNCS, vol. 5023, pp. 371–388. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-68164-9-25
- 25. Papamanthou, C., Shi, E., Tamassia, R.: Signatures of correct computation. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 222–242. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36594-2_13
- 26. Ristretto contributors (2023). The Ristretto Group
- 27. Setty, S.: Spartan: efficient and general-purpose zkSNARKs without trusted setup. Cryptology ePrint Archive, Report 2019/550 (2019). https://eprint.iacr.org/2019/550
- 28. Setty, S.: Spartan: efficient and general-purpose zkSNARKs without trusted setup. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12172, pp. 704–737. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_25
- 29. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-69053-0_18
- 30. Thaler, J.: (2015–2023). Proofs, Arguments, and Zero-Knowledge
- 31. Xie, T., et al.: zkBridge: trustless cross-chain bridges made practical. In: ACM CCS 2022. ACM Press (2022)
- 32. Zhang, J., Xie, T., Zhang, Y., Song, D.: Transparent polynomial delegation and its applications to zero knowledge proof. In: 2020 IEEE Symposium on Security and Privacy. IEEE Computer Society Press (2020)

- 33. Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: A zero-knowledge version of vSQL. Cryptology ePrint Archive, Report 2017/1146 (2017). https://eprint.iacr.org/2017/1146
- 34. zk Harness contributors (2023). zk-Harness