

Batching-Efficient RAM using Updatable Lookup Arguments

Moumita Dutta¹, Chaya Ganesh¹, Sikhar Patranabis², Shubh Prakash¹, and Nitin Singh²

¹ Indian Institute of Science

{moumitadutta,chaya,shubhprakash}@iisc.ac.in

² IBM Research India

sikhar.patranabis@ibm.com,nitisin1@in.ibm.com

Abstract. RAM (random access memory) is an important primitive in verifiable computation. In this paper, we focus on realizing RAM with efficient batching property, i.e, proving a batch of m updates on a RAM of size N while incurring a cost that is sublinear in N . Classical approaches based on Merkle-trees or address ordered transcripts to model RAM correctness are either concretely inefficient, or incur linear overhead in the size of the RAM. Recent works explore cryptographic accumulators based on unknown-order groups (RSA, class-groups) to model the RAM state. While recent RSA accumulator based approaches offer significant improvement over classical methods, they incur linear overhead in the size of the accumulated set to compute witnesses, as well as prohibitive constant overheads.

We realize a batching-efficient RAM with superior asymptotic and concrete costs as compared to existing approaches. Towards this: (i) we build on recent constructions of lookup arguments to allow efficient lookups even in presence of table *updates*, and (ii) we realize a variant of sub-vector relation addressed in prior works, which we call *committed index lookup*. We combine the two building blocks to realize batching-efficient RAM with sublinear dependence on size of the RAM. Our construction incurs an amortized proving cost of $\tilde{O}(m \log m + \sqrt{mN})$ for a batch of m updates on a RAM of size N . Our results also benefit the recent arguments for sub-vector relation, by enabling them to be efficient in presence of updates to the table. We believe that this is a contribution of independent interest.

We implement our solution to evaluate its concrete efficiency. Our experiments show that it offers significant improvement over existing works on batching-efficient accumulators/RAMs, with a substantially reduced resource barrier.

Table of Contents

1	Introduction	3
1.1	Our Contribution	3
1.2	Techniques	4
2	Related Work	5
3	Preliminaries	7
3.1	Succinct Arguments of Knowledge	8
3.2	Security Model	8
3.3	KZG Commitment Scheme	9
3.4	Lookup Arguments	9
4	Technical Overview	10
4.1	Batching-Efficient RAM: Blueprint	10
4.2	Batching-Efficient RAM: Components	11
5	Memory Consistency for RAM	13
5.1	Correctness of RAM Update	13
5.2	Consistency Check via Transcripts	13
6	Improved Batching-Efficient RAM	14
6.1	Committed Index Lookup	14
6.2	Almost Identical RAM States	14
6.3	Batching-Efficient RAM: Combined Protocol	17
7	Fast Lookups from Approximate Pre-Processing	17
7.1	Base + Cache approach	17
7.2	Amortized Sublinear Batching	21
8	Experiments	21
A	More Preliminaries	25
A.1	Polynomial Commitment Scheme	25
A.2	Succinct Argument of Knowledge	26
B	Argument for RAM From Polynomial Protocols	27
B.1	Polynomial Encoding	27
B.2	Relations over Polynomial Encodings	27
C	Succinct Argument for Verifiable RAM	28
D	Proof of Lemma 5	31
D.1	Computational Algebra Preliminaries	32
D.2	Some Useful Results	33
D.3	Proof of Lemma 5	36
E	Committed Index Lookup from Caulk+	39

1 Introduction

General purpose *Succinct Non-interactive Arguments of Knowledge* (SNARKs) enable one to generate succinct proofs of membership of a statement in an NP relation expressed as an arithmetic circuit. These proofs are extremely cheap to verify, which makes them useful for *Verifiable Computation* (VC), where a resource-constrained client (e.g., a mobile phone), can outsource an expensive computation to an untrusted server, and later verify the correctness of the computation at a minimal cost.

