

On amortization techniques for FRI-based SNARKs

Albert Garreta¹, Hayk Hovhanissyan², Aram Jivanyan², Ignacio Manzur¹, Isaac Villalobos¹, and Michał Zając¹

¹Nethermind Research, {albert,ignacio,isaac,michal}@nethermind.io

²Yerevan State University, aram@skycryptor.com,hayk.9816@gmail.com

Abstract

We present two techniques to improve the computational and/or communication costs of STARK proofs: packing and modular split-and-pack.

Packing allows to generate a single proof of the satisfiability of several constraints. We achieve this by packing the evaluations of all relevant polynomials in the same Merkle leaves, and combining all DEEP FRI functions into a single randomized validity function. Our benchmarks show that packing reduces the verification time and proof size compared to individually proving the satisfiability of each witness, while only increasing the prover time moderately.

Modular split-and-pack is a proof acceleration technique where the prover divides a witness into smaller sub-witnesses. It then uses packing to prove the simultaneous satisfiability of each sub-witness. Compared to producing a proof of the original witness, splitting improves the prover time and memory usage, while increasing the verifier time and proof size. Ideas similar to modular split-and-pack seem to be used throughout the industry, but 1) generally execution traces are split by choosing the first k rows, then the next k rows, and so on; and 2) full recursion is used to prove the simultaneous satisfiability of the sub-witnesses, usually combined with a final wrapper proof (typically a Groth16 proof). We present a different way to split the witness that allows for an efficient re-writing of Plonkish-type constraints. Based on our benchmarks, we believe this approach (together with a wrapper proof) can improve upon existing splitting methods, resulting in a faster prover at essentially no cost in proof size and verification time.

Both techniques apply to popular FRI-based proof systems such as ethSTARK [Sta21], Plonky2/3 [Tea22], RISC Zero [BGT23], Boojum [zks23], and others.

Contents

1	Introduction	2
1.1	Our contribution	2
1.1.1	Packing	3
1.1.2	Modular Split-and-pack	5
1.2	Related work	8
1.3	Structure of the paper	8
2	Preliminaries	8
2.1	Plonkish arithmetization	8
2.2	Grand product arguments	9
3	The packing technique	11
3.1	Description of the STARKPack IOP	11
3.2	Adding <code>proofup</code> to the STARKPack IOP	14
3.3	Benchmarks	14
3.4	Soundness Analysis	16

4	Modular split-and-pack	19
4.1	The split Plonkish relation	20
4.2	An IOP for split-and-pack modulo k	23
4.3	Performance	26
4.4	Security analysis	27
5	References	30
A	Detailed benchmarks	32
A.1	Packed prover times	32
A.2	Packed verifier times	33
A.3	Packed proof size	34
A.4	Splitting in Winterfell	36

1 Introduction

STARKs are non-interactive proof systems known for their efficient prover and verifier characteristics. These systems are now crucial in decentralized environments, especially regarding blockchain scalability. Projects like StarkWare, Polygon, and RISC Zero use STARKs to scale Ethereum. At the same time, these companies have created frameworks that can be used to generate computational integrity proofs with multiple applications.

However, because these proof systems use Merkle tree commitments, they have not benefited from the exciting breakthrough of folding schemes [KST21], which require homomorphic commitments. Instead, STARKs have continued to rely on their sublinear verifier costs to implement recursion. This recursive proof generation method serves both as an IVC/PCD construction in the STARK setting, and as the preferred method for batch-proof generation. Batch-proof generation is now becoming crucial in scaling L2 solutions. For instance SHARP ("SHARed Prover") is a fundamental part of the Starkware tech stack that allows for the batch-generation and verification of multiple Cairo programs. This increases transaction throughput and reduces both latency and gas fees on Starknet, Starkware's decentralized zk-rollup.

Another use case of recursion comes from a technique that seems to be used throughout the industry: splitting long execution traces into smaller traces. When doing such a "splitting", one needs to carefully express constraints that span across trace chunks, so as to enforce the same exact constraints as before the splitting. Then, a proof is generated for each chunk, and the prover produces a recursive proof of the correctness of all the generated proofs. As we mentioned, this seems to be a widespread technique, and a concrete example of this technique being used by RISC Zero can be found in a zkSummit talk [Gaf]. As far as we are aware, this technique has not been formally analyzed yet, and further, it has not been fully exploited. Indeed, the common way teams are splitting execution traces is by selecting the first k rows, then the next k rows, and so on. We call this approach *sequential splitting*. In this work, we present a more efficient variation of this idea.

Overall, this outlines two interesting research directions for optimizing STARK systems: improving batch/recursive proof generation, and finding ways to improve the split-and-prove method.

1.1 Our contribution

We present two ideas to improve STARK proof generation and/or verification: *packing* and *modular split-and-pack*. We will describe the two techniques in the context of the *Plonkish* relation and what we call the *STARKish* IOP.

In this work, by *Plonkish relation* we refer to a relation which consists of an AIR instance [Sta21] with copy constraints [GWC19]. This relation is essentially the relation found in Plonky2/3 [Tea22] (see also [BGK⁺23]). By *STARKish IOP* we refer to the Plonky2-based

IOP as described in [BGK+23]. This is essentially the ethSTARK protocol [Sta21] with added support for permutation arguments. The two techniques described in this paper still apply in the simpler setting, where constraints are just AIRs, and the underlying IOP is the ethSTARK IOP. Similarly, the techniques still apply when we add lookups to the Plonkish relation and support for lookup arguments in the STARKish IOP.

At a high level, the two techniques utilize the fact that recent STARKs are constructed from a " δ -correlated IOP" [BGK+23]. These are roughly IOPs in which the verifier's final decision procedure is as follows: 1) the verifier checks that a polynomial equation holds at a single point (DEEP QUERY procedure), and 2) the verifier checks that certain functions are polynomials of low degree by using FRI (DEEP FRI procedure).

1.1.1 Packing

Packing is a batch-proof generation method. With packing, a prover can create a proof for the simultaneous satisfiability of witnesses to different Plonkish (or AIR, or RAP, and so on) instances. These witnesses can be thought of as execution traces (matrices of dimension $n \times r$) representing the application of a certain program during n steps to r registers. The correctness of the trace representing a certain program is enforced by using polynomial constraints between its entries. As mentioned, one can reduce the integrity of a trace to two checks: that a certain polynomial equation holds at a random point, and that certain functions are low-degree polynomials. The second check is the most expensive of the two. For our purposes, the batched-FRI proximity test is used to carry out this check, which guarantees that a list of functions $f_1, \dots, f_n : D \rightarrow \mathbb{F}$ are (close to) low degree polynomials¹. Here D is certain subset of a finite field \mathbb{F} .

Batched-FRI simply works by applying the FRI low-degree test to a random linear combination of the maps f_1, \dots, f_n . Hence if we need to test more functions we can simply linearly combine them with the functions that we already combined. This allows us to batch all functions to be tested for the low-degree property across different execution traces (representing different programs), and apply a single batched-FRI test. Our second observation was already presented in [MAGABMMT23]. During the batched-FRI protocol, we often have to reveal the evaluations of functions at the same point, and therefore when using Merkle tree commitments we may batch all the evaluations at the same point in the same Merkle leaf. We do this across different execution traces, while [MAGABMMT23] was limited to the same execution trace. These two improvements theoretically reduce the Prover work and the proof size, because the Prover performs essentially one FRI protocol and computes less Merkle commitment roots. The Verifier has a bit more work to do to hash Merkle leaves, but performs much less hash operations.

As we discuss later, we believe packing to be a technique that, in combination with recursion, can improve the cost of proving (multiple) large and complex statements. Further, packing is used as a subroutine in the second technique we present in this paper, and which we call modular-split-and-pack.

Asymptotic and concrete improvements of packing. The theoretical improvement brought by packing is shown in Table 1.

The improvements concern mainly the communication complexity and the Verifier complexity. More precisely, for the proof size we see that the Prover sends less Merkle commitments to witness functions, as evaluations of witness functions are packed in the same leaves across instances. Additionally, the Prover sends less Merkle commitments and decommitment information during the FRI protocol, as it performs only one FRI protocol. The Prover sends exactly the same amount of evaluations when packing proofs. On the Verifier side, the improvement is contained in the hash operations it performs. When packing proofs, the

¹In fact, batched-FRI shows the stronger property that the maps f_i have correlated agreement in a Reed-Solomon code

	Sequential, N instances	Packed, N instances
Proof size	$(4 + \sum_{i \in [N]} \log(n_i))N \cdot \mu$ bits $+ \sum_{i \in [N]} (4r_i + \nu_i + s_i + \ell_i + 1) \cdot \log(\mathbb{K})$ bits $+ O(s_{FRI} \sum_{i \in [N]} \log(\rho^{-1} \mathbf{n}_i) \log(\mathbf{n}_i)) \cdot \mu$ bits	$(4 + \log(\mathbf{n})) \cdot \mu$ bits $+ \sum_{i \in [N]} (4r_i + \nu_i + s_i + \ell_i + 1) \cdot \log(\mathbb{K})$ bits $+ O(s_{FRI} \log(\rho^{-1} \mathbf{n}) \log(\mathbf{n})) \cdot \mu$ bits
Verifier work	$O\left(\sum_{i \in [N]} \left(s_i + \log(n_i) + \sum_{k \in \mathcal{P}_i} \text{fops}(P_{i,k})\right)\right) \mathbb{K}$ $+ 2s_{FRI} \mathbb{H}$ of size $r_1 \log(\mathbb{K})$ $+ 2s_{FRI} \mathbb{H}$ of size $s_1 \log(\mathbb{K})$ $+ 2s_{FRI} \mathbb{H}$ of size $\nu_1 \log(\mathbb{K})$ \vdots $+ 2s_{FRI} \mathbb{H}$ of size $r_N \log(\mathbb{K})$ $+ 2s_{FRI} \mathbb{H}$ of size $s_N \log(\mathbb{K})$ $+ 2s_{FRI} \mathbb{H}$ of size $\nu_N \log(\mathbb{K})$ $+ O(s_{FRI} \sum_{i \in [N]} \log(\rho^{-1} \mathbf{n}_i) \log(\mathbf{n}_i)) \mathbb{H}$	$O\left(\sum_{i \in [N]} \left(s_i + \log(\mathbf{n}_i) + \sum_{k \in \mathcal{P}_i} \text{fops}(P_{i,k})\right)\right) \mathbb{K}$ $+ 2s_{FRI} \mathbb{H}$ of size $\sum_{i \in [N]} r_i \log(\mathbb{K})$ $+ 2s_{FRI} \mathbb{H}$ of size $\sum_{i \in [N]} s_i \log(\mathbb{K})$ $+ 2s_{FRI} \mathbb{H}$ of size $\sum_{i \in [N]} \nu_i \log(\mathbb{K})$ $+ O(s_{FRI} \log(\rho^{-1} \mathbf{n}) \log(\mathbf{n})) \mathbb{H}$

Table 1: Impact of packing N proofs of the satisfiability of Plonkish instances $\text{PL}_i = (\mathcal{P}_i, \mathcal{Q}_i, \sigma_i, \text{PI}_i, r'_i, \ell_i, \mathbf{n}_i, \mathbb{F}, \mathbb{K})$. Here, \mathbf{n}_i and r_i is the number of rows and columns in the execution trace, respectively; σ_i is the permutation specifying the permutation constraints; r'_i is the number of columns that are subject to permutation constraints; \mathcal{P}_i are the constraint polynomials; PI are the public inputs; \mathcal{Q}_i are “selector vectors” and $\ell_i = |\mathcal{Q}_i|$.

On the left, we specify the proof size and verifier work associated to applying the STARKish IOP to each witness separately, and on the right we apply the STARKPack IOP. We set $\mathbf{n} = \max_{i \in [N]}(n_i)$, and recall that ρ is the rate used in the IOP (we may assume it is constant). The hash function is assumed to be a function $\{0, 1\}^* \rightarrow \{0, 1\}^\mu$. In the table, for the verifier work, \mathbb{K} represents the cost of a field operation over \mathbb{K} , and \mathbb{H} represents the cost of a hash function application. Whenever we write “of size” for a hash computation, this means that the hash function is applied to a bitstring of the corresponding size (e.g., $r_1 \log(\mathbb{K})$). If no size is mentioned, it is assumed that the hash is applied to a bitstring of length μ . The quantity s_{FRI} is the number of times the QUERY phase of FRI is repeated. The FRI COMMIT phase is assumed to be carried out until polynomials are supposedly constant. The quantities s_i, ν_i are the quantities appearing in the STARKish IOP (packed and not packed), see for example Section 3.1. The quantity $\text{fops}(P_{i,k})$ represents the maximum number of \mathbb{K} field operations necessary to evaluate the $2r_i + \ell_i$ -variate polynomial $P_{i,k}$ at any point in $\mathbb{H}_i^{2r_i + \ell_i}$ (\mathbb{H}_i is the evaluation domain).

Verifier hashes bigger leaves for witness functions Merkle trees, but there are less of these. Finally when packing proofs, the Verifier only checks Merkle decommitment information for one application of FRI. In the case of proving multiple instances of the same program, more improvements are possible, e.g., we can make multiple proof instances share common parameters such as public input parameter columns, fixed selector columns, or lookup tables.

Our benchmarks can be found in Section 3.3 and Appendices A.1, A.2 and A.3. We use the Winterfell and RISC Zero libraries, and focus on proving execution traces of simple programs. Precisely, for the Winterfell library, the program checks that $x_{i+1} = x_i^3 + 42 \pmod{p}$ for some prime p , for $i \in [0, 2^{16})$. The execution trace has 2^{16} rows. For the RISC Zero library, the program proves the knowledge of the prime factors of a number, yielding, again a trace of 2^{16} rows. The number of columns in RISC Zero is fixed to 275 (it cannot be customized because it is determined by the r0vm circuit²). We tested different column sizes in Winterfell.

For 9 traces of 275 columns, the packed RISC Zero verifier is around 2.2 times faster than the unpacked RISC Zero verifier. Further, the packed Winterfell verifier is 2.4 times faster than the unpacked Winterfell verifier. Both the unpacked RISC Zero and Winterfell verifier checked each proof for the 9 traces sequentially.

The packed RISC Zero proof is around 2.5 times smaller than the proof produced by the unpacked RISC Zero scheme, and the packed Winterfell proof is 1.4 times smaller than the proof created by the unpacked Winterfell scheme³. We also note that as the number of traces packed together increases, the improvements in both proof size, shown in Appendix A.3, and verifier time, shown in Appendix A.2, become more significant as can be seen in Fig. 1 and Fig. 2.

In all our Winterfell benchmarks, the prover time was improved slightly, but remained

²https://docs.rs/risc0-circuit-rv32im-sys/latest/risc0_circuit_rv32im_sys/

³Both the RISC Zero and Winterfell proofs are the proofs for the 9 traces generated one after the other.

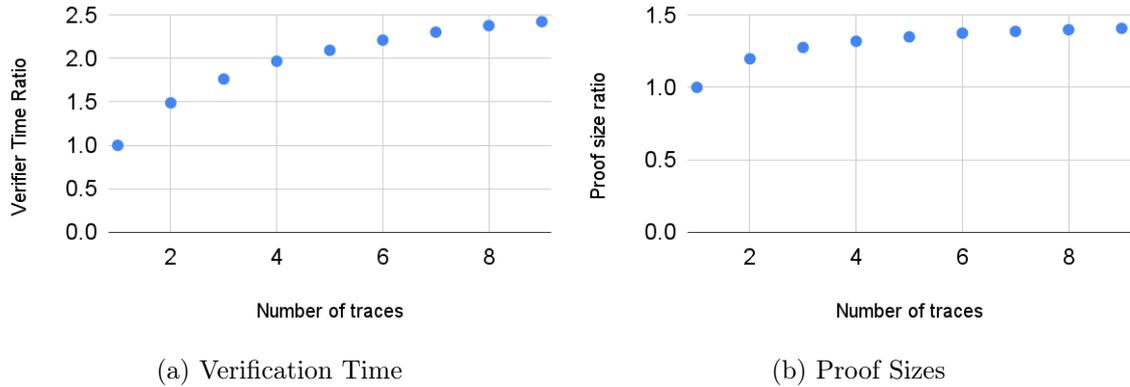


Figure 1: Improvements for Winterfell with 275 columns

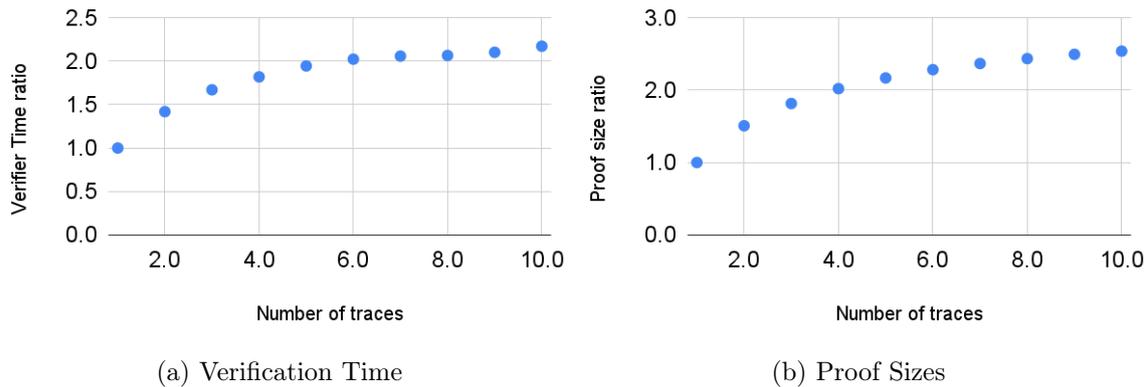


Figure 2: Improvements for Risc Zero

comparable to the prover that individually proves the satisfiability of all witnesses. This was especially true when the number of columns got larger (for 275 columns, the improvement in prover time was around 5%). For our RISC Zero benchmark, the prover time was slightly worse as can be seen in Appendix A.1.

1.1.2 Modular Split-and-pack

Informally, trace splitting is a proof acceleration method where the prover informally splits the execution trace of size $n \times r$ into chunks of size $n/k \times r$ and then applies the original constraints to the entries in the split witnesses. Since the chunks have less rows, we can interpolate their columns with smaller degree polynomials over smaller domains. The constraints for the entries in the chunks need to be translated into polynomial equations involving the polynomials interpolating the new, smaller columns. The new constraints must stay consistent with the previous constraints. The way to do this translation will depend on the splitting, and in general it is hard to do efficiently.

