

SubIonK: Sublinear Prover PIonK

Arka Rai Choudhuri ^{*} Sanjam Garg [†] Aarushi Goel [‡] Sruthi Sekar [§] Rohit Sinha [¶]

Abstract

We propose **SubIonK** — a new zero-knowledge succinct non-interactive argument of knowledge (zkSNARK). **SubIonK** builds on **PIonK** [EPRINT’19], a popular state-of-the-art practical zkSNARK. Our new construction preserves all the great features of **PIonK**, i.e., it supports constant size proofs, constant time proof verification, a circuit-independent universal setup, as well as support for custom gates and lookup gates. Moreover, **SubIonK** achieves improved prover running time over **PIonK**. In **PIonK**, the prover runtime grows with circuit size. Instead in **SubIonK**, the prover runtime grows with the size of the “active part” of the circuit. For instance, consider circuits encoding conditional execution, where only a fraction of the circuit is exercised by the input. For such circuits, the prover runtime in **SubIonK** grows only with the exercised execution path.

As an example, consider the zkRollup circuit. This circuit involves executing one of n code segments k times. For this case, using **PIonK** involves proving a circuit of size $n \cdot k$ code segments. In **SubIonK**, the prover costs are close to proving a **PIonK** proof for a circuit of size roughly k code segments. Concretely, based on our implementation, for parameter choices derived from rollup contracts on Ethereum, $n = 8, k = \{2^{10} \dots 2^{16}\}$, the **SubIonK** prover is approximately $4.6\times$ faster than the **PIonK** prover. Proofs in **SubIonK** are 2.4KB, and can be verified in under 50ms.

^{*}NTT Research arkarai.choudhuri@ntt-research.com.

[†]UC Berkeley and NTT Research sanjamg@berkeley.edu. The author is supported in part by DARPA under Agreement No. HR00112020026, AFOSR Award FA9550-19-1-0200, NSF CNS Award 1936826, and research grants by the Sloan Foundation, and Visa Inc.

[‡]NTT Research aarushi.goel@ntt-research.com.

[§]UC Berkeley. sruthi@berkeley.edu. The author is supported in part by DARPA under Agreement No. HR00112020026, AFOSR Award FA9550-19-1-0200, NSF CNS Award 1936826, and research grants by the Sloan Foundation, and Visa Inc.

[¶]Swirls Labs. sinharog@gmail.com.

Contents

1	Introduction	3
1.1	Our Contributions	4
1.2	Our Techniques	5
1.3	Example Applications	7
1.4	Related Works	8
2	Preliminaries	10
2.1	Algebraic Group Model	10
2.2	Preliminary Definitions	11
2.3	Batched KZG Commitments	11
2.4	Pre-processing SNARKs in the Algebraic Group Model	11
2.5	Preliminary Lemmas	12
2.6	Our Model	13
2.7	Background on $\mathcal{P}\text{Ion}\mathcal{K}$	13
3	Segment-Lookup Argument	15
3.1	Overview	15
3.2	Definition	18
3.3	Building Blocks	18
3.4	Tables with Repeated Elements	20
3.5	Protocol	21
3.6	Adding Zero-Knowledge	26
4	Sublonk: Segment Lookup + $\mathcal{P}\text{Ion}\mathcal{K}$	27
4.1	Pre-Processing Layered Branching Circuit	27
4.2	Post-Processing F_Y Polynomials	29
4.3	$\text{Sub}\text{Ion}\mathcal{K}$ Protocol Description	30
5	Implementation and Evaluation	32
5.1	Evaluated Applications	32
5.2	Prover Time	33
5.3	Proof Size	34
5.4	Verification Cost	34
6	Conclusion	34
A	Caulk Sub-protocol: Multi-unity Proof	39

1 Introduction

Zero-knowledge succinct non-interactive arguments of knowledge (zkSNARK) [Mic94, BCC⁺17] are cryptographic primitives, that allow a prover to generate a *small certificate* of correctness of a potentially expensive computation. Furthermore, these certificates are cheap to verify and hide secrets that the prover may have used in performing the computation. Over the past few years, realizing efficient zkSNARKs have been a topic of extensive research, including numerous applications (e.g., see [BCG⁺14, ZkR21, KGC⁺18, RPX⁺22, GHH⁺23, ZFZS20]) and real-world deployments (e.g., see [Gro16, GWC19, BBHR18a]).

Constant Proof Size, Universal Setup zkSNARKs. An influential line of work on zkSNARKs [Gro10, GGPR13, PHGR13] has focused on realizing zkSNARKs with constant proof size and constant time verification. In his seminal paper, Groth [Gro16] gave a zkSNARK with a 3 group element proof. These short proof sizes make such SNARKs quite appealing for several applications, e.g. ones involving on-chain verification. However, these first-generation constant-size zkSNARK constructions relied on a circuit-specific trusted setup — making real-world deployments challenging. The next generation of constant-size zkSNARKs removed this obstacle by realizing universal setup zkSNARKs [GKM⁺18, MBKM19]. In particular, these zkSNARKs relied on a circuit-independent trusted setup that only needed to be done once. Furthermore, following a circuit-specific pre-processing, these constructions also achieved constant time verification.

The $\mathcal{P}\text{Ion}\mathcal{K}$ Proof System. Improving upon prior work [MBKM19], Gabizon, Williamson, and Ciobotaru [GWC19] introduced a practical universal setup zkSNARK with constant proof size (about 400 bytes in practice) and constant time verification. This construction has found widespread real-world deployments. A significant reason for the success of $\mathcal{P}\text{Ion}\mathcal{K}$ is its easy adaptability. For example, $\mathcal{P}\text{Ion}\mathcal{K}$ supports *custom gates* — gates other than $+$ and \times — that significantly enhance concrete performance. plookup [GW20] and $\mathcal{P}\text{Ion}\mathcal{K}\text{up}$ [PFM⁺22] augment $\mathcal{P}\text{Ion}\mathcal{K}$ to add support for *lookup gates* — a gate checking that its input is from a pre-defined *lookup list*.

Limitations of $\mathcal{P}\text{Ion}\mathcal{K}$ Proof System. A key limitation of $\mathcal{P}\text{Ion}\mathcal{K}$, and all zkSNARKs with constant size proofs, is that proof generation is expensive — especially when compared to SNARKs with larger proof size, such as STARKs [BBHR18a]. Thus, improving proof generation times for $\mathcal{P}\text{Ion}\mathcal{K}$, or constant size SNARKs, continues to be an important problem of significant interest. A recent line of work on lookup arguments [ZBK⁺22a, PK22, GK22, ZGK⁺22, EFG22] improves $\mathcal{P}\text{Ion}\mathcal{K}$ proof generation for the special case of lookup gates to (essentially) independent of the lookup list. Targeting use cases where larger proof sizes are acceptable, several works [KPV22, plo, CBBZ23a] accelerate $\mathcal{P}\text{Ion}\mathcal{K}$ proof generation at the cost of increasing proof size.

Going Sublinear. We observe that $\mathcal{P}\text{Ion}\mathcal{K}$ and other similar SNARKs can be wasteful. For example, consider the following code snippet: *if X then evaluate, C_0 else evaluate C_1 .* When evaluating this code snippet only the “active part,” i.e., either C_0 or C_1 depending on the conditional X needs to be evaluated. In fact, it is often the case that in a computation only a fraction of the circuit is “active.”

More generally, we consider circuits C that can be divided into \bar{k} layers, where each layer has the same *branch* of n s -sized circuits C_1, \dots, C_n . These layers are interleaved in C by activation layers that select a single *active* circuit to execute in each layer - the choice of the circuit to be activated in each layer may depend on the input to C . Thus, for any input to C , the total size of the *executed/active* sub-circuit is $O(\bar{k}s)$ (independent of n). We denote circuits that satisfy the above structure as *layered branching circuit*.¹

¹While the applications we consider can naturally be cast into this model, our model is general enough to handle a larger class of circuits but may require more work to be viewed in this framework. Such transformations for more general circuits are

In certain scenarios, the input to C may determine the total number of layers $k \leq \bar{k}$ that is activated - we handle this by adding a special identity circuit C_{id} (that passes input unchanged to the output), and requiring that for every layer $\geq k$, C_{id} is activated. For such inputs, we say that the *effective* number of layers is k . For the remainder of the introduction and overview, we will refer to the activated sub-circuit as *effective* activated sub-circuit of size $O(ks)$.

Rollups, for instance, naturally map to our model of computation, where at each step, the prover executes one of several transaction types. For example, in a typical decentralized exchange (DEX) smart contract (e.g., Loopring [loo]), which allows users to create one of several types of transactions: deposits, spot trades, transfers, withdrawals, etc.

Naïvely, the prover runtime for SNARKs for the aforementioned circuits grow with the size of the entire circuit (i.e., grow with $O(nks)$). However, there has been a line of work [WSR⁺15, ZGK⁺18, KPPS20, BBHR18b, lib18, gen20, hod21, GPR21, MAGABMMT23, CFQ19, Lip16, CFH⁺15, san23, KS22, KS23, BC23] addressing this issue by additionally requiring an “a la carte” cost profile from the prover, where the cost of proving should grow only with the size of the executed sub-circuit (of size $O(ks)$) rather than the entire circuit. Unfortunately, prior works either fail to achieve a constant proof size or resort to using cryptographic hash functions in a non-black-box manner², which is undesirable given the overhead caused.³ See Section 1.4 for further discussion on this.

Specifically, in our model, we allow the prover to perform a one-time pre-processing that depends on the circuit C . The stored pre-processed material may in fact depend on $O(nks)$, but the online proof generation cost grows only with the size of the *active* sub-circuit (i.e. $O(ks)$). Note that this pre-processing is separate from the pre-processing required for the verifier to achieve constant time verification.

1.1 Our Contributions

The current state of affairs thus motivates us to construct SNARKs for layered branching circuit in the pre-processing model where: (i) the proof is of constant size; and (ii) the overhead in the prover cost in selecting the active circuit at each layer is $O(1)$. Our contributions are as follows:

1. We present **SubIonK**, building on a popular SNARK system **PIonK**. Our new construction preserves all the great features of **PIonK**, i.e., it supports constant size proofs, constant time proof verification, a circuit-independent universal setup, as well as support for custom gates and lookup gates. Additionally, **SubIonK** proof generation time grows only with the size of the active sub-circuit. Previously, **PIonK** proof generation grew with the size of the entire circuit.
2. We provide an implementation of **SubIonK** in Rust, and evaluate it on circuits modeling a popular rollup application.
3. We demonstrate the practical improvements with **SubIonK**. For instance, in our rollup application, we demonstrate improvements in prover time of up to 4.6×, and we also show potential far greater speedups for general-purpose programs which only exercise a small fraction of the entire logic in any execution. Proofs are 2.4KB in size, and verification requires 50 ms on a commodity machine or 716.6K EVM gas units to verify on-chain.

orthogonal to our work.

²I.e., proving a statement about the hash function by representing it as a set of constraints.

³Non-black-box use of cryptography inherently induces a non-constant overhead when selecting the active circuit at each layer.

1.2 Our Techniques

We now discuss the key technical ideas underlying SubIonK . As discussed earlier, our goal is to design pre-processing SNARKs for layered branching circuit, where the online cost to prove computation grows only with the active sub-circuit.

Recall that the pre-processing phase for the verifier in pre-processing SNARKs outputs a short summary - typically a commitment - of all the constraints in the circuit. These SNARKs then allow for the verifier to determine whether to accept the proof based on the commitment, thereby allowing the verifier to run in time sub-linear in the size of the circuit (the verifier no longer has to parse the entire circuit).

Core Idea: Note that once the prover executes the layered branching circuit C on input x , it induces the active sub-circuit \tilde{C}_x of size $O(ks)$ (with the non-active circuits from each layer removed). Thus, with the induced sub-circuit \tilde{C}_x , the prover can generate the proof in time proportional to $O(ks)$ as desired. Unfortunately, such a proof is not useful to the verifier since it cannot verify the proof without a commitment to the constraints specified only by \tilde{C}_x , whereas the verifier has a commitment to the constraints specified only by C . The first main idea is to enable the verifier to derive the commitment to \tilde{C}_x only given the commitment to C .

As a starting point, note that this is not something that can be addressed in the pre-processing phase since \tilde{C}_x depends on the input x . Further, having the prover simply send over the claimed commitment to \tilde{C}_x will not work either since the verifier needs to be convinced that the commitment was correctly generated. Thus we will require the prover to additionally prove that the new commitment is indeed a commitment to a valid induced sub-circuit of C , which the verifier can check given the commitment to C . But note that since the prover is now proving validity with respect to the original circuit of size $O(nks)$, care must be taken that the prover cost for this does not become proportional to $O(nks)$.

A natural approach for implementing this idea is to compute a Merkle hash of all the constraints in the pre-processing phase. Later in the online phase, the prover can generate a proof to convince the verifier that the new claimed commitment to \tilde{C}_x is a commitment to a subset of the leaf nodes (constraints) in this Merkle tree. This simple approach requires making a non-black-box use of hash functions, which adds significant computational overheads. The inefficiency of this type of approach is well established - in fact, a recent relevant line of work moved away from the non-black-box use of hash functions for breakthrough results on efficient sub-linear time *lookup arguments* [ZBK⁺22a, ZGK⁺22, PK22, GK22, EFG22].

Lookup Arguments. As we shall see shortly, lookup arguments will, in fact, be central to our work, and here we provide an informal description of the requirements in such works. Specifically, given a table T of size N to which the verifier only has access via an untrusted pre-processed commitment, *lookup arguments* allow one to prove in time proportional to $O(m)$ (independent of N) that the values of a committed polynomial of size m are contained within the table T . This is achieved by allowing a one-time prover pre-processing on the table T , taking time proportional to $O(N)$, and re-usable across multiple proofs. The online proof generation time grows only with $O(m)$. In fact, the m values allow for *repeated* elements from the table T , a property we will crucially leverage. Finally, the verification for the lookup arguments we consider is, in fact, $O(1)$ time. The specific lookup protocol that we build on is *cached quotients* (or *cq*), introduced by Eagen, Fiore and Gabizon [EFG22], which fits well with the $\mathcal{P}\text{IonK}$ proof system that we will use.

SubIonK Template: Recall, that in the $\mathcal{P}\text{IonK}$ proof system [GWC19], to achieve $O(1)$ -time verification for a circuit C , there is an untrusted verifier pre-processing phase that outputs $O(1)$ sized commitment on input C . We refer to the pre-processed verifier commitment as a $\mathcal{P}\text{IonK}$ commitment to C . Tying in with our previous discussion on *lookup arguments*, we have the following high-level template for SubIonK for

layered branching circuit C : (i) generate a lookup table that *appropriately* encodes information about the layered branching circuit C - this will also require prover to compute a one-time prover pre-processing of the lookup table; (ii) once the induced circuit \tilde{C}_x is fixed, use the lookup arguments to derive on-the-fly the $\mathcal{P}\text{Ion}\mathcal{K}$ commitment to C in time proportional to $O(ks)$ and prove that the derivation was done correctly. While this is indeed the template we follow in this work, there are several challenges in implementing its details that necessitate new ideas, as we illustrate below.

We begin by describing how we populate the initial table T given the layered branching circuit C in the context of cq . We use the implicit representation of C and store in T the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints for each circuit branch C_1, \dots, C_n .⁴ The exact nature of the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints are not important for this discussion. Since the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints for an s sized circuit can be represented in $O(s)$ constraints, the table consists of $O(ns)$ entries, where each entry is only a single field element (from an appropriate field).

In the online phase, once the induced circuit \tilde{C}_x is fixed, the prover can compute the polynomial commitment com to the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints for \tilde{C}_x as a concatenation of the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints for each active circuit in the k layers (since \tilde{C}_x itself is a concatenation of k circuits)⁵. Since each of these constraints are present in the table T , the prover can simply run the lookup argument protocol in time proportional to $O(ks)$ to generate proof that the constraints are contained in T . Unfortunately, the only guarantee provided by cq , or any lookup argument for that matter, is that each element in the committed polynomial com is contained in the table T . While this is certainly necessary, it is not a sufficient condition in our setting. For instance, the relative ordering of the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints is crucial for us to rely on the $\mathcal{P}\text{Ion}\mathcal{K}$ argument system since the $\mathcal{P}\text{Ion}\mathcal{K}$ security analysis assumes that the pre-processing is done correctly. This motivates us to extend the notion of lookup arguments to *segment* lookup arguments that we detail next.

Segment Lookup. In order to address the specific needs of our application, we extend the cq protocol to achieve a notion of *segment lookup*. The initial table T of size $O(ns)$ in a segment lookup protocol is sub-divided into n segments, each consisting of $O(s)$ contiguous elements in T (starting with the first element). The prover provides a commitment to a polynomial that encodes values in T , and proves that the committed values additionally satisfy *segment granularity*. Specifically, for $O(ks)$ values committed to via the polynomial, each of the k segments of size $O(s)$ (starting with the first element) must correspond exactly to a segment in T , maintaining relative ordering with the segment. It is easy to see that if each segment corresponds to a $\mathcal{P}\text{Ion}\mathcal{K}$ constraint, a segment lookup protocol will indeed provide the necessary guarantees to ensure that the polynomial commitment sent by the prover is, in fact, a $\mathcal{P}\text{Ion}\mathcal{K}$ commitment to \tilde{C}_x (for some x), where the validity of the choice of the segments will be checked separately by the $\mathcal{P}\text{Ion}\mathcal{K}$ proof system.

Unfortunately, we are not able to use existing lookup protocols in a black-box manner to achieve segment lookup. We extend the ideas present in the cq protocol to construct a new segment lookup argument, where the prover costs grow with the size of the polynomial that is committed.⁶

Putting It Together. The table T containing the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints for the circuit branches in C is pre-processed and provided to the verifier. Once the input x is fixed, the prover uses the induced circuit \tilde{C}_x to compute the $\mathcal{P}\text{Ion}\mathcal{K}$ verifier pre-processing for \tilde{C}_x , and subsequently, the corresponding segment lookup

⁴We handle the activation layer constraints by embedding it within each circuit branch, such that the circuit branch activated in the j -th layer also outputs the circuit branch to be activated in the $j + 1$ -th layer, and thus the activation layer can be ignored for the purposes of our discussion.

⁵It should be noted that this description of a concatenation of $\mathcal{P}\text{Ion}\mathcal{K}$ constraints is not fully accurate and written here as such for simplicity, and we handle this in our technical sections

⁶One could choose an appropriately large field to encode an entire segment into a single field element to use cq in a black-box manner, but the overhead would be too large for this approach to be meaningful in our setting.

proof to it, in time $O(ks)$. Thus, the SubIonK proof consists of (i) the $\mathcal{P}\text{IonK}$ verifier pre-processing for \widetilde{C}_x ; (ii) a segment lookup proof that the verifier pre-processing was correctly derived from T ; and (iii) $\mathcal{P}\text{IonK}$ proof for \widetilde{C}_x to be verified using the verifier pre-processing sent by the prover. All the communication, and verification can be done in $O(1)$, thus satisfying our efficiency requirements.