Modeling RAM in Verifiable Computation. It turns out that arithmetic circuit-based representations are inefficient in expressing relations involving the result of a program execution on memory/state. Such relations frequently arise in the context of verifiable computation, in scenarios that require proving the correctness of query execution against a database, inference from a decision tree, or updates on a table of account balances (e.g., when a batch of transactions, such as account transfers, is applied to the table).

In the aforementioned examples, objects such as database tables, decision trees, and accounts tables can be naturally modeled as instances of addressable memory, or more generally, random access memory (RAM), where one needs to prove that the RAM has been accessed/updated in accordance with the correct execution of the computation. There exists a rich and expanding body of work on efficiently modeling abstractions of RAM in verifiable computation. While a complete treatment of this vast body of work is beyond the scope of this paper (a fairly recent survey in [WB15] is a good starting point), we mention two additional properties that are often demanded of the RAM primitive: *persistence* – the ability to persist the RAM state across several computations, and *batching* – where verifiable update of the RAM state is required for small batches of updates. These properties are also the focus of this work.

Application to Blockchain Rollups. Batching-efficient RAM is especially relevant in the context of blockchain *rollups* [bar], an umbrella term for recent efforts to scale blockchains by moving expensive computation off the blockchain to the so-called *layer two* (or L2) chains. The blockchain only needs to verify succinct proofs attesting to the correctness of the off-chain computation. This approach is popularly called *rollup* as it allows verifying the result of several (rolled-up) transactions modifying the L2 state, as part of one transaction verified on the main chain. This simultaneously improves scalability and lowers the cost (e.g., gas fees) per transaction due to succinct verification. We consider improving efficiency of rollups an important motivation for our work, but avoid precise details of a smart-contract based instantiation of our solution.

1.1 Our Contribution

We present batching-efficient RAM construction, which advances the efforts towards achieving *verifiable outsourcing of state update* such as in [BFR⁺13b] and more recently in [OWWB20,CFH⁺22]. The most popular approaches to succinctly represent state involve accumulators based on Merkle-trees [Mer88], or ones based on groups of unknown order (e.g. RSA, class-groups) [CL02,BBF19,OWWB20,CFH⁺22]. The updates to the state are effected by insertions or deletions in the accumulated set. In this work, we model the state as an addressable memory (RAM) described by vector \mathbf{T} , which stores value v_i at address i . We denote this as $\mathbf{T}[i] = v_i$. The RAM supports two operations, viz, *loads* expressed as $v_i := \mathbf{T}[i]$, and *stores* expressed as $\mathbf{T}[i] = v_i$. We think of addresses $i \in [0, N]$ for some $N \in \mathbb{Z}$ while the values $v_i \in \mathbb{F}$ for some finite field \mathbb{F} . In our construction, we represent both the RAM and operations on it as polynomials, and use appropriate polynomial commitment schemes to obtain succinct commitments (digests) to them. In this paper, we do not require commitments to be *hiding*, as our focus is on succinctness. We consider privacy as an orthogonal goal, one we believe is easily achievable via small adaptations to our construction.

We summarize our contributions below.

- As our first contribution, we propose *update friendly lookup arguments*, which addresses the strict dependence of recent constructions on table-specific pre-processing parameters. Our innovation extends the utility of table-specific parameters to enable efficient lookups from tables, which are within certain Hamming distance of the pre-processed table.
- We construct *committed index lookup arguments* via black-box reduction to sub-vector arguments that use homomorphic commitments. A committed index lookup involves three committed vectors \mathbf{t} , \mathbf{a} and \mathbf{v} satisfying $v_i = t_{a_i}$ for all i . Similar definition is also used in recent multi-variate lookup

arguments in [STW23], where a similar reduction to sub-vector arguments is obtained under a more restrictive assumption about the elements of the table.

- We crucially employ the above two contributions to construct a *batching-efficient* RAM, which can prove a batch of m updates with an amortized prover complexity of $O(m \log m + \sqrt{mN})$, with N being the size of the RAM. Our dependence on the RAM size is sublinear, in contrast to the linear complexity inherent in recent works on batching-efficient RAM using RSA accumulators [OWWB20,CFH⁺22] or using generic memory checking techniques [WSR⁺15,BCTV14b,BCG⁺13,ZGK⁺18]. All of our protocols are public-coin, and can be made non-interactive using standard techniques [FS87].
- We implement our scheme in Rust³. Experimentally, we show that our scheme performs significantly better than prior works, and is eminently deployable on a commodity hardware.