As far as we understand, ideas similar to splitting are used throughout the industry. For instance, from [Gaf] we see that companies like RISC Zero split long execution traces into sequential chunks (with constraints that may overlap from one chunk to the other). This means that if the execution trace is represented as a matrix of size $n \times r$, they split it by choosing the first k rows, then the next k rows, and so on. Recursion is usually used to prove the simultaneous satisfiability of each chunk. Splittings other than sequential do not seem to be explored or considered. We present a new type of splitting modulo k , where rows are selected according to the residue of their index modulo k , where k divides n . We show that for Plonkish-type constraints that involve consecutive rows, our modular-split-and-pack technique allows to write k constraints for each original constraint (which is optimal). For typical constraints, a sequential splitting requires more, essentially because constraints can

Proving Action	% of Overall Proof Time
Executing trace low-degree extension (LDE) polynomial computation and commitment	45% – 50%
Constraint evaluation over LDE domain	25% – 30%
Constraint polynomial factoring into lower-degree polynomial components	6% – 10%
Constraint polynomial evaluation and commitment	\approx 5%
DEEP composition polynomial construction and evaluation	10% – 12%
FRI folding and query proofs	\approx 1%

(a) Modulus Labs benchmarks on Winterfell’s prover time repartition

Proving Action	% of Overall Proof Time
Computing and committing to the wire polynomials, i.e. the intermediate computed values	60% – 70%
Computing and committing to permutation polynomials and permutation argument product polynomials	5% – 8%
Computing and committing to the vanishing/quotient polynomial	6% – 10%
Computing opening proofs for the verifier’s challenge point (Plonky2 uses a multi-opening version of FRI)	10%

(b) Modulus Labs benchmarks on Plonky2’s prover time repartition

Figure 3: Modulus Labs benchmarks on Prover time spread for the Winterfell (a) and Plonky2 (b) libraries.

involve the last row from a chunk and the first from the next one.

It is not immediately clear that splitting traces in other ways could lead to improvements in the number and degree of the resulting constraints. It just so happens that if constraints have some structure, then there are splittings which are more efficient in the rewriting of the constraints over the smaller chunks.

For instance, for Plonkish-type constraints (see Definition 2.1), we show that a splitting "modulo k " leads to better performance than the sequential splitting. Splitting modulo k splits rows in the execution trace according to the residue of their index modulo k . We show in Example 4.4 that this allows the Plonkish-type constraints to be efficiently translated into polynomial constraints. Once this splitting and constraint translation is done, the Prover can use a technique similar to packing to prove the simultaneous satisfiability of each chunk. This means that the evaluations of the witness polynomials at the same point in the evaluation domain are batched in the same Merkle leaf. It also means that all functions involved in all batched-FRI test are linearly combined into a single validity function.

We call the combination of the modular splitting with the packing technique *modular split-and-pack*.

Asymptotic and concrete improvements of modular split-and-pack. In Table 2, we see the impact of modular split-and-pack for a Plonkish instance. Precisely, the splitting is done modulo k . The operation is used on a Plonkish instance with n rows.

We see that modular split-and-pack reduces the cost of number-theoretic transforms from $O(n \log(n))$ to $O(n \log(n/k))$. The leaves of the Merkle trees corresponding to the witness functions are around k times larger, but the Prover performs $\sim k$ times less hash operations. Note the leaves of all other Merkle trees do not grow in size.

To understand the significance of these cost improvements, it is important to understand the different profile costs of creating STARKish proofs. For this, we follow the benchmarks (c.f. Fig. 3) run by the Modulus Labs team [Lab23]. We see that the NTTs and Merkle tree commitments to witness functions represent a significant part of the total computation. For the Winterfell library [KGL⁺20], these operations represent around 50% of the overall proving time. For Plonky2 [Tea22], they comprise around 70% of the proving cost. Therefore,

	No splitting	Split mod k
Prover work	$ \begin{aligned} &O((r + \lceil \frac{r}{u} \rceil)n \log(n))\mathbb{F} \\ &+O((r + \frac{r}{u} + 1)\rho^{-1}n \log(\rho^{-1}n) + M'\rho^{-1}n)\mathbb{K} \\ &+O(\rho^{-1}n \sum_{j \in [\mathcal{P}]} \text{fops}(P_j))\mathbb{K} \\ &+(\rho^{-1}n)\mathbb{H} \text{ of size } r \cdot \log(\mathbb{K}) \\ &+(\rho^{-1}n)\mathbb{H} \text{ of size } (\lceil r/u \rceil + 1) \cdot \log(\mathbb{K}) \\ &+(\rho^{-1}n)\mathbb{H} \text{ of size } \nu \cdot \log(\mathbb{K}) \\ &+O(\rho^{-1}n)\mathbb{H} \end{aligned} $	$ \begin{aligned} &O((r + \lceil \frac{r}{u} \rceil)n \log(n/k))\mathbb{F} \\ &+O((r + \frac{r}{u} + 1)\rho^{-1}n \log(\rho^{-1}n/k) + M\rho^{-1}n/k)\mathbb{K} \\ &+O(\rho^{-1}n \sum_{j \in [\mathcal{P}]} \text{fops}(P_j))\mathbb{K} \\ &+(\rho^{-1}n/k)\mathbb{H} \text{ of size } kr \cdot \log(\mathbb{K}) \\ &+(\rho^{-1}n/k)\mathbb{H} \text{ of size } (\lceil kr/u \rceil + 1) \cdot \log(\mathbb{K}) \\ &+(\rho^{-1}n/k)\mathbb{H} \text{ of size } \nu \cdot \log(\mathbb{K}) \\ &+O(\rho^{-1}n/k)\mathbb{H} \end{aligned} $
Proof size	$ \begin{aligned} &(4 + \log(n))\mu \text{ bits} \\ &+M' \cdot \log(\mathbb{K}) \text{ bits} \\ &+O(\log(\rho^{-1}n) \log(n))\mu \text{ bits} \end{aligned} $	$ \begin{aligned} &(4 + \log(n/k))\mu \text{ bits} \\ &+M \cdot \log(\mathbb{K}) \text{ bits} \\ &+O(\log(\rho^{-1}n/k) \log(n/k))\mu \text{ bits} \end{aligned} $
Verifier work	$ \begin{aligned} &O(s + \log(n) + \sum_{j \in [\mathcal{P}]} \text{fops}(P_j))\mathbb{K} \\ &+2s_{FRI}\mathbb{H} \text{ of size } r \cdot \log(\mathbb{K}) \\ &+2s_{FRI}\mathbb{H} \text{ of size } (\lceil r/u \rceil + 1) \cdot \log(\mathbb{K}) \\ &+2s_{FRI}\mathbb{H} \text{ of size } \nu \cdot \log(\mathbb{K}) \\ &+O(\log(\rho^{-1}n) \log(n))\mathbb{H} \end{aligned} $	$ \begin{aligned} &O(s + \log(n/k) + k \cdot \sum_{j \in [\mathcal{P}]} \text{fops}(P_j))\mathbb{K} \\ &+2s_{FRI}\mathbb{H} \text{ of size } kr \cdot \log(\mathbb{K}) \\ &+2s_{FRI}\mathbb{H} \text{ of size } (\lceil kr/u \rceil + 1) \cdot \log(\mathbb{K}) \\ &+2s_{FRI}\mathbb{H} \text{ of size } \nu \cdot \log(\mathbb{K}) \\ &+O(\log(\rho^{-1}n/k) \log(n/k))\mathbb{H} \end{aligned} $

Table 2: Impact of modular split-and-pack for a Plonkish instance $\text{PL} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K})$ "modulo k " (c.f. Example 4.4) into a split Plonkish instance, when proving the satisfiability of witnesses to each instance. Recall that ρ is the rate used in the IOP (we may assume it is constant). The hash function is assumed to be a function $\{0, 1\}^* \rightarrow \{0, 1\}^\mu$. In the table, for the prover and verifier work, \mathbb{K} represents the cost of a field operation over \mathbb{K} (similarly for \mathbb{F}), and \mathbb{H} represents the cost of a hash function application. Whenever we write "of size" for a hash computation, this means that the hash function is applied to a bitstring of the corresponding size (e.g., $r \log(\mathbb{K})$). If no size is mentioned, it is assumed that the hash is applied to a bitstring of length μ . The quantity s_{FRI} is the number of times the QUERY phase of FRI is repeated, we assume it is a small constant and remove it from asymptotic estimates for simplicity. The FRI COMMIT phase is assumed to be carried out until polynomials are supposedly constant. The quantities s, ν are the quantities appearing in the STARKish IOP (split and not split), see for example Section 4. The quantity M is equal to $(3k + 1)r + \lceil kr/u \rceil + k\ell + \nu + 1$, and M' is equal to $4r + \lceil r/u \rceil + \nu + \ell + 1$. The quantity $\text{fops}(P_j)$ represents the maximum number of \mathbb{K} field operations necessary to evaluate the $2r + \ell$ -variate polynomial P_j at any point in $\mathbb{H}^{2r+\ell}$ (\mathbb{H} is the evaluation domain).

reducing the amount of prover computation in those operations is crucial to have an impact on the overall proving time. This is addressed by the modular-split-and-pack technique.

However, this should be nuanced by saying that the prover needs to send around k times more evaluations in the proof, and the verifier performs k times more field operations to compute the FRI validity functions. This dependency on k means that one should generally split into a small number of chunks (see the benchmarks below), and use a wrapper proof (with, say, Groth16 [Gro16]) to decrease the proof size and verification time. We believe that modular-split-and-pack could improve upon current methods for proving (multiple) long execution traces. We believe it could lead to a technique that is faster for the prover, while having a marginal impact on the proof size and verifier work.

Our benchmarks (cf. Section 4.3) compare modular split-and-pack to running a STARK on the whole execution trace without splitting. They were run with the Winterfell library [KGL⁺20] (which uses ethSTARK [Sta21]) for an execution trace representing the same simple program we used for the packing benchmarks. In particular, there were no copy constraints or lookups. For this program, and a trace of a typical dimensions ($2^{16} \times 275$) the reduction in the Prover time was around **20%** when splitting modulo 4 (into 4 chunks of dimension $2^{14} \times 275$). The Verifier was around 2.2 times slower, and the proof around 3 times larger. When splitting modulo 32, the prover time is reduced by around **40%**, but the proof is ~ 11 times larger, and the verifier ~ 6 times slower as can be seen in Fig. 15 and Fig. 16. The prover's performance exhibits an exponential improvement when the number of splits is relatively low. However, as the number of splits increases, the prover's improvement quickly reaches an asymptotic limit. In contrast, the verifier time and proof size demonstrate rapid improvements for a small number of splits, followed by a linear growth as the number of splits continues to increase. This behavior suggests the existence of an optimal "sweet spot" in the lower range of the number of splits, where the trade-off between the prover's performance and the verifier time and proof size is most favorable.

1.2 Related work

Over the past few years there has been a rapid development of folding techniques [KST21][KS22][KS23][ZGGX23][BC23][EG23]. These methods build upon the idea of post-processing for some predicates in SNARKs, which first appeared in [BGH19] and was formalized in [BCMS20]. Most rely on additively homomorphic commitments, which STARKs do not use (there is some recent work that explores IVC and PCD in the context of non-homomorphic vector commitments [BMNW24]). The packing technique can be seen as an attempt to transpose the idea of post-processing to the STARK setting, by batching FRI instances into a single execution. Of course, the resulting technique is weaker, as the witness size increases linearly in the number of packed instances. Furthermore, packing does not lead to primitives like IVC (and PCD).

The idea of batching relevant evaluations of witness functions at the same point in the same Merkle leaf was already present in [MAGABMMT23]. That work did not leverage the full power of the batched FRI protocol to extend this Merkle commitment optimization to multiple instances.

1.3 Structure of the paper

Section 2 will introduce all the necessary preliminary notions, such as the description of the Plonkish arithmetization and permutation arguments. In Section 3, we introduce the packing technique and the STARKPack IOP, and formally prove the security of the resulting proof system. In Section 4, we introduce the splitting technique, and show how to construct an IOP for a particular type of splitting "modulo k ". The performance of both methods is discussed in Sections 3 and 4, and detailed benchmarks are presented in Appendix A.

2 Preliminaries

2.1 Plonkish arithmetization

We first recall the definition of Plonkish arithmetization. We then see that witnesses for the simultaneous satisfiability of Plonkish instances are equivalent to a witness for a *packed* Plonkish relation, which we describe. We closely follow [BGK⁺23]⁴. For a vector $x \in \mathbb{F}^n$, we denote by x_i its i -th entry. We can also write $x[i]$ for the i -th entry, typically if the vector already has an index in its notation.

Definition 2.1 (Plonkish relation). A Plonkish instance $\text{PL} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{PI}, r, r', \ell, n, \mathbb{F}, \mathbb{K})$ consists of:

- A list \mathcal{P} of $(2r + \ell)$ -variate polynomials $P_i(X_1, \dots, X_r, Y_1, \dots, Y_r, S_1, \dots, S_\ell)$ for all $i \in [|\mathcal{P}|]$.
- A list \mathcal{Q} of ℓ vectors $(\mathbf{q}_1, \dots, \mathbf{q}_\ell) \in (\mathbb{F}^n)^\ell$. These are called selector vectors.
- $r' \leq r$. Intuitively, witnesses for Plonkish instances are traces of size $n \times r$, and r' is the number of columns whose entries are subject to a permutation constraint.
- $\sigma : [r'n] \rightarrow [r'n]$ is a permutation.
- $\text{PI} \subset [rn]$ indicates, intuitively speaking, the location of the public inputs in the witness traces.

An execution trace for a Plonkish instance is a vector $\mathbf{w} \in \mathbb{F}^{nr}$. It is helpful to consider \mathbf{w} as a matrix of size $n \times r$ whose entry (i, j) is given by $\mathbf{w}_{(i-1) \cdot r + j}$. Such a trace is said to satisfy PL if the following conditions hold:

⁴Here we refer to the "Oplonky" relation [BGK⁺23] as "Plonkish"

1. (Circuit constraints satisfaction). For all $j \in [|\mathcal{P}|]$ and all $i \in [n-1]$,

$$P_j(\mathbb{w}_{(i-1) \cdot r+1}, \dots, \mathbb{w}_{i \cdot r}, \dots, \mathbb{w}_{(i+1) \cdot r}, \mathbf{q}_1[i], \dots, \mathbf{q}_\ell[i]) = 0$$

2. (Copy constraints satisfaction). For all $i \in [r'n]$,

$$\mathbb{w}_i = \mathbb{w}_{\sigma(i)}$$

We denote by \mathbb{w}^{Pl} the vector formed by the entries $\{\mathbb{w}_i \mid i \in \text{Pl}\}$. We call this vector the public part of \mathbb{w} . The Plonkish relation is then defined as:

$$\mathbf{R}_{\text{Plonkish}} := \left\{ (\text{PL}, \mathbf{x}, \mathbb{w}) \left| \begin{array}{l} \text{PL} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K}), \\ \mathbf{x} \in \mathbb{F}^{|\text{Pl}|}, \\ \mathbf{x} = \mathbb{w}^{\text{Pl}}, \\ \mathbb{w} \in \mathbb{F}^{nr} \text{ satisfies PL} \end{array} \right. \right\}$$

Remark 2.2. The original Plonkish relation in [Tea22, BGK⁺23] includes further a so-called “repetition parameter”. This parameter controls the number certain checks are performed. Repeating certain checks enables working over a small field, such as the Goldilocks field [Tea22]. However, recent trends in industry seem to have departed from this approach. Instead of repeating some checks, one simply samples challenges from a sufficiently large extension field [BGT23]. In this work, for simplicity, we also assume there are no repeated checks.

2.2 Grand product arguments

In the STARKish proof system, copy constraints and lookup checks are implemented by using what is called “grand product” arguments [GWC19, GW20, Gab]. In our techniques (splitting and modular-split-and-pack) these subprotocols are the ones that will undergo the most significant changes. Because of these, we describe them in detail here.

Protocol: multiset check. In a multiset check, we want to check that two vectors $a = (a_1, \dots, a_n)$, $b = (b_1, \dots, b_n) \in \mathbb{F}^n$ have the same elements counting repetition, possibly in a different order. With high probability over $\gamma \in \mathbb{F}$ chosen uniformly at random, this is true if and only if:

$$\prod_{i \in [n]} (a_i + \gamma) = \prod_{i \in [n]} (b_i + \gamma)$$

This fact allows one to construct the following IOP for multiset check.

Assumption: There is $D = \langle g \rangle \subset \mathbb{F}^*$ such that $|D| = n$.

Inputs: Two vectors $a = (a_1, \dots, a_n)$, $b = (b_1, \dots, b_n) \in \mathbb{F}^n$. We denote \mathbf{a}, \mathbf{b} the polynomial interpolants of a and b , respectively, over the set D . That is, for all $i \in [n]$, $\mathbf{a}(g^i) = a_i$, and similarly for \mathbf{b} .

Protocol:

1. The verifier chooses $\gamma \in \mathbb{F}$ uniformly at random.
2. The prover gives oracle access to the verifier to the polynomial $Z \in \mathbb{F}_{<n}[X]$ such that $Z(g) = 1$, and for all $i \in \{2, \dots, n\}$, $Z(g^i) = \prod_{1 \leq j < i} (a_j + \gamma) / (b_j + \gamma)$.
3. The verifier checks that for all $x \in D$, $L_1(x) \cdot (Z(x) - 1) = 0$, and $Z(x) \cdot (\mathbf{a}(x) + \gamma) = (\mathbf{b}(x) + \gamma) \cdot Z(x \cdot g)$.

Here L_1 is the Lagrange polynomial for g , i.e. $L_1(g) = 1$ and for all $x \in D \setminus \{g\}$, $L_1(x) = 0$. Note in turn that the last two conditions can be expressed as polynomial constraints: the multiset check holds if and only if the polynomials $L_1(X) \cdot (Z(X) - 1)$ and $Z(X) \cdot (a(X) + \gamma) - (b(X) + \gamma) \cdot Z(X \cdot g)$ are divisible by $\prod_{y \in D} (X - y)$. The quantities $Z(g^i)$ are the partial products that are included in the accumulator column.