While this overview captures the main ideas, we refer the reader to the relevant technical sections for details. Specifically, we present our segment lookup protocol in Section 3 and show how it can be combined with $\mathcal{P}\text{IonK}$ to get SubIonK in Section 4.

1.3 Example Applications

SubIonK has the potential to improve prover run-time in nearly all applications of SNARKs, where the active part of the circuit during execution is not the entire circuit. In this section, we explore some example applications where it could be particularly beneficial and yield substantial computational savings.

1. **Rollups:** Rollups are becoming increasingly popular due to their potential to address the scaling issue of modern layer 1 blockchains. Consider a typical decentralized exchange (DEX) smart contract (e.g. Loopring [loo]), which allows users to create one of several types of transactions: deposits, spot trades, transfers, withdrawals, etc. The logic within each of these transaction types can be encoded as a circuit (typically under 60K arithmetic gates for each transaction). A single instance of a rollup transaction that is submitted to a layer 1 blockchain can batch together over hundreds of these DEX transactions, along with a single proof attesting to the validity of the state transition (from having applied all of the above DEX transactions on the state prior to the rollup transaction). Rollups naturally map to our model of computation, where at each step, the prover executes one of several transaction types (which map to segments in SubIonK). Specifically, if there are n different DEX transaction types, and a rollup batches together k such DEX transactions (each of size s), then we expect SubIonK to operate in roughly $O(ks \log(ks))$ time, whereas the $\mathcal{P}\text{IonK}$ prover would operate in $O(nks \log(nks))$ time. Our experiments in section 5 show significant speedups for rollups, for parameter values inspired by Loopring [loo].
2. **Smart Contracts:** Smart contracts support general computation (beyond rollups discussed above), but these can include arbitrary conditional statements, thereby often resulting in the active circuit only comprising a small fraction of the entire logic. For instance, consider a program that is a nested sequence of conditional statements - which can be represented as a complete tree in our graph-based model of computation. In such a setting, if each code segment is roughly the same size, the fraction of the executed path is exponentially smaller than the total size of the program. Specifically in the above example, if the nested conditional statements resulted in $O(n)$ segments each of size s , the run time of a $\mathcal{P}\text{IonK}$ prover on any input would grow with $O(ns)$. However, since the execution path along the tree would only execute $O(\log(n))$ code segments, the SubIonK prover running time would only grow with $O(\log(n)s)$. In section 5, we provide some data points that indicate significant concrete speedups for the above.
3. **Proving Existence of Bugs in Large Codebase:** Exploitation attacks pose a significant risk to large and critical software systems, leading to the emergence of bug bounty programs. These programs involve independent research teams auditing deployed software and revealing vulnerabilities in exchange for monetary incentives.

Recent works [HK20a, HK20b, GHH⁺23] have explored the idea of using zero-knowledge proof systems as a means for vulnerability research teams to substantiate to bug bounty program managers that they have successfully detected a critical exploit. This guarantees that they obtain their reward without disclosing the exploit prematurely.

Although the relation circuit for these proofs grows with the size of the software system, the execution path needed to prove the existence of a bug is expected to be much smaller than the entire software system. Having the proof generation time depend on the size of the entire software system could be very costly, particularly for complex systems. For programs that can be cast as layered branching circuit, `SublonK` is well-suited for such scenarios.

4. **Combating Disinformation.** Naveh and Tromer [NT16] recently demonstrated that zero-knowledge proofs can be used to verify that images featured in media have undergone a pre-approved set of modifications since their creation. This capability is especially valuable as it enables journalists to conceal sensitive content while simultaneously establishing the image’s authenticity.

The complete list of pre-approved edits determines the size of the relation circuit for generating proofs, while the execution path only considers the edits that are applied to the image. `SublonK` could help significantly improve proof generation times in this application.

1.4 Related Works

There are several prior works that focus on building SNARKs where the prover cost grows only with the size of the program execution. We summarize the most relevant works in Figure 1 and give a detailed description below.

A La Carte Cost Profile. There is a sequence of works including Buffet[WSR⁺15], vRAM[ZGK⁺18] and Mirage [KPPS20] that consider an “a la carte” cost profile for the provers where the prover cost for proving a step of computation (akin to layers in our setting) grow only with the size of the circuit representing the instruction invoked on that step, i.e. independent of the number of branches. Mirage [KPPS20] achieves a constant proof size using the universal circuit approach, where the trusted setup is run for the universal circuit (setup is indicated by a * in Figure 1) and the executed circuit is passed as input to this universal circuit. Since the prover knows the executed sub-circuit, it can simply provide this as input to the universal circuit. But to achieve constant verification time, one needs to pre-process the circuit that is passed to the universal circuit and since the sub-circuit is input dependent, it results in an input dependent pre-processing to achieve an “a la carte” cost profile. vRAM[ZGK⁺18] handles the issue of conveying the executed sub-circuit/instructions by only conveying the multiplicity of each instruction and appropriately encoding the constraints into the proof system, while fully black-box in the use of cryptography vRAM does not achieve a constant proof size. The proof generation in these works requires the entire transcript of the program execution in order to compute the proof, making them inherently non-incremental.

Incremental Proofs. A recent line of work; Sangria [san23], SuperNova [KS22], HyperNova [KS23], ProtoStar [BC23] address the lack of incremental property in the aforementioned works. These works build on the novel folding technique introduced by Nova [KST22] for designing IVCs (incrementally verifiable computation). Sangria, SuperNova, HyperNova and ProtoStar generalize the notion of IVCs to non-uniform IVCs, where at each step of the computation one out of a pre-determined set of instructions is executed. While the prover cost in these works only grows with the size of executed instructions at each step, they inherently rely on making non-black-box use of cryptography. This is because all these works follow the same high-level approach of designing an efficient folding argument and then efficiently compiling it into a non-uniform IVC using proof recursion. The proof-size of the resulting non-uniform IVC depends on the underlying SNARK used in this compilation. In Figure 1, we quote the proof sizes mentioned in the respective papers. However, we note that most of these schemes are compatible with and hence can be adapted to work with a variety of existing SNARKs (to further reduce their proof size).

	Constant Proof Size	Input Independent Verifier Preprocessing	Black-Box in Cryptography	Incremental Proof	Setup
vRAM [ZGK ⁺ 18]	✗	✓	✓	✗	urs
Mirage [KPPS20]	✓	✗	✓	✗	urs*
Sangria [san23]	✓	✓	✗	✓	urs
SuperNova [KS22]	✗	✓	✗	✓	trans
HyperNova [KS23]	✗	✓	✗	✓	trans
ProtoStar [BC23]	✗	✓	✗	✓	trans
eSTARK [MAGABMMT23]	✗	✓	✗	✗	trans
Geppetto [CFH ⁺ 15]	✗	✗	✗	✓	srs
Our Work	✓	✓	✓	✗	urs

Figure 1: We use ✓ to denote that a certain property is satisfied and ✗ to denote that it is not. When we say that a scheme does not have constant proof size (by constant we mean that the proof is a constant number of group or field elements), it means that they have a size that depends on the program execution size. In the Setup column, we refer to a circuit-dependent setup by srs, the circuit-independent universal setup by urs, and the untrusted transparent setup by trans. By urs* we mean that while the setup is circuit-dependent, the particular scheme is defined for universal circuits.

Transparent Setup. The works on building zkSTARKs [BBHR18b, lib18, gen20, hod21, GPR21, MAGABMMT23] use a transparent (i.e. untrusted) setup. All these schemes use the algebraic intermediate representation (AIR), which only encodes the step-by-step trace of the program execution. This inherently leads to the a la carte prover cost since it will only grow with the AIR size, which in turn grows only with the number of steps of the program that are executed. However, all these constructions have two shortcomings - non-constant proof size and non-black-box use of cryptographic hash functions. To ensure scalability while keeping the verifier pre-processing input independent, it is crucial that a hash of the computation trace is given along with proof that the hash computations were done correctly (uses AIR for “STARK-friendly” hashes). We summarize these properties for the most recent STARK [MAGABMMT23] in figure 1.

Commit and Prove SNARKs (CP-SNARKs). Several works [CFQ19, Lip16, CFH⁺15] build CP-SNARKs, which rely on proving statements of the form “ $C_{ck}(x)$ contains x such that $R(x, w)$ ”, where $C_{ck}(x)$ is a commitment. Such CP-SNARKs are shown to be useful [CFQ19] in proving the correctness of different parts of computation using different representations and proof systems (e.g. a QAP-based scheme may be used to prove one component, while a GKR [GKR08]-based scheme may be used for another). LegoSNARK [CFQ19] builds a general framework for CP-SNARKs that would help in linking such different components and also build CP-SNARKs for some existing SNARKs (Groth [Gro16], Pinocchio [Vee17] and zk-vSQL [ZGK⁺17]). This framework requires the prover to prove the knowledge of a valid opening for the commitment corresponding to the component used. Moreover, the only methods shown [CFH⁺15] to combine the proofs of different components involve a bounded bootstrapping (giving proof of a proof), making them non-black-box. We summarize these properties corresponding to the a la carte CP-SNARK, Geppetto [CFH⁺15], in fig. 1.

In a recent work [PFM⁺22], Pearson et al. also consider the idea of integrating $\mathcal{P}\text{IonK}$ with a lookup argument [GK22]. However, their goal was very different from ours. They propose an extension of $\mathcal{P}\text{IonK}$ that enables faster proof generation for relation circuits that include lookup gates, without having to encode the lookup relation as an arithmetic circuit.

2 Preliminaries

In this section, we present our model, establish notation and present an overview of $\mathcal{P}\text{IonK}$.

Notation. We denote our field by \mathbb{F} . We use $\mathbb{F}_{<d}[X]$ to denote the ring of univariate polynomials over \mathbb{F} with a degree smaller than d . We denote our security parameter by λ . For a polynomial $P \in \mathbb{F}[X]$, and a subgroup $\mathbb{H} \subset \mathbb{F}$, we denote the evaluations of P at \mathbb{H} by $P|_{\mathbb{H}}$. We use the additive notations for groups \mathbb{G}_1 and \mathbb{G}_2 , and denote their corresponding group elements by $[x]_1 := x.g_1$ and $[x]_2 := x.g_2$, where g_1 and g_2 are the generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively. $[n]$ and $[k, n]$ are used to denote the sets of integers $\{1, \dots, n\}$ and $\{k, \dots, n\}$, respectively.

Lagrange, and Vanishing Polynomials. For a subgroup containing n -th roots of unity, i.e., $\mathbb{H} = \{1, \omega, \dots, \omega^{n-1}\} \subset \mathbb{F}$, we denote the vanishing polynomial corresponding to \mathbb{H} by $Z_{\mathbb{H}}(X) \in \mathbb{F}[X]$, defined as $Z_{\mathbb{H}}(X) := \prod_{i=1}^n (X - \omega^i)$. Furthermore, for each $i \in [n]$, we denote the i -th Lagrange polynomial corresponding to \mathbb{H} by $\psi_i^{\mathbb{H}}(X) := \frac{Z'_{\mathbb{H}}(X)}{Z'_{\mathbb{H}}(X - \omega^i)}$, where $Z'_{\mathbb{H}}$ is the derivative of the polynomial $Z_{\mathbb{H}}$.

Bilinear Groups. Let $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ be cyclic groups of prime order q with generators $g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2$. $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ be an efficiently computable and non-degenerate pairing, such that $e(h_1^\alpha, h_2^\beta) = e(h_1, h_2)^{\alpha\beta}$, for all $\alpha, \beta \in \mathbb{F}_q$, and all $h_1 \in \mathbb{G}_1$ and $h_2 \in \mathbb{G}_2$.

2.1 Algebraic Group Model

We use the same terminology as prior works [FKL18, GWC19, EFG22]. An *algebraic adversary* \mathcal{A} in an SRS-based protocol is a $\text{poly}(\lambda)$ -time algorithm which satisfies the following: For each $i \in \{1, 2\}$, whenever \mathcal{A} outputs a group element $G \in \mathbb{G}_i$, it also outputs a vector v over \mathbb{F} such that $G = \langle v, \text{srs}_i \rangle$. srs_i is said to have degree d if all elements of srs_i are of the form $[F(\tau)]_i$ for $F \in \mathbb{F}_{<d}[X]$ and uniform $\tau \in \mathbb{F}$. In the following, it is assumed that a degree d SRS is used. Let $F_{i,j}$ denote the corresponding polynomial for the j th element of srs_i .

We require the following d -DLOG assumption to ensure that the srs in our protocol hides the random $\tau \in \mathbb{F}$ from \mathcal{A} .

Definition 1 (d -DLOG Assumption [FKL18, EFG22]). *Fix an integer d . The d -DLOG assumption for $(\mathbb{G}_1, \mathbb{G}_2)$ ensures that given $[1]_1, [\tau]_1, \dots, [\tau^d]_1, [1]_2, [\tau]_2, \dots, [\tau^d]_2$ for a uniformly random $\tau \in \mathbb{F}$, the probability of \mathcal{A} outputting τ is $\text{negl}(\lambda)$.*

Furthermore, we require the following lemma from [EFG22] which ensures that if the “*real*” pairing checks in our protocol are guaranteed to pass, then so would the “*ideal*” pairing check, where essentially the algebraic adversary \mathcal{A} also gives a field vector corresponding to the group elements in the pairing check (as described above).

Lemma 1 ([EFG22]). *Assume d -DLOG assumption for $(\mathbb{G}_1, \mathbb{G}_2)$. Given an algebraic adversary \mathcal{A} , and the d -degree srs from our protocol, the probability of a real pairing check passing is larger than the corresponding ideal check, by at most an additive factor of $\text{negl}(\lambda)$. The real and ideal pairing checks are described below.*

Real Pairing Check: For \mathbb{F} -vectors a, b , whose $\mathbb{G}_1, \mathbb{G}_2$ encodings are given by \mathcal{A} during the protocol execution, the real check is of the form: $(a \cdot T_1) \cdot (T_2 \cdot b) = 0$, for some matrices T_1, T_2 over \mathbb{F} . Such a check is done efficiently given the encoded group elements using the pairing function $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_t$.

Ideal Pairing Check: Since \mathcal{A} is algebraic, for each group encoding $[a_j]_i$ corresponding to a_j in the vector a , it also outputs a vector v such that $a_j = \sum v_\ell F_{i,\ell}(\tau) = R_{i,j}(\tau)$, for $R_{i,j}(X) := \sum v_\ell F_{i,\ell}(X)$. For $i \in \{1, 2\}$, we denote the vector of polynomials here by $R_i = (R_{i,j})_j$. The ideal check is of the form: $(R_1 \cdot T_1) \cdot (T_2 \cdot R_2) \equiv 0$.

2.2 Preliminary Definitions

In this section, we provide a formal definition of KZG polynomial commitments [KZG10] and pre-processing SNARKs in the algebraic group model, as defined in [GWC19].

2.3 Batched KZG Commitments

Definition 2 (Batched KZG Commitment [GWC19]). *A d -polynomial commitment scheme consists of*

- $\text{KZG.Gen}(d)$: randomized algorithm that outputs srs_{kzg} .
- $\text{KZG.Commit}(P, \text{srs}_{\text{kzg}})$: given a polynomial $P \in \mathbb{F}_{<d}[X]$, outputs a commitment com to P .
- A public coin protocol between a prover and a verifier. The prover gets $P_1, \dots, P_t \in \mathbb{F}_{<d}[X]$. Both the parties get the inputs $t = \text{poly}(\lambda)$, $z_1, \dots, z_t \in \mathbb{F}$, $\text{com}_1, \dots, \text{com}_t$, the alleged commitments to P_1, \dots, P_t , and $s_1, \dots, s_t \in \mathbb{F}$, the alleged correct openings $P_1(z_1), \dots, P_t(z_t)$. At the end of the protocol the verifier outputs accept or reject.

such that the following properties hold

- **Completeness:** Fix integer t , $z_1, \dots, z_t \in \mathbb{F}$, $P_1, \dots, P_t \in \mathbb{F}_{<d}[X]$. Suppose that for each $i \in [t]$, $\text{com}_i = \text{KZG.Commit}(P_i, \text{srs}_{\text{kzg}})$. Then if the prover and verifier honestly run the protocol with inputs $t, \{\text{com}_i, z_i, s_i = P_i(z_i)\}_{i \in [t]}$, the verifier accepts w.p. 1.
- **Knowledge Soundness in the algebraic group model:** The probability of any efficient algebraic adversary \mathcal{A} winning the following game is $\text{negl}(\lambda)$.
 - Given srs_{kzg} , \mathcal{A} outputs $t, \text{com}_1, \dots, \text{com}_t$ along with polynomials $P_1, \dots, P_t \in \mathbb{F}_{<d}[X]$, corresponding to the commitments.
 - \mathcal{A} outputs $z_1, \dots, z_t, s_1, \dots, s_t \in \mathbb{F}$.
 - \mathcal{A} (as a prover) and the verifier run the commitment protocol with inputs $\text{com}_1, \dots, \text{com}_t, z_1, \dots, z_t, s_1, \dots, s_t$.
 - \mathcal{A} wins if and only if the verifier accepts and for some $i \in [t]$, $s_i \neq P_i(z_i)$.

2.4 Pre-processing SNARKs in the Algebraic Group Model

In this section, we provide a formal definition of the pre-processing SNARKs in the algebraic group model, as defined in [GWC19].

Definition 3 (SNARKs in the Algebraic Group Model [GWC19]). *Let \mathcal{R} be a relation generator that given a security parameter λ in unary returns a polynomial time decidable binary relation R . For pairs $(\phi, w) \in R$ we call ϕ the statement and w the witness. We define \mathcal{R}_λ to be the set of possible relations R that the relation generator may output given 1^λ . We will in the following for notational simplicity assume λ can be deduced from the description of R . The relation generator may also output some side information, an auxiliary input z , which will be given to the adversary. An efficient prover publicly verifiable non-interactive argument for \mathcal{R} is a quadruple of probabilistic polynomial algorithms (Setup, Prove, Ver) defined as:*

- $\text{srs} \leftarrow \text{Setup}(1^\lambda, R)$: The setup takes the security parameter λ and the relation R as input and produces a structured reference string srs .
- $\pi \leftarrow \text{Prove}(\text{srs}, R, \phi, w)$: The prover algorithm takes as input srs and $(\phi, w) \in R$ and returns an argument π .
- $0/1 \leftarrow \text{Ver}(\text{srs}, R, \phi, \pi)$: The verification algorithm takes as input srs , a statement ϕ and an argument π and returns 0 (reject) or 1 (accept).