1.2 Techniques

We present a brief summary of our techniques below. A more detailed technical overview appears in Section 4.

Update-friendly Lookup Arguments. Our starting point is the recent line of works on lookup arguments which prove that a vector of size m appears as a sub-vector in a large fixed vector (table) of size N with succinct proof sizes and verification, but most notably ensuring that prover runs in time sublinear in the size of the table (N). The pioneering work [ZBK⁺22] obtained prover complexity of $O(m^2 + m \log N)$, which was improved in subsequent works to $O(m^2)$ [PK22], $O(m \log^2 m)$ [ZGK⁺22], and $O(m \log m)$ [EFG22,CFF⁺24]. However, the sublinear prover complexity requires table-dependent $O(N \log N)$ pre-processing and $O(N)$ storage. This table-dependent pre-processing implies that while the aforementioned lookup arguments can be used to obtain efficient ROM (read only memory) semantics they cannot be used as is for RAM (which supports update operations). Moreover, an update involving even a single index renders the entire $O(N)$ pre-processing unusable for further lookups, thus necessitating entire $O(N \log N)$ re-computation. This work is the first effort towards mitigating this rigid dependence, thereby increasing the applicability of the recent lookup arguments. An important contribution we make here is a new method for computing “encoded quotients” used in several recent lookup constructions such as [ZBK⁺22,PK22,EFG22,CFF⁺24]. Our approach for computing these quotients from pre-computed parameters remains efficient even when the table is updated, and it directly applies to all the aforementioned constructions. For a table δ -hamming distance away from the pre-processed one, we incur $(m + \delta) \log^2(m + \delta)$ additional overhead for proving m lookups. To achieve such a quasi-linear overhead in both m and δ , we rely on novel algebraic algorithms described in Section 7. We informally summarize our contribution in this regard below, whereas Theorem 4 states the precise result.

Theorem 1 (Informal). *There exists a deterministic $O(N \log N)$ time algorithm $\text{Preprocess}(\mathbf{T}) \rightarrow \text{pp}_T$ which on input $\mathbf{T} \in \mathbb{F}^N$, outputs parameters pp_T of size $O(N)$ such that: Given pp_T , vectors $\mathbf{T}' \in \mathbb{F}^N$, $\mathbf{t} \in \mathbb{F}^m$ with \mathbf{t} being a sub-vector of \mathbf{T}' an argument of knowledge for the same can be computed in time $O((m + \delta) \log^2(m + \delta) + f(m))$ where $\delta = \Delta(\mathbf{T}, \mathbf{T}')$ is the Hamming distance between \mathbf{T} and \mathbf{T}' while $f(m)$ depends on the specific lookup protocol.*

For the constructions based on [ZBK⁺22,PK22], we set $f(m) = m^2$ in the above, while for [EFG22,CFF⁺24], we have $f(m) = m \log m$.

Committed Index Lookup: We augment the sub-vector relation in prior lookup arguments which considers whether each entry of a given vector appears in the target vector to one that also identifies the precise positions where the given vector appears in the target vector. When this relation is checked over commitments of the respective vectors; given vector, the target vector and the position vector, we call it *committed index lookup*. The relation we consider is similar to the one considered in [STW23]. For lookup arguments with homomorphic commitment schemes, we show that committed index lookup can be obtained using a sub-vector lookup argument (Lemma 2, Section 6.1). Such a construction was also considered in [STW23], but under a more restrictive assumption that the size of the elements in the table have to be within a certain bound. Lemma 2 yields a construction of committed index lookup that uses (a single instance of) the underlying sub-vector protocol in a black-box manner. This immediately implies efficient constructions of arguments for committed index lookups from [ZBK⁺22,PK22,ZGK⁺22,EFG22,CFF⁺24]. In Appendix E, we also present an explicit

³ <https://github.com/nitsatiisc/caulk/tree/updateable-ram>

(non-black-box) adaptation of [PK22] to obtain a committed index lookup, which again incurs costs comparable to a single instance of the underlying sub-vector protocol.