The multiset check is usually used to perform a permutation check between polynomials $f, h \in \mathbb{F}_{<n}[X]$, meaning that for a permutation $\sigma : [n] \rightarrow [n]$, we want to check that $h(g^i) = f(g^{\sigma(i)})$ for all i . The way this is done is by applying the multiset check to the vectors

$$(f(g) + \beta, \dots, f(g^n) + \beta \cdot n), (h(g) + \beta \cdot \sigma(1), \dots, h(g^n) + \beta \cdot \sigma(n))$$

for $\beta \in \mathbb{F}$ chosen uniformly at random. It can be extended further to a permutation check for $f_1, \dots, f_k, h_1, \dots, h_k \in \mathbb{F}_{<n}[X]$ given a permutation $\sigma : [nk] \rightarrow [nk]$. This is done by applying the permutation check to the vectors in \mathbb{F}^{nk} constructed by setting

$$f_{(i-1) \cdot k + j} = f_j(g^i)$$

(and similarly for $h_{(i-1) \cdot k + j}$). In this case, for $j \in [k]$ one defines the polynomials

$$\begin{aligned} S_{\text{ID}_j}(\mathbf{g}^i) &= (i-1) \cdot k + j \\ S_{\sigma_j}(\mathbf{g}^i) &= \sigma((i-1) \cdot k + j) \end{aligned}$$

Then the polynomials $f'_j(X) = f_j(X) + \beta S_{\text{ID}_j}(X)$ and $h'_j(X) = h_j(X) + \beta S_{\sigma_j}(X)$ correctly interpolate the value of the vectors $(f_j(\mathbf{g}) + \beta j, \dots, f_j(\mathbf{g}^n) + \beta((n-1)k + j))$ and $(h_j(\mathbf{g}) + \beta \sigma(j), \dots, h_j(\mathbf{g}^n) + \beta \sigma((n-1)k + j))$ respectively. Then one can apply the grand product argument to $f' = \prod f'_j$ and $h' = \prod h'_j$.

Finally, these extended permutation checks can be applied to enforce copy constraints. For a partition $\mathcal{T} = \{T_1, \dots, T_s\}$ of $[kn]$, a set of polynomials $f_1, \dots, f_k \in \mathbb{F}_{<n}[X]$ is said to copy-satisfy \mathcal{T} if, for $f_{(j-1) \cdot n + i}$ as defined above, we have $f_\ell = f_{\ell'}$ whenever ℓ, ℓ' are in the same block T_i . This can be checked with an extended permutation check on $f_1, \dots, f_k, f_1, \dots, f_k$ (with polynomials repeated twice), by defining a permutation $\sigma(\mathcal{T}) : [kn] \rightarrow [kn]$ whose cycle decomposition is given by the blocks T_i .

Protocol: `lookup`. Lookups are a powerful primitive which allows to verify that a certain subset of witness values are contained in the values of predefined (lookup) tables. We present and use `lookup` as it is the lookup argument used in STARKish, and it is also used in the RISC Zero proof system.

Suppose for simplicity that we want to check that the values taken by $f \in \mathbb{F}_{<n}[X]$ over $D = \{g, \dots, g^n = 1\} \subset \mathbb{F}^*$ are included in the table $t \in \mathbb{F}^n$ (if the table is smaller, we pad it by repeating its last element until it has length n). Let s be the concatenation (f, t) , and sort it by t (meaning that elements in s appear in the same order as they do in t). It is shown in [GW20] that it now suffices to do a multiset check between s' and $((1 + \beta)f, t')$, where:

- $\beta \in \mathbb{F}$ is chosen uniformly at random
- $s' \in \mathbb{F}^{2n}$ and $t' \in \mathbb{F}^n$ are vectors defined by $s'_i = s_i + \beta s_{i+1}$ and $t'_i = t_i + \beta t_{i+1}$

This also extends to vector lookups and multiple tables, as seen in [GW20].

Remark 2.3. Note again that the multiset check can be translated into polynomial constraints involving the polynomials representing the vectors of values we are checking, and the polynomial representing the partial product quantities. All permutation, copy constraint, and lookup checks are multiset checks. This means that if the prover commits to the partial product polynomials, all these checks can be translated into polynomial constraints involving the witness polynomials.

3 The packing technique

In this section we describe how the packing technique applies to the STARKish IOP. Recall that the STARKish IOP is the Plonky2 IOP [Tea22, BGK⁺23], this is an IOP that can prove the satisfiability of a witness for a Plonkish instance which is essentially an AIR with copy constraints [GWC19]. In the literature, this arithmetization –with small variations in their formulations– is also referred to as RAPs or TurboPlonk. When adding lookups to the relation and support for lookup arguments in the IOP, it is also referred to as UltraPlonk. The packing technique applied to the STARKish IOP results in an IOP, which we call *STARKPack*, that can prove the simultaneous satisfiability of witnesses for different Plonkish instances. We then explain how to extend that argument to support **lookup**, which we still call the STARKPack IOP, or STARKPack with lookups. Then, we provide soundness proofs for the STARKPack IOP.

Definition 3.1 (Packed Plonkish relation). Informally, the packed Plonkish relation is a conjunction of Plonkish relations (cf Definition 2.1). Formally, it is defined as:

$$\mathbf{R}_{\text{PPlonkish}} := \left\{ (N, \text{PL}_1, \dots, \text{PL}_N, \vec{\mathbf{x}}, \vec{\mathbf{w}}) \left| \begin{array}{l} N \text{ is an integer } \geq 1, \\ \text{PL}_i = (\mathcal{P}_i, \mathcal{Q}_i, \sigma_i, \text{Pl}_i, r_i, r'_i, \ell_i, \mathbf{n}_i, \mathbb{F}, \mathbb{K}) \\ \quad \text{is a Plonkish relation } \forall i \in [N], \\ \vec{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^N) \in \prod_{i \in [N]} (\mathbb{F}^{\text{Pl}_i}), \\ \vec{\mathbf{w}} = (\mathbf{w}^1, \dots, \mathbf{w}^N) \in \prod_{i \in [N]} (\mathbb{F}^{\mathbf{n}_i r_i}), \\ \text{For all } i \in [N], (\mathbf{w}^i)^{\text{Pl}_i} = \mathbf{x}^i, \\ \text{For all } i \in [N], \mathbf{w}^i \text{ satisfies PL}_i \end{array} \right. \right\}$$

The STARKPack IOP is an IOP constructed from the STARKish IOP that can prove the simultaneous satisfiability of Plonkish instances, or equivalently, the satisfiability of a Packed Plonkish relation.

3.1 Description of the STARKPack IOP

The STARKPack IOP, denoted by $\Pi_{\text{STARKPack}}$, mainly follows the STARKish protocol. We fix an integer $N \geq 1$, and a finite field extension \mathbb{K} of \mathbb{F} . We fix N Plonkish instances, for each $i \in [N]$ their parameters will be denoted by $\text{PL}_i = (\mathcal{P}_i, \mathcal{Q}_i, \sigma_i, \text{Pl}_i, r_i, r'_i, \ell_i, \mathbf{n}_i, \mathbb{F}, \mathbb{K})$. For each i , we assume that there is a multiplicative subgroup $\text{D}_i = \langle \mathbf{g}_i \rangle \subset \mathbb{F}^\times$ of size \mathbf{n}_i . Its vanishing polynomial is denoted by $Z_{\text{D}_i}(X) \in \mathbb{F}_{<\mathbf{n}_i}[X]$. Further, we assume that if $\mathbf{n}_i \leq \mathbf{n}_{i'}$, then D_i is a multiplicative subgroup of $\text{D}_{i'}$ (in which case we necessarily have $\mathbf{n}_i \mid \mathbf{n}_{i'}$).

We also fix $\mathbf{n} = \max_i \mathbf{n}_i$. The integer \mathbf{n} is the largest trace length among all instances. It corresponds to a trace evaluation domain $\text{D} = \langle \mathbf{g} \rangle \subset \mathbb{F}^\times$ of size \mathbf{n} . We fix an evaluation domain $\text{H} \subset \mathbb{K}^\times$ which contains a non-trivial coset of D . When we say that the prover gives oracle access to a function, we mean that the prover interpolates a vector with entries in \mathbb{F} over a certain subgroup of \mathbb{F}^\times ; and then provides oracle access to this interpolant over H . We fix a witness (or execution trace) $\mathbf{w} \in \prod_{i \in [N]} (\mathbb{F}^{\mathbf{n}_i r_i})$ and public part \mathbf{x} for the Packed Plonkish instance $(N, \text{PL}_1, \dots, \text{PL}_N)$. Finally, we fix auxiliary parameters $\text{aux} = (\vec{t}, s)$ for the Batched-FRI protocol. The parameter $\vec{t} = (t_1, \dots, t_m) \in \mathbb{N}$ indicates that we will use t_i -to-1 maps in the i -th COMMIT phase of FRI, and s is the number of repetitions of the QUERY phase of FRI.

Preprocessing. The prover and verifier compute the selector polynomials as:

$$\forall i \in [N], \forall j \in [\ell_i], \mathbf{q}_j^i(X) = \sum_{k=1}^{\mathbf{n}_i} \mathbf{q}_j^i[k] \cdot L_k^i(X)$$

where the L_k^i are the Lagrange polynomials for the basis $\{\mathbf{g}_i, \dots, \mathbf{g}_i^{n_i} = 1\}$. Additionally, for all $i \in [N]$ and all $j \in [r'_i]$, the prover and verifier agree on the splitting parameter u_i for the permutation argument (see below), and compute the permutation polynomials $\mathbf{S}_{\mathbb{D}_j}^i, (\mathbf{S}_{\sigma_i}^i)_j$ described in Section 2.2.

Additionally, define the FRI degree correction terms $\mathbf{e}_i := (n-1) - (n_i-2)$ for all $i \in [N]$.

Round 1. The prover computes the wire polynomials as follows:

$$\forall i \in [N], \forall j \in [r_i], \mathbf{a}_j^i(X) = \sum_{k=1}^{n_i} \mathbb{w}_{(j-1) \cdot n_i + k} \cdot L_k^i(X)$$

The prover then gives oracle access to the wire polynomials $\mathbf{a}_j^i : \mathbb{H} \rightarrow \mathbb{K}$ to the verifier. The verifier samples uniform randomness used in the permutation argument $\beta, \gamma \in \mathbb{K}$. The prover will use the same randomness for all copy constraint arguments and lookup arguments (see next section for lookups) across instances.

Round 2. The prover sends oracle access to all permutation and partial product polynomials. Here the terminology does not concord with what we have presented in Section 2.1, so we explain it. For simplicity fix $i \in [N]$. The prover is supposed to provide oracle access to a polynomial $z^i \in \mathbb{F}_{<n_i}[X]$ such that:

$$z^i(\mathbf{g}X) \cdot g^i(X) = f^i(X) \cdot z^i(X) \quad (1)$$

where:

$$f^i(X) = \prod_{l=1}^{[r_i]} (\mathbf{a}_l^i(X) + \beta \cdot \mathbf{S}_{\mathbb{D}_l}^i(X) + \gamma) := \prod_{l=1}^{[r_i]} f_l^i(\beta, \gamma, X)$$

$$g^i(X) = \prod_{l=1}^{[r_i]} (\mathbf{a}_l^i(X) + \beta \cdot (\mathbf{S}_{\sigma_i}^i)_l(X) + \gamma) := \prod_{l=1}^{[r_i]} g_l^i(\beta, \gamma, X)$$

We have called z^i the partial product polynomial before. However, if r_i is large, then so is the degree of the polynomial constraint in Eq. (1). To reduce its degree, it is split into many polynomial constraints. For the integer u_i determined in preprocessing, the prover splits the products as follows. Define $s_i = \lceil r_i / u_i \rceil$, then for each $k \in [s_i]$:

$$\overline{f}_k^i(\beta, \gamma, X) = \overline{f}_k^i(X) := \prod_{l=(u_i-1) \cdot k + 1}^{u_i \cdot k} f_l^i(\beta, \gamma, X)$$

$$\overline{g}_k^i(\beta, \gamma, X) = \overline{g}_k^i(X) := \prod_{l=(u_i-1) \cdot k + 1}^{u_i \cdot k} g_l^i(\beta, \gamma, X)$$

Now the prover will compute polynomials $\pi_1^i, \dots, \pi_{s_i-1}^i \in \mathbb{F}_{<n_i}[X]$ such that for all $m \in [n_i]$:

$$\pi_1^i(\mathbf{g}^m) = z^i(\mathbf{g}^m) \overline{f}_1^i(\mathbf{g}^m) \overline{g}_1^i(\mathbf{g}^m)^{-1}$$

$$\pi_k^i(\mathbf{g}^m) = \pi_{k-1}^i(\mathbf{g}^m) \overline{f}_k^i(\mathbf{g}^m) \overline{g}_k^i(\mathbf{g}^m)^{-1}, \text{ for } k = 2, \dots, s_i - 1$$

These equalities can then be translated into polynomial constraints of the following form:

$$z^i(X) \overline{f}_1^i(X) - \pi_1^i(X) \overline{g}_1^i(X) \equiv 0 \pmod{Z_{\mathbb{D}_i}(X)}$$

$$\pi_{k-1}^i(X) \overline{f}_k^i(X) - \pi_k^i(X) \overline{g}_k^i(X) \equiv 0 \pmod{Z_{\mathbb{D}_i}(X)}, \text{ for } k = 2, \dots, s_i - 1$$

$$\pi_{s_i-1}^i(X) \overline{f}_{s_i}^i(X) - z^i(\mathbf{g}X) \overline{g}_{s_i}^i(X) \equiv 0 \pmod{Z_{\mathbb{D}_i}(X)}$$

$$L_1^i(X) (z^i(X) - 1) \equiv 0 \pmod{Z_{\mathbb{D}_i}(X)}$$

The π_k^i will now be called the partial product polynomials, while the polynomial z^i will be called the permutation polynomial. The prover sends oracle access to the permutation and partial product polynomials $(z^i)_{i \in [N]}, (\pi_k^i)_{i \in [N], k \in [s_i-1]}$.

The verifier samples uniform randomness $\alpha \in \mathbb{K}$.

Round 3. The prover computes polynomials $(\mathbf{u}^i, \mathbf{d}^i, \mathbf{quot}^i)_{i \in [N]}$ as follows:

$$\begin{aligned} \mathbf{u}^i(X) &:= (z^i(X)\overline{f_1^i}(X) - \pi_1^i(X)\overline{g_1^i}(X))\alpha \\ &\quad + \sum_{k=2}^{s_i} (\pi_{k-1}^i(X)\overline{f_k^i}(X) - \pi_k^i(X)\overline{g_k^i}(X))\alpha^k \\ &\quad + (\pi_{s_i-1}^i(X)\overline{f_{s_i}^i}(X) - z^i(X)\overline{g_{s_i}^i}(X))\alpha^{s_i+1} \\ &\quad + L_1^i(X)(z^i(X) - 1)\alpha^{s_i+2} \\ \mathbf{d}^i(X) &:= \sum_{k \in [\mathcal{P}]} \alpha^{k-1} P_k^i(\mathbf{a}_1^i(X), \dots, \mathbf{a}_{r_i}^i(X), \mathbf{q}_1^i, \dots, \mathbf{q}_{\ell_i}^i) \\ &\quad + \alpha^{|\mathcal{P}|} \mathbf{u}^i(X) \\ \mathbf{quot}^i(X) &:= \mathbf{d}^i(X) / Z_{D_i}(X) \end{aligned}$$

Then the prover splits each $\mathbf{quot}^i(X)$ into degree n_i polynomials $\mathbf{quot}_1^i, \dots, \mathbf{quot}_\nu^i$ such that $\mathbf{quot}^i(X) = \sum_l X^{n_i \cdot (l-1)} \mathbf{quot}_l^i(X)$.

The prover gives oracle access to the maps $\mathbf{quot}_1^i, \dots, \mathbf{quot}_\nu^i : \mathbb{H} \rightarrow \mathbb{K}$. The verifier samples uniform randomness $\mathfrak{z} \in \mathbb{K} \setminus (\mathbb{H} \cup \mathbf{g}^{-1}\mathbb{H})$.

Round 4. The prover sends the evaluations

$$\mathbf{eval} = (\mathbf{a}_j^i(\mathfrak{z}), \mathbf{a}_j^i(\mathbf{g}\mathfrak{z}), \mathbf{quot}_l^i(\mathfrak{z}), \pi_m^i(\mathfrak{z}), z^i(\mathfrak{z}), z^i(\mathbf{g}\mathfrak{z}), \mathbf{q}_c^i(\mathfrak{z}), \mathbf{S}_{\text{ID}_j}^i(\mathfrak{z}), (\mathbf{S}_{\sigma_i}^i)_j(\mathfrak{z}))$$

for all $i \in [N], j \in [r_i], l \in [\nu], m \in [s_i - 1], c \in [\ell_i]$. The verifier samples uniform FRI randomness $\eta \in \mathbb{K}$.

Batched-FRI. The prover and verifier engage in a Batched-FRI protocol with auxiliary parameters \mathbf{aux} to check proximity to the Reed-Solomon code $\text{RS}(\mathbb{K}, \mathbb{H}, \mathbf{n})$ of the function defined as follows. Consider the vector of functions defined as:

$$\left(\frac{\mathbf{a}_j^i(X) - \mathbf{a}_j^i(\mathfrak{z})}{X - \mathfrak{z}}, \frac{\mathbf{a}_j^i(X) - \mathbf{a}_j^i(\mathbf{g}\mathfrak{z})}{X - \mathbf{g}\mathfrak{z}}, \frac{\mathbf{quot}_l^i(X) - \mathbf{quot}_l^i(\mathfrak{z})}{X - \mathfrak{z}}, \dots, \frac{(\mathbf{S}_{\sigma_i}^i)_j(X) - (\mathbf{S}_{\sigma_i}^i)_j(\mathfrak{z})}{X - \mathfrak{z}} \right) \quad (2)$$

for all $i \in [N], j \in [r_i], l \in [\nu], m \in [s_i - 1], c \in [\ell_i]$ (it contains all relevant quotients of functions whose evaluations appear in \mathbf{eval}). For simplicity, relabel this vector as $\vec{F} = (F_1(X), F_2(X), \dots, F_M(X))$ (for some integer M). The function we apply FRI to is the linear combination of the F_j 's multiplied by the factors $(\eta^j + \eta^{j+2} X^{e(F_j)})$. Here \mathbf{e} is defined by the fact that if F_j is supposedly of degree at most $n_i - 2$, then $\mathbf{e}(F_j) = \mathbf{e}_i$ (the FRI degree correction term defined in Preprocessing).