We say that $\Sigma = (\text{Setup}, \text{Prove}, \text{Ver})$ is a SNARK in the algebraic group model if it satisfies the following properties.

- **Completeness**: Given any true statement, an honest prover should be able to convince an honest verifier. For all $\lambda \in \mathbb{N}$, $R \in \mathcal{R}_\lambda$, $(\phi, w) \in R$

$$\Pr \left[\text{srs} \leftarrow \text{Setup}(1^\lambda, R); \pi \leftarrow \text{Prove}(\text{srs}, R, \phi, w) : \text{Ver}(\text{srs}, R, \phi, w) = 1 \right] \geq 1 - \text{negl}(\lambda).$$

- **Knowledge soundness in the algebraic group model**: The probability of an efficient algebraic adversary winning in the following game should be $\text{negl}(\lambda)$.
 - Given the srs and R , \mathcal{A} outputs (ϕ, π) along with the corresponding witness w .
 - \mathcal{A} wins if and only if the verifier accepts and $(\pi, w) \notin R$.
- **Succinctness**. A non-interactive argument of knowledge where the verifier runs in polynomial time in $\lambda + |\phi| + \log(|R|)$ and the proof size is polynomial in $\lambda + \log(|R|)$ is called a pre-processing SNARK. If we also restrict the srs to be polynomial in $\lambda + \log(|R|)$ we say that the non-interactive argument is a fully succinct SNARK.

2.5 Preliminary Lemmas

We require several properties of polynomials over fields for our protocol, which are described below. We require the following lemma on the uniqueness of fractional representations from [EFG22].

Lemma 2. [EFG22, Lemma 4] Let \mathbb{F} be an arbitrary field and $m_1, m_2 : \mathbb{F} \rightarrow \mathbb{F}$ be any functions. Then $\sum_{z \in \mathbb{F}} \frac{m_1(z)}{X-z} = \sum_{z \in \mathbb{F}} \frac{m_2(z)}{X-z}$ in the rational field $\mathbb{F}(X)$, if and only if, $m_1(z) = m_2(z)$ for every $z \in \mathbb{F}$.

Sumcheck Lemma. We require the following sumcheck lemma from [EFG22, BCR⁺19].

Lemma 3. Let $H \subset \mathbb{F}$ be a multiplicative subgroup of size t . For $F \in \mathbb{F}_{<t}[X]$, we have

$$\sum_{a \in H} F(a) = t \cdot F(0).$$

We require the following pre-processing lemma that combines [EFG22, Lemma 3.1] and [EFG22, Theorem 1].

Lemma 4. Let $G \in \mathbb{F}_{<ns}[X]$ and a subgroup $\mathbb{W} \subset \mathbb{F}$ of size ns . Further, suppose the \mathbb{G}_1 elements $\{[\tau^i]_1\}_{i \in [0, ns-1]}$ are given. Then, the following two computations are possible:

1. For each $i \in [ns]$, it is possible to compute the elements $g_i := [Q_i(\tau)]_1$, where $Q_i(X) \in \mathbb{F}[X]$ is such that

$$\psi_i^W(X) \cdot G(X) = g_i \cdot \psi_i^W(X) + Z_{\mathbb{W}}(X) \cdot Q_i(X),$$

for $g_i = G(\omega^i)$, in $O(ns \log(ns)) \mathbb{G}_1$ operations.

2. Furthermore, there is an algorithm, that takes as input, the q_i 's and $[\psi_i^{\mathbb{W}}(\tau)]_1$ for each $i \in [ns]$ from step 1 above, and a ks -sparse polynomial $H(X) \in \mathbb{F}_{<ns}[X]$ given in the sparse representation, and does the following. It takes $O(ks)$ \mathbb{F} -operations and ks \mathbb{G}_1 -operations to compute $[Q_A(\tau)]_1$ and $[R_A(\tau)]_1$, where $Q_A(X), R_A(X) \in \mathbb{F}_{<ns}[X]$ are such that

$$H(X)G(X) = Q_H(X)Z_{\mathbb{W}}(X) + R_H(X).$$

2.6 Our Model

We discuss below the model of computation used in this work. At a high level, we will consider layered circuits C with \bar{k} layers, where each layer supports a branch of n circuits $\{C_1, \dots, C_n\}$. The circuits are connected via interleaving *activation layers* where the $j - 1$ -th activation layer specifies the circuit C_i that will be the *active branch* in the j -th layer. We allow the output of each activation layer can depend on the input to the circuit.

Further, we note that our definition follows closely to that of works that consider (i) incremental verifiable computation (IVC) with non-uniform circuits at every step [BC23]; (ii) the ‘‘a la carte’’ cost profile of program execution [ZGK⁺18, WSR⁺15, KS22].

We formally define the model below by specifying an efficiently computable function ξ that on circuit input specifies a vector $I \in [1, n]^{\bar{k}}$. For notational simplicity, we denote by ξ_x the vector $\xi(x)$, which is indexed at the j -th location by $\xi_x[j]$.

Definition 4. Let $s, \bar{k}, n, m \in \mathbb{N}$, and for every $i \in [n]$, C_i is an arithmetic circuit of size s . Then, a C is a $(s, \bar{k}, n, \{C_i\}_{i=1}^n)$ -layered branching circuit is a circuit such that there exists an efficiently computable function $\xi : \mathbb{F}^m \rightarrow [1, n]^{\bar{k}}$ such that for every input $x \in \mathbb{F}^m$, $C(x) = \tilde{C}_\xi(x)$, where $\tilde{C}_\xi := C_{\xi_x[1]} \parallel \dots \parallel C_{\xi_x[\bar{k}]}$.

Remark 1. To handle scenarios where the input determines the number of layers $k \leq \bar{k}$ that are activated, we can include a ‘special’ circuit C_{id} which implements the identity (i.e., just passes the input through unchanged). We then say that for any $k \leq \bar{k}$, the input x has effective layer k if the circuits active in the final $\bar{k} - k$ layers are all C_{id} , i.e. $\forall j \in [k, \bar{k}] C_{\xi_x[j]} = C_{\text{id}}$. In such a situation, one can in fact require a stronger requirement - that the prover cost grows with $O(ks)$ rather than $O(\bar{k}s)$. For the rest of this work, we will indeed consider the effective layer k .

The above formalism allows us to contrast between an explicit representation of the layered branching circuit, which is of size $O(kns)$, and the specific sub-sequence of k circuits each of size s that are *activated* for a specific input. Looking ahead, we will utilize the aforementioned property of layered circuits to construct a proof system where the prover cost grows with cost of the executed circuits $O(ks)$.

We note that while the example applications discussed in Section 1.3 can be naturally cast into this model, our model is general enough to handle a larger class of circuits. For instance, a generic program consisting of nested conditional statements of depth d can be recast in our above formulation by appropriately choosing $n = 2^d$ circuits, where the executed path will indeed consist of $d = O(\log n)$ circuits.

2.7 Background on $\mathcal{P}\text{Ion}\mathcal{K}$

$\mathcal{P}\text{Ion}\mathcal{K}$ [GWC19] is a popular state-of-the-art, pre-processing zkSNARK with a constant-sized proof. As discussed earlier, our work builds on $\mathcal{P}\text{Ion}\mathcal{K}$ to design a zkSNARK where the proof generation time grows only with the size of the active sub-circuit. In this section, we provide the relevant background on the $\mathcal{P}\text{Ion}\mathcal{K}$ proof system that will be useful for understanding our construction. Some of the text in this section is taken verbatim from [GWC19].

$\mathcal{P}\text{IonK}$ Constraint System. The $\mathcal{P}\text{IonK}$ constraint system is meant to capture fan-in two arithmetic circuits of unlimited fan-out with n gates and m wires, but is more general. It is defined as $\mathcal{C} = (\mathcal{V}, \mathcal{Q})$, where:

- \mathcal{V} is of the form $\mathcal{V} = (\mathbf{a}, \mathbf{b}, \mathbf{c}) \in [m]^{3n}$, which implicitly describe a *permutation* (to be explained shortly) on $[3n]$.
- $\mathcal{Q} = (\mathbf{q}_L, \mathbf{q}_R, \mathbf{q}_O, \mathbf{q}_M, \mathbf{q}_C) \in (\mathbb{F}^n)^5$ where $\mathbf{q}_L, \mathbf{q}_R, \mathbf{q}_O, \mathbf{q}_M, \mathbf{q}_C \in \mathbb{F}^n$ are the *selector vectors*.

We say $\mathbf{x} \in \mathbb{F}^m$ satisfies \mathcal{C} if for each $\ell \in [n]$,

$$(\mathbf{q}_L)_\ell \cdot \mathbf{x}_{a_\ell} + (\mathbf{q}_R)_\ell \cdot \mathbf{x}_{b_\ell} + (\mathbf{q}_O)_\ell \cdot \mathbf{x}_{c_\ell} + (\mathbf{q}_M)_\ell \cdot (\mathbf{x}_{a_\ell} \mathbf{x}_{b_\ell}) + (\mathbf{q}_C)_\ell = 0.$$

This lets us define relation $\mathcal{R}_{\mathcal{C}}$ which is a set of pairs $\mathbf{x} := (\mathbf{x}, \mathbf{w})$, where \mathbf{x} satisfies \mathcal{C} .

From Arithmetic Circuits to $\mathcal{P}\text{IonK}$ Constraint System. As an example, we demonstrate how a fan-in two arithmetic circuit with n gates (each one is either an addition or a multiplication gate) and m wires (since every gate is assumed to have 2 input wires and 1 output wire associated with it, $m = 3n$.) can be captured by the $\mathcal{P}\text{IonK}$ constraint system. For each gate $\ell \in [n]$

- Let a_ℓ, b_ℓ and c_ℓ denote the index of the left, right and output wire of the ℓ^{th} gate. Set $(\mathbf{q}_C)_\ell = 0$.⁷
- Set $(\mathbf{q}_L)_\ell = 0, (\mathbf{q}_R)_\ell = 0, (\mathbf{q}_O)_\ell = -1, (\mathbf{q}_M)_\ell = 1$, when the ℓ^{th} gate is a multiplication gate.
- Set $(\mathbf{q}_L)_\ell = 1, (\mathbf{q}_R)_\ell = 1, (\mathbf{q}_O)_\ell = -1, (\mathbf{q}_M)_\ell = 0$, when the ℓ^{th} gate is an addition gate.

A circuit constraint system needs to ensure (1) *Correct Gate Evaluation*: given the left and right input wires, each gate is evaluated correctly. This is checked by choosing appropriate entries in the selector vectors based on the gate types (as described in the above example). (2) *Consistency of Wire Values*: if a wire is “split” (for instance as input to multiple gates or as the output of one gate and input to another), all the split wires must indeed contain the same wire value. This is done in the $\mathcal{P}\text{IonK}$ proof system via the *copy-check constraints* implemented by a permutation $\sigma : [3n] \rightarrow [3n]$. Specifically, σ is a collection of cycles (possibly of length 1), where each cycle is over all wires that are required to contain the same value as a consequence of the aforementioned “split”.

We note that the $\mathcal{P}\text{IonK}$ constraint system can be further generalized to handle custom gates and gates with arbitrary fan-in. However, for simplicity of presentation, we only work with the above simple variant.

Verifier Pre-Processing in $\mathcal{P}\text{IonK}$. The $\mathcal{P}\text{IonK}$ protocol for the above constraint system is defined over a multiplicative subgroup $\mathbb{W} = \{1, \omega, \dots, \omega^{n-1}\}$ of size n . Let k_1 and k_2 be picked such that $k_1 \cdot \mathbb{W}$ and $k_2 \cdot \mathbb{W}$ are distinct cosets of \mathbb{W} .⁸ Let $\mathbb{W}' = \mathbb{W} \cup (k_1 \cdot \mathbb{W}) \cup (k_2 \cdot \mathbb{W})$. Identify $[3n]$ with \mathbb{W}' via $\ell \rightarrow \omega^\ell, \ell + n \rightarrow k_1 \cdot \omega^\ell, \ell + 2n \rightarrow k_2 \cdot \omega^\ell$. Finally, let σ^* denote the mapping from $[3n]$ to \mathbb{W}' derived from applying σ (as described above) and then this injective mapping into \mathbb{W}' .

The $\mathcal{P}\text{IonK}$ protocol requires the following universal trusted setup (needed for computing KZG commitments [KZG10] throughout the protocol): $(\tau \cdot [1]_1, \dots, \tau^{n+5} \cdot [1]_1)$, for a randomly chosen τ . In addition, in order to keep the verifier cost low, the $\mathcal{P}\text{IonK}$ protocol pre-processes the constraint system to produce

⁷We remark that the above is only an example. The $\mathcal{P}\text{IonK}$ constraint system is quite general and can be used to enforce other types of constraints as well (e.g., checking if some wire value is equal to a public input by setting the corresponding entry in \mathbf{q}_C to that public value, or whether a wire value is a boolean value etc.).

⁸Since further details regarding the cosets are not relevant to our discussion, we do not elaborate here and refer the reader to [GWC19] for details.

the following pre-processed input, where ψ_i correspond to the Lagrange polynomials over multiplicative sub-group \mathbb{W} .

$$\begin{aligned}
& n, (q_{Mi}, q_{Li}, q_{Ri}, q_{Oi}, q_{Ci})_{i=1}^n, \sigma^*, \\
q_M(X) &= \sum_{\ell=1}^n q_{Mi} \psi_\ell(X), \quad q_L(X) = \sum_{\ell=1}^n q_{Li} \psi_\ell(X), \\
q_R(X) &= \sum_{\ell=1}^n q_{Ri} \psi_\ell(X), \quad q_O(X) = \sum_{\ell=1}^n q_{Oi} \psi_\ell(X), \\
q_C(X) &= \sum_{\ell=1}^n q_{Ci} \psi_\ell(X), \quad S_{\sigma_1}(X) = \sum_{\ell=1}^n \sigma^*(i) \psi_i(X), \\
S_{\sigma_2}(X) &= \sum_{\ell=1}^n \sigma^*(n + \ell) \psi_\ell(X), \quad S_{\sigma_3}(X) = \sum_{\ell=1}^n \sigma^*(2n + \ell) \psi_\ell(X)
\end{aligned}$$

Here the KZG commitment to these polynomials are sent to the verifier as $[q_M(\tau)]_1, [q_L(\tau)]_1, [q_R(\tau)]_1, [q_O(\tau)]_1, [q_C(\tau)]_1, [S_{\sigma_1}(\tau)]_1, [S_{\sigma_2}(\tau)]_1, [S_{\sigma_3}(\tau)]_1$. The rest of the protocol proceeds by assuming that this pre-processing was done honestly, and that the verifier has access to these commitments, while the prover has access to the above polynomials.⁹

Observe that the size of the above polynomials (and hence the prover work in $\mathcal{P}\text{Ion}\mathcal{K}$) depends on the “entire” circuit. Looking ahead, we adopt the following high-level approach to reduce the prover work. Given “some” input-independent verifier pre-processing (based on the entire circuit), we will allow the prover to *efficiently derive* a “smaller” verifier pre-processing material (of the above form) that only depends on the activated sub-circuit and have the prover send KZG commitments to this derived pre-processing material to the verifier, along with a proof that certifies that these were computed honestly. Finally, given this “smaller” derived pre-processing, the rest of our protocol will work exactly as $\mathcal{P}\text{Ion}\mathcal{K}$. A majority of the rest of this paper is dedicated towards describing our approach that allows the prover to efficiently derive and convince the verifier that the derived verifier pre-processing for the activated sub-circuit was honestly computed.

3 Segment-Lookup Argument

In this section, we present an efficient SNARK for *segment-lookup*. Looking ahead, in Section 4 we show how this protocol can be combined with $\mathcal{P}\text{Ion}\mathcal{K}$ to obtain $\text{SubIon}\mathcal{K}$.

3.1 Overview

We start with an overview of the techniques in our construction. As discussed in the introduction, the first step towards designing $\text{SubIon}\mathcal{K}$ involves designing a protocol for segment lookup. Since our protocol borrows techniques developed in cq [EFG22], we first briefly recall the cq protocol. This will also allow us to pinpoint the shortcomings that necessitate a *segment* lookup protocol, which will allow us to naturally define the properties required from a segment lookup protocol.

Overview of cq [EFG22]. Lookup arguments are succinct proof systems, where given a commitment to a large lookup table (of size n), the prover wants to convince the verifier that a commitment to a vector of

⁹We omit the discussion on how $\mathcal{P}\text{Ion}\mathcal{K}$ works given the above pre-processing, since it is not relevant for understanding our techniques.

k values are all contained in the large lookup table (the vector is allowed to repeat values from the table). `cq` allows the prover to pre-process the table and generate such proofs in time proportional to k , where the proofs themselves are constant-sized.

The lookup table in `cq` is encoded using a polynomial $T(X)$ of degree at most n , and the vector encoded using a polynomial $F(X)$. Both these polynomials are committed via the KZG [KZG10] polynomial-commitment. The pre-processing for the prover involves pre-computing succinct “quotient” commitments based on T . These quotient commitments help the prover generate a proof in the online phase in time that is proportional to $O(k)$.

In more detail, `cq` relies on the following log-derivative lemma from [Hab22], which essentially says that the values encoded using polynomial $F(X)$ are a subset of the values encoded using polynomial $T(X)$ if and only if for some $m \in \mathbb{F}^n$

$$\sum_{i \in [n]} \frac{m_i}{X + t_i} = \sum_{i \in [k]} \frac{1}{X + f_i},$$

where t_i indicates the i -th entry of the table T , and correspondingly f_i indicates the i -th entry of F . Here m_i essentially encode the multiplicity of the i -th element T in F .

In [EFG22], this identity is checked by letting the verifier evaluate it at a random β by requiring the prover to send polynomial commitments to the following polynomials: (1) $M(X)$, which is an encoding of m , and (2) polynomials $A(X)$ and $B(X)$, which are encodings of the summands on the left and right hand sides of the above equation respectively.

While $M(X)$ and $A(X)$ are both degree $n - 1$ polynomials, the number of non-zero evaluations of these polynomials is at most k ($B(X)$ is only a sum of size $O(k)$). Hence, given pre-processed commitments to Lagrange polynomials, it is possible for the prover to generate a commitment to M and A simply using $O(k)$ operations. The only remaining step in enabling the verifier to check the above equality is to convince them that the commitment to polynomial $A(X)$ is well-formed with respect to $T(X)$. This is be done by providing the verifier with a commitment to a quotient polynomial $Q_A(X)$ and letting them check if

$$A(X) \cdot (T(X) + \beta) - M(X) \stackrel{?}{=} Q_A(X) \cdot Z_{\mathbb{W}}(X).$$

Even though polynomials $A(X)$, $M(X)$ have sparse representations, computing the $Q_A(X)$ still requires $O(n)$ operations. To reduce this overhead, `cq` introduced the idea of cached quotients. They show that if one pre-process commitments to quotients $\{Q_i(X)\}_{i \in [n]}$, of the form

$$\psi_i^{\mathbb{W}}(X) \cdot T(X) = Q_i(X) \cdot Z_{\mathbb{W}}(X) + R_i(X),$$

then it is possible to compute a commitment to $Q_A(X)$ using just $O(k)$ operations.