Batching-Efficient RAM from Lookup Arguments. Memory checking methods based on *address ordered transcripts* [WSR⁺15,BCTV14b,BCG⁺13,ZGK⁺18], which are popularly used in efficient RAM abstractions, incur a cost linear in the size of the RAM. This is prohibitive for efficient batching. As a key idea in this work, we invoke committed index lookup on the large RAMs, to verifiably extract smaller sub-RAMs, which correspond to indices actually involved in the batch update. Then, we use the linear time memory-checking techniques to argue the consistency of these smaller sub-RAMs.

The idea needs to work through some more details, such as showing that the larger RAMs are identical on positions not referenced by the batch of updates (considered in Section 6.2). The overall idea is illustrated in Figure 1. We also note that the extracted sub-RAMs can have duplicate records, corresponding to multiple updates referencing the same RAM index; however, memory checking methods can be easily adapted to handle such cases. Finally, we would still hit the “rigidity” of lookup arguments in realizing this plan; once the table has changed, lookups are no longer efficient from it. To circumvent this, we use our first contribution on extending the utility of table-specific parameters to defer parameter re-computation optimally while still availing efficient lookups. More specifically, if we choose to re-compute the full table-specific parameters after k batches (of m updates each), the average cost per batch is $O(N \log N/k + mk \log^2(mk) + f(m))$. Here, $f(m)$ as earlier denotes complexity of the non-updatable base protocol. Setting $k \approx \sqrt{N/m}$ yields the average cost of m updates as $\tilde{O}(f(m) + \sqrt{mN})$, which scales sublinearly with the size of the RAM. While the preceding analysis considers the worst case, in specific applications (such as account transactions, where few accounts contribute a large volume of transactions), it may be possible to further delay the computation of table-specific parameters. Thus we have:

Theorem 2 (Informal). *Given $m, N \in \mathbb{N}$, there exists an argument for verifiable RAM which proves updates of batch size m on RAM of size N with amortized prover complexity of $\tilde{O}(f(m) + \sqrt{mN})$.*

Polynomial Protocol for RAM. There are several ways to implement the ordered transcript based memory consistency check on the smaller $O(m)$ -sized RAMs, for example by expressing the same as an arithmetic circuit. However, for completeness, we also present an argument for RAM as an interactive *polynomial protocol* [GWC19b], which is then compiled into an argument of knowledge using the KZG [KZG10] commitment scheme in the algebraic group model (AGM) [FKL18]. This construction appears in Appendix B.

2 Related Work

Efficient modeling of the RAM primitive is a widely studied problem in verifiable computation (VC), due to its inherent usefulness in modeling several computations of interest. Encapsulating RAM semantics in VC circuits is also challenging; since (i) arithmetic/boolean circuits do not adequately model random access, and (ii) incorporating entire memory as gates in a circuit is prohibitive.

Several novel techniques have been proposed to work around the above limitations of circuit based representation of RAM. Among them, Merkle tree-based accumulators to model the RAM state are popular [BFR⁺13b,BFR⁺13a,BCTV14a] as they can efficiently prove updates to the state, without modeling the entire memory in the arithmetic circuit. Other approaches based on *address ordered time-scripts* avoid the concrete costs of Merkle tree-based approaches by letting the prover provide inputs and outputs of RAM operations in a non-deterministic manner, which are then checked to satisfy consistency of *loads* and *stores*. Several works such as [WSR⁺15,BCTV14b,BCG⁺13,ZGK⁺18] implement and improve variants of the aforementioned approach. Most transcript-based realizations of RAM only consider it to be transient, i.e, its state is useful only during the execution of a program, and do not consider *persistence* of the RAM state across several executions.