Decision.

1. **Circuit and copy constraint check.** The verifier uses \mathbf{eval} to verify that for all $i \in [N]$, $\mathbf{d}^i(\mathfrak{z}) = \mathbf{quot}^i(\mathfrak{z})Z_{D_i}(\mathfrak{z})$. If this check fails, the verifier rejects.
2. **Batched FRI check.** If the prover fails the Batched-FRI protocol with auxiliary parameters \mathbf{aux} , the verifier rejects.

If all checks pass, the verifier accepts the proof.

Commitments in the STARKPack IOP. When we instantiate this IOP protocol with Merkle tree commitments, in Round 1 the prover needs to provide a commitment to the evaluations of the wire polynomials $\mathbf{a}_j^i : \mathbf{H} \rightarrow \mathbb{K}$ coming from different instances. There is a way to compactly pack leaves so that they contain the necessary evaluations of all relevant functions. To do this, the prover performs an IFFT over \mathbf{D}_i , and then an FFT over \mathbf{H} taking $O(n_i \log(n_i) + |\mathbf{H}| \log(|\mathbf{H}|))$ field operations per function. Looking at Eq. (2), when querying the Batched-FRI function at $x_0 \in \mathbf{H}$, we need to provide all the evaluations of the wire polynomials at x_0 . Therefore, the prover can pack all the evaluations at a point $x_0 \in \mathbf{H}$ of all wire polynomials (from all Plonkish instances) into one leaf of the Merkle tree, and compute one root for all functions.

The same holds for other functions computed by the prover during the protocol. This idea was also presented in [MAGABMMT23].

3.2 Adding `lookup` to the STARKPack IOP

As we mentioned in Section 2.2, `lookup` allows one to verify that $f \in \mathbb{F}^n$ is included in $t \in \mathbb{F}^d$ (meaning that $\{f\} \subset \{t\}$). Because we may pad t by repeating its last value, we may assume $d = n$ without loss of generality. The `lookup` protocol applies a multiset check to $s', ((1 + \beta)f, t') \in \mathbb{F}^{2n}$ for a uniformly random $\beta \in \mathbb{F}$, where:

- $s \in \mathbb{F}^{2n}$ is (f, t) sorted by t (elements in s appear in the same order as they do in t)
- $s'_i = s_i + \beta s_{i+1}$ and $t'_i = t_i + \beta t_{i+1}$

We make no difference in notation between the vectors and their polynomial interpolants over the domain $\mathbf{D} = \langle \mathbf{g} \rangle \subset \mathbb{F}^\times$ of size n , except that we denote by $h_1, h_2 \in \mathbb{F}_{<n}[X]$ the polynomials such that $h_1(\mathbf{g}^i) = s_i$ and $h_2(\mathbf{g}^i) = s_{n+i}$. For uniformly random $\beta, \gamma \in \mathbb{F}$, denote by $Z \in \mathbb{F}_{<n}[X]$ the polynomial such that:

$$\begin{aligned} Z(\mathbf{g}) &= 1 \\ Z(\mathbf{g}^i) &= \frac{(1 + \beta)^{i-1} \prod_{j < i} (f_j + \gamma) \cdot \prod_{1 \leq j < i} (\gamma(1 + \beta) + t_j + \beta t_{j+1})}{\prod_{1 \leq j < i} (\gamma(1 + \beta) + s_j + \beta s_{j+1}) (\gamma(1 + \beta) + s_{n+j} + \beta s_{n+j+1})}, \text{ for } 2 \leq i < n \\ Z(\mathbf{g}^n) &= 1 \end{aligned}$$

The relevant multiset check translates into the polynomial constraints:

$$\begin{aligned} L_i(X)(Z(X) - 1) &\equiv 0 \pmod{Z_{\mathbf{D}}(X)}, \text{ for } i = 1, n \\ Z(X) \cdot (1 + \beta)(f(X) + \gamma)(\gamma(1 + \beta) + t(X) + \beta t(\mathbf{g}X)) &\equiv \\ Z(\mathbf{g}X) \cdot (\gamma(1 + \beta) + h_1(X) + \beta h_1(\mathbf{g}X))(\gamma(1 + \beta) + h_2(X) + \beta h_2(\mathbf{g}X)) &\pmod{Z_{\mathbf{D}}(X)} \end{aligned}$$

This means that if the lookup table t is encoded as one of the wire polynomials, then the verifier can use the same randomness sampled in Round 1 to implement `lookup`. The prover will need to additionally commit to h_1, h_2, Z over the evaluation domain. If the degree of the polynomial constraints gets too large, we can also split it into smaller degree constraints exactly as for the permutation argument. Then the polynomial constraints will be included in the \mathbf{d} polynomials in Round 3. The rest of the protocol remains the same. We can modify the Plonkish instances slightly to include lookup tables in the public data, we still refer to these instances as Plonkish instances. By implementing `lookup` in the STARKPack IOP as we described, we obtain an IOP to prove the simultaneous satisfiability of Plonkish instances with lookups, still denoted by $\Pi_{\text{STARKPack}}$.

3.3 Benchmarks

In this section we present benchmarks comparing the performance of STARKPack and the protocol whereby each proof is created separately. For a detailed description of the costs of STARKPack we refer to Table 1 in the introduction of the paper.

In general, STARKPack optimizes both the computational and communication complexities of STARK proof generation. For N programs, there is a single FRI proof. This reduces the joint proof size and the verification time. All N proofs share a single Merkle commitment and decommitment step, and rely on a single composition polynomial evaluation. This also helps reduce both the proving and verification time, and shrinks the proof size. Furthermore, the verifier samples the same randomness for all N instances, which reduces the communication complexity.

We recall the simple programs used for benchmarks: for Winterfell, checking that $x_{i+1} = x_i^3 + 42 \pmod{p}$ for some prime p until the rows of the trace were 2^{16} ; and for RISC Zero, the program proved the knowledge of the prime factors of a number over a trace with 2^{16} rows.

A note about the RISC Zero and Winterfell proof systems. The arithmetization used in RISC Zero [BGT23] is a RAP. This is another way to formalize the STARKish/TurboPlonk/UltraPlonk arithmetization. There are selector columns (which act as the selector polynomials), and the prover and verifier engage in a randomized preprocessing phase to compute the values of the relevant permutation/partial product polynomials in additional trace columns called accumulator columns. Therefore, this arithmetization is also amenable to permutation arguments and lookups. The IOP used in RISC Zero is similar to ethSTARK [Sta21] with support for permutation and lookup arguments, and closely follows the blueprint we described in Section 3.1 when the number of traces is $N = 1$. For this reason, the same ideas that we applied to obtain the STARKPack IOP can be applied in this setting. In this way, we obtain a proof system that can show the simultaneous satisfiability of multiple RAP instances.

Winterfell uses ethSTARK for AIRs, and the ideas apply in the same way (AIRs are RAPs without selector or accumulator columns).

Benchmarks. For our benchmarks, we applied the STARKPack ideas for the same program run many times with different inputs. Fig. 1 shows the STARKPack performance improvement when implemented in RISC Zero. The acceleration for the proving time has been small, around 5% for the benchmarked instances: this is due to the large number of columns. The Winterfell benchmarks presented below for traces with fewer columns show how the proving time optimization depends on the number of columns (Fig. 5).

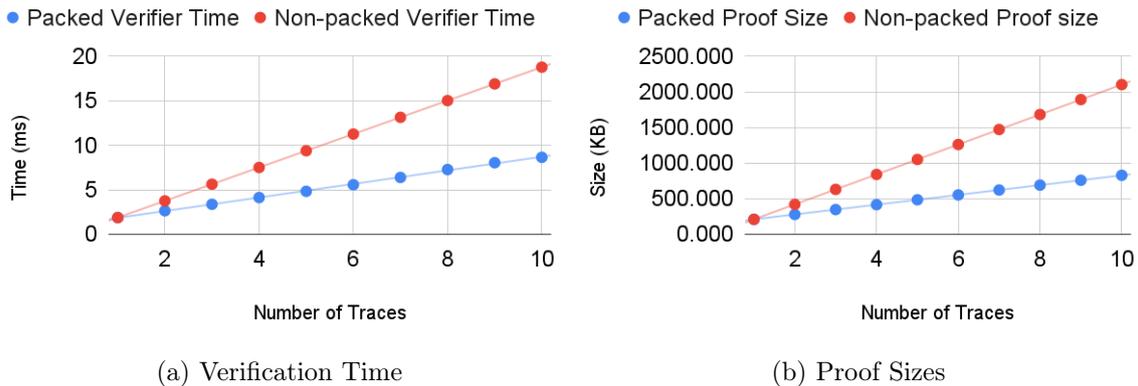


Figure 4: RISC Zero STARKPack Implementation Performance

The diagrams below show the STARKPack performance improvement implemented in Winterfell for AIR traces comprised of 10, 100 and 275 columns.

Remark 3.2. When we run our benchmarks on the same program with different inputs, further optimizations are possible. For instance, all programs will use the same selector columns and lookup tables. We can use these columns only once for all RAPs. In this extension, only

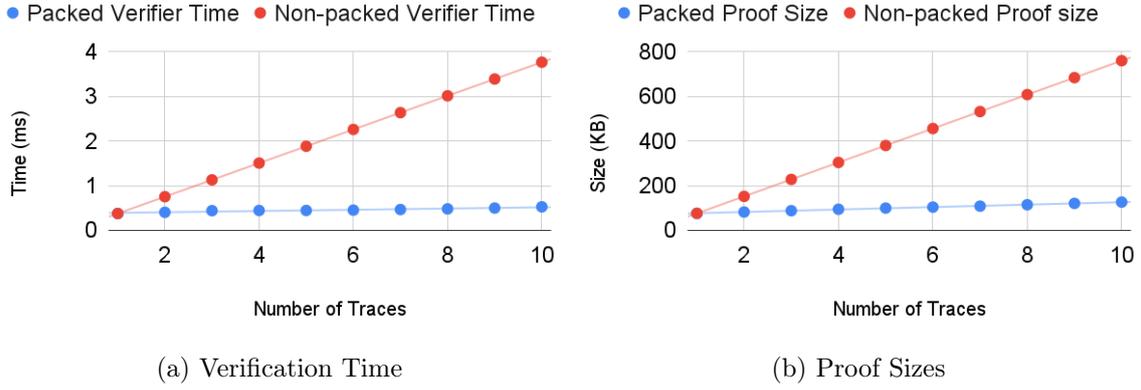


Figure 5: Winterfell STARKPack Implementation Performance for AIR traces of 10 columns

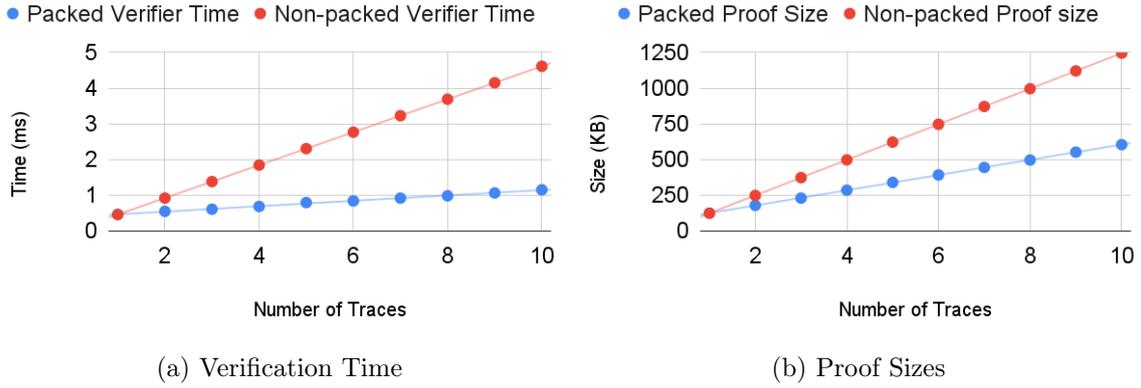


Figure 6: Winterfell STARKPack Implementation Performance for AIR traces of 100 columns

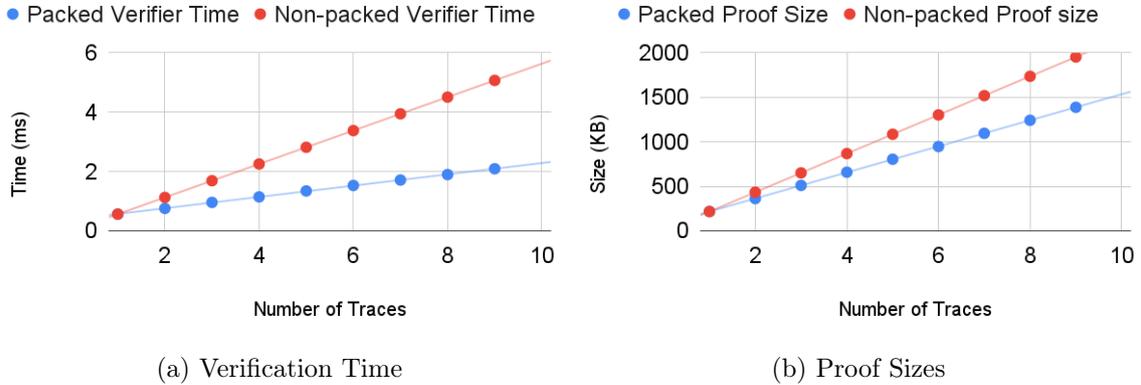


Figure 7: Winterfell STARKPack Implementation Performance for AIR traces of 275 columns

one RAP has selector and lookup columns, while all others only have accumulator columns. Constraints can span across different RAPs to enforce the selector constraints across RAPs.

3.4 Soundness Analysis

In [BGK⁺23, Sta21] it was proved that both the ethSTARK and the STARKish (called OPlonky in [BGK⁺23]) IOPs are round-by-round knowledge sound, and, consequently, special sound [BGTZ23]. These are strong soundness properties that imply knowledge soundness after compiling the IOP via Merkle tree vector commitments and the Fiat-Shamir transformation. As a consequence, the ethSTARK and Starkish protocols, in their final

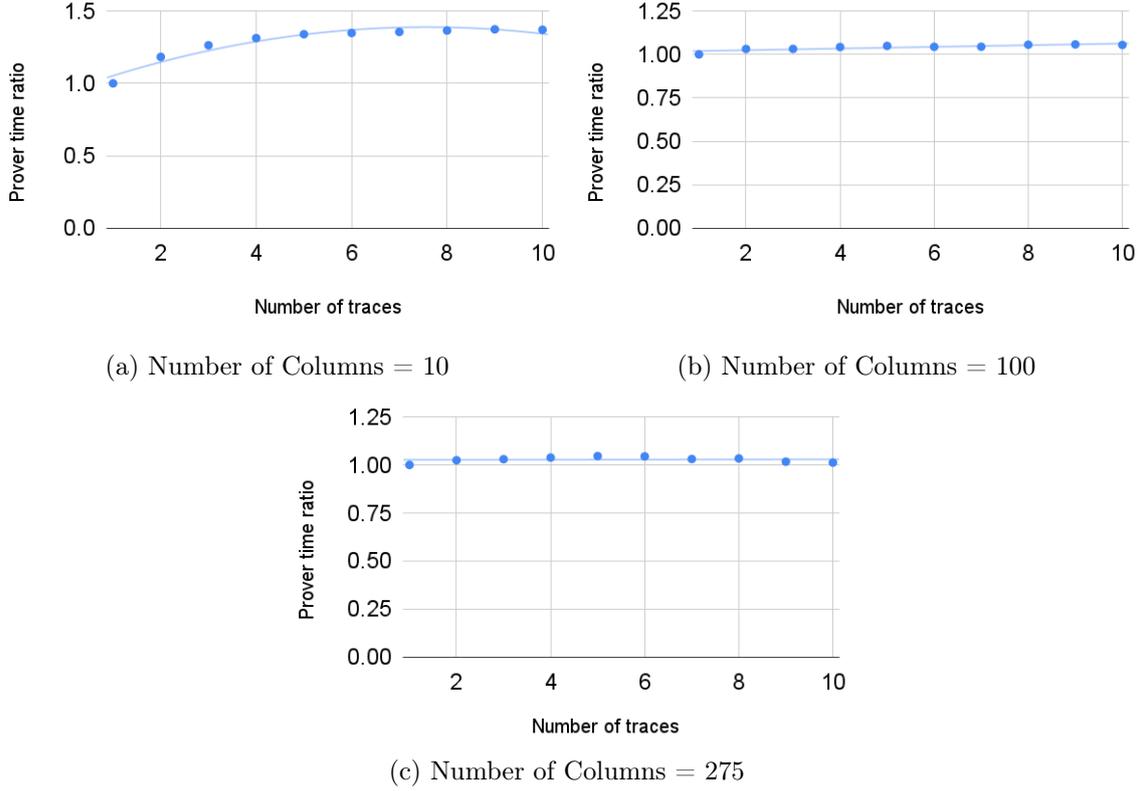


Figure 8: Winterfell: Proving Time Optimizations for various AIR traces

non-interactive and succinct form, are knowledge sound.

We argue that the same properties hold for the STARKPack IOP $\Pi_{\text{STARKPack}}$ and its Merkle-tree+Fiat-Shamir transformed final form. To show this, we first prove the following result, which states, essentially, that an instance of the Packed Plonkish relation is equivalent to an instance of the standard Plonkish relation. We use this fact to later derive the knowledge soundness of the STARKpack IOP from the knowledge soundness of STARKish IOP.

Lemma 3.3. *Let $N \geq 1$ and let*

$$\text{PACK} = (N, \text{PL}_1, \dots, \text{PL}_N, \mathbb{x}_1, \dots, \mathbb{x}_N) \quad (3)$$

be an instance of the Packed Plonkish relation $\mathbf{R}_{\text{PPlonkish}}$, with

$$\text{PL}_i = (\mathcal{P}_i, \mathcal{Q}_i, \sigma_i, \text{Pl}_i, r_i, r'_i, \ell_i, n_i, \mathbb{F}, \mathbb{K})$$

for $i \in [N]$. Then there exists a (deterministically and polynomially computable) instance $(\text{PL}_{\text{PACK}}, \mathbb{x}_{\text{PACK}})$ for the Plonkish relation $\mathbf{R}_{\text{Plonkish}}$, with

$$\text{PL}_{\text{PACK}} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K}),$$

with the following properties: 1) $r = \sum_{i \in [N]} r_i$; 2) there is an efficient deterministic algorithm that, given a valid witness $(\mathbb{w}_1, \dots, \mathbb{w}_N) \in \prod_{i \in [N]} \mathbb{F}^{n_i r_i}$ for PACK , outputs a valid witness \mathbb{w} for PL_{PACK} .