As evidenced from the above equations, the `cq` protocol only provides guarantees that the element encoded via $F(X)$ are contained in $T(X)$. As described in the introduction, our application for lookup protocol requires stronger properties that we describe below.

Segment-Lookup. We propose a variant of the standard lookup problem that we refer to as the *segment-lookup* problem. At a high level, the look-up table of size ns is partitioned into n segments consisting of s elements each. Now the polynomial $F(X)$ encoding ks elements must contain k segments from $T(X)$ in its entirety, ensuring that the relative order of elements within each segment is maintained. From our aforementioned discussion of `cq`, it does not guarantee such a property. We formally define our requirements now.

Let $\mathbb{W} = \{\omega^0, \dots, \omega^{ns-1}\}$ be a set of ns^{th} roots of unity and $\mathbb{V} = \{v^0, \dots, v^{ks-1}\}$ denote the set of ks^{th} roots of unity. Given commitments to $ns - 1$ degree polynomial $T(X)$, for each $i \in [0, n - 1]$, the i -th segment is $\{T(\omega^{is}), \dots, T(\omega^{(i+1)s-1})\}$. Thus, given a degree $ks - 1$ polynomial $F(X)$, the prover in a segment-lookup

protocol wants to convince the verifier that for each $i \in [0, k-1]$, $F(v^{is+0}), \dots, F(v^{(i+1)s-1})$ represents one of the original n segments in $T(X)$. In other words, the prover wants to prove that $F(X)$ is well-formed, and each segment embedded in $F(X)$ is taken from $T(X)$.

With the above requirement in mind, we want to design an efficient pre-processing SNARK with a constant proof size and constant time verification, where the prover work only grows with ks .

Our Approach. We now describe our main ideas for designing such a protocol. For simplicity of notation, for any polynomial $P(X)$, we will denote by p_i the evaluation at ω^i (the i -th power of the corresponding set of roots of unity), i.e. $p_i := P(\omega^i)$. Note that the range of i will vary depending on the corresponding roots of unity P is defined over.

Similar to **cq**, we start by defining a polynomial $M(X)$ that is used to indicate which elements of $T(X)$ are included in $F(X)$, and how many times. i.e. $m_i = \# \text{times } t_i \text{ appears in } (f_1, \dots, f_{ks})$. The number of non-zero entries in M is bounded above by $\min\{ks, ns\}$.

Since we want the “granularity” of elements selected to correspond to an *entire segment*, unlike **cq**, we want to enforce some additional constraints on $M(X)$.

- **Constraint I:** The first constraint we enforce is that the value $\{m_i\}_{i \in [js, (j+1)s-1]}$ in $M(X)$ corresponding to each segment, must all be *equal*. We capture this by comparing each pair of consecutive values in M , except the first value in each segment since each segment can have distinct values. The test is then described as, $\forall i \in [ns]$ s.t. $s \nmid i, m_i = m_{i-1}$. Or in polynomial terms,

$$\forall x \in \mathbb{W}, (x^n - 1)(M(x) - M(x/\omega)) = 0.$$

By the fact that \mathbb{W} consists of ks -th roots of unity, first term, $(x^n - 1)$ is 0 if and only if $x = \omega^{js}$ for some j , which ensures: (i) the equation is trivially satisfied by the starting index of each segment, which by description is of the form js , i.e. encodes $s \nmid i$; and (ii) for all other indices i , it must be the case that $m_i = m_{i-1}$ to ensure the condition holds. This check is encoded by the polynomial check in Step 2 of Round 1.

- **Constraint II:** A consequence of how the table lookup protocol works in **cq** is that the relative ordering of elements in $F(X)$ need not be consistent with the relative ordering of elements in $T(X)$. Since our segment will encode a circuit, it is imperative that we maintain the relative ordering of the elements *within* a segment.

We encode this test by first defining a function, $L : [ks] \mapsto [ns]$ which we will encode as a polynomial $L(X) : \mathbb{K} \mapsto \mathbb{W}$ (by overloading notation).¹⁰ L maps the indices from F to their corresponding location within T . For relative ordering, we perform a check akin to that of M , the indices for consecutive elements in F should be consecutive in T (except for the start of the segment), $\forall i \in [ks]$ s.t. $s \nmid i, \ell_{i+1} = \ell_i + 1$. In polynomial terms, since we have ks -th roots of unity, we can rewrite this as,

$$\forall x \in \mathbb{V}, (x^k - 1)(L(xv) - \omega L(x)) = 0.$$

$L(X)$ has degree at most ks .

- **Constraint III:** Unfortunately, the above checks are not yet sufficient to achieve the desired segment granularity. While the tests do ensure that relative ordering is maintained within the segment, it does not enforce that the segments in $F(X)$ indeed *start* at the specified location, i.e. we want to ensure that ℓ_{js} for each $j \in [k]$ must map to *an* index is for some $i \in [n]$ in T . This will ensure that each segment in F indeed corresponds to a segment in T .

¹⁰Since the same segment can be invoked multiple times, the function is not injective, and thus L^{-1} is not well defined.

We define a polynomial D , that selects $\{\ell_{js}\}_j$, i.e. $\forall i \in [k], d_{is} = \ell_{is}$. Finally, we need to check that $\forall i \in [k], s \mid d_{is}$. In polynomial terms, this translates to checking if all of the elements in $\{d_{is}\}_{i \in [k]}$ are n -th roots of unity. To perform this check, we invoke a sub-protocol from [ZBK⁺22a].

Most of the remaining protocol follows the cq template, except that since we need to enforce the above additional constraints, the log-derivative lemma used in cq does not suffice in our setting. We work with the following modified lemma: Each segment embedded in F is taken from T if and only if for some M and L as defined above,

$$\sum_{i \in [ns]} \frac{m_i}{X + t_i + Y\omega^i} = \sum_{i \in [ks]} \frac{1}{X + f_i + Y\ell_i}$$

The verifier picks random field elements to replace X, Y . The prover provides commitments to polynomials $A(X)$ and $B(X)$ that encode the summands in the LHS and RHS, respectively. Similar to cq, the verifier then verifies if the above equality holds. Throughout the protocol, when computing commitments to quotient polynomials of degree $ns - 1$, we use the idea of cached quotients to reduce online prover work at the cost of some additional pre-processing.

3.2 Definition

We begin by formally defining Segment-Lookup, secure against algebraic adversaries.

Definition 5 (Segment-Lookup). *(n, k, s) -Segment-Lookup is a pair (gen, segmentLookup) such that:*

- *gen(n, k, s, T) : This is a PPT algorithm that takes as input integers n, k, s and a polynomial $T \in \mathbb{F}[X]$ of degree $ns - 1$. It outputs a string srs of \mathbb{G}_1 and \mathbb{G}_2 elements. This algorithm is run in the pre-processing phase.*
- *segmentLookup(com, srs, T, F, \mathbb{V}) : This is an interactive public coin protocol between the Prover and the Verifier, where the prover has a private input $F \in \mathbb{F}[X]$ of degree $ks - 1$, and both the parties have access to T, com and $\text{srs} := \text{gen}(n, k, s, T)$, such that it satisfies the following properties of completeness and knowledge soundness in the algebraic group model.*
 - **Completeness.** *If $\text{com} = [F(\tau)]_1$, and if for each $i \in [0, k - 1]$, there exists $j \in [0, n - 1]$ such that, for each $q \in [0, s - 1]$, $F(v^{is+q}) = T(\omega^{js+q})$, then the verifier accepts with probability 1.*
 - **Knowledge Soundness in the algebraic group model.** *The probability of any efficient algebraic adversary \mathcal{A} winning the following game is $\text{negl}(\lambda)$.*
 1. \mathcal{A} chooses integer parameters n, k, s and the polynomial $T(X) \in \mathbb{F}[X]$ of degree $ns - 1$.
 2. Compute $\text{srs} := \text{gen}(n, k, s, T)$.
 3. \mathcal{A} sends a message com and polynomial $F(X) \in \mathbb{F}_{<ks}[X]$ such that $\text{com} = [F(\tau)]_1$. Here, for $\text{max} = \max(k, n)$, all \mathbb{G}_1 elements in the srs are linear combinations of $\{[\tau^i]_1\}_{i \in [0, \text{max}-1]}$.
 4. \mathcal{A} (as the prover) and the verifier run the interactive protocol $\text{segmentLookup}(\text{com}, \text{srs}, T, F, \mathbb{V})$, where $\mathbb{V} \subset \mathbb{F}$ is a subgroup of order ks generated by v .
 5. \mathcal{A} wins if and only if the verifier accepts and there exists some $i \in [0, k - 1]$, such that for each $j \in [0, n - 1]$ there exists at least one $q \in [0, s - 1]$ such that $F(v^{is+q}) \neq T(\omega^{js+q})$.

3.3 Building Blocks

We require two main building blocks for our segment-lookup protocol, which we state as two lemmas below. The first lemma gives a protocol that helps in checking if the evaluations of some polynomial are all n -th roots of unity, which can be instantiated using Caulk [ZBK⁺22a]. The second lemma is a modification of the log-derivative lemma from [EFG22, Hab22].

Lemma 5. *There exists a knowledge-sound argument for $\mathcal{R}_{\text{unity}} := \{(\text{srs}, [D]_1, \{v^{is} : i \in [0, k-1]\}, \{\mu^i : i \in [0, n-1]\}) : [D]_1 = [D(\tau)] \text{ for } D(x) \text{ s.t. } \forall i \in [0, k-1], D(v^{is}) = \mu^j \text{ for some } j \in [0, n-1]\}$, where μ is some n -th root of unity, under the algebraic group model. Moreover, the prover uses $O(ks \log n)$ \mathbb{G}_1 - and \mathbb{F} -operations, while the proof size and verifier cost are constant.*

We instantiate this lemma using the Multi-unity sub-protocol from Caulk [ZBK⁺22a, Theorem 4], an overview of which is given below. The aim of the protocol is to check if the evaluations of some polynomial are all roots of unity, i.e., prove that for each $i \in [0, k-1]$, $D(v^{is}) \in \{\mu^0, \dots, \mu^{n-1}\}$, where $\mu = \omega^s$ and $\{\mu^0, \dots, \mu^{n-1}\}$ are the n^{th} roots of unity. To do this, the prover defines a vector $\vec{u}_0 = (D(1), D(v^s), \dots, D(v^{(k-1)s}))$, and iteratively defines $\vec{u}_j = \vec{u}_{j-1} \circ \vec{u}_{j-1}$, for all $j \in [\log n]$. This means the following must be proved:

1. \vec{u}_0 consists of $(D(1), D(v^s), \dots, D(v^{(k-1)s}))$.
2. $\vec{u}_j = \vec{u}_{j-1} \circ \vec{u}_{j-1}$ for all $j \in [\log n - 1]$.
3. $\vec{u}_{\log n - 1} \circ \vec{u}_{\log n - 1} = \vec{1}$.

Proving these three statements would imply that the initial statement is proved to hold. To prove steps 1-3 in terms of the polynomial encodings, set $U_0(X) = D(X)$, the polynomial with evaluations \vec{u}_0 on \mathbb{K} , $U_{\log n}(X) = \text{id}(X)$, where id is the polynomial evaluating to 1 at all points of \mathbb{V} , and $U_j(X)$'s to encode the intermediate vectors \vec{u}_j 's. Then, it needs to be proved that the following equations hold.

$$\begin{aligned} U_0(X)U_0(X) - U_1(X) &\equiv Z_{\mathbb{V}}(X)Q_{U,1}(X) \\ U_1(X)U_1(X) - U_2(X) &\equiv Z_{\mathbb{V}}(X)Q_{U,2}(X) \\ &\vdots \\ U_{\log n - 1}(X)U_{\log n - 1}(X) - \text{id}(X) &\equiv Z_{\mathbb{V}}(X)Q_{U, \log n}(X) \end{aligned}$$

All the above checks are aggregated into one equation, using lagrange polynomials $\{\Delta_i\}_{i \in [\log n]}$ over roots of unity $\mathbb{S} = \{1, \phi, \dots, \phi^{\log n - 1}\}$ as follows.

$$\left(U_0^2(X)\Delta_1(Y) + \sum_{j=2}^{\log n} U_{j-1}^2(X) \cdot \Delta_j(Y) \right) - \left(\sum_{j=1}^{\log n - 1} U_j(X) \cdot \Delta_j(Y) + \text{id}(X) \cdot \Delta_{\log n}(Y) \right) = Z_{\mathbb{V}}(X)Q_2(X, Y)$$

for some polynomial $Q_2(X, Y)$. It needs to be shown that the above equation holds at some random point (ζ, ξ) using bivariate KZG. We refer the reader to Appendix A for the details of the above subprotocol.

Log-derivative Method. We use a variant of the log-derivative method from [Hab22, EFG22] that is tailored for our segment lookup. We state the lemma below, whose proof is similar to the log-derivate lemma from [Hab22].

Lemma 6. *Let \mathbb{F} be a field with characteristic $p > \max\{ks, ns\}$. Given three sequences of field elements $f = (f_i = F(v^i))_{i \in [ks]}$, $\ell = (\ell_i = L(v^i))_{i \in [ks]}$, and $t = (t_i = T(\omega^i))_{i \in [ns]}$. Consider segments of s -consecutive elements in f and t , resulting in k and n segments of f and t , respectively. We have that for each $i \in [0, k-1]$, the i -th segment of f corresponds to the ℓ_i -th segment of t , i.e., $\forall \delta \in \mathbb{F}$, for each $i \in [ks]$, there exists some $j \in [ns]$ such that $f_i + \delta \ell_i = t_j + \delta \omega^j$, if and only if, for some $m \in \mathbb{F}^{ns}$, the following identity of rational functions holds.*

$$\sum_{i \in [ns]} \frac{m_i}{X + t_i + \delta \omega^i} = \sum_{i \in [ks]} \frac{1}{X + f_i + \delta \ell_i} \quad (1)$$

Proof. Denote by $m_{f,\ell}(z)$, the multiplicity of field element z in the sequence $(f_i + \delta\ell_i)_{i \in [ks]}$. Since $p > \max\{ks, ns\}$, the multiplicities will be non-zero elements in \mathbb{F} . Suppose that for each $i \in [ks]$, there exists some $j \in [ns]$ such that for each $\delta \in \mathbb{F}$, it holds that $f_i + \delta\ell_i = t_j + \delta\omega^j$. Then, set $m_j = m_{f,\ell}(t_j + \delta\omega^j)$ for the $j \in [ns]$ corresponding to the i 's above, and to be 0 for the remaining j 's in $[ns]$. Clearly, for these choice of m_j 's, equation 1 holds.

Conversely, suppose equation 1 holds for each $\delta \in \mathbb{F}$. By collecting repeating terms of each of the summands in the field, we get:

$$\begin{aligned} \sum_{i \in [ns]} \frac{m_i}{X + t_i + \delta\omega^i} &= \sum_{z \in \mathbb{F}} \frac{m'(z)}{X + z}, \\ \sum_{i \in [ks]} \frac{1}{X + f_i + \delta\ell_i} &= \sum_{z \in \mathbb{F}} \frac{m_{f,\ell}(z)}{X + z}. \end{aligned}$$

where $m'(z)$ is $m_t(z)$ multiplied with the corresponding m_i 's. Since, $p > \max\{ks, ns\}$, for each $z \in \{f_i + \delta\ell_i\}$, we have $m_{f,\ell}(z) \neq 0$. By uniqueness of fractional representations (Lemma 2), $m_f(z) = m'(z)$ for each $z \in \{f_i\}$. Hence, for each $z \in \{f_i + \delta\ell_i\}$, there must exist some $j \in [ns]$ such that $z = t_j + \delta\omega^j$. \square

3.4 Tables with Repeated Elements

An astute reader may have noted that the Lemma 6 only works if within each *segment* the table T consists of unique values, i.e. $\forall i \in [0, n-1] \nexists j, j' \in [s]$ such that $T(\omega^{is+j}) = T(\omega^{is+j'})$. At a high level, this is implicit in the proof of Lemma 6. Particularly, if there is a repeated element in a segment, i.e., 2 elements in the table correspond to the same t_j , then the multiplicity count argument for proving equation 1 in the proof will no longer hold.

Looking ahead, in our application of the segment lookup protocol, we will indeed deal with tables that will contain repeated elements with a segment. We now describe how to generically start with any table $T(X)$, to (i) convert it into a table $T'(X)$ during the pre-processing step such that each segment consists of unique values - the prover will run the protocol using $F'(X)$ derived from $T'(X)$; and (ii) have the verifier perform a single step to *recover* $F(X)$ from $F'(X)$. The guarantee of our transformation will be that the *witness polynomial* is identical in both cases. Specifically, our segment lookup protocol when run on T' and F' ensures that \exists a function $\xi : [k] \rightarrow [N]$ such that $\forall i \in [0, k], j \in [s], F'(v^{is+j}) = T'(v^{\xi(i)s+j})$, and our transformation will ensure that $\forall i \in [0, k], j \in [s], F(v^{is+j}) = T(v^{\xi(i)s+j})$ for the *same* function ξ .

Let $T_{\max} := \max_i |T(\omega^i)|$, and we shall set $\forall i \in [0, n-1], j \in [s], T'(\omega^{is+j}) := T(\omega^{is+j}) + 2 \cdot j \cdot (T_{\max} + 1)$. This is done by further defining two polynomials, $E(X) = \sum_{i=0}^{n-1} \sum_{j \in [s]} 2 \cdot j \cdot (T_{\max} + 1) \psi_{is+j}^{\mathbb{W}}(X)$ and $D(X) = \sum_{i=0}^{k-1} \sum_{j \in [s]} 2 \cdot j \cdot (T_{\max} + 1) \psi_{is+j}^{\mathbb{V}}(X)$. We define the following function `makeUnique`:

`makeUnique(T, E):`

 Output $T'(X) := T(X) + E(X)$.

Within the context of our segment lookup protocol, $[T'(\tau)]$ and $[D(\tau)]$ are computed as a part of pre-processing by running `makeUnique(T, E)`, when given $T(X)$ as input - the rest of the pre-processing remains unchanged. The prover sets $F'(X)$ to be $F'(X) = F(X) + D(X)$, and the verifier sets $[F'(\tau)] = [F(\tau)] + [D(\tau)]$ to run `segmentLookup(com, srs, T', F', \mathbb{V})`, and verifier setting the “output” of the protocol to be $[F(\tau)]$. For correctness, we rely on the following claims.