Another feature, which has only been considered in recent works [OWWB20,CFH⁺22] is *batching*, where a verifiable update of RAM state is required for a batch of m updates, with m being much smaller than the RAM size. Both the Merkle tree-based approaches and the transcript-based approaches are inefficient with respect to batching. While constructions using Merkle tree-based accumulators (realized from collision-resistant hash functions) suffer from high concrete costs and poor ability to batch proofs, those based on checking consistency using transcripts incur a linear overhead in the RAM size.

Batching-Efficient RAM. There have been recent efforts [OWWB20,CFH⁺22] on batching-efficient realization of the RAM primitive (see Table 1 for a summary of these schemes and the associated efficiency parameters). This is a natural setting in applications of verifiable computation, most notably in the context of blockchain *rollups*. Here, one is required to show that a batch of m transactions correctly updates the state of a table of account balances, which is maintained off-chain by the rollup provider. Here the batch-size m ranges from few hundreds to few thousands, whereas the table itself could contain several million accounts. Similar to prior work on batching-efficient RAMs [OWWB20,CFH⁺22], our work is also motivated by the problem of enabling more efficient rollup for tables, which are naturally modeled as RAMs.

While the aforementioned works substantially mitigate disadvantages of both the Merkle-tree based approaches and transcript-based approaches by using RSA accumulators to model the state, they still incur large prover costs and memory requirements even for modest sized batches. The approaches in [OWWB20,CFH⁺22] encode complex modular arithmetic over RSA groups and *hash to prime* functions as arithmetic circuits, which results in a fixed overhead of around 10 million R1CS constraints at batch sizes of $m = 1000$ (this overhead is significantly larger for [OWWB20]). This is already prohibitive on a modest hardware. In addition, the witness computation for each update incurs cost linear in the size of accumulated set. The prior works [OWWB20,CFH⁺22] seek to mitigate this through pre-computation and parallel/distributed processing. However, the issue of maintaining pre-computed parameters in sync with dynamic accumulator state has not been adequately addressed in [OWWB20,CFH⁺22]. For example, in RSA-based accumulators, generating $O(N)$ non-membership witnesses straightforwardly requires $O(N^2)$ time; however, these witnesses become stale once the accumulator state changes, and hence cannot be used as-is for subsequent update proofs.

Our approach considers both the cost of online proof generation as well as the offline cost of maintaining pre-computed parameters. In addition, our solution is almost “circuit-free”. Our entire RAM operation is modeled as a polynomial protocol, which is readily transformed into an argument of knowledge using a polynomial commitment scheme. In an application of our primitive to rollups, the only part of the statement that would need to be expressed as a circuit is the verification of digital signatures on transactions (which is around 500 constraints per verification for EDDSA signatures). By suitably balancing the online and offline costs, we can prove a batch of 1000 updates on a RAM of size 1 million in an average of 90 seconds on a commodity laptop with a single-threaded implementation. Our performance can be substantially improved using a parallel implementation. See Sections 4 and 8 for more detailed discussions on the efficiency of our scheme.

Lookup Arguments. There are several recently proposed constructions of lookup arguments which enable proving that a vector of size m is a sub-vector of a larger (predetermined) vector of size N (see Table 1 for a summary of these schemes and the associated efficiency parameters). Of these, the very initial schemes [BGH19,GW20,BGH20] incur proving costs linear in N . Starting with Caulk [ZBK⁺22], many lookup arguments were proposed with proving costs that are (largely) independent of N . Broadly, these schemes can be divided into categories based on how they achieve proving costs independent of N . The lookup arguments in the first category [ZBK⁺22,PK22,GK22,EFG22,ZGK⁺22,CFF⁺24,ZSG24] use polynomial protocols in conjunction with the KZG polynomial commitment scheme, where the lookup efficiency relies crucially on *pre-computed* KZG opening proofs for the polynomial encoding the predetermined N -sized vector. We point out that naïvely adapting these lookup arguments for updatable tables/RAM is challenging since even small updates to the table require re-computing all of the opening proofs, which is prohibitively expensive (requires $\tilde{O}(N)$ computation). To overcome the “rigidity”, we propose a new method of constructing lookup arguments which allows re-using the pre-computed opening proofs across several batch updates, thus avoiding the need for re-computing after each batch (see Sections 4 and 7 for more details).