Proof. Assume first that $r'_i = r_i$ for all $i \in [N]$ (i.e. there are no advice wires) and $n_i = n_j$ for all $i, j \in [N]$. The idea in this case is that the Packed Plonkish instance is equivalent to a Plonkish instance where N execution traces are “glued together horizontally”, and the permutations in each of the Plonkish instances are combined into a single permutation of

the glued trace. Formally, we define $\text{PL}_{\text{PACK}} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K})$ as follows. First, for each $i \in [N]$, let

$$\mathbf{X}_i := (X_{i1}, \dots, X_{ir_i}), \mathbf{Y}_i := (Y_{i1}, \dots, Y_{ir_i}), \mathbf{S}_i := (S_{i1}, \dots, S_{i\ell_i})$$

be tuples of variables. Intuitively, $\mathbf{X}_i, \mathbf{Y}_i$ are variables reserved for the values of the " i -th chunk" of r_i columns of the glued execution trace, and \mathbf{S}_i are variables that will take the values in the selector vectors \mathcal{Q}_i .

Next we let $r = \sum_{i \in [N]} r_i$, $n = n_i$ for all i (as we are assuming all n_i are the same), $r' = r$, and $\ell := \sum_{i \in [N]} \ell_i$. Further, we define \mathcal{P} to be the set of polynomials

$$\mathcal{P} := \{P(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{S}_i) \mid i \in [N], P \in \mathcal{P}_i\}.$$

Formally, we view each of these as a polynomial on the $2r + \ell$ variables

$$(\mathbf{X}_1, \dots, \mathbf{X}_N, \mathbf{Y}_1, \dots, \mathbf{Y}_N, \mathbf{S}_1, \dots, \mathbf{S}_N).$$

We further let

$$\mathcal{Q} := (\mathcal{Q}_1, \dots, \mathcal{Q}_N).$$

To make the upcoming definitions clearer, we identify the sets $[n] \times [r_i]$ and $[nr_i]$ ($i \in [N]$) via the bijection $(j-1) \cdot n + k \mapsto (j, k+1)$, for $j \geq 1$ and $0 \leq k < n$. Further, we identify $[n] \times [r]$ and $[nr]$ via an analogous bijection. For the rest of the proof, we look at all sets $[nr_i]$ (and similarly for $[nr]$) as $[n] \times [r_i]$, unless otherwise stated.

Now, we let

$$\text{Pl} := (\text{Pl}_1^{r_1}, \text{Pl}_2^{r_1, r_2}, \dots, \text{Pl}_N^{r_1, \dots, r_N}),$$

where $\text{Pl}_1^{r_1} := \text{Pl}$ and, for $i > 1$, $\text{Pl}_i^{r_1, \dots, r_i}$ denotes Pl_i after translating the first component of the indices in Pl_i by $\sum_{j < i} r_j$. Further, we define the input \mathbf{x} for PL_{PACK} to be $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$.

Next, we define $\sigma : [n] \times [r] \rightarrow [n] \times [r]$ as the result of composing the permutations $\sigma_1, \dots, \sigma_N$ in the following way:

$$\sigma := \sigma_1^{\text{Ext}} \dots \sigma_N^{\text{Ext}} \quad (4)$$

where $\sigma_i^{\text{Ext}} : [n] \times [r] \rightarrow [n] \times [r]$ is the permutation that acts as σ_i on the set $[n] \times \{1 + \sum_{j < i} r_j, \dots, \sum_{j < (i+1)} r_j\}$ (which is in bijective correspondence with $[n] \times [r_i]$), and fixes the rest of the elements in $[n] \times [r]$. In (4), concatenation stands for permutation composition.

Next, we describe how to, given a valid witness $(\mathbf{w}_1, \dots, \mathbf{w}_N)$ for PACK , construct a witness \mathbf{w} for PL_{PACK} . Following our previous convention, we look at each \mathbf{w}_i as a $n \times r_i$ matrix, and let \mathbf{w} be the $n \times r$ matrix obtained by concatenating the matrices \mathbf{w}_i . Note that, by construction, \mathbf{w} satisfies the copy constraints of PL_{PACK} . Further, since $\mathbf{w}_i^{\text{Pl}_i} = \mathbf{x}_i$ for all $i \in [N]$, we have $\mathbf{w}^{\text{Pl}} = \mathbf{x}$. Further, \mathbf{w} satisfies the circuit constraints of PL_{PACK} , since, given $i \in [N]$ and

$$P = P(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = P(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i) = P(X_{i1}, \dots, X_{ir_i}; Y_{i1}, \dots, Y_{ir_i}; S_{i1}, \dots, S_{i\ell_i}) \in \mathcal{P}_i,$$

we have that for all row index $j \in [n]$ (below, we look at \mathbf{w} again as a vector from \mathbb{F}^{nr} , rather than a matrix from $\mathbb{F}^{n \times r}$),

$$\begin{aligned} &P(\mathbf{w}_{(j-1)r+1}, \dots, \mathbf{w}_{(j-1)r+r-1}; \mathbf{w}_{jr+1}, \dots, \mathbf{w}_{jr+r-1}; (\mathbf{q}[j] \mid \mathbf{q} \in \mathcal{Q})) \\ &= P(\mathbf{w}_{(j-1)r+i+1}, \dots, \mathbf{w}_{(j-1)r+r_i-1}; \mathbf{w}_{jr+i+1}, \dots, \mathbf{w}_{jr+r_i-1}; (\mathbf{q}[j] \mid \mathbf{q} \in \mathcal{Q}_i)). \end{aligned}$$

This expression is zero by construction of \mathbf{w} .

This completes the proof of the first part of the lemma under the assumption that $r'_i = r_i$ for all $i \in [N]$ and $n_i = n_j$ for all $i, j \in [N]$. Assume now that not all n_i 's are the same. The key claim we make in this scenario is that for any Plonkish instance of the form $\text{PL} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K})$, and any power of two $\tilde{n} > n$, there is a Plonkish instance

of the form $\widetilde{\text{PL}} = (\mathcal{P}, \widetilde{\mathcal{Q}}, \widetilde{\sigma}, \text{Pl}, r, r', \ell, \widetilde{\mathbf{n}}, \mathbb{F}, \mathbb{K})$ an efficient algorithm that transforms valid witness for PL to valid witnesses for $\widetilde{\text{PL}}$, and viceversa. Once this claim is proved, we can assume without loss of generality that $n_i = \max_{j \in [N]} n_j$ for all $i \in [N]$, by replacing PL_i with $\widetilde{\text{PL}}_i$ for all $i \in [N]$ such that $n_i < \max_{j \in [N]} n_j$.

Finally, the case where $r'_i \neq r_i$ for some $i \in [N]$ can be dealt with with analogous arguments as the ones given so far. In a bit more detail, the proof proceeds in the same way, except that, formally speaking, we cannot have \mathbf{w} just be the concatenation of the witness matrices $\mathbf{w}_1, \dots, \mathbf{w}_N$. This is because, again formally, the witness (or execution trace) of a Plonkish relation has all its advice wires (i.e. columns) on the rightmost part of the trace. Accordingly, in the presence of advice wires, we construct \mathbf{w} by concatenating $\mathbf{w}_1, \dots, \mathbf{w}_N$ and then moving the advice wires of each submatrix \mathbf{w}_i to the right-most columns of \mathbf{w} . Then, $\text{Pl}, \mathbf{x}, \sigma, \mathbf{X}, \mathbf{Y}, \mathbf{Z}$ need to be modified in order to accomodate this rearrangement of the columns in \mathbf{w} . \square

Now Lemma 3.3 and [BGK⁺23, Sta21] yield:

Corollary 3.4. *The STARKPack IOP $\Pi_{\text{STARKPack}}$ has knowledge soundness error, given an instance PACK as in (3),*

$$\varepsilon_{\text{ks}} = \frac{1}{2\eta\sqrt{\rho}} \cdot \max \left\{ \left(\frac{3n(r' + u)}{|\mathbb{F}|} \right), \left(\frac{|\mathcal{P}| + s + 1}{|\mathbb{F}|} \right), \frac{n \cdot \max\{u + 1, d_{\max}\}}{|\mathbb{K} \setminus \mathbf{D}|}, 2\eta\sqrt{\rho} \cdot \varepsilon_{\text{FRI}} \right\},$$

where $n = \max_{i \in [N]} n_i$, \mathbf{D} is a subgroup of size n , $r' = \sum_{i \in [N]} r'_i$, $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_N)$, $d_{\max} = \max_{P \in \mathcal{P}} \deg(P)$, u is a parameter used during the permutation argument of the STARKish IOP (analogous to the splitting parameters u_i in Section 3.1), $s = \lceil r'/u \rceil$, $\rho = (d_{\max} + 1)/n$, η is a parameter used in the batched FRI protocol (cf. [BGK⁺23]), and ε_{FRI} is the round-by-round soundness error of the batched FRI protocol used at the end of the STARKish IOP.

Furthermore, the Succinct Non-interactive Argument of Knowledge obtained by compiling the $\Pi_{\text{STARKPack}}$ IOP with Merkle tree commitments and the Fiat-Shamir transform is knowledge sound with error

$$Q \cdot \varepsilon_{\text{ks}} + O(Q^2/2^\kappa),$$

where Q is the number of random oracle queries with κ -bits of output that an adversary is capable of making.

Proof. Let P_{PACK}^* be a malicious prover for the StarkPack IOP (respectively, P_{PACK}^* is a malicious prover for the SNARK compiled from the StarkPack IOP using Merkle trees and the FS transform, which we call compiled StarkPack). We build a prover P_0^* for the STARKish IOP (resp. compiled STARKish) as follows. The prover P_0^* only accepts Plonkish instances of the form PL_{PACK} . In its execution, it first recovers the packed instance PACK, and then it runs an extractor Ext_{PACK} for the Starkpack IOP (resp. compiled IOP) on the instance PACK and the prover P_{PACK}^* . If the extractor outputs a valid witness $(\mathbf{w}_1, \dots, \mathbf{w}_N)$ for the instance PACK, then P_0^* uses Lemma 3.3 to efficiently construct a valid witness \mathbf{w} for PL_{PACK} . This shows that the knowledge error of the StarkPack IOP (resp. compiled SNARK) is at most the knowledge error of the STARKish IOP (resp. compiled IOP) for instances of the form PL_{PACK} . Now the corollary follows by applying the results from [BGK⁺23] in order to bound the knowledge error of the STARKish IOP (resp. compiled IOP) for an instance of the form PL_{PACK} . \square

4 Modular split-and-pack

The idea of splitting is quite simple, in essence it is about breaking execution traces into smaller pieces. We describe any type of splitting first, and then move on to modular splitting. When interpolating and evaluating the relevant polynomials, the prover performs more number theoretic transforms but over smaller domains. As we will see this is cheaper overall.

The cost of committing to these polynomials is amortized by batching the evaluations at the same point in the same Merkle leaf. There will be less leaves for committing to witness polynomials, but they will be bigger. Computing the combined constraint for the splitting (the polynomial we have called `quot`) is comparable—if not cheaper—than when there is no splitting. Finally, the cost of FRI will be asymptotically comparable (and even a bit cheaper) to when there is no splitting. We expect that the splitting operation therefore leads to an improvement in the proof time. The trade-off is at the level of the proof size (as the Prover is forced to send more evaluations) and the Verifier time (as the Verifier now needs to check more evaluations and hash bigger leaves).

Formalizing the splitting operation can be cumbersome in its full generality. In this paper, we focus on a simpler and natural case where we split witnesses into chunks with the same column size, by selecting entire rows out of the original witness. Let $\text{PL} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K})$ be a Plonkish instance. Choose a divisor k of n . Instead of considering a witness $w \in \mathbb{F}^{nr}$ for PL which we consider as a matrix whose (i, j) -th entry is $w[(i-1)r + j]$, we split it into vectors $w^0, \dots, w^{k-1} \in \mathbb{F}^{(n/k) \times r}$ by selecting rows out of w . We consider these vectors as matrices in the same way. It should hold that for any $(i, j) \in [n] \times [r]$ there are unique $(s, i') \in \{0, \dots, k-1\} \times [n/k]$ such that:

$$w[(i-1)r + j] = w^s[(i'-1)r + j] \quad (5)$$

We should still ask that the constraints in \mathcal{P} hold when evaluated at any of the corresponding entries in the newly created vectors w^s . More specifically, for any $(i, j) \in [n] \times [r]$, let $(s(i), i'(i))$ be the indices such that (5) holds, emphasizing the dependence on i . The condition:

$$\forall P \in \mathcal{P}, \forall i \in [n-1], P(w[(i-1)r + 1], \dots, w[ir], \dots, w[(i+1)r], q_1[i], \dots, q_\ell[i]) = 0$$

must translate into the condition:

$$\forall P \in \mathcal{P}, \forall i \in [n-1], P(w^{s(i-1)}[i'(i-1)r + 1], \dots, w^{s(i-1)}[i'(i-1)r + r], \dots, w^{s(i)}[i'(i)r + r], q_1[i], \dots, q_\ell[i]) = 0$$

And a similar condition must hold for copy constraints. The way this translates into polynomial constraints is hard to systematize when the splitting is general. We will see that in some special cases, like when splitting over cosets of a multiplicative subgroup, the constraints can be expressed very naturally and efficiently. Note from the previous equation that the constraints may now involve much more than two consecutive rows of the same w^s , and may even involve values from different vectors w^s . For this reason, after we formally define what a splitting of a Plonkish instance is, we need to define a new *split Plonkish relation*. We do this next.

4.1 The split Plonkish relation

We can think informally of the w^s as buckets in which we place rows of the original witness w successively without replacement. Once we have chosen n/k rows of each w^s , we are done. A natural way to formalize this is by using a permutation $\mathfrak{S} : \{0, \dots, n-1\} \rightarrow \{0, \dots, n-1\}$. The first⁵ row of w^0 will be the row with index $\mathfrak{S}^{-1}(0)$ in w , and so on. We formalize this next.

Definition 4.1 (Splitting of a Plonkish instance). Let $\text{PL} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K})$ be a Plonkish instance. A splitting of PL is the additional data of a strict divisor k of n , together with a permutation $\mathfrak{S} : \{0, \dots, n-1\} \rightarrow \{0, \dots, n-1\}$.

⁵Note that the entries in the first row of w , that is, entries with the form $(1, j)$ are given by $w[j]$. This means that rows are indexed from 0.

We also define the following auxiliary functions $\mathfrak{T} : \{0, \dots, n-1\} \rightarrow \{0, \dots, k-1\}$ and $\gamma : \{0, \dots, n-1\} \rightarrow \{0, \dots, n/k-1\}$ by the following relations:

$$\begin{aligned}\mathfrak{T}(i) &= \left\lfloor \frac{k\mathfrak{S}(i)}{n} \right\rfloor \\ \gamma(i) &= \mathfrak{S}(i) - \mathfrak{T}(i)\frac{n}{k}\end{aligned}$$

Informally, $\mathfrak{T}(i)$ tells us the s such that row i in \mathbf{w} will be sent to \mathbf{w}^s , and $\gamma(i)$ tells us the index that row will have in \mathbf{w}^s .

We now define the split Plonkish relation.

Definition 4.2 (Split Plonkish Relation). A split Plonkish instance

$$\text{PL}_{\text{split}} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, k, \mathfrak{S}, \mathbb{F}, \mathbb{K})$$

is a Plonkish instance with the additional data of a splitting. Recall from Definition 4.1 that this is given by a strict divisor k of n and a permutation $\mathfrak{S} : \{0, \dots, n-1\} \rightarrow \{0, \dots, n-1\}$.

A set of execution traces for PL_{split} is a set of k vectors $\mathbf{w}^0, \dots, \mathbf{w}^{k-1} \in \mathbb{F}^{(n/k) \times r}$. It will be helpful to consider these as matrices whose entry $(i, j) \in [n/k] \times [r]$ is given by $\mathbf{w}^s[(i-1)r+j]$ for $s = 0, \dots, k-1$. Such a set of traces is said to satisfy PL_{split} if the following conditions hold:

1. (Circuit constraint satisfaction). For all $P \in \mathcal{P}$, and all $i \in [n-1]$,

$$P(\mathbf{w}^{\mathfrak{T}(i-1)}[\gamma(i-1)r+1], \dots, \mathbf{w}^{\mathfrak{T}(i-1)}[\gamma(i-1)r+r], \dots, \mathbf{w}^{\mathfrak{T}(i)}[\gamma(i)r+r], \mathbf{q}_1[i], \dots, \mathbf{q}_\ell[i]) = 0$$

where \mathfrak{T} and γ are defined as in Definition 4.1.

2. (Copy constraint satisfaction). Let $S \subset [r]$ be the subset of column indices subject to a copy constraint. By definition $|S| = r'$. For all $(i, j) \in [n] \times S$, let $(\zeta(i, j), \xi(i, j)) \in [n] \times S$ be such that $\sigma((i-1)r+j) = (\zeta(i, j)-1)r + \xi(i, j)$. Then the following should hold for all $(i, j) \in [n] \times S$:

$$\mathbf{w}^{\mathfrak{T}(i-1)}[(i-1)r+j] = \mathbf{w}^{\mathfrak{T}(\zeta(i,j)-1)}[(\zeta(i,j)-1)r + \xi(i, j)]$$

We denote by $(\mathbf{w}^0, \dots, \mathbf{w}^{k-1})^{\text{Pl}}$ the vector formed by the entries $\{\mathbf{w}^{\mathfrak{T}(i-1)}[\gamma(i-1)r+j] \mid (i-1)r+j \in \text{Pl}\}$. We call this vector the public part of $\mathbf{w}^0, \dots, \mathbf{w}^{k-1}$. The split Plonkish relation is then defined as:

$$\mathbf{R}_{\text{SPlonkish}} := \left\{ (\text{PL}_{\text{split}}, \mathbf{x}, \mathbf{w}^0, \dots, \mathbf{w}^{k-1}) \left| \begin{array}{l} \text{PL}_{\text{split}} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, k, \mathfrak{S}, \mathbb{F}, \mathbb{K}), \\ \mathbf{x} \in \mathbb{F}^{|\text{Pl}|}, \\ \mathbf{x} = (\mathbf{w}^0, \dots, \mathbf{w}^{k-1})^{\text{Pl}}, \\ \mathbf{w}^0, \dots, \mathbf{w}^{k-1} \in \mathbb{F}^{(n/k) \times r} \text{ satisfy } \text{PL}_{\text{split}} \end{array} \right. \right\}$$

Remark 4.3. We have explained what a witness for the split Plonkish relation is, and how it satisfies a split instance. However, we have not explained how to turn the conditions in Point 1 and 2 in Definition 4.2 into polynomial constraints. As we mentioned, for circuit constraint satisfaction this might depend on the particular splitting. There are cases in which it translates seamlessly into polynomial constraints (see the next Example 4.4). However, for copy constraints there is a natural way to express the permutation argument, which we explain next.