Claim 1. *If $T_{\max} \cdot (2 \cdot s + 1) < |\mathbb{F}|$, then $\forall i \in [N] \nexists j, j' \in [s]$ such that $T'(\omega^{is+j}) = T'(\omega^{is+j'})$.*

Proof Sketch. If not, for any distinct $j, j' \in [s]$ consider $T'(\omega^{is+j}) - T'(\omega^{is+j'}) = T(\omega^{is+j}) - T(\omega^{is+j'}) + 2(T_{\max} + 1)(j - j')$. Since by definition, $T(\omega^{is+j}) - T(\omega^{is+j'}) \in [-2T_{\max}, 2T_{\max}]$, we have that $T'(\omega^{is+j}) - T'(\omega^{is+j'}) > 0$ if $T_{\max} \cdot (2 \cdot s + 1) < |\mathbb{F}|$, and thus $T'(\omega^{is+j}) \neq T'(\omega^{is+j'})$. \square

Note that we are assuming that $T_{\max} \cdot (2 \cdot s + 1) < |\mathbb{F}|$ to avoid any “wrap-around” in the field. This restriction will not hinder the applications we consider.

Claim 2. \exists a function $\xi : [k] \rightarrow [N]$ such that $\forall i \in [0, k], j \in [s], F'(v^{is+j}) = T'(v^{\xi(i)s+j})$, then $\forall i \in [0, k], j \in [s], F(v^{is+j}) = T(v^{\xi(i)s+j})$.

The proof of the above claim follows from the definition of the relevant polynomials.

For the rest of this work, we will assume that we invoke `segmentLookup`, we are running the above modified version of the segment lookup protocol that deals with repeated elements within a segment. Thus, we ignore this issue henceforth.

3.5 Protocol

We now give a formal description of our segment-lookup protocol.

gen(n, k, s, T): Given n, k, s and the polynomial $T(X) \in \mathbb{F}[X]$ of degree $ns - 1$, the pre-processing information is computed as follows:

The pre-processing involves computing all the powers of a random τ (step 1) which is the only step of pre-processing requiring a trusted setup. In the remaining steps, these trusted powers of τ are used to generate the commitments of the relevant vanishing polynomials, the Lagrange polynomials, and some quotient polynomials for reducing prover and verifier work. These quotient polynomial computations are key to reducing the online prover cost.

1. Choose a random $\tau \in \mathbb{F}$. Let $\max = \max(k, n)$. Compute $\{[\tau^i]_1\}_{i \in [0, \max \cdot s - 1]}$ and $\{[\tau^i]_2\}_{i \in [0, \max \cdot s - 1]}$.
2. Compute $[Z_{\mathbb{W}}(\tau)]_2, [Z_{\mathbb{V}}(\tau)]_2$ and $[Z_{\mathbb{K}}(\tau)]_2$.
3. Compute and output $[T(\tau)]_2$.
4. For $i \in [ns]$, use lemma 4 to compute:
 - (a) $q_{i,1} = [Q_{i,1}(\tau)]_1$ and $q_{i,2} = [Q_{i,2}(\tau)]_1$ such that
$$\psi_i^{\mathbb{W}}(X) \cdot T(X) = t_i \cdot \psi_i^{\mathbb{W}}(X) + Z_{\mathbb{W}}(X) \cdot Q_i(X)$$

$$\psi_i^{\mathbb{W}}(X) \cdot X = \omega^i \cdot \psi_i^{\mathbb{W}}(X) + Z_{\mathbb{W}}(X) \cdot Q_{i,2}(X)$$
 - (b) $[\psi_i^{\mathbb{W}}(\tau)]_1$.
 - (c) $[\frac{\psi_i^{\mathbb{W}}(\tau) - \psi_i^{\mathbb{W}}(0)}{\tau}]_1$.
5. For $i \in [ks]$, compute the commitments $[\psi_i^{\mathbb{V}}(\tau)]_1, [\psi_i^{\mathbb{V}}(\tau\nu)]_1$, for all $i \in [ks]$, corresponding to the set \mathbb{V} .
6. For $i \in [ns]$, use lemma 4 to compute:
 - (a) $q_{i,3} = [Q_{i,3}(\tau)]_1$ and $q_{i,4} = [Q_{i,4}(\tau)]_1$ such that:
$$\psi_i^{\mathbb{W}}(X) \cdot (X^n - 1) = g_i \psi_i^{\mathbb{W}}(X) + Z_{\mathbb{W}}(X) \cdot Q_{i,3}(X)$$

$$\psi_i^{\mathbb{W}}(X/\omega) \cdot (X^n - 1) = g_i \psi_i^{\mathbb{W}}(X/\omega) + Z_{\mathbb{W}}(X) \cdot Q_{i,4}(X),$$

where $g_i := (\omega^{in} - 1)$.
 - (b) $[\psi_i^{\mathbb{W}}(\tau/\omega)]_1$.

Output srs: $\{[\tau^i]_1, [\tau^i]_2\}_{i \in [0, \max \cdot s - 1]}$, $[Z_{\mathbb{W}}(\tau)]_2, [Z_{\mathbb{V}}(\tau)]_2, [Z_{\mathbb{K}}(\tau)]_2, [T(\tau)]_2, q_{i,1}, q_{i,2}, q_{i,3}, q_{i,4}, [\psi_i^{\mathbb{W}}(\tau)]_1, [\psi_i^{\mathbb{W}}(\tau/\omega)]_1$, and $[\frac{\psi_i^{\mathbb{W}}(\tau) - \psi_i^{\mathbb{W}}(0)}{\tau}]_1$ for each $i \in [ns]$, and $[\psi_i^{\mathbb{V}}(\tau)]_1, [\psi_i^{\mathbb{V}}(\tau\nu)]_1$ for each $i \in [ks]$.

segmentLookup(com, srs, T, F, \mathbb{V}): This protocol proceeds as follows:

In Round 1, the prover first commits to the multiplicity vector m from eq. 1. Then, the prover computes the vector ℓ , such that ℓ_i corresponds to the segment of t that matches the i -th segment of f , and generates commitments of two polynomials encoding $\ell_i: L$, which helps in checking that within each of the i -th segment, the consecutive elements of f correspond to consecutive elements of t in the ℓ_i -th segment, and D , which helps in checking that L is well-formed, i.e., the first entries of each segment form a set of n -th roots of unity. Note here that all the commitments on τ (which the prover does not know) are computed by taking a linear combination of the corresponding Lagrange commitments, all of which are given in the srs.

Round 1 (Prover \rightarrow Verifier). The prover does the following:

1. Computes a polynomial M of degree $ns - 1$ as follows: For each $i \in [0, n - 1]$, if i^{th} segment in t is executed y times in $F|_{\mathbb{V}}$, then for each $j \in [0, s - 1]$, $M(\omega^{is+j}) = y$, and sends $[M(\tau)]_1$ and $[M(\tau/\omega)]_1$ to the verifier.
2. Computes and sends $[Q_M(\tau)]_1$ using the srs and lemma 4, where $Q_M(X)$ is such that¹¹

$$(X^n - 1)(M(X) - M(X/\omega)) = Z_{\mathbb{W}}(X)Q_M(X) \quad (2)$$

where $Z_{\mathbb{W}}(X)$ denotes the vanishing polynomial corresponding to the set \mathbb{W} .

3. Computes $L(X)$ of degree $ks - 1$ as follows: For each $i \in [0, k - 1]$, if the segment executed at the i^{th} step (i.e., the i^{th} segment in $f = F|_{\mathbb{V}}$) is the j^{th} segment in t , then for each $q \in [0, s - 1]$, $L(v^{is+q}) = \omega^{js+q}$, and sends $[L(\tau)]_1$ and $[L(\tau v)]_1$ to the verifier.
4. Computes $Q_L(X)$ such that

$$(X^k - 1)(L(Xv) - \omega L(X)) = Z_{\mathbb{V}}(X)Q_L(X) \quad (3)$$

where $Z_{\mathbb{V}}(X)$ denotes the vanishing polynomial corresponding to the set \mathbb{V} , and sends $[Q_L(\tau)]_1$ to the verifier.

5. Computes a degree $k - 1$ polynomial $D(X)$ such that for each $i \in [0, k - 1]$, $D(v^{is}) = L(v^{is})$, and sends $[D(\tau)]_1$ to the verifier.
6. Computes Q_D such that

$$L(X) - D(X) = Z_{\mathbb{K}}(X)Q_D(X) \quad (4)$$

where $Z_{\mathbb{K}}(X)$ denotes the vanishing polynomial corresponding to the set \mathbb{K} , and sends $[Q_D(\tau)]_1$ to the verifier.

In Round 2, the verifier checks validity of the equation 2 at τ . For soundness, it is important that the eqs 3 and 4 corresponding to L and D in Round 1 are checked at random challenges. For efficiency reasons, we combine these checks with the other checks in rounds 14-15.

¹¹Given $q_{i,3}, q_{i,4}, [\psi_i^{\mathbb{W}}(\tau)]_1, [\psi_i^{\mathbb{W}}(\tau/\omega)]_1$, for each $i \in [ns]$, we can slightly modify the step 2 of lemma 4 to compute a linear combination and obtain $[Q_M(\tau)]_1$ and the commitment of the remainder polynomials corresponding to the ks -sparse polynomial $H(X) := (M(X) - M(X/\omega))$.

Round 2 (Verifier). The verifier proceeds to check the following:

$$e([M(\tau)]_1 - [M(\tau/\omega)]_1, [\tau^n]_2 - [1]_2) = e([Q_M(\tau)]_1, [Z_{\mathbb{W}}(\tau)]_2)$$

In Rounds 3-8, the prover and verifier run the sub-protocol from [ZBK⁺ 22a, Figure 5], as described in Section 3.3 to check that the first entries of each segment $i \in [0, k-1]$ of f , form the n -th roots of unity v^0, \dots, v^{n-1} , for $v = \omega^s$. Recall, this was needed to guarantee the correctness of L .

Round 3-8 (Prover \longleftrightarrow Verifier). Using the instantiation of Lemma 5, the prover and verifier engage in a protocol to prove that the polynomial L is well-formed, i.e., for each $i \in [0, k-1]$, $D(v^{is}) \in \{\mu^0, \dots, \mu^{n-1}\}$, where $\mu = \omega^s$ and $\{\mu^0, \dots, \mu^{n-1}\}$ are the n^{th} roots of unity.

In Rounds 8-10, the prover computes the polynomials A and B corresponding to the two summations in equation 1, and sends the needed commitments for the verifier to check the correctness of A at τ , and to perform degree-check of the polynomial that has the least degree amongst A and B at τ . No such degree check is needed if $k = n$.

Round 9 (Verifier). The verifier sends random $\beta, \delta \in \mathbb{F}$ to the prover.

Round 10 (Prover \rightarrow Verifier). The prover does the following:

1. The prover computes $A(X)$ of degree $ns - 1$ such that for each $i \in [0, ns - 1]$,

$$A(\omega^i) = \frac{M(\omega^i)}{\beta + T(\omega^i) + \delta\omega^i}$$

and sends $[A(\tau)]_1$ to the verifier.

2. The prover computes $[Q_A(\tau)]_1$ using the srs and step 2 of lemma 4, where $Q_A(X)$ is such that¹²

$$A(X)(\beta + T(X) + \delta X) - M(X) = Z_{\mathbb{W}}(X)Q_A(X).$$

3. The prover computes $B(X)$ of degree $ks - 1$ such that for each $i \in [0, ks - 1]$,

$$B(v^i) = \frac{1}{\beta + F(v^i) + \delta L(v^i)}$$

and sends $[B(\tau)]_1$ to the verifier.

4. The prover computes $Q_B(X)$ such that

$$B(X)(\beta + F(X) + \delta L(X)) - 1 = Z_{\mathbb{V}}(X)Q_B(X),$$

and sends $[Q_B(\tau)]_1$ to the verifier.

5. The prover computes $B_0(X) = \frac{B(X) - B(0)}{X}$, $A_0(X) = \frac{A(X) - A(0)}{X}$, and sends $[A_0(\tau)]_1, [B_0(\tau)]_1$ to the verifier.

6. **For degree check**, which is needed only for $k \neq n$,

if $n > k$: the prover computes $P_B(X) = B_0(X) \cdot X^{ns - (ks+1)}$ and sends $[P_B(\tau)]_1$ to the verifier;

else if $k > n$: the prover computes $P_A(X) = A_0(X) \cdot X^{ks - (ns+1)}$ and sends $[P_A(\tau)]_1$ to the verifier.

¹²Given $q_{i,1}, q_{i,2}$, and $[\psi_i^{\mathbb{W}}(\tau)]_1$, for each $i \in [ns]$, we can slightly modify the step 2 of lemma 4 to obtain $[Q_A(\tau)]_1$ and the commitment of the remainder polynomials: compute a linear combination of the quotients and remainders corresponding to the division of $A(X)T(X)$ and $A(X)X$ by $Z_{\mathbb{W}}(X)$.

Round 11 (Verifier). The verifier proceeds as follows:

1. Checks that A encodes the correct values:

$$e([A(\tau)]_1, [T(\tau)]_2 + \delta[\tau]_2) = e([Q_A(\tau)]_1, [Z_{\mathbb{W}}(\tau)]_2) \cdot e([M(\tau)]_1 - \beta[A(\tau)]_1, [1]_2)$$

2. **Degree Check:** If $n > k$, check

$$e([B_0(\tau)]_1, [\tau^{ns-ks-1}]_2) = e([P_B(\tau)]_1, [1]_2)$$

else if $k > n$, check

$$e([A_0(\tau)]_1, [\tau^{ks-ns-1}]_2) = e([P_A(\tau)]_1, [1]_2)$$

3. Samples random $\gamma \in \mathbb{F}$ and sends them to the prover.

In Rounds 12-15, the prover gives a correctness proof of B at a random challenge. As we mentioned before, we combine this check with the correctness checks of eqs 3 and 4 to reduce the number of pairing checks and give the KZG-opening proofs corresponding to all the involved polynomials evaluated at the random challenge.

Round 12 (Prover). The prover sends $b_{0,\gamma} = B_0(\gamma)$, $f_\gamma = F(\gamma)$, $\ell_\gamma = L(\gamma)$, $a_0 = A(0)$, $\ell_{\gamma,v} = L(v\gamma)$, $q_{\gamma,L} = Q_L(\gamma)$, $d_\gamma = D(\gamma)$, $q_{\gamma,D} = Q_D(\gamma)$. to the verifier¹³.

Round 13 (Verifier). Verifier samples a random $\eta \in \mathbb{F}$ and sends to the prover.

Round 14 (Prover). The prover computes the following:

$$P(X) = L(Xv) + \eta L(X) + \eta^2 Q_L(X) + \eta^3 D(X) + \eta^4 Q_D(X) + \eta^5 B_0(X) + \eta^6 F(X) + \eta^7 Q_B(X) \text{ and}$$

$$H_P(X) = \frac{P(X) - p_\gamma}{X - \gamma}$$

where $p_\gamma = \ell_{\gamma,v} + \eta \ell_\gamma + \eta^2 q_{\gamma,L} + \eta^3 d_\gamma + \eta^4 q_{\gamma,D} + \eta^5 b_{0,\gamma} + \eta^6 \cdot f_\gamma + \eta^7 q_{B,\gamma}$, for $b_\gamma = b_{0,\gamma} \cdot \gamma + B(0)$ and

$$q_{B,\gamma} = \frac{b_\gamma(f_\gamma + \beta + \delta \ell_\gamma) - 1}{Z_V(\gamma)}.$$

The prover then sends $[H_P(\tau)]_1$ and $[P(\tau)]_1$ to the verifier.

Round 15 (Verifier). The verifier proceeds as follows:

1. Sets $b_0 = ns \cdot a_0 / ks$.

2. As part of checking the correctness of B , it computes $Z_V(\gamma) = \gamma^{ks} - 1$, $b_\gamma = b_{0,\gamma} \cdot \gamma + b_0$ and

$$q_{B,\gamma} = \frac{b_\gamma(f_\gamma + \beta + \delta \ell_\gamma) - 1}{Z_V(\gamma)}$$

3. Computes $p_\gamma = \ell_{\gamma,v} + \eta \ell_\gamma + \eta^2 q_{\gamma,L} + \eta^3 d_\gamma + \eta^4 q_{\gamma,D} + \eta^5 b_{0,\gamma} + \eta^6 \cdot \text{com} + \eta^7 q_{B,\gamma}$.

¹³Given the sparse representation of A , the prover can compute $A(0)$ in time $O(ks)$, given the pre-computed Lagrange evaluations at 0.

4. Checks the following:

$$\begin{aligned}
e([H_p(\tau)]_1, [\tau]_2) &= e([P(\tau)]_1 - p_\gamma + \gamma[H_p(\tau)]_1, [1]_2) \\
e([A(\tau)]_1 - [a_0]_1, [1]_2) &= e([A_0(\tau)]_1, [\tau]_2) \\
(\gamma^k - 1)(\ell_{\gamma, \nu} - \omega \ell_\gamma) - Z_{\mathbb{V}}(\gamma)q_{\gamma, L} &= 0 \\
\ell_\gamma - d_\gamma - Z_{\mathbb{K}}(\gamma)q_{\gamma, D} &= 0
\end{aligned}$$

Proof. It is easy to check that our protocol is complete, assuming the completeness of the caulk sub-protocol.

Proof of Knowledge Soundness. Suppose \mathcal{A} is an efficient algebraic adversary attacking the knowledge soundness of our protocol, as in Definition 5. Since \mathcal{A} is algebraic, it will send a polynomial $F(X) \in \mathbb{F}_{<ks}[X]$ corresponding to com in Step 3) of the knowledge soundness game. Furthermore, corresponding to all the commitments in the protocol, \mathcal{A} sends the corresponding polynomials of appropriate degree. Let Win denote the event that \mathcal{A} wins the knowledge soundness game and Acc denote the event that the verifier accepts. Then, $\text{Win} \subset \text{Acc}$ and Acc implies that all the pairing checks in our protocol verify. By Lemma 1, this means that the corresponding ideal checks will also verify, except with a $\text{negl}(\lambda)$ -probability. Assuming this, we want to prove that for each $i \in [0, k-1]$, there exists $j \in [0, n-1]$ such that, for each $q \in [0, s-1]$, $F(\nu^{is+q}) = T(\omega^{js+q})$.

If all the ideal pairing checks verify, then the following holds.