Lasso. The second category of lookup arguments is exemplified by the recently proposed Lasso scheme [STW23,AST24], which enables efficient lookups for tables with a decomposable structure. Informally, a table $\mathbf{T} : \{0, 1\}^n \rightarrow \{0, 1\}^k$ is said to have a decomposable structure if there exists a decomposition of the table \mathbf{T} into c sub-tables $\mathbf{T}_1, \dots, \mathbf{T}_c : \{0, 1\}^{n/c} \rightarrow \{0, 1\}^\ell$ and a succinctly computable function f such that for $x = x_1 \parallel \dots \parallel x_c$ where $x_i \in \{0, 1\}^{n/c}$, we have

$$\mathbf{T}[x] = f(\mathbf{T}_1(x_1), \dots, \mathbf{T}_c(x_c))$$

A simple example of such a function is a bit-wise AND (we refer to [STW23,AST24] for a more detailed exposition). Lasso crucially leverages this decomposability of the table to reduce a lookup into a table

Scheme	Setup	Proof Size	Prover Work	Verifier Work
Lookup Arguments for Static Tables				
Plookup [GW20]	Updatable	$5\mathbb{G}_1, 9\mathbb{F}$	$O(N)\mathbb{G}_1, O(N \log N)\mathbb{F}$	$2P$
LogUp [Hab22,PH23]		$O(1)\mathbb{G}_1, O(1)\mathbb{F}$	$O(N)\mathbb{G}_1, O(N \log N)\mathbb{F}$	$O(1)P$
Halo2 [BGH19,BGH20]		$6\mathbb{G}_1, 5\mathbb{F}$	$O(N)\mathbb{G}_1, O(N \log N)\mathbb{F}$	$2P$
Caulk [ZBK ⁺ 22]		$14\mathbb{G}_1, 1\mathbb{G}_2, 4\mathbb{F}$	$O(m^2)\mathbb{G}_1, O(m^2 + m \log N)\mathbb{F}$	$4P$
Caulk+ [PK22]		$7\mathbb{G}_1, 1\mathbb{G}_2, 2\mathbb{F}$	$O(m^2)\mathbb{G}_1, O(m^2)\mathbb{F}$	$3P$
Flookup [GK22]		$7\mathbb{G}_1, 1\mathbb{G}_2, 2\mathbb{F}$	$O(m)\mathbb{G}_1, O(m \log^2 m)\mathbb{F}$	$3P$
Baloo [ZGK ⁺ 22]		$12\mathbb{G}_1, 1\mathbb{G}_2, 4\mathbb{F}$	$14m\mathbb{G}_1, O(m \log^2 m)\mathbb{F}$	$5P$
CQ [EFG22]		$8\mathbb{G}_1, 3\mathbb{F}$	$8m\mathbb{G}_1, O(m \log m)\mathbb{F}$	$5P$
CQ+ [CFE ⁺ 24]		$7\mathbb{G}_1, 1\mathbb{F}$	$8m\mathbb{G}_1, O(m \log m)\mathbb{F}$	$5P$
Locq [ZSG24]		Trusted	$4\mathbb{G}_1, 1\mathbb{G}_2$	$6m\mathbb{G}_1, m\mathbb{G}_2, O(m \log m)\mathbb{F}$
Lasso [STW23]	Transparent	$O(1)\mathbb{G}, O(\log m)\mathbb{F}$	$o(cm + cN^{1/c})\mathbb{G}_1, O(cm)\mathbb{F}$	$O(\log m)\mathbb{F}, O(1)\mathbb{G}$
Lookup Arguments for Updatable Tables/Batching-Efficient RAM				
OWWB20 [OWWB20]	Trusted	$O(1)\mathbb{G}$	$\tilde{O}(m)\mathbb{F}, \mathbb{G}, \tilde{O}(N)\mathbb{G}'$	$1P$
B-INS-ARISA [CFH ⁺ 22]		$O(1)\mathbb{G}$	$\tilde{O}(m)\mathbb{F}, \mathbb{G}, \tilde{O}(N)\mathbb{G}'$	$1P, O(1)\mathbb{G}'$
Our work (c.f. Table 2)	Updatable	$65\mathbb{G}_1, 1\mathbb{G}_2, 43\mathbb{F}$	$\tilde{O}(\sqrt{mN})\mathbb{F}, \mathbb{G}_1$	$9P$