For simplicity, assume that all wires are routable (i.e. $r = r'$). We first want to find a permutation $\tilde{\sigma} : \{0, \dots, k-1\} \times \{0, \dots, n/k-1\} \times [r] \rightarrow \{0, \dots, k-1\} \times \{0, \dots, n/k-1\} \times [r]$ that expresses the action of σ on the $\mathbf{w}^s[(i'-1)r+j']$. Let us temporarily use the notation

$E = \{0, \dots, k-1\} \times \{0, \dots, n/k-1\} \times [r]$ and $F = \{0, \dots, n-1\} \times [r]$. The permutation $\tilde{\sigma}$ is given by the composition of the following three maps:

$$\begin{aligned} E &\longrightarrow F \\ (s, i', j') &\longmapsto (\mathfrak{S}^{-1}(i' + s\frac{n}{k}), j') \\ \\ F &\xrightarrow{\sigma} F \\ (i, j) &\longmapsto (\zeta(i+1, j) - 1, \xi(i+1, j)) \\ \\ F &\longrightarrow E \\ (i, j) &\longmapsto (\tau(i), \gamma(i), j) \end{aligned}$$

where ζ, ξ are defined as in Definition 4.2. All in all, this means that for all $(s, i', j') \in E$:
 $\tilde{\sigma}(s, i', j') = (\tau(\zeta(\mathfrak{S}^{-1}(i' + s\frac{n}{k}) + 1, j') - 1), \gamma(\zeta(\mathfrak{S}^{-1}(i' + s\frac{n}{k}) + 1, j') - 1), \xi(\mathfrak{S}^{-1}(i' + s\frac{n}{k}) + 1, j'))$

Define a vector $W \in \mathbb{F}^{nr}$ by setting, for all $(s, i, j) \in \{0, \dots, k-1\} \times [n/k] \times [r]$:

$$W[(k(i-1) + s)r + j] = w^s[(i-1)r + j]$$

The vector W is constructed by putting the matrices w^s side-by-side. We can make the permutation $\tilde{\sigma}$ act on $[nr]$ in the same way we can make the permutation σ act on $[nr]$. For all $(s, i, j) \in \{0, \dots, k-1\} \times [n/k] \times [r]$, if $(\bar{s}, \bar{i}, \bar{j}) = \tilde{\sigma}(s, i-1, j)$, then

$$\tilde{\sigma}((k(i-1) + s)r + j) = (k\bar{i} + \bar{s})r + \bar{j}$$

One checks that with these definitions, Point 2 in in Definition 4.2 corresponds precisely to the requirement that for all $i \in [nr]$:

$$W_{\tilde{\sigma}(i)} = W_i$$

So now define the polynomials $S_{ID_j}^s$ and $(S_{\tilde{\sigma}}^s)_j$ for all $s \in \{0, \dots, k-1\}$ and $j \in [r]$ by the following:

$$\begin{aligned} \forall i \in \{0, \dots, n/k-1\}, S_{ID_j}^s((g^k)^i) &:= (k(i-1) + s)r + j \\ (S_{\tilde{\sigma}}^s)_j((g^k)^i) &:= \tilde{\sigma}((k(i-1) + s)r + j) \end{aligned}$$

If we set \mathbf{a}_j^s to be the polynomial such that for all $i \in \{0, \dots, n/k-1\}$, $\mathbf{a}_j^s((g^k)^i) = w^s[(i-1)r + j]$, then in order to check that copy constraints hold it suffices to exhibit a polynomial $z(X) \in \mathbb{F}_{<n/k}[X]$ such that:

$$z(g^k X) \cdot f(X) = g(X) \cdot z(X)$$

where:

$$\begin{aligned} f(X) &:= \prod_{s=0}^{k-1} \prod_{j=1}^r (\mathbf{a}_j^s(X) + \delta \cdot S_{ID_j}^s(X) + \eta) \\ g(X) &:= \prod_{s=0}^{k-1} \prod_{j=1}^r (\mathbf{a}_j^s(X) + \delta \cdot (S_{\tilde{\sigma}}^s)_j(X) + \eta) \end{aligned}$$

and δ, η are randomly chosen by the verifier. This can also be split into constraints of smaller degree as we have seen previously. Note that the degree in the variable X of the polynomials f and g is smaller or equal to $r(n-k)$ (a product of rk polynomials of degree smaller or equal to $n/k-1$). However, the degree in δ and η (considered as formal variables) increases with respect to the permutation argument without splitting, which results in a small soundness loss.

The computation of $\tilde{\sigma}$ and the polynomials $S_{ID_j}^s, (S_{\tilde{\sigma}}^s)_j$ will be done in the preprocessing phase.

Example 4.4 (Splitting along cosets modulo k). We can perform a splitting by selecting rows depending on the residue of their index modulo k . Choose k a strict divisor of n , and set $\mathfrak{S}(i) = \lfloor i/k \rfloor + n/k \cdot (i \bmod k)$. One checks that \mathfrak{S} is a permutation, and that:

$$\begin{aligned}\mathfrak{T}(i) &:= \left\lfloor \frac{k\mathfrak{S}(i)}{n} \right\rfloor = (i \bmod k) \\ \gamma(i) &:= \mathfrak{S}(i) - \mathfrak{T}(i) \frac{n}{k} = \lfloor i/k \rfloor\end{aligned}$$

Here, the circuit constraints can be translated easily to polynomial constraints. Let $((\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, k, \mathfrak{S}, \mathbb{F}, \mathbb{K}), \mathbf{x})$ be a split Plonkish instance where \mathfrak{S} is given by the permutation we defined above, and let w^0, \dots, w^{k-1} be a set of execution traces for it. For simplicity focus on one circuit constraint P . We ask that for all $i \in [n-1]$,

$$P(w^{i-1} [k] [\lfloor (i-1)/k \rfloor r + 1], \dots, w^{i [k]} [\lfloor i/k \rfloor r + r], \mathbf{q}_1[i], \dots, \mathbf{q}_\ell[i]) = 0$$

where we have used the notation $i [k] = (i \bmod k)$. We notice that the constraints always involve w^s and w^{s+1} , unless $s = k-1$ in which case it involves w^{k-1} and w^0 . Further, the constraints always involve the same rows of w^s and w^{s+1} if $(i-1) [k] = s < k-1$, namely the row with index $\lfloor (i-1)/k \rfloor = \lfloor i/k \rfloor$. And if $(i-1) [k] = s = k-1$, the constraint involves row $\lfloor (i-1)/k \rfloor$ of w^{k-1} and row $\lfloor i/k \rfloor = \lfloor (i-1)/k \rfloor + 1$ of w^0 . This means that we can translate the circuit constraint into:

$$\forall s \in \{0, \dots, k-2\}, \forall i \in \{0, \dots, n/k-1\},$$

$$P(w^s [ir + 1], \dots, w^{s+1} [ir + r], \mathbf{q}_1[\mathfrak{S}^{-1}(i + n(s+1)/k)], \dots, \mathbf{q}_\ell[\mathfrak{S}^{-1}(i + n(s+1)/k)]) = 0$$

$$\forall i \in \{0, \dots, n/k-2\},$$

$$P(w^{k-1} [ir + 1], \dots, w^0 [(i+1)r + r], \mathbf{q}_1[\mathfrak{S}^{-1}(i+1)], \dots, \mathbf{q}_\ell[\mathfrak{S}^{-1}(i+1)]) = 0$$

Set \mathbf{a}_j^s to be the polynomials such that $\mathbf{a}_j^s((\mathbf{g}^k)^i) = w^s[(i-1)r + j]$ for all $(s, i, j) \in \{0, \dots, k-1\} \times \{0, \dots, n/k-1\} \times [r]$. For $0 \leq s < k-1$ and $1 \leq l \leq \ell$ set \mathbf{q}_l^s to be the polynomial such that $\mathbf{q}_l^s((\mathbf{g}^k)^i) = \mathbf{q}_l[\mathfrak{S}^{-1}(i + n(s+1)/k)]$, and set $\mathbf{q}_l^{k-1}((\mathbf{g}^k)^i) = \mathbf{q}_l[\mathfrak{S}^{-1}(i+1)]$. Then in order to check the circuit constraint satisfaction, we may check that all polynomials:

$$\begin{aligned}P(\mathbf{a}_1^0(X), \dots, \mathbf{a}_r^0(X), \dots, \mathbf{a}_r^1(X), \mathbf{q}_1^0(X), \dots, \mathbf{q}_\ell^0(X)) \\ \vdots \\ P(\mathbf{a}_1^{k-1}(X), \dots, \mathbf{a}_r^{k-1}(X), \dots, \mathbf{a}_r^0(\mathbf{g}^k X), \mathbf{q}_1^{k-1}(X), \dots, \mathbf{q}_\ell^{k-1}(X))\end{aligned}$$

are divisible by the vanishing polynomial of $\mathbf{D}^k = \langle \mathbf{g}^k \rangle$. Note that the degree of these polynomial constraints is less than $n \deg(P)/k - n/k$. This is k times less than the degree of the circuit constraints had we not performed the splitting. This may look like it increases the soundness of the protocol, but the fact that we now have k times more constraints will balance out the gain from having smaller degree constraints. The newly created selector polynomials should also be computed during preprocessing.

It is clear that a similar phenomenon happens for any splitting which has the property that the circuit constraint involves the same pair of rows of the same k pairs of traces w^0, \dots, w^{k-1} . The splitting modulo k happens to be a splitting with this property.

4.2 An IOP for split-and-pack modulo k

In the previous section, we have seen that copy constraints over the trace evaluation domain \mathbf{D} can always be expressed as a polynomial permutation argument over the domain \mathbf{D}^k , provided that we perform the computation of the induced permutation $\tilde{\sigma}$ and the polynomials $S_{ID_j}^s, (S_{\tilde{\sigma}}^s)_j$ correctly (see Remark 4.3). We have also seen that when splitting along cosets

modulo k (see Example 4.4), circuit constraints over the trace evaluation domain \mathbb{D} can be expressed efficiently as k polynomial constraints over the domain \mathbb{D}^k , provided that we perform the computation of the selectors $\mathbf{q}_1^s, \dots, \mathbf{q}_\ell^s$ correctly. In this section, we present an IOP for proving the satisfiability of a witness to a split Plonkish instance where the splitting is along cosets modulo k , which we denote by Π_{split}^k . Of course, one can extrapolate the ideas to more general splittings.

We fix a split Plonkish instance $(\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, k, \mathfrak{S}, \mathbb{F}, \mathbb{K})$, where $\mathfrak{S}(i) = \lfloor i/k \rfloor + n/k \cdot (i \bmod k)$. For simplicity, we assume that $r = r'$ (all wires are routable). The trace evaluation domain is denoted by $\mathbb{D} = \langle \mathbf{g} \rangle$, and the subgroup generated by $\mathbf{g}_k := \mathbf{g}^k$ is denoted by \mathbb{D}^k . Its vanishing polynomial is denoted by $Z_k(X) := \prod_{x \in \mathbb{D}^k} (X - x) = \prod_{i=0}^{n/k-1} (X - \mathbf{g}_k^i)$. We fix an evaluation domain $\mathbb{H} \subset \mathbb{K}^\times$, a multiplicative subgroup containing a non-trivial coset of \mathbb{D}^k . When we say that the prover gives oracle access to a function, we mean that the prover interpolates a vector with entries in \mathbb{F} over \mathbb{D}^k ; and then provides oracle access to this interpolant over \mathbb{H} . We fix a set of execution traces $\mathbf{w}^0, \dots, \mathbf{w}^{k-1} \in \mathbb{F}^{(n/k)r}$ and public part \mathbf{x} for the split Plonkish instance. Finally, we fix auxiliary parameters $\text{aux} = (\vec{t}, s)$ for the Batched-FRI protocol. The parameter $\vec{t} = (t_1, \dots, t_m) \in \mathbb{N}$ indicates that we will use t_i -to-1 maps in the i -th COMMIT phase of FRI, and s is the number of repetitions of the QUERY phase of FRI.

Preprocessing. The prover and verifier compute the induced permutation $\tilde{\sigma} : [nr] \rightarrow [nr]$ as explained in Remark 4.3. They compute the polynomials $S_{ID_j}^s, (S_{\tilde{\sigma}}^s)_j \in \mathbb{F}_{<n/k}[X]$ for all $s \in \{0, \dots, k-1\}$ and $j \in [r]$ as:

$$\begin{aligned} \forall i \in \{0, \dots, n/k-1\}, S_{ID_j}^s(\mathbf{g}_k^i) &:= (k(i-1) + s)r + j \\ (S_{\tilde{\sigma}}^s)_j(\mathbf{g}_k^i) &:= \tilde{\sigma}((k(i-1) + s)r + j) \end{aligned}$$

They also compute the selectors $\mathbf{q}_1^s, \dots, \mathbf{q}_\ell^s \in \mathbb{F}_{<n/k}[X]$ for all $s \in \{0, \dots, k-1\}$ as:

$$\begin{aligned} \forall i \in \{0, \dots, n/k-1\}, \mathbf{q}_l^s(\mathbf{g}_k^i) &:= \mathbf{q}_l[\mathfrak{S}^{-1}(i + n(s+1)/k)] \\ \mathbf{q}_l^{k-1}(\mathbf{g}_k^i) &:= \mathbf{q}_l[\mathfrak{S}^{-1}(i+1)] \end{aligned}$$

Finally, they agree on the splitting parameter u for the permutation argument.

Round 1. The prover computes the polynomials $\mathbf{a}_j^s \in \mathbb{F}_{<n/k}[X]$ such that $\mathbf{a}_j^s(\mathbf{g}_k^i) = \mathbf{w}^s[(i-1)r + j]$ for all $(s, i, j) \in \{0, \dots, k-1\} \times \{0, \dots, n/k-1\} \times [r]$. The prover gives oracle access to the wire polynomials $\mathbf{a}_j^s : \mathbb{H} \rightarrow \mathbb{K}$ to the verifier. The verifier samples uniform randomness $\delta, \eta \in \mathbb{K}$ to be used in the permutation argument.

Round 2. The prover sends oracle access to the permutation and partial product polynomials $z, (\pi_l)_{l \in [s-1]} \in \mathbb{F}_{<n/k}[X]$, where $s = \lceil kr/u \rceil$. We recall their definition quickly. First recall the definition of the f and g polynomials:

$$\begin{aligned} f(X) &:= \prod_{s=0}^{k-1} \prod_{j=1}^r (a_j^s(X) + \delta \cdot S_{ID_j}^s(X) + \eta) \\ g(X) &:= \prod_{s=0}^{k-1} \prod_{j=1}^r (a_j^s(X) + \delta \cdot (S_{\tilde{\sigma}}^s)_j(X) + \eta) \end{aligned}$$

Then $z(\mathbf{g}_k) = 1$, and for all $i \in \{2, \dots, n/k\}$:

$$z(\mathbf{g}_k^i) := \prod_{1 \leq j < i} f(\mathbf{g}_k^j) / g(\mathbf{g}_k^j)$$

We can arrange each of the kr factors appearing in the definition of f and g above in lexicographic order with respect to the indices $(s, j) \in \{0, \dots, k-1\} \times [r]$. For each $l \in [s]$, one defines the polynomials $\overline{f_l}(\delta, \eta, X) = \overline{f_l}(X)$ (resp. $\overline{g_l}(\delta, \eta, X) = \overline{g_l}(X)$) as the product of the factors indexed from $(u-1)l+1$ to $\min(ul, kr)$ appearing in the definition of f (resp. g), according to that lexicographic order. Then for all $i \in [n/k]$:

$$\begin{aligned}\pi_1(\mathbf{g}_k^i) &:= z(\mathbf{g}_k^i) \overline{f_1}(\mathbf{g}_k^i) \overline{g_1}(\mathbf{g}_k^i)^{-1} \\ \pi_l(\mathbf{g}_k^i) &:= \pi_{l-1}(\mathbf{g}_k^i) \overline{f_l}(\mathbf{g}_k^i) \overline{g_l}(\mathbf{g}_k^i)^{-1}, \text{ for } l = 2, \dots, s-1\end{aligned}$$

The verifier samples uniform randomness $\alpha \in \mathbb{K}$.

Round 3. The prover computes the polynomials $\mathbf{u}, \mathbf{d}, \mathbf{quot}$ as follows:

$$\begin{aligned}\mathbf{u}(X) &:= (z(X) \overline{f_1}(X) - \pi_1(X) \overline{g_1}(X)) \alpha \\ &\quad + \sum_{l=2}^s (\pi_{l-1}(X) \overline{f_l}(X) - \pi_l(X) \overline{g_l}(X)) \alpha^l \\ &\quad + (\pi_{s-1}(X) \overline{f_s}(X) - z(\mathbf{g}_k X) \overline{g_s}(X)) \alpha^{s+1} \\ &\quad + L_1(X) (z(X) - 1) \alpha^{s+2} \\ \mathbf{d}(X) &:= \sum_{j \in [|\mathcal{P}|]} \sum_{l=0}^{k-2} \alpha^{(j-1)k+l} P_j(\mathbf{a}_1^l(X), \dots, \mathbf{a}_r^{l+1}(X), \mathbf{q}_1^l(X), \dots, \mathbf{q}_\ell^l(X)) \\ &\quad + \sum_{j \in [|\mathcal{P}|]} \alpha^{(j-1)k+k-1} P_j(\mathbf{a}_1^{k-1}(X), \dots, \mathbf{a}_r^0(\mathbf{g}_k X), \mathbf{q}_1^{k-1}(X), \dots, \mathbf{q}_\ell^{k-1}(X)) \\ &\quad + \alpha^{(|\mathcal{P}|k-1)} \mathbf{u}(X) \\ \mathbf{quot}(X) &:= \mathbf{d}(X) / Z_k(X)\end{aligned}$$

where L_1 is the first Lagrange basis polynomial with respect to the set $\{\mathbf{g}_k, \dots, \mathbf{g}_k^{n/k}\}$. Then the prover splits each $\mathbf{quot}(X)$ into degree n/k polynomials $\mathbf{quot}_1, \dots, \mathbf{quot}_\nu$ such that $\mathbf{quot}(X) = \sum_l X^{n(l-1)/k} \mathbf{quot}_l(X)$. The prover gives oracle access to the maps $\mathbf{quot}_1, \dots, \mathbf{quot}_\nu : \mathbf{H} \rightarrow \mathbb{K}$. The verifier samples uniform randomness $\mathfrak{z} \in \mathbb{K} \setminus (\mathbf{H} \cup \mathbf{g}^{-1}\mathbf{H})$.