- By the security of KZG commitments, the pairing checks in Round 15 implies that except with probability $ks/|\mathbb{F}|$ over $\gamma, \eta \in \mathbb{F}$, it holds that $\ell_{\gamma, \nu} = L(\gamma\nu)$, $\ell_\gamma = L(\gamma)$, $q_{\gamma, L} = Q_L(\gamma)$, $d_\gamma = D(\gamma)$, $q_{\gamma, D} = Q_D(\gamma)$, $b_{0, \gamma} = B_0(\gamma)$, $f_\gamma = F(\gamma)$ and $q_{B, \gamma} = Q_B(\gamma)$.
- This implies that the checks in Round 15, Step 4) at the random point γ implies that the following equations hold:

$$\begin{aligned}
B(X)(\beta + F(X) + \delta L(X)) - 1 &= Z_{\mathbb{V}}(X)Q_B(X) \\
(X^k - 1)(L(X\nu) - \omega L(X)) &= Z_{\mathbb{V}}(X)Q_L(X) \\
L(X) - D(X) &= Z_{\mathbb{K}}(X)Q_D(X)
\end{aligned}$$

- By the knowledge soundness of the caulk sub-protocol (c.f. [ZBK⁺22a, Theorem 4]), Rounds 3-8 in our protocol guarantee that L is well-formed, i.e., for each $i \in [0, k-1]$, $D(\nu^{is}) \in \{\mu^0, \dots, \mu^{n-1}\}$, for $\mu = \omega^s$.
- Since the verifier obtains $[\tau^{ks-ns-1}]_2, [\tau^{ns-ks-1}]_2$ from the srs, the degree pairing checks in Round 11 imply the the following:
 - If $n > k$, then $\deg(B_0) \leq ks - 2$, and if $k > n$, then $\deg(A_0) \leq ns - 2$.
 - In either case above, the other polynomial also satisfies the degree check because the commitment of the corresponding power of τ will be the highest \mathbb{G}_1 -power in srs. This is also why no degree check is needed if $k = n$. Furthermore, since $B(X) = B_0(X)X + b_0$ for $b_0 = nsa_0/ks$, $\deg(B) < ks$.
- In Round 11, we check the following equation at τ :

$$A(X)(\beta + T(X) + \delta X) - M(X) = Z_{\mathbb{W}}(X)Q_A(X)$$

This implies that $A(\omega^i) = \frac{M(\omega^i)}{\beta + T(\omega^i) + \delta \omega^i}$ for each $i \in [ns]$. Furthermore, since B is well-formed, we know that $\sum_{i \in [ks]} B(\nu^i) = \sum_{i \in [ks]} \frac{1}{\beta + F(\nu^i) + \delta L(\nu^i)}$. Further, by Lemma 3, $\sum_{i \in [ks]} B(\nu^i) = ks \cdot b_0$ and

$\sum_{i \in [ns]} A(\omega^i) = ns \cdot a_0$. Since, in the protocol b_0 is set such that $ks \cdot b_0 = ns \cdot a_0$, we have that $\sum_{i \in [ns]} \frac{M(\omega^i)}{\beta + T(\omega^i) + \delta\omega^i} = \sum_{i \in [ks]} \frac{1}{\beta + F(v^i) + \delta L(v^i)}$. This implies that except with probability $(ks \cdot ns)/|\mathbb{F}|$ over $\beta \in \mathbb{F}$, the following will hold:

$$\sum_{i \in [ns]} \frac{M(\omega^i)}{X + T(\omega^i) + \delta\omega^i} = \sum_{i \in [ks]} \frac{1}{X + F(v^i) + \delta L(v^i)} \quad (5)$$

This implies that the following two equation checks on τ in Rounds 2 and 15,

$$\begin{aligned} (X^n - 1)(M(X) - M(X/\omega)) &= Z_{\mathbb{W}}(X)Q_M(X) \\ A(X) - a_0 &= A_0(X) \cdot X \end{aligned}$$

will guarantee that these equations hold and that M is guaranteed to be well-formed¹⁴. Thus, by Lemma 6, eq. 5 implies that for each $i \in [ks]$, there exists some $j \in [ns]$ such that $F(v^i) + \delta L(v^i) = T(\omega^j) + \delta\omega^j$.

- Thus, by the definition of polynomial L , this implies that for each $i \in [0, k-1]$, there exists a $j \in [0, n-1]$ such that for each $q \in [0, s-1]$, $F(v^{is+q}) + \delta\omega^{js+q} = T(\omega^{js+q}) + \delta\omega^{js+q}$, i.e., $F(v^{is+q}) = T(\omega^{js+q})$.

Thus, the event Acc that the verifier accepts implies that all the ideal pairing checks will verify except with a $\text{negl}(\lambda)$ probability. By the above implications, this in turn implies that for each $i \in [0, k-1]$, there exists a $j \in [0, n-1]$ such that for each $q \in [0, s-1]$, $F(v^{is+q}) = T(\omega^{js+q})$, except with a $\text{negl}(\lambda)$ probability. This proves the knowledge soundness of our protocol.

Efficiency. The efficiency of gen follows from Lemma 4. The additional multiplicative log-factor is due to the use of FFTs to compute the quotient polynomials. By Lemma 4 it is guaranteed that the online prover cost of our protocol is dominated by the following two costs: first, the instantiation of Lemma 5 uses $O(ks \log n)$ \mathbb{G}_1 - and \mathbb{F} -operations; second, the FFT computation for computing the quotient polynomials, on \mathbb{V} requires $O(ks \log ks)$ \mathbb{F} -operations. It is easy to see that the proof size and verifier cost of our protocol are $O(1)$. □

3.6 Adding Zero-Knowledge

Our segment-lookup protocol as described in the previous section, does not achieve *zero-knowledge* (i.e., it does not hide the segments encoded using polynomial $F(X)$). However, as is the case with most existing efficient SNARKs [GWC19, EFG22, CHM⁺20], our protocol can be easily modified to achieve zero-knowledge with the caveat that it *leaks the effective layer k of the input*. The main idea is to use randomized polynomial encodings (as opposed to unique polynomial encodings) when computing commitments to witness-dependent vectors. The reason why this simple modification suffices for ensuring zero-knowledge is because, throughout the protocol, the verifier only sees evaluations of these polynomial encodings at a few random points. When using higher-degree randomized polynomial encodings, these polynomial evaluations do not leak any information about the encoded vector.

In more detail, if for instance the verifier receives α evaluations of polynomial $F(X)$ in our segment-lookup protocol, then it suffices for the prover to commit to and work with a randomized polynomial $\hat{F}(X) = F(X) + R(X) \cdot Z_{\mathbb{V}}(X)$ instead of $F(X)$, where $R(X)$ is a random polynomial of degree $\alpha - 1$. Our segment-lookup protocol will achieve zero-knowledge if all witness-dependent polynomials in the

¹⁴if M was not correctly formed, and the pairing check of round 2 was set to be true by \mathcal{A} , then the equality guaranteed by eq. 5 will not hold.

above protocol are randomized in a similar manner. For simplicity of presentation, we chose to present the protocol in the previous section without the zero-knowledge property. Since the primary focus of our work was to reduce the prover run-time, this allowed us to highlight the main technical ideas without having to deal with unnecessarily complex notation.

4 Sublonk: Segment Lookup + $\mathcal{P}\text{IonK}$

In this section, we will utilize the segment-lookup protocol described in Section 3 in conjunction with the $\mathcal{P}\text{IonK}$ proof system to construct a proof system where, as described in the introduction, the prover cost grows with the length of the execution path, rather than the entire circuit. We elaborate on this requirement now. Consider a layered branching circuit C with at most \bar{k} layers, where each layer has the same branch of n s -sized circuits C_1, \dots, C_n . On a given input, let it be the case that a sequence of $k \leq \bar{k}$ circuit branches $\tilde{C} := (\tilde{C}^{(i)})_{i \in [k]}$ (where k , as well as the sequence, is possibly dependent on the input), are executed in order. We want the prover to cost to grow with the size of the total *executed* sub-circuit of size ks , rather than kns . For simplicity of notation, throughout the rest of this section, we will assume that the effective number of layers k is fixed in advance (and hence even the pre-processing algorithm will take k as input). At the end of Section 4.3, we discuss how our protocol can be generalized to handle any $k \leq \bar{k}$ and hence the pre-processing algorithm of SubIonK only needs to take \bar{k} as input.

The core idea underlying our solution is to store the encoding of the circuit constraints for each sub-circuit in a segment-lookup table such that the prover only needs to access the constraints determined by the sequence \tilde{C} , and further prove to the verifier that the appropriate constraints were selected from the table. In more detail, as discussed in Section 2.7, the $\mathcal{P}\text{IonK}$ proof system encodes circuit constraints via a number of polynomials that are pre-processed to their KZG commitment (see Section 2.3) and provided to the verifier. Since the output of our segment-lookup protocol in Section 3 is indeed a KZG commitment to the polynomial defined by the selected segments, this gives the following natural approach to achieving a sub-linear prover:

1. We represent each of the circuit branches C_1, \dots, C_n using the $\mathcal{P}\text{IonK}$ constraint system. Generate a table T where the i -th segment contains the $\mathcal{P}\text{IonK}$ constraints for C_i , i.e. T consists of n segments.
2. Depending on the sequence of activated circuit branches, the prover can use our segment-lookup protocol (from Section 3) to generate the KZG commitment of the k constraints corresponding to the k activated circuit branches along with a proof of correctness.
3. The prover post-processes the above derived KZG commitment so that it has the same form as the KZG commitments received by a $\mathcal{P}\text{IonK}$ verifier in the pre-processing phase for \tilde{C} .
4. Given the above, we can directly rely on the $\mathcal{P}\text{IonK}$ protocol for \tilde{C} .

Unfortunately, while this is indeed the template for our protocol, the details do not work out in such a straightforward manner. We present the details in this section.

4.1 Pre-Processing Layered Branching Circuit

In this section, we describe how the $\mathcal{P}\text{IonK}$ constraints for each of the circuit branches C_1, \dots, C_n is stored in a table T consisting of n segments. We start by establishing some notation.

Ensuring Correct Sequence of Activated Circuit Branches. The *activated sub-circuit* \tilde{C} (of C on input x) of size ks is specified by a sequence of k activated circuit branches $(\tilde{C}^{(i)})_{i \in [k]}$. Further, as discussed in Section 2.6, \exists a function $\xi : \mathbb{F}^m \rightarrow [1, n]^k$ such that $\forall i \in [k]$, $\tilde{C}^{(i)} = C_{\xi_x(i)}$ (where we use ξ_x to denote ξ_x). We assume that each circuit branch when executed, specifies the next circuit branch to be executed in the next layer. To ensure that the circuit branches are executed in order, we assume that one of the output wires for $\tilde{C}^{(i)}$ outputs $\xi_x(i+1)$, and each circuit $C^{(i)}$ has hardcoded within it the index i to check whether the aforementioned incoming wire into $\xi_x(i+1)$ indeed has the hardwired value $\xi_x(i+1)$.

$\mathcal{P}\text{Ion}\mathcal{K}$ constraints. Let us denote the n $\mathcal{P}\text{Ion}\mathcal{K}$ constraint systems for the n circuit branches $\{C_i\}_{i \in [n]}$ as $\{\mathcal{C}^{(i)} = (\mathcal{V}^{(i)}, \mathcal{Q}^{(i)})\}_{i \in [n]}$. The corresponding pre-processed polynomials (refer to Section 2.7 for notation) are denoted by $(q_M^{(i)}(X), q_L^{(i)}(X), q_R^{(i)}(X), q_O^{(i)}(X), q_C^{(i)}(X), S_{\sigma_1}^{(i)}(X), S_{\sigma_2}^{(i)}(X), S_{\sigma_3}^{(i)}(X))$. These polynomials will be relevant to our subsequent discussion.

Let $\{\mathbf{x}_\ell\}_{\ell \in [3ks]}$ be the extended witness/satisfying assignment for the $\mathcal{P}\text{Ion}\mathcal{K}$ constraints such that for each $j \in [k]$, $\{\mathbf{x}_{js+i}, \mathbf{x}_{ks+js+i}, \mathbf{x}_{2ks+js+i}\}_{i \in [s]}$ correspond to the extended-witnesses for branching circuits $\tilde{C}^{(j)} (= C^{(\xi_x(j))})$.

Checking Wire-Consistency or Copy-Check Constraints. Recall from Section 2.7, that the consistency of wire values in a circuit is checked in $\mathcal{P}\text{Ion}\mathcal{K}$ using the copy-check constraints. In our setting, in addition to checking these constraints within each individual circuit branch, we need an additional copy-constraint. In particular, in our setting, the j -th output wire of a circuit $\tilde{C}^{(i)}$ must have the same value as the j -th input wire of $\tilde{C}^{(i+1)}$. We ensure this by assuming that the output wires of $\tilde{C}^{(i)}$ are *not* a part of any cycles induced by the copy-check constraints $\sigma^{(i)}$ within $\tilde{C}^{(i)}$. For instance, this requirement can be satisfied by considering “output gates” in the final layer of the circuit that pass values unchanged, but ensure the input and output wire remain in different cycles within the permutation. Finally, since the (relative) index of the j -th output wire of $\tilde{C}^{(i)}$ is fixed and known to $\tilde{C}^{(i+1)}$, the copy-constraint for $\tilde{C}^{(i+1)}$ is constructed ensuring that the j -th output wire of $\tilde{C}^{(i)}$ and j -th input wire of $\tilde{C}^{(i+1)}$ are a part of the same cycle in $\sigma^{(i+1)}$. Thus in each circuit branch C_i , if there are p input and output wires, the permutation within the circuit-branch is defined to be $\sigma^{(i)} : [-3p : 3(s-p)] \rightarrow [-3p : 3(s-p)]$.

This avoids any potential for cycles across circuit branches, except for the ones we discussed, and thus \tilde{C} consists of disjoint cycles, and thus $\{\sigma^{(i)}\}_{i \in [n]}$ remains a permutation. For simplicity of notation, for the remainder of the section, we treat $\sigma^{(i)}$ to be $\sigma^{(i)} : [3s] \rightarrow [3s]$.

Parallel Repeated Version of the Segment-Lookup Protocol. Before the next discussion, we remark that in our final protocol, we will run 9 copies of our segment-lookup protocol `segmentLookup`. In particular, we will encode each of the 8 types of polynomials $Y \in \{q_M, q_L, q_R, q_O, q_C, S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}\}$ for all the circuit branches in a separate table (to be described shortly) and run our segment lookup protocol on each of these 8 tables individually. Let $\{T_i\}_{i \in [8]}$ denote the 8 tables. Each of these T_i s has identical sizes and an identical number of segments. Further, even though we run separate instances of `segmentLookup` on each of them, we require that the *same witness polynomial* $M(X)$ polynomial is used across each copy of `segmentLookup`(T_i) execution. Note that this can be done naively by running the base `segmentLookup` protocol many times and have the prover sending a *single* round 1 message (which includes the KZG commitment to $[M(\tau)]_1$), which is then subsequently shared across each execution of the protocol.

But we utilize the fact that our segment-lookup protocol inherits linear homomorphism from `cq`. Specifically, once the prover sends the commitments to the various $F(X)$ polynomials, say $F_1(X)$ and $F_2(X)$ for tables $T_1(X)$ and $T_2(X)$, the verifier samples a random challenge γ , the prover and verifier can run the segment-lookup protocol on $F_1(X) + \gamma F_2(X)$ for the table $T_1(X) + \gamma T_2(X)$. With this observation, we only need to run one copy of the `segmentLookup` protocol internally.

Polynomials for Segment-Lookup. We are finally ready to describe how we encode the constraints for each circuit branch $\{C_i\}_{i \in [n]}$ in the 8 tables. We define the following 8 *table* polynomials of degree $ns - 1$ in their value representation form - for $Y \in \{q_M, q_L, q_R, q_O, q_C, S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}\}$,

$$\forall i \in [0, n - 1], j \in [s], T_Y(\omega^{is+j}) := Y^{(i)}(\eta^j)$$

where the set $\{0, \eta, \eta^2, \dots, \eta^{(s-1)}\} := \{0, \omega^n, \omega^{2n}, \dots, \omega^{(s-1)n}\}$ are the s -th roots of unity. We use \mathbb{W} and \mathbb{V} as defined in Section 3 to be the set of ns and ks roots of unity respectively. If we view T_Y as consisting of n *segments* each of size s , the above equation indicates that the i -th segment consists of $\{Y_j^{(i)}\}_{j \in [s]}$. Similarly, we define the following 8 $ks - 1$ -degree polynomials, for $Y \in \{q_M, q_L, q_R, q_O, q_C, S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}\}$,

$$\forall i \in [0, k - 1], j \in [s], F_Y(v^{is+j}) := Y^{(\xi(i))}(\eta^j)$$

For each $Y \in \{q_M, q_L, q_R, q_O, q_C, S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}\}$, we run instances of `segmentLookup` as described above. For clarity of exposition, we abuse notation and denote this as the P and V executing `segmentLookup` ^{$\otimes \ell$} ($\{T_i\}_{i \in [8]}$) (shortening from `segmentLookup` ^{$\otimes \ell$} ($\{\text{com}_i, T_i, F_i\}_{i \in [8]}, \text{srs}, \mathbb{V}\}$), where the verifier output is $\{[F_i(\tau)]_1\}_i$ if the proof accepts.

In Section 4.2, we describe how the F_Y polynomials, can be easily *post-processed* to obtain a verifier pre-processing identical to that in $\mathcal{P}\text{IonK}$ for the activated sub-circuit \tilde{C} .

4.2 Post-Processing F_Y Polynomials

In this section, we describe how the F_Y polynomials, can be *post-processed* to obtain a verifier pre-processing identical to that in $\mathcal{P}\text{IonK}$ for the activated sub-circuit \tilde{C} . Let $\{\widetilde{q}_M, \widetilde{q}_L, \widetilde{q}_R, \widetilde{q}_O, \widetilde{q}_C, \widetilde{S}_{\sigma_1}, \widetilde{S}_{\sigma_2}, \widetilde{S}_{\sigma_3}\}$ denote the $\mathcal{P}\text{IonK}$ constraints for \tilde{C} . We prove the following claims.

Selector Polynomials. For $Y \in \{q_M, q_L, q_R, q_O, q_C\}$, we observe that the corresponding F_Y polynomials are exactly the selector polynomials $\{\widetilde{q}_M, \widetilde{q}_L, \widetilde{q}_R, \widetilde{q}_O, \widetilde{q}_C\}$.

Claim 3. For $Y \in \{q_M, q_L, q_R, q_O, q_C\}$, $F_Y(X) = \widetilde{Y}(X)$.