Table 1: Comparison with state-of-the-art lookup arguments. Here, N denotes the size of the RAM and m denotes the number of updates (typically $m \ll N$). We use $(\mathbb{F}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2, g_t)$ to denote a bilinear group (\mathbb{G} denotes either \mathbb{G}_1 or \mathbb{G}_2 , and is used in cases where the exact group is unspecified), and \mathbb{G}' to denote an RSA group. P denotes a pairing evaluation. For Lasso, we report the overheads considering the polynomial commitment scheme Sona and assuming *structured* tables (here, c denotes an arbitrary positive integer). For our scheme, we report performance for the CQ-based instantiation.

of size $N = 2^n$ into c lookups, each into a table of size $N^{1/c}$. While this strategy works elegantly for tables with special structure, it is not compatible with arbitrary tables/updates, which is the focus of our work (in particular, in applications such as rollups, we need the ability to handle updates to arbitrary tables).

To summarize, existing lookup arguments achieve efficiency either by leveraging table-specific pre-processing or exploiting special structure, both of which do not naively extend to arbitrary dynamic tables. We focus on handling batch updates for arbitrary tables, and our techniques can be viewed as enabling the utility of table-specific pre-processing even *across batch updates*.

3 Preliminaries

This section presents notations and preliminary background material used in the rest of the paper.

Notation. Throughout the paper, we assume a bilinear group generator BG which on input λ outputs parameters for the protocols. Specifically $\text{BG}(1^\lambda)$ outputs $(\mathbb{F}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2, g_t)$ where:

- $\mathbb{F} = \mathbb{F}_p$ is a prime field of super-polynomial size in λ , with $p = \lambda^{\omega(1)}$.
- $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T are groups of order p , and e is an efficiently computable non-degenerate bilinear pairing $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$.
- Generators g_1, g_2 are uniformly chosen from \mathbb{G}_1 and \mathbb{G}_2 respectively and $g_t = e(g_1, g_2)$.

We write groups \mathbb{G}_1 and \mathbb{G}_2 additively, and use the shorthand notation $[x]_1$ and $[x]_2$ to denote group elements $x \cdot g_1$ and $x \cdot g_2$ respectively for $x \in \mathbb{F}$. We implicitly assume that all the setup algorithms for the protocols invoke BG to generate descriptions of groups and fields over which the protocol is instantiated. We use $[n]$ to denote the set of integers $\{1, \dots, n\}$.

Lagrange Polynomials. We denote the N th root of unity by ξ and define the subgroup \mathbb{H} as $\mathbb{H} = \{\xi, \dots, \xi^N\}$. Let $\{\mu_i(X)\}_{i=1}^N$ be the associated Lagrange basis polynomials over the set \mathbb{H} ; that is, $\mu_i(X) = \prod_{j \neq i} \frac{X - \xi^j}{\xi^i - \xi^j}$. We denote by $Z_{\mathbb{H}}$ the vanishing polynomial of \mathbb{H} ; $Z_{\mathbb{H}}(X) = X^N - 1$.

Formal Derivatives of Polynomials. For a polynomial $f(X) = \sum_{i=0}^d a_i X^i \in \mathbb{F}[X]$, we define its formal derivative to be the polynomial $f'(X) = \sum_{i=1}^d i a_i X^{i-1}$.

3.1 Succinct Arguments of Knowledge

Let \mathcal