Round 4. The prover sends the evaluations

$$\mathbf{eval} = (\mathbf{a}_j^s(\mathfrak{z}), \mathbf{a}_j^0(\mathbf{g}_k \mathfrak{z}), \mathbf{quot}_l(\mathfrak{z}), \pi_m(\mathfrak{z}), z(\mathfrak{z}), z(\mathbf{g}_k \mathfrak{z}), \mathbf{q}_c^s(\mathfrak{z}), \mathbf{S}_{\text{ID}_j}^s(\mathfrak{z}), (\mathbf{S}_\sigma^s)_j(\mathfrak{z}))$$

for all $s \in \{0, \dots, k-1\}, j \in [r], l \in [\nu], m \in [s-1], c \in [\ell]$. The verifier samples uniform FRI randomness $\xi \in \mathbb{K}$.

Batched-FRI. The prover and verifier engage in a Batched-FRI protocol with auxiliary parameters \mathbf{aux} to check proximity to the Reed-Solomon code $\text{RS}(\mathbb{K}, \mathbf{H}, n)$ of the function defined as follows. Consider the vector of functions defined as:

$$\left(\frac{\mathbf{a}_j^s(X) - \mathbf{a}_j^s(\mathfrak{z})}{X - \mathfrak{z}}, \frac{\mathbf{a}_j^0(X) - \mathbf{a}_j^0(\mathbf{g}_k \mathfrak{z})}{X - \mathbf{g}_k \mathfrak{z}}, \frac{\mathbf{quot}_l(X) - \mathbf{quot}_l(\mathfrak{z})}{X - \mathfrak{z}}, \dots, \frac{(\mathbf{S}_\sigma^s)_j(X) - (\mathbf{S}_\sigma^s)_j(\mathfrak{z})}{X - \mathfrak{z}} \right) \quad (6)$$

for all $s \in \{0, \dots, k-1\}, j \in [r], l \in [\nu], m \in [s-1], c \in [\ell]$ (it contains all relevant quotients of functions whose evaluations appear in \mathbf{eval}). For simplicity, relabel this vector as $\vec{F} = (F_1(X), F_2(X), \dots, F_M(X))$ (for some integer M). The function we apply FRI to is the function $\sum_{i=1}^M \xi^{i-1} F_i(X)$.

Decision.

1. **Circuit and copy constraint check.** The verifier uses `eval` to verify that $d(\mathfrak{z}) = \text{quot}(\mathfrak{z})Z_k(\mathfrak{z})$. If this check fails, the verifier rejects.
2. **Batched FRI check.** If the prover fails the Batched-FRI protocol with auxiliary parameters `aux`, the verifier rejects.

If all checks pass, the verifier accepts the proof.

Remark 4.5. Our remark about Merkle commitments still applies in this case, where we can pack the evaluations of all relevant polynomials at the same evaluation point at each commitment step.

Remark 4.6 (Adding `proofup` support for splitting). For the kind of splittings we have been considering, i.e., those where we split the execution trace by selecting entire rows, one can see that a lookup argument involving a column of the execution trace can be expressed equivalently as k lookup arguments involving the same column of a set of execution traces for a splitting. This is due to the fact that if $f, t \in \mathbb{F}^n$, and f_0, \dots, f_{k-1} is a partition of f , then $\{f\} \subset \{t\} \Leftrightarrow (\{f_0\} \subset \{t\}) \wedge \dots \wedge (\{f_{k-1}\} \subset \{t\})$.

This translates into polynomial constraints that we include in the `d` polynomial in Round 3. of the IOP. Even though there are k times as much polynomial constraints, their degree is divided at least by k (because all of the polynomials involved are of smaller degree except the t polynomial, but the vanishing polynomials Z of D^k are now of degree n/k).

This still holds for more general splittings, however in this case we might be forced to perform more than k permutation arguments.

4.3 Performance

Benchmarks. To evaluate the performance implications of splitting, we conducted benchmarks on the Winterfell library. Our analysis focuses on the effects of splitting the trace into more columns and fewer rows via the `mod k` approach. Figures 9-16 illustrate the performance trends observed for a trace with an increasing number of columns. Both Prover time and RAM peak usage decrease upon applying splitting techniques. However, the proof size and Verifier time exhibit initial reductions before increasing linearly as the number of splits grows.

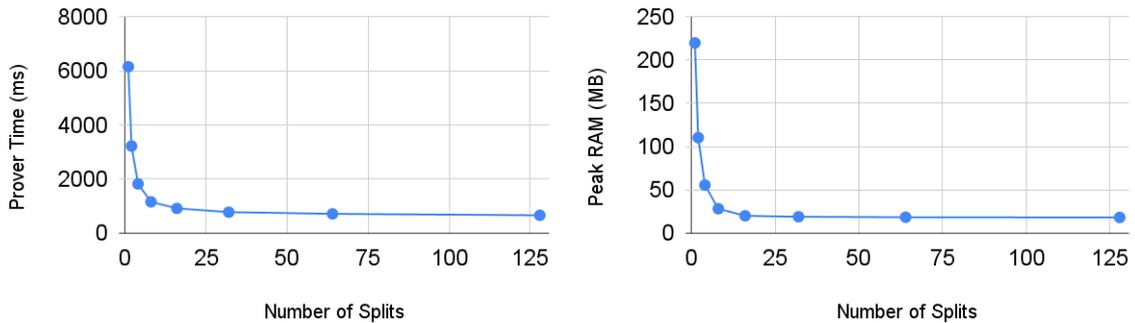


Figure 9: Prover time and Peak RAM splitting performance for 1 column

These findings suggest that splitting offers substantial gains in Prover time and RAM peak usage, particularly for scenarios involving a limited number of columns and a reasonable number of splittings. The improvement is less sharp whenever any of these two parameters is large.

The proof sizes and Verifier times are asymptotically linear in the number of chunks the original trace is split into. The Verifier time is still less than a few hundred ms at worst, however the proof gets large. This is especially true as the number of columns and

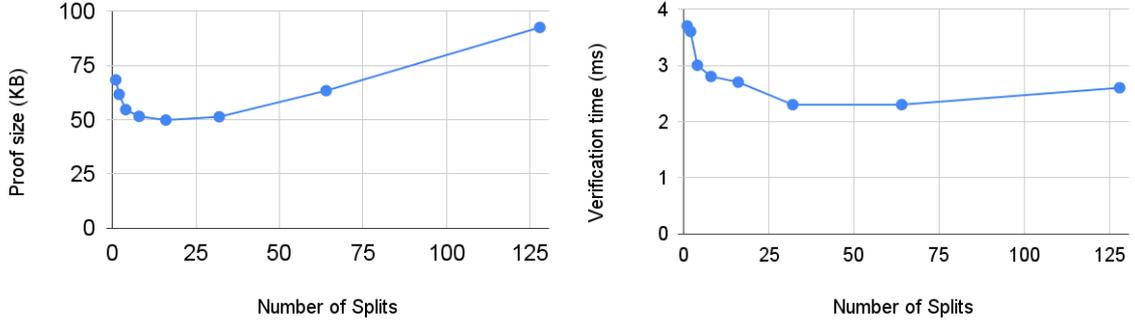


Figure 10: Proof size and Verifying time splitting performance for 1 column

split chunks gets large (the proof is around $10000kB$ for 275 columns and 64 chunks). This means that one should limit the splitting into a small number of chunks before the proof gets prohibitively large. However, because the Prover time decreases so sharply before attaining a sort of plateau, it means that we can get significant Prover time improvements without making the proof too large, or the verifier too slow.

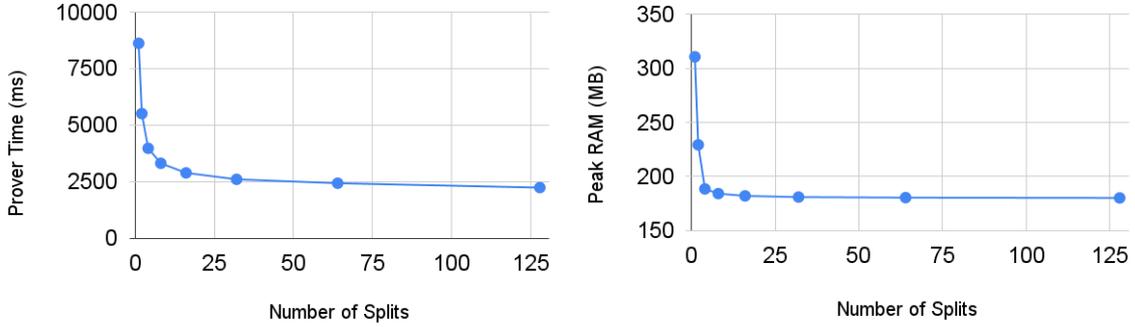


Figure 11: Prover time and Peak RAM splitting performance for 10 column

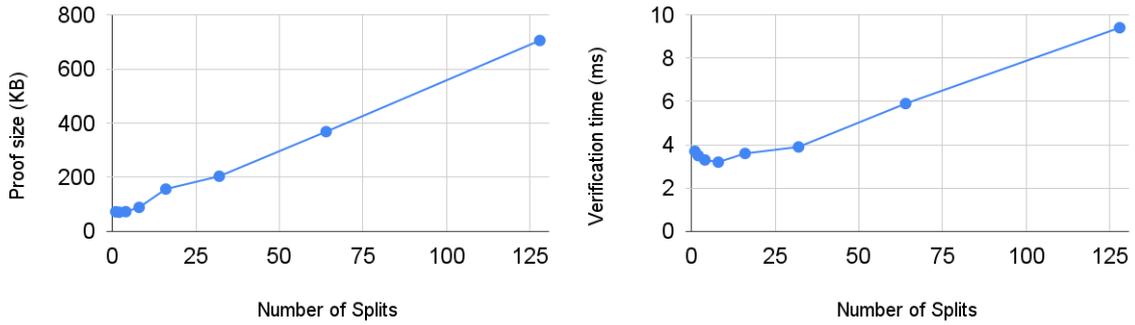


Figure 12: Verifier time and Proof size splitting performance for 10 columns

4.4 Security analysis

In this section we obtain knowledge soundness results for the IOP Π_{split}^k . Following a strategy we employed when proving knowledge soundness of the STARKPack IOP, we show that witnesses to a splitting of a Plonkish instance are in 1-to-1 correspondende with witnesses to the Plonkish instance. Because of this general correspondence, we have no doubt that these ideas extend to IOPs for more general splittings.

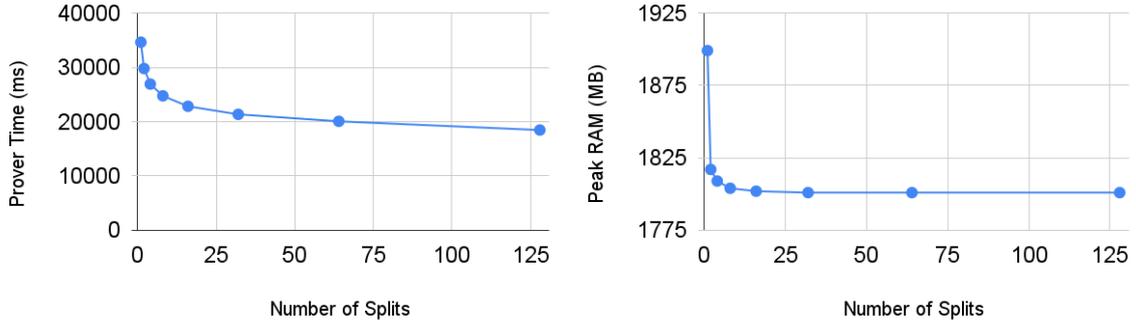


Figure 13: Prover time and Peak RAM splitting performance

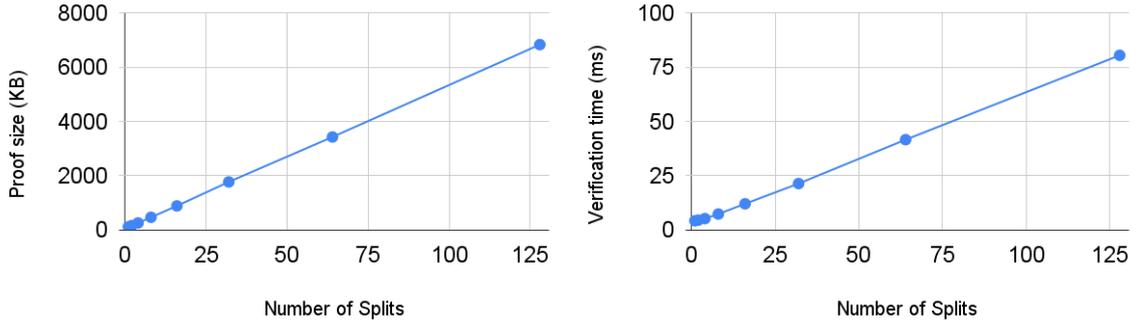


Figure 14: Verifier time and Proof size splitting performance for 100 columns

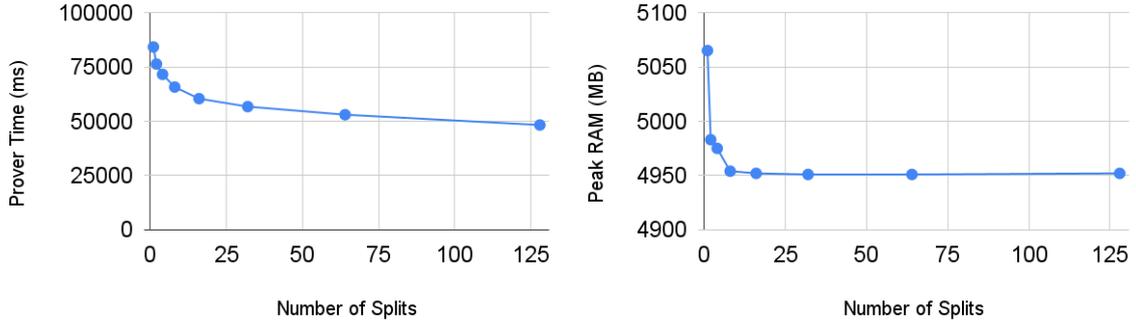


Figure 15: Prover time and Peak RAM splitting performance

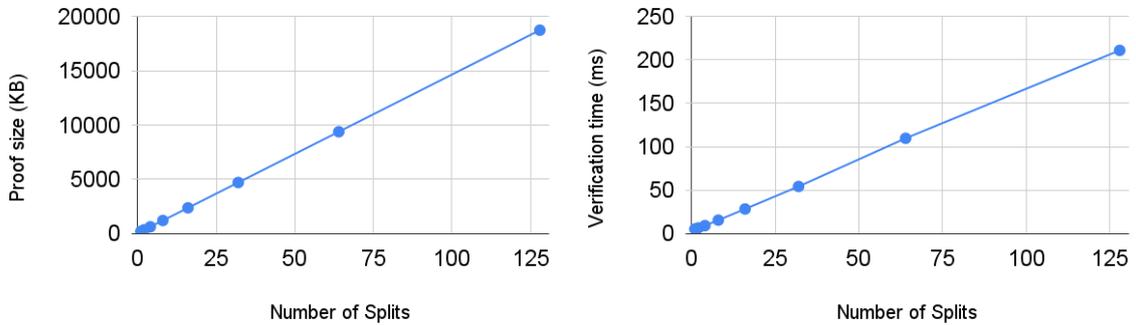


Figure 16: Verifier time and Proof size splitting performance for 275 columns

Lemma 4.7. Let $PL_{\text{split}} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, k, \mathfrak{G}, \mathbb{F}, \mathbb{K})$ be a split Plonkish instance. Let $PL = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, \mathbb{F}, \mathbb{K})$ be the corresponding Plonkish instance. Then there is a bijective (polynomially computable) correspondence between satisfying execution traces $w \in$

\mathbb{F}^{nr} for PL, and satisfying sets of execution traces $\mathbf{w}^0, \dots, \mathbf{w}^{k-1} \in \mathbb{F}^{(n/k)r}$ for PL_{split} .

Proof. This is almost by construction. Given a satisfying execution trace $\mathbf{w} \in \mathbb{F}^{nr}$ for PL, set for all $(i, j) \in \{0, \dots, n-1\} \times [r]$:

$$\mathbf{w}[ir + j] = \mathbf{w}^{\mathfrak{T}(i)}[\gamma(i)r + j]$$

Recall the definition of \mathfrak{T}, γ from Definition 4.1. This gives us k vectors $\mathbf{w}^0, \dots, \mathbf{w}^{k-1} \in \mathbb{F}^{(n/k)r}$. The assumption that \mathbf{w} is satisfying means that:

$$\forall i \in [n-1], \forall P \in \mathcal{P}, P(\mathbf{w}[(i-1)r+1], \dots, \mathbf{w}[ir+r], \mathbf{q}_1[i], \dots, \mathbf{q}_\ell[i]) = 0$$

and for all $i \in [nr]$, $\mathbf{w}_i = \mathbf{w}_{\sigma(i)}$. The former condition translates to:

$$\forall i \in [n-1], \forall P \in \mathcal{P}, P(\mathbf{w}^{\mathfrak{T}(i-1)}[\gamma(i-1)r+1], \dots, \mathbf{w}^{\mathfrak{T}(i)}[\gamma(i)r+r], \mathbf{q}_1[i], \dots, \mathbf{q}_\ell[i]) = 0$$

We set again $(\zeta(i, j), \xi(i, j)) \in [n] \times S$ to be such that $\sigma((i-1)r+j) = (\zeta(i, j)-1)r + \xi(i, j)$ for all $(i, j) \in [n] \times [r]$. The latter condition becomes $\mathbf{w}[(i-1)r+j] = \mathbf{w}[(\zeta(i, j)-1)r + \xi(i, j)]$ for all i, j . In turn, this becomes:

$$\forall (i, j) \in [n] \times [r], \mathbf{w}^{\mathfrak{T}(i-1)}[\gamma(i-1)r+j] = \mathbf{w}^{\mathfrak{T}(\zeta(i, j)-1)}[\gamma(\zeta(i, j)-1)r + \xi(i, j)]$$

So that $\mathbf{w}^0, \dots, \mathbf{w}^{k-1}$ are satisfying for PL_{split} .