Proof Sketch. We sketch the proof for $\widetilde{q}_M(X)$ and $F_{q_M}(X)$, and the proof extends identically to the other polynomials in the claim. For any $i \in [0, k - 1], j \in [s]$, it is the case that $\widetilde{q}_M(v^{is+j}) = q_M^{(\xi(i))}(\eta^j) = F_{q_M}(v^{is+j})$, where the first equality follows from the fact that $\tilde{C} = (C^{(\xi(i))})_{i \in [0, k-1]}$ and the definition of $q_M^{(\xi(i))}(X)$, and the second follows from the definition of $F_{q_M}(X)$. □

Permutation Polynomials. For $Y \in \{S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}\}$, we observe that the corresponding F_Y polynomials are a function of the selector polynomials $\{\widetilde{S}_{\sigma_1}, \widetilde{S}_{\sigma_2}, \widetilde{S}_{\sigma_3}\}$. More formally,

Claim 4. $\forall a \in [3], \forall i \in [0, k - 1], j \in [s], F_{S_{\sigma_a}}(v^{is+j}) \cdot v^{is} = \widetilde{S}_{\sigma_a}(v^{is+j})$

Proof Sketch. For any $i \in [0, k - 1], j \in [s]$, consider $\widetilde{S}_{\sigma_a}(v^{is+j})$. As described, the permutation $\tilde{\sigma}$ for C can be split into disjoint permutations $\{\tilde{\sigma}^{(i)}\}_i$, where by the definition of $\mathcal{P}\text{IonK}$, $\tilde{\sigma}^{(i)} : I_i \rightarrow I_i$ for $I_i := [is + 1, (i + 1)s] \cup [ks + is + 1, ks + (i + 1)s] \cup [2ks + is + 1, 2ks + (i + 1)s]$.

This in turn implies that $\widetilde{S}_{\sigma_a}(v^{is+j}) \in \{v^{is+\ell}, v^{is+\ell} \cdot k_1, v^{is+\ell} \cdot k_2\}$, where $\ell \in [s]$ and k_1 and k_2 are such that $k_1\mathbb{V}$ and $k_2\mathbb{V}$ are disjoint cosets of \mathbb{V} (see $\mathcal{P}\text{IonK}$ [GWC19] for details regarding the cosets).

Further, ℓ is determined by $\sigma^{\xi(i)}(j)$, the permutation polynomial for $C^{(\xi(i))}$. Therefore, we have for any $i \in [0, k-1], j \in [s]$,

$$\begin{aligned}\widetilde{S}_{\sigma_a}(v^{is+j}) &= \widetilde{\sigma}^{*(i)}(j) \\ &= v^{is} \sigma^{*(\xi(i))}(j) \\ &= v^{is} S_{\sigma_a}^{(\xi(i))}(\eta^j) \\ &= v^{is} F_{S_{\sigma_a}}(v^{is+j})\end{aligned}$$

□

Post-Processing the $\mathcal{P}\text{Ion}\mathcal{K}$ Permutation Polynomial. From our previous two claims, it is clear that while the (commitment of the) selector polynomials output from the segment-lookup protocol work as is, the same is *not* true of the permutation polynomial. But we will utilize Claim 4 to make the prover send the commitment to the correct permutation polynomial $\widetilde{S}_{\sigma_a}(X)$, and prove the relation to $F_{S_{\sigma_a}}(X)$. To this end, we define a new polynomial

$$U(X) := \sum_{i=0}^{k-1} v^i \sum_{j=1}^s \psi_{is+j}^{\mathbb{V}}(X).$$

From Claim 4, we have that $i \in [0, k-1], j \in [s]$

$$F_{S_{\sigma_a}}(v^{is+j}) \cdot U(v^{is+j}) - \widetilde{S}_{\sigma_a}(v^{is+j}) = 0.$$

Since the above holds for all $v \in \mathbb{V}$, it can be represented by the following polynomial check,

$$F_{S_{\sigma_a}}(X) \cdot U(X) - \widetilde{S}_{\sigma_a}(X) = Z_{\mathbb{V}}(X) Q_{S_{\sigma_a}}(X)$$

Given commitments to the polynomials above, the verifier will check that the equation is satisfied at a random point, and then proceed to invoking $\mathcal{P}\text{Ion}\mathcal{K}$ using the commitment to $\widetilde{S}_{\sigma_a}(X)$. We note that the commitment to $U(X)$ will be produced as a part of the pre-processing since it is input independent.

4.3 SubIonK Protocol Description

We finally describe our full $\text{SubIon}\mathcal{K}$ protocol in Figure 2.

We state our theorem relative to the costs of the segment-lookup and $\mathcal{P}\text{Ion}\mathcal{K}$ protocols. Specifically, $\text{Time}_p^{\text{segmentLookup}}(n, k, s)$ denotes the prover time for the segment-lookup protocol when parameterized by n, k and s . Similarly $\text{Time}_p^{\mathcal{P}\text{Ion}\mathcal{K}}(ks)$ denotes the prover time when $\mathcal{P}\text{Ion}\mathcal{K}$ is run on a circuit of size $O(ks)$. Verifier time and proof size are denoted in an analogous manner.

Theorem 1. *Given a $(s, \bar{k}, n, \{C_i\}_{i=1}^n)$ layered branching circuit C such that can be partitioned into n sub-circuits of size s such that the execution path has length $O(ks)$, the above protocol is a pre-processing SNARK in the Algebraic Group Model for \mathcal{C} induced by C such that the following properties hold:*

Prover Time: $\text{Time}_p^{\text{segmentLookup}}(n, k, s) + \text{Time}_p^{\mathcal{P}\text{Ion}\mathcal{K}}(ks) + O(ks \log(ks)) \mathbb{G}_1, \mathbb{G}_2$ and \mathbb{F} operations

Verifier Time: $\text{Time}_v^{\text{segmentLookup}}(n, k, s) + \text{Time}_v^{\mathcal{P}\text{Ion}\mathcal{K}}(ks) + 3$ Pairings.

Proof Size: $\text{Size}_v^{\text{segmentLookup}}(n, k, s) + \text{Size}_v^{\mathcal{P}\text{Ion}\mathcal{K}}(ks) + 13 \mathbb{G}_1$ elements.

Except with negligible probability (over the choice of γ), given that the verifier check in Step 7 succeeds, the commitments sent by the prover in Step 4 must be the correct polynomial. Thus the security of the scheme follows directly from the security of the underlying schemes.

Plugging in the asymptotic costs of the underlying protocol, we have the following corollary for Theorem 1 for Theorem 1.

PreProcess($C = (k, \{C^{(i)}\}_{i \in [n]})$)

1. Define polynomials $(T_{q_M}(X), T_{q_L}(X), T_{q_R}(X), T_{q_O}(X), T_{q_C}(X), T_{S_{\sigma_1}}(X), T_{S_{\sigma_2}}(X), T_{S_{\sigma_3}}(X))$ as described in Section 4.1.
2. Run pre-processing for the segment-lookup, i.e. for each $Y \in \{q_M, q_L, q_R, q_O, q_C, S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}\}$,

$$\text{pre}_Y = \text{gen}(n, k, s, T_Y).$$
3. Compute $U(X) := \sum_{i=0}^{k-1} v^i \sum_{j=1}^s \psi_{is+j}^{\mathbb{V}}(X)$
4. Output $\text{pre} := \{\text{pre}_Y\}_Y, [U(\tau)]_2$.

SubIonK($\text{pre}, \phi, C, \mathbb{V}, \mathbb{W}$)

1. P parses the execution path $\widetilde{C} = \{\widetilde{C}^{(i)}\}_{i \in [k]}$, and computes the function $\xi : \mathbb{F}^m \rightarrow [1, n]^k$ (on statement and witness pair $x = (\phi, w)$), that maps $\widetilde{C}^{(i)}$ to $C_{\xi_x(i)}$.
2. P uses ξ to compute $F_Y(X)$ for all $Y \in \{q_M, q_L, q_R, q_O, q_C, S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}\}$.
3. P and V run $\text{segmentLookup}^{\otimes \ell}(T_{q_M}, T_{q_L}, T_{q_R}, T_{q_O}, T_{q_C}, T_{S_{\sigma_1}}, T_{S_{\sigma_2}}, T_{S_{\sigma_3}})$ with verifier output $([F_{q_M}(\tau)]_1, [F_{q_L}(\tau)]_1, [F_{q_R}(\tau)]_1, [F_{q_O}(\tau)]_1, [F_{q_C}(\tau)]_1, [F_{S_{\sigma_1}}(\tau)]_1, [F_{S_{\sigma_2}}(\tau)]_1, [F_{S_{\sigma_3}}(\tau)]_1)$.
4. P sends $[\widetilde{S}_{\sigma_1}(\tau)]_1, [\widetilde{S}_{\sigma_2}(\tau)]_1, [\widetilde{S}_{\sigma_3}(\tau)]_1$ to V.
5. V sends a random $\gamma \in \mathbb{F}$ to P.
6. P computes for each $a \in [3]$, $[Q_{\sigma_a}(\tau)]_1$ such that

$$U(X) \cdot F_{S_{\sigma_a}}(X) - \widetilde{S}_{\sigma_a}(X) = \mathbb{Z}_{\mathbb{V}}(X) \cdot Q_{\sigma_a}(X)$$

7. V checks if the following holds for each $a \in [3]$,

$$\begin{aligned} e([F_{S_{\sigma_a}}(\tau)]_1, [U(\tau)]_2) \cdot e([\widetilde{S}_{\sigma_a}(\tau)]_1, [1]_2)^{-1} \\ = e([Q_{\sigma_a}(\tau)]_1, [\mathbb{Z}_{\mathbb{V}}(\tau)]_2) \end{aligned}$$

8. P and V run $\mathcal{P}\text{IonK}$ ($P_{\mathcal{P}\text{IonK}}, V_{\mathcal{P}\text{IonK}}$) with pre-processed inputs $([F_{q_M}(\tau)]_1, [F_{q_L}(\tau)]_1, [F_{q_R}(\tau)]_1, [F_{q_O}(\tau)]_1, [F_{q_C}(\tau)]_1, [\widetilde{S}_{\sigma_1}(\tau)]_1, [\widetilde{S}_{\sigma_2}(\tau)]_1, [\widetilde{S}_{\sigma_3}(\tau)]_1)$
9. V accepts if all checks pass.

Figure 2: The SubIonK Protocol

Corollary 1. *The asymptotic costs for the protocol in Theorem 1 is $O(ks \cdot (\log(ks) + \log(n)))$ for prover cost, and $O(1)$ for verifier cost and proof size.*

Allowing Arbitrary Number of Effective Layers. So far in this section, we assume that the effective number of layers k is fixed in advance. However, as discussed in Section 2.6, k depends on the input to the circuit C . Our protocol can be easily generalized to handle this case. In particular, let \bar{k} be the maximum

number of layers in C . We can easily modify the pre-processing algorithm in SubIonK to take as input \bar{k} , instead of k as follows: The modified pre-processing algorithm does the computation dependent on n , as described in Figure 2. For the remaining computation in the pre-processing phase that depends on k , we compute this for every $k' < \bar{k}$. In the online phase, once the prover learns the value of k upon evaluating the circuit on a given statement and witness, he can communicate this k to the verifier. For the rest of the protocol, both the prover and verifier work with the pre-processing corresponding to $k' = k$ and ignore the rest.

5 Implementation and Evaluation

We implement the SubIonK proof system in Rust, and are working towards an open-source release. The implementation relies on the BLS12-377 pairing-based curve as implemented in arkworks [ac22]. Moreover, we rely on the implementation of PionK available at [jel].

All experiments are run on a Macbook Pro with M1 Pro 3.2 Ghz chip and 32 GB RAM. We also report EVM gas costs¹⁵ for publishing and verifying signatures on-chain.

5.1 Evaluated Applications

DEX Rollups. We demonstrate the improvement of SubIonK on rollup applications. Consider a typical decentralized exchange (DEX) smart contract (e.g. Loopring [loo]) which allows users to create one of several types of transactions: deposits, spot trades, transfers, withdrawals, etc. The logic within each of these transaction types is encoded as a circuit, typically ranging from 30K to 60K arithmetic gates – we round up the circuit size to the nearest power of 2 when modeling as a segment, so we use 2^{16} gates¹⁶. A single instance of a rollup transaction that is submitted to a layer 1 blockchain (e.g. Ethereum) can batch together hundreds of such transactions, along with a single proof attesting to the validity of the next state transition (i.e., the final state of the DEX contract is attained by correctly evaluating the hundreds of transactions starting from the initial state that is recorded on the layer 1 chain). Rollups naturally map to our layered branching circuit model, where at each step, the prover executes one of several transaction types (segments).

Additional Rollup Applications In addition to application-specific rollups, as discussed above with DEX, we experiment with generic rollup solutions which allow for sequentially composing multiple different applications. For each of the following applications (borrowed from [CBBZ23b]), we compose circuits for that application with other similar-sized mock circuits.

- Proof of knowledge of exponent: each circuit segment encodes the statement that the witness, which is the scalar x , maps to the group element g^x , which is part of the input. We use the BLS12-377 curve (group \mathbb{G}_1). The segment has size 2^{12} gates.
- ZCash: each circuit segment encodes the proof circuit for a ZCash [HBHW16] (Sapling) transaction. Here, each segment has size 2^{17} gates.
- Nested conditional branching: we consider a generic program containing nested conditional statements, producing a control flow graph comprising a number of possible segments that can be exer-

¹⁵Our calculation uses the pre-compiled gas costs for the BLS12-381 curve as defined in EIP-2537 [eip]: \mathbb{G}_1 additions cost 600, \mathbb{G}_1 multiplications cost 12000, \mathbb{G}_2 additions cost 800, \mathbb{G}_2 multiplications cost 45000, and n pairings cost $115000 + n \cdot 23000$.

¹⁶This is not a technical limitation, as we could further decompose a segment into a set of smaller segments, each of size a power of two. We have not implemented this idea.

cised by a particular input – for instance, we report performance for an artificial program with 256 total segments, each of size 2^{16} gates, where an active path only exercises 16 segments.

Note that unlike [CBBZ23a], we do not use custom gates in our circuit encoding, so it results in larger circuits (comparable to the R1CS encoding).

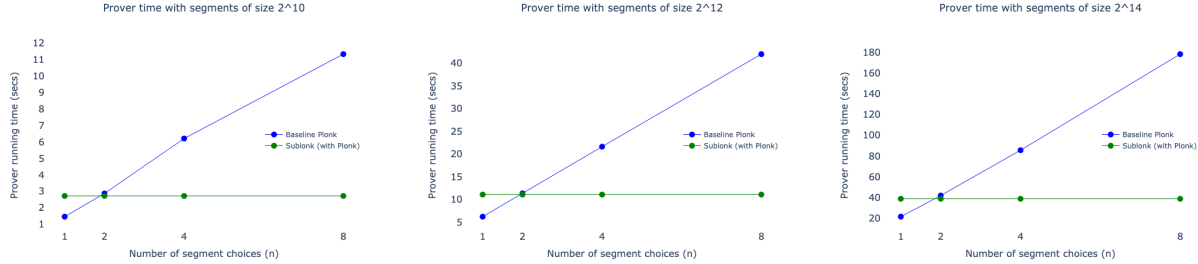


Figure 3: Prover runtime with $k = 128$ segments, with segment choices $n \in \{1, 2, 4, 8\}$, and segments of size 2^{10} , 2^{12} , and 2^{14}

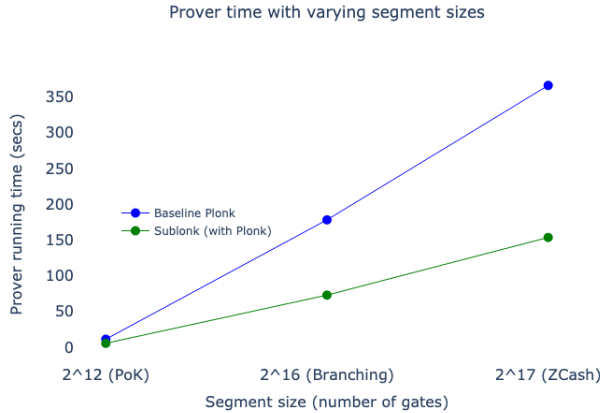


Figure 4: Prover runtime for varying segment sizes

5.2 Prover Time

In our first experiment, we encode the DEX rollup computation as a sequence of $k = 128$ steps, where each step has a conditional choice between $n \in \{1, 2, 4, 8\}$ segments, each of size 2^{16} gates; note that n is typically between 4 and 8 for a typical DEX rollup (see [loo]), however we also demonstrate **Sublonk** on smaller values of n to better illustrate its performance tradeoffs.

Figure 3 reports the prover time for both **Sublonk** and the baseline **Plonk** systems. The baseline refers to the invocation of **Plonk** using the conventional encoding where the circuit representation of all n transaction types are “stitched” together for each of the 128 steps; i.e., the circuit size is $128 \times n \times 2^{16}$ gates. **Sublonk** with **Plonk** refers to our proof system, which includes our segment-lookup protocol followed by an invocation of **Plonk**. The overhead incurred by the segment-lookup protocol can be computed by subtracting the **Sublonk** prover time from the baseline prover time for $n = 1$. Not surprisingly, as we increase n , we find larger speedups from **Sublonk**, with over 4.8x speedup for $n = 8$ and segment size 2^{16} . In addition to improvements in the prover time, we also believe that **Sublonk** enables better scaling

as it reduces the degree of the polynomials that are provided to the $\mathcal{P}\text{IonK}$ sub-procedure, thus reducing the memory needs (and likely to execute on smaller machines).

More generally, programs with conditional branches observe a large gap between the total number of segments and the number of executed segments, which makes SubIonK 's improvement even more significant. As an example, consider a control flow graph of a program with 256 possible segments, each of size 2^{16} . Now consider the an execution, where the active path only exercised 16 segments. The baseline $\mathcal{P}\text{IonK}$ prover requires roughly 175 seconds to produce a proof, whereas SubIonK requires 20.04 seconds; that marks an improvement in prover time of $8.7\times$.

Finally, we report prover times for the additional rollup applications. In Figure 4, we vary the segment size while fixing $k = 64$ and $n = 4$ – i.e., for each application, the rollup consists of mock circuits of size equal to the segment size, where the rollup prover can choose between 4 possible statements to prove at each of the 64 steps.

5.3 Proof Size

Our proof consists of 42 \mathbb{G}_1 elements and 12 scalar \mathbb{F} elements, of which 9 \mathbb{G}_1 and 6 \mathbb{F} elements arise from the proof generated by the $\mathcal{P}\text{IonK}$ subprocedure. Concretely, this makes our proof 2.4 KB in size. While the proof is substantially larger compared to the $\mathcal{P}\text{IonK}$ baseline, which is 624 bytes, we find this to be a valuable performance tradeoff when considering the reductions in prover time.

5.4 Verification Cost

The primary operations of the verifier include 23 pairings, 26 \mathbb{G}_1 operations, and 1 \mathbb{G}_2 operation – if one wishes to strictly enforce that the verifier must be constant time, then one could have the prover compute the opening to the vanishing polynomials, but we find the distinction to be minor in practice due to the efficiency of computing $\log(ns)$ field multiplications. A SubIonK proof requires 716.6K EVM gas units to verify on-chain and 50 ms to verify on a Macbook Pro laptop, where the computation is dominated by the cost to compute pairings. In comparison, verifying a $\mathcal{P}\text{IonK}$ proof requires 2 pairings and 18 \mathbb{G}_1 multiplications, and costs 377K EVM gas units to verify on-chain.

6 Conclusion

We propose SubIonK , significantly reducing the proof generation costs of $\mathcal{P}\text{IonK}$. Our improvements are most dramatic when only a small fraction of the circuit is active.

References

- [ac22] arkworks contributors. arkworks zksnark ecosystem. <https://arkworks.rs>, 2022. 32
- [BBHR18a] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *ICALP 2018*, volume 107 of *LIPICs*, pages 14:1–14:17. Schloss Dagstuhl, July 2018. 3
- [BBHR18b] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Report 2018/046, 2018. <https://eprint.iacr.org/2018/046>. 4, 9

- [BC23] Benedikt Bünz and Binyi Chen. Protostar: Generic efficient accumulation/folding for special sound protocols. Cryptology ePrint Archive, Paper 2023/620, 2023. 4, 8, 9, 13
- [BCC⁺17] Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinfeld, and Eran Tromer. The hunting of the SNARK. *J. Cryptol.*, 30(4):989–1066, 2017. 3
- [BCG⁺14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014. 3
- [BCR⁺19] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for R1CS. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of LNCS, pages 103–128. Springer, Heidelberg, May 2019. 12
- [CBBZ23a] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. HyperPlonk: Plonk with linear-time prover and high-degree custom gates. In *EUROCRYPT 2023, Part II*, LNCS, pages 499–530. Springer, Heidelberg, June 2023. 3, 33
- [CBBZ23b] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In *Advances in Cryptology–EUROCRYPT 2023: 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23–27, 2023, Proceedings, Part II*, pages 499–530. Springer, 2023. 32
- [CFH⁺15] Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17–21, 2015*, pages 253–270. IEEE Computer Society, 2015. 4, 9
- [CFQ19] Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2075–2092. ACM Press, November 2019. 4, 9
- [CHM⁺20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of LNCS, pages 738–768. Springer, Heidelberg, May 2020. 26
- [EFG22] Liam Eagen, Dario Fiore, and Ariel Gabizon. cq: Cached quotients for fast lookups. Cryptology ePrint Archive, Report 2022/1763, 2022. <https://eprint.iacr.org/2022/1763>. 3, 5, 10, 12, 15, 16, 18, 19, 26
- [eip] EIP-2537: Precompile for BLS12-381 curve operations. 32
- [FKL18] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of LNCS, pages 33–62. Springer, Heidelberg, August 2018. 10

- [gen20] genSTARK. <https://github.com/guildofweavers/genstark>, 2020. 4, 9
- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013. 3
- [GHH⁺23] Matthew Green, Mathias Hall-Andersen, Eric Hennenfent, Gabriel Kaptchuk, Benjamin Perez, and Gijs Van Laer. Efficient proofs of software exploitability for real-world processors. *Proc. Priv. Enhancing Technol.*, 2023(1):627–640, 2023. 3, 7
- [GK22] Ariel Gabizon and Dmitry Khovratovich. flookup: Fractional decomposition-based lookups in quasi-linear time independent of table size. Cryptology ePrint Archive, Report 2022/1447, 2022. <https://eprint.iacr.org/2022/1447>. 3, 5, 10
- [GKM⁺18] Jens Groth, Markulf Kohlweiss, Mary Maller, Sarah Meiklejohn, and Ian Miers. Updatable and universal common reference strings with applications to zk-SNARKs. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 698–728. Springer, Heidelberg, August 2018. 3
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 113–122. ACM Press, May 2008. 9
- [GPR21] Lior Goldberg, Shahar Papini, and Michael Riabzev. Cairo – a turing-complete STARK-friendly CPU architecture. Cryptology ePrint Archive, Report 2021/1063, 2021. <https://eprint.iacr.org/2021/1063>. 4, 9
- [Gro10] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, Heidelberg, December 2010. 3
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016. 3, 9
- [GW20] Ariel Gabizon and Zachary J. Williamson. plookup: A simplified polynomial protocol for lookup tables. Cryptology ePrint Archive, Report 2020/315, 2020. <https://eprint.iacr.org/2020/315>. 3
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. <https://eprint.iacr.org/2019/953>. 3, 5, 10, 11, 13, 14, 26, 29
- [Hab22] Ulrich Haböck. Multivariate lookups based on logarithmic derivatives. page 1530, 2022. 16, 18, 19
- [HBHW16] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification. *GitHub: San Francisco, CA, USA*, 4:220, 2016. 32

- [HK20a] David Heath and Vladimir Kolesnikov. A 2.1 KHz zero-knowledge processor with BubbleRAM. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 2055–2074. ACM Press, November 2020. 7
- [HK20b] David Heath and Vladimir Kolesnikov. Stacked garbling for disjunctive zero-knowledge proofs. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 569–598. Springer, Heidelberg, May 2020. 7
- [hod21] hodori. <https://github.com/matter-labs/hodori>, 2021. 4, 9
- [jel] Jellyfish: A Rust Implementation of the PLONK ZKP System and Extensions. 32
- [KGC⁺18] Harry A. Kalodner, Steven Goldfeder, Xiaoqi Chen, S. Matthew Weinberg, and Edward W. Felten. Arbitrum: Scalable, private smart contracts. In William Enck and Adrienne Porter Felt, editors, *USENIX Security 2018*, pages 1353–1370. USENIX Association, August 2018. 3
- [KPPS20] Ahmed E. Kosba, Dimitrios Papadopoulos, Charalampos Papamanthou, and Dawn Song. MIRAGE: Succinct arguments for randomized algorithms with applications to universal zk-SNARKs. In Srdjan Capkun and Franziska Roesner, editors, *USENIX Security 2020*, pages 2129–2146. USENIX Association, August 2020. 4, 8, 9
- [KPV22] Assimakis A. Kattis, Konstantin Panarin, and Alexander Vlasov. RedShift: Transparent SNARKs from list polynomial commitments. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 1725–1737. ACM Press, November 2022. 3
- [KS22] Abhiram Kothapalli and Srinath Setty. SuperNova: Proving universal machine executions without universal circuits. Cryptology ePrint Archive, Report 2022/1758, 2022. <https://eprint.iacr.org/2022/1758>. 4, 8, 9, 13
- [KS23] Abhiram Kothapalli and Srinath Setty. Hypernova: Recursive arguments for customizable constraint systems. Cryptology ePrint Archive, Paper 2023/573, 2023. 4, 8, 9
- [KST22] Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla. Nova: Recursive zero-knowledge arguments from folding schemes. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 359–388. Springer, Heidelberg, August 2022. 8
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, December 2010. 11, 14, 16
- [lib18] libSTARK. <https://github.com/elibensasson/libstark>, 2018. 4, 9
- [Lip16] Helger Lipmaa. Prover-efficient commit-and-prove zero-knowledge SNARKs. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 16*, volume 9646 of *LNCS*, pages 185–206. Springer, Heidelberg, April 2016. 4, 9
- [loo] Loopring: A Decentralized Token Exchange Protocol. 4, 7, 32, 33
- [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. 4, 9

- [MBKM19] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2111–2128. ACM Press, November 2019. 3
- [Mic94] Silvio Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, November 1994. 3
- [NT16] Assa Naveh and Eran Tromer. Photoproof: Cryptographic image authentication for any set of permissible transformations. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22–26, 2016*, pages 255–271. IEEE Computer Society, 2016. 8
- [PFM⁺22] Luke Pearson, Joshua Fitzgerald, Héctor Masip, Marta Bellés-Muñoz, and Jose Luis Muñoz-Tapia. PlonKup: Reconciling PlonK with plookup. Cryptology ePrint Archive, Report 2022/086, 2022. <https://eprint.iacr.org/2022/086>. 3, 10
- [PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013. 3
- [PK22] Jim Posen and Assimakis A. Kattis. Caulk+: Table-independent lookup arguments. Cryptology ePrint Archive, Report 2022/957, 2022. <https://eprint.iacr.org/2022/957>. 3, 5
- [plo] Plonky2: Fast Recursive Arguments with PLONK and FRI. 3
- [RPX⁺22] Deevashwer Rathee, Guru Vamsi Policharla, Tiancheng Xie, Ryan Cottone, and Dawn Song. Zebra: Anonymous credentials with practical on-chain verification and applications to kyc in defi. Cryptology ePrint Archive, Paper 2022/1286, 2022. <https://eprint.iacr.org/2022/1286>. 3
- [san23] Sangria: A Folding Scheme for PLONK, 2023. 4, 8, 9
- [Vee17] Meilof Veeningen. Pinocchio-based adaptive zk-SNARKs and secure/correct adaptive function evaluation. In Marc Joye and Abderrahmane Nitaj, editors, *AFRICACRYPT 17*, volume 10239 of *LNCS*, pages 21–39. Springer, Heidelberg, May 2017. 9
- [WSR⁺15] Riad S. Wahby, Srinath T. V. Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *NDSS 2015*. The Internet Society, February 2015. 4, 8, 13
- [ZBK⁺22a] Arantxa Zapico, Vitalik Buterin, Dmitry Khovratovich, Mary Maller, Anca Nitulescu, and Mark Simkin. Caulk: Lookup arguments in sublinear time. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 3121–3134. ACM Press, November 2022. 3, 5, 18, 19, 23, 25
- [ZBK⁺22b] Arantxa Zapico, Vitalik Buterin, Dmitry Khovratovich, Mary Maller, Anca Nitulescu, and Mark Simkin. Caulk: Lookup arguments in sublinear time. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7–11, 2022*, pages 3121–3134. ACM, 2022. 39

- [ZFZS20] Jiaheng Zhang, Zhiyong Fang, Yupeng Zhang, and Dawn Song. Zero knowledge proofs for decision tree predictions and accuracy. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 2039–2053. ACM Press, November 2020. 3
- [ZGK⁺17] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. A zero-knowledge version of vsql. *IACR Cryptol. ePrint Arch.*, page 1146, 2017. 9
- [ZGK⁺18] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. vRAM: Faster verifiable RAM with program-independent preprocessing. In *2018 IEEE Symposium on Security and Privacy*, pages 908–925. IEEE Computer Society Press, May 2018. 4, 8, 9, 13
- [ZGK⁺22] Arantxa Zapico, Ariel Gabizon, Dmitry Khovratovich, Mary Maller, and Carla Ràfols. Baloo: Nearly optimal lookup arguments. *Cryptology ePrint Archive*, Report 2022/1565, 2022. <https://eprint.iacr.org/2022/1565>. 3, 5
- [ZkR21] ZkRollups. An incomplete guide to rollups. <https://vitalik.ca/general/2021/01/05/rollup.html>, 2021. 3

A Caulk Sub-protocol: Multi-unity Proof

For the sake of completeness, we give the full protocol description of [ZBK⁺22b, Figure 5] whose informal description was given in Section 3.3. We require the following variant of bivariate KZG commitment from [ZBK⁺22b] for the sub-protocol.

KZG for Bivariate Polynomials. For a bivariate polynomial $P(X, Y)$ with degree up to $d_1 - 1$ in X and $d_2 - 1$ in Y , this protocol requires a universal setup with $d_1 d_2$ powers. For this, the setup corresponding to the univariate KZG itself can be used. To commit to $P(X, Y)$, using the commit algorithm of univariate KZG, one can commit to $[P(\tau^{d_2}, \tau)]_1$. Opening this commitment requires two steps: first, partially open $P(X, Y)$ at some $X = \alpha$ to a commitment $[P(\alpha, \tau)]_1$. The partial proof is given by $[H_\alpha(\tau^{d_2}, \tau)]_1$, where $H_\alpha(X, Y) = \frac{P(X, Y) - P(\alpha, Y)}{X - \alpha}$; second, fully evaluate $P(\alpha, Y)$ at $Y = \beta$ via standard univariate KZG proof with a degree bound of $d_2 - 1$ on $[P(\alpha, \tau)]_1$.

The main protocol takes as input the commitment $[u_0]_1$ and polynomial $U_0(X)$, and proves that the commitment indeed corresponds to $U_0(X)$ that encodes $(D(1), D(\nu^s), \dots, D(\nu^{(k-1)s}))$, and checks that for each $i \in [0, k - 1]$, $D(\nu^{is}) \in \{\mu^0, \dots, \mu^{n-1}\}$, for $\mu = \omega^s$. As explained in the informal description, this is equivalent to proving that the following aggregated equation holds:

$$\begin{aligned} & (U_0^2(X)\Delta_1(Y) + \sum_{j=2}^{\log n} U_{j-1}^2(X)\Delta_j(Y)) - \left(\sum_{j=1}^{\log n-1} U_j(X) \right. \\ & \left. \cdot \Delta_j(Y) + \text{id}(X)\Delta_{\log n}(Y) \right) = Z_V(X)Q_2(X, Y) \end{aligned}$$

for some polynomial $Q_2(X, Y)$. We describe the details of the protocol below, which helps in checking the above equation at a random point (α, β) . On a high level each round of the protocol does the following:

- Checking the above equation for some polynomial $Q_2(X, Y)$ at (α, β) is equivalent to showing that the following polynomial evaluates to 0 at $Y = \beta$:

$$\begin{aligned}
P(Y) &= (U_0^2(\alpha)\Delta_1(\beta) + \sum_{\ell=2}^{\log n} U_{\ell-1}^2(\alpha)\Delta_\ell(\beta)) \\
&+ Z_{\mathbb{S}}(\beta)(-Q_1(\beta) + Q_1(Y)) - \left(\sum_{\ell=1}^{\log n-1} U_\ell(\alpha)\Delta_\ell(\beta) \right) \\
&+ \text{id}(\alpha)\Delta_{\log n}(\beta) - Z_{\mathbb{V}}(\alpha)Q_2(\alpha, Y)
\end{aligned}$$

for some polynomial $Q_1(Y)$. For this, the prover sends commitments and values needed to reconstruct $[P(\tau)]_1$ and provides a proof for opening P at 0 and β .

- The other observation used is that since Δ_ℓ 's take 0/1 values, for each $Y \in \mathbb{S}$, it holds that:

$$\begin{aligned}
&(U_0^2(X)\Delta_1(Y) + \sum_{\ell=2}^{\log n} U_{\ell-1}^2(X)\Delta_\ell(Y)) \\
&= (U_0(X)\Delta_1(Y) + \sum_{\ell=2}^{\log n} U_{\ell-1}(X)\Delta_\ell(Y))^2
\end{aligned}$$

In the protocol below, we denote $\bar{U}(X, Y) = \sum_{\ell=2}^{\log n} U_{\ell-1}(X)\Delta_\ell(Y)$ and $U(X, Y) = \bar{U}(X, Y) + U_0(X)\Delta_1(Y)$.

- Now, the prover takes the univariate polynomials corresponding to \bar{U} and Q_2 and sends the commitments $[\bar{U}(\tau^{\log n}, \tau)]_1, [Q_2(\tau^{\log n}, \tau)]_1$ to the verifier. This is intended to be a bivariate commitment, but we want to reuse the srs containing only powers of τ .
- For the remaining rounds, the prover essentially sends the necessary commitments and opening proofs for the remaining polynomials corresponding to the equation above. Additional to the equation check, one additional check is done for the correctness of \bar{U} and U (check $\bar{U}(X, 1) = 0$, and then enforce the degree check).

The protocol uses the srs from our protocol of section 3.5. The common inputs for the prover and verifier is $[U_0(\tau)]_1$.

Round 1 (Prover \rightarrow Verifier). The prover takes the input srs and $U_0(X)$ and samples $t_1, \dots, t_{\log n} \leftarrow \mathbb{F}$ to compute:

1. For $\ell = 1, \dots, \log n$, $U_\ell(X) = \sum_{j=1}^{ks} (\mu^{ij})^{2^\ell} \Delta_j(X) + t_\ell Z_{\mathbb{V}}(X)$.
2. $U(X, Y) = \sum_{\ell=1}^{\log n} U_{\ell-1}(X)\Delta_\ell(Y)$.
3. $\bar{U}(X, Y) = U(X, Y) - U_0(X)\Delta_1(Y)$.
4. $Q_2(X, Y) = \sum_{\ell=1}^{\log n} \Delta_\ell(Y)Q_{2,\ell}(X)$, for $Q_{2,\ell}(X) = (U_{\ell-1}^2(X) - U_\ell(X))/Z_{\mathbb{V}}(X)$.

The prover sends $[\bar{U}(\tau^{\log n}, \tau)]_1, [Q_2(\tau^{\log n}, \tau)]_1$ to the verifier.

Round 2 (Verifier). The verifier sends a challenge $\alpha \in \mathbb{F}$.

Round 3 (Prover \rightarrow Verifier). The prover computes $Q_1(Y) = (U^2(\alpha, Y) - \sum_{\ell=1}^{\log n} U_{\ell-1}^2(\alpha)\Delta_{\ell}(Y))/Z_{\mathbb{S}}(Y)$ and sends $[Q_1(\tau)]_1$ to the verifier.

Round 4 (Verifier). Sends challenge $\beta \in \mathbb{F}$.

Round 5 (Prover \rightarrow Verifier). The prover does the following:

1. Computes $P(Y) = (U^2(\alpha, \beta) - \bar{U}(\alpha, \beta\phi) + \text{id}(\alpha)\Delta_{\log n}(\beta)) - Z_{\mathbb{V}}(\alpha)Q_2(\alpha, Y)$.
2. Computes and sends $U_0(\alpha)$, $[\bar{U}(\alpha, \tau)]_1$, $[Q_2(\alpha, \tau)]_1$, and their KZG opening proofs at α : π_1, π_2, π_3 , respectively.
3. Computes and sends $\bar{U}(\alpha, 1) = 0$, $\bar{U}(\alpha, \beta)$, $\bar{U}(\alpha, \beta\phi)$ along with their KZG opening proof at $(1, \beta, \beta\phi)$: π_4 .
4. Computes and sends $P(\beta) = 0$, along with its opening proof at β : π_5 .

Round 6 (Verifier). The verifier computes $[P(\tau)]_1 = (U_0(\alpha)\Delta_1(\beta) + \bar{U}(\alpha, \beta))^2 - [Q_1(\tau)]_1 Z_{\mathbb{S}}(\beta) - (\bar{U}(\alpha, \beta\phi) + \text{id}(\alpha)\Delta_{\log n}(\beta)) - Z_{\mathbb{V}}(\alpha)[Q_2(\alpha, \tau)]_1$, and accepts if and only if all the KZG opening proofs π_1, \dots, π_5 verify.