Conversely, let $\mathbf{w}^0, \dots, \mathbf{w}^{k-1} \in \mathbb{F}^{(n/k)r}$ be satisfying for PL_{split} . For all $(s, i, j) \in \{0, \dots, k-1\} \times \{0, \dots, n/k-1\} \times [r]$, let

$$\mathbf{w}^s[ir + j] = \mathbf{w}[\mathfrak{S}^{-1}(i + sn/k)r + j]$$

This gives a vector $\mathbf{w} \in \mathbb{F}^{nr}$. The relation $\gamma(i) = \mathfrak{S}(i) - n\mathfrak{T}(i)/k$ implies that $\mathfrak{S}^{-1}(\gamma(i) + n\mathfrak{T}(i)/k) = i$. Therefore, the circuit constraint satisfaction conditions

$$\forall i \in [n-1], \forall P \in \mathcal{P}, P(\mathbf{w}^{\mathfrak{T}(i-1)}[\gamma(i-1)r+1], \dots, \mathbf{w}^{\mathfrak{T}(i)}[\gamma(i)r+r], \mathbf{q}_1[i], \dots, \mathbf{q}_\ell[i]) = 0$$

immediately give:

$$\forall i \in [n-1], \forall P \in \mathcal{P}, P(\mathbf{w}[(i-1)r+1], \dots, \mathbf{w}[ir+r], \mathbf{q}_1[i], \dots, \mathbf{q}_\ell[i]) = 0$$

It is clear by construction that the copy constraints on $\mathbf{w}^0, \dots, \mathbf{w}^{k-1}$ are equivalent to the condition that $\mathbf{w}_i = \mathbf{w}_{\sigma(i)}$, for all $i \in [nr]$. \square

Corollary 4.8. *The IOP for split-and-pack modulo k from Section 4.2 has knowledge soundness error, given an instance $\mathbf{R}_{\text{SPlonkish}} = (\mathcal{P}, \mathcal{Q}, \sigma, \text{Pl}, r, r', \ell, n, k, \mathfrak{S}, \mathbb{F}, \mathbb{K})$,*

$$\varepsilon_{\text{ks}} = \frac{1}{2\eta\sqrt{\rho}} \cdot \max \left\{ \left(\frac{3n(r' + u)}{|\mathbb{F}|} \right), \left(\frac{|\mathcal{P}| + s + 1}{|\mathbb{F}|} \right), \frac{n \cdot \max\{u + 1, d_{\text{max}}\}}{|\mathbb{K} \setminus \mathcal{D}|}, 2\eta\sqrt{\rho} \cdot \varepsilon_{\text{FRI}} \right\},$$

where \mathcal{D} is a subgroup of size n , u is a parameter used during the permutation argument of the STARKish IOP (analogous to the splitting parameters u_i in Section 3.1), $s = \lceil r'/u \rceil$, $\rho = (d_{\text{max}} + 1)/n$, η is a parameter used in the batched FRI protocol (cf. [BGK⁺23]), and ε_{FRI} is the round-by-round soundness error of the batched FRI protocol used at the end of the STARKish IOP.

Furthermore, the Succinct Non-interactive Argument of Knowledge obtained by compiling the IOP with Merkle tree commitments and the Fiat-Shamir transform is knowledge sound with error

$$Q \cdot \varepsilon_{\text{ks}} + O(Q^2/2^\kappa),$$

where Q is the number of random oracle queries with κ -bits of output that an adversary is capable of making.

Proof. The proof is analogous to that of Corollary 3.4, using Lemma 4.7 instead of Lemma 3.3. \square

5 References

- [BC23] Benedikt Bünz and Binyi Chen. Protostar: Generic efficient accumulation/folding for special sound protocols. Cryptology ePrint Archive, Paper 2023/620, 2023. <https://eprint.iacr.org/2023/620>.
- [BCMS20] Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data from accumulation schemes. Cryptology ePrint Archive, Paper 2020/499, 2020. <https://eprint.iacr.org/2020/499>.
- [BGH19] Sean Bowe, Jack Grigg, and Daira Hopwood. Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Paper 2019/1021, 2019. <https://eprint.iacr.org/2019/1021>.
- [BGK⁺23] Alexander R. Block, Albert Garreta, Jonathan Katz, Justin Thaler, Pratyush Ranjan Tiwari, and Michał Zając. Fiat-shamir security of fri and related snarks. Cryptology ePrint Archive, Paper 2023/1071, 2023. <https://eprint.iacr.org/2023/1071>.
- [BGT23] Jeremy Bruestle, Paul Gafni, and RISC Zero Team. Risc zero zkvm: Scalable, transparent arguments of risc-v integrity, 2023. <https://dev.risczero.com/proof-system-in-detail.pdf>.
- [BGTZ23] Alexander R. Block, Albert Garreta, Pratyush Ranjan Tiwari, and Michał Zając. On soundness notions for interactive oracle proofs. Cryptology ePrint Archive, Paper 2023/1256, 2023. <https://eprint.iacr.org/2023/1256>.
- [BMNW24] Benedikt Bünz, Pratyush Mishra, Wilson Nguyen, and William Wang. Accumulation without homomorphism. Cryptology ePrint Archive, Paper 2024/474, 2024. <https://eprint.iacr.org/2024/474>.
- [EG23] Liam Eagen and Ariel Gabizon. Protogalaxy: Efficient protostar-style folding of multiple instances. Cryptology ePrint Archive, Paper 2023/1106, 2023. <https://eprint.iacr.org/2023/1106>.
- [Gab] Ariel Gabizon. Multiset checks in plonk and plookup. <https://hackmd.io/@arielg/ByFgSDA7D>. Accessed: 2023-11-27.
- [Gaf] Paul Gafni. Zk10: Fri-based recursion with a groth16 wrapper. https://www.youtube.com/watch?v=wkIBN2CGJdc&list=PLj80z0cJm8QFnY6VLVa84nr-21DNvjWH7&index=17&ab_channel=ZeroKnowledge. Accessed: 2024-3-20.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. Cryptology ePrint Archive, Paper 2016/260, 2016. <https://eprint.iacr.org/2016/260>.
- [GW20] Ariel Gabizon and Zachary J. Williamson. plookup: A simplified polynomial protocol for lookup tables. Cryptology ePrint Archive, Paper 2020/315, 2020. <https://eprint.iacr.org/2020/315>.
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Paper 2019/953, 2019. <https://eprint.iacr.org/2019/953>.

- [KGL⁺20] Irakliy Khaburzaniya, François Garillot, Kevin Lewi, Konstantinos Chalkias, and Jasleen Malvai. Winterfell, 2020. <https://github.com/facebook/winterfell>.
- [KS22] Abhiram Kothapalli and Srinath Setty. Supernova: Proving universal machine executions without universal circuits. Cryptology ePrint Archive, Paper 2022/1758, 2022. <https://eprint.iacr.org/2022/1758>.
- [KS23] Abhiram Kothapalli and Srinath Setty. Hypernova: Recursive arguments for customizable constraint systems. Cryptology ePrint Archive, Paper 2023/573, 2023. <https://eprint.iacr.org/2023/573>.
- [KST21] Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla. Nova: Recursive zero-knowledge arguments from folding schemes. Cryptology ePrint Archive, Paper 2021/370, 2021. <https://eprint.iacr.org/2021/370>.
- [Lab23] Modulus Labs. The cost of intelligence: Proving machine learning inference with zero-knowledge, 2023. <https://drive.google.com/file/d/1tylpowpaqc0hKQtYolPlqvx6R2Gv4IzE/view>.
- [MAGABMMT23] Héctor Masip-Ardevol, Marc Guzmán-Albiol, Jordi Baylina-Melé, and Jose Luis Muñoz-Tapia. estark: Extending starks with arguments. Cryptology ePrint Archive, Paper 2023/474, 2023. <https://eprint.iacr.org/2023/474>.
- [Sta21] StarkWare. ethSTARK documentation. Cryptology ePrint Archive, Paper 2021/582, 2021. <https://eprint.iacr.org/2021/582>.
- [Tea22] PolygonZero Team. Plonky2: Fast recursive arguments with plonk and fri, 2022. <https://github.com/0xPolygonZero/plonky2/blob/main/plonky2/plonky2.pdf>.
- [ZGGX23] Tianyu Zheng, Shang Gao, Yu Guo, and Bin Xiao. Kilonova: Non-uniform pcd with zero-knowledge property from generic folding schemes. Cryptology ePrint Archive, Paper 2023/1579, 2023. <https://eprint.iacr.org/2023/1579>.
- [zkS23] zkSync. Boojum upgrade: zkSync era’s new high-performance proof system for radical decentralization, 2023. Accessed: 2023-7-17.

A Detailed benchmarks

A.1 Packed prover times

Table 3: Comparison of Packed and Non-packed(Winterfell) Prover Times with 10 columns traces

Number of Traces	Packed Prover Time (s)	Non-packed Prover Time (s)	Prover Time Ratio
1	1.0173	1.0173	1
2	1.6527	2.0346	1.23107642
3	2.3065	3.0519	1.32317364
4	2.9295	4.0692	1.389042499
5	3.5535	5.0865	1.431405656
6	4.2125	6.1038	1.448973294
7	4.8912	7.1211	1.455900393
8	5.5078	8.1384	1.477613566
9	6.1971	9.1557	1.477416856
10	6.7718	10.173	1.50225937

Table 4: Comparison of Packed and Non-packed(Winterfell) Prover Times with 100 columns traces

Number of Traces	Packed Prover Time (s)	Non-packed Prover Time (s)	Prover Time Ratio
1	5.0106	5.0106	1
2	9.1251	10.0212	1.098201664
3	13.141	15.0318	1.143885549
4	17.242	20.0424	1.162417353
5	21.843	25.053	1.146957835
6	25.453	30.0636	1.181141712
7	29.344	35.0742	1.195276718
8	33.329	40.0848	1.202700351
9	37.469	45.0954	1.203538926
10	42.099	50.106	1.190194541

Table 5: Comparison of Packed and Non-packed(Winterfell) Prover Times with 275 columns traces

Number of Traces	Packed Prover Time (s)	Non-packed Prover Time (s)	Prover Time Ratio
1	11.568	11.568	1
2	22.292	23.136	1.037861116
3	33.271	34.704	1.043070542
4	44.044	46.272	1.050585778
5	54.995	57.84	1.051731976
6	66.506	69.408	1.043635161
7	76.914	80.976	1.052812232
8	88.239	92.544	1.048787951
9	99.898	104.112	1.042183027

Table 6: Comparison of Packed and Non-packed(RISC Zero) Prover Times

Number of Traces	Packed Prover Time (s)	Non-packed Prover Time (s)	Prover Time Ratio
1	4.0001	4.0001	1
2	8.2615	8.0002	0.9683713611
3	12.404	12.0003	0.9674540471
4	16.621	16.0004	0.962661693
5	21.523	20.0005	0.92926172
6	25.012	24.0006	0.9595634096
7	29.996	28.0007	0.9334811308
8	34.47	32.0008	0.9283666957
9	39.646	36.0009	0.9080588206
10	44.005	40.001	0.9090103397

A.2 Packed verifier times

Table 7: Comparison of Packed and Non-packed(Winterfell) Verifier Times(ms) with 10 columns traces

No. Traces	Packed Verifier Time	Non-packed Verifier Time	Verifier Time Ratio
1	0.37661	0.37661	1
2	0.40327	0.75322	1.867780891
3	0.43987	1.12983	2.568554346
4	0.4445	1.50644	3.389066367
5	0.4449	1.88305	4.232524163
6	0.45203	2.25966	4.998916001
7	0.46565	2.63627	5.661483947
8	0.48044	3.01288	6.271084839
9	0.49884	3.38949	6.794743806
10	0.52587	3.7661	7.161655923

Table 8: Comparison of Packed and Non-packed(Winterfell) Verifier Times(ms) with 100 columns traces

Number of Traces	Packed Verifier Time	Non-packed Verifier Time	Verifier Time Ratio
1	0.46116	0.46116	1
2	0.54923	0.92232	1.67929647
3	0.60912	1.38348	2.271276596
4	0.68608	1.84464	2.688666045
5	0.7959	2.3058	2.897097625
6	0.84213	2.76696	3.285668483
7	0.92047	3.22812	3.50703445
8	0.98587	3.68928	3.742156674
9	1.0649	4.15044	3.897492722
10	1.154	4.6116	3.996187175

Table 9: Comparison of Packed and Non-packed(Winterfell) Verifier Times(ms) with 275 columns traces

Number of Traces	Packed Verifier Time	Non-packed Verifier Time	Verifier Time Ratio
1	0.5633	0.5633	1
2	0.75701	1.1266	1.488223405
3	0.95869	1.6899	1.762717875
4	1.145	2.2532	1.967860262
5	1.3456	2.8165	2.093118312
6	1.5293	3.3798	2.210030733
7	1.7129	3.9431	2.302002452
8	1.896	4.5064	2.376793249
9	2.0925	5.0697	2.422795699

Table 10: Comparison of Packed and Non-packed(RISC Zero) Verifier Times(ms)

Number of Traces	Packed Verifier Time	Non-packed Verifier Time	Verifier Time Ratio
1	1.8774	1.8774	1
2	2.6476	3.7548	1.418190059
3	3.3735	5.6322	1.669542019
4	4.1306	7.5096	1.818040963
5	4.8293	9.387	1.943759965
6	5.5719	11.2644	2.021644322
7	6.3896	13.1418	2.056748466
8	7.2754	15.0192	2.06438134
9	8.0432	16.8966	2.100731052
10	8.6492	18.774	2.170605374

A.3 Packed proof size

Table 11: Comparison of Packed and Non-packed(Winterfell) Proof Sizes in KB for 10 columns traces

Number of Traces	Packed Proof Size	Non-packed Proof Size	Proof Size Ratio
1	76	76	1
2	81.9	152	1.855921856
3	87.9	228	2.593856655
4	94.3	304	3.223753977
5	99.6	380	3.815261044
6	103.9	456	4.388835419
7	108.3	532	4.912280702
8	114.5	608	5.310043668
9	120.6	684	5.671641791
10	126.4	760	6.012658228

Table 12: Comparison of Packed and Non-packed(Winterfell) Proof Sizes in KB for 100 columns traces

Number of Traces	Packed Proof Size	Non-packed Proof Size	Proof Size Ratio
1	124.5	124.5	1
2	179	249	1.391061453
3	230.5	373.5	1.620390456
4	285.8	498	1.742477257
5	340.2	622.5	1.829805996
6	392.3	747	1.904154983
7	445	871.5	1.958426966
8	496.8	996	2.004830918
9	552.1	1120.5	2.029523637
10	604.6	1245	2.059212703

Table 13: Comparison of Packed and Non-packed(Winterfell) Proof Sizes in KB for 275 columns traces

Number of Traces	Packed Proof Size	Non-packed Proof Size	Proof Size Ratio
1	217	217	1
2	362.5	434	1.197241379
3	510.5	651	1.275220372
4	658.6	868	1.317947161
5	804.6	1085	1.348496147
6	947.4	1302	1.374287524
7	1095.8	1519	1.386201862
8	1242	1736	1.397745572
9	1387.9	1953	1.407161899

Table 14: Comparison of Packed and Non-packed(RISC Zero) Proof Sizes in KB

Number of Traces	Packed Proof Size	Non-packed Proof Size	Proof Size Ratio
1	210.223	210.223	1
2	278.941	420.445	1.507
3	347.660	630.668	1.814
4	416.378	840.891	2.020
5	485.098	1051.113	2.167
6	553.816	1261.336	2.278
7	622.535	1471.559	2.364
8	691.254	1681.781	2.433
9	759.973	1892.004	2.490
10	828.691	2102.227	2.537

A.4 Splitting in Winterfell

Table 15: Splitting benchmarks for a trace of 2^{16} rows with 1 column.

Number of Splits	Prover Time (ms)	Proof Size (KB)	Verification Time (ms)
1	6147	68.3	3.7
2	3219	61.6	3.6
4	1823	54.6	3
8	1160	51.5	2.8
16	919	49.8	2.7
32	780	51.3	2.3
64	717	63.3	2.3
128	665	92.4	2.6

Table 16: Splitting benchmarks for a trace of 2^{16} rows with 10 columns.

Number of Splits	Prover Time (ms)	Proof Size (KB)	Verification Time (ms)
1	8618	72.8	3.7
2	5514	71.1	3.5
4	3981	73.3	3.3
8	3320	89.1	3.2
16	2896	156.6	3.6
32	2612	203.9	3.9
64	2440	368.4	5.9
128	2247	704.8	9.4

Table 17: Splitting benchmarks for a trace of 2^{16} rows with 100 columns.

Number of Splits	Prover Time (ms)	Proof Size (KB)	Verification Time (ms)
1	34605	118.5	4.3
2	29766	166	4.6
4	26906	264.3	5.3
8	24738	471.5	7.4
16	22816	891.6	12.1
32	21337	1773.3	21.4
64	20055	3427.3	41.6
128	18451	6825	80.4

Table 18: Splitting benchmarks for a trace of 2^{16} rows with 275 columns.

Number of Splits	Prover Time (ms)	Proof Size (KB)	Verification Time (ms)
1	84143	212.2	5.4
2	76286	353.2	6.6
4	71546	636.2	9.3
8	65711	1215.1	15.6
16	60372	2378.6	28.4
32	56712	4708.8	54.3
64	52987	9379.5	109.7
128	48255	18724.3	210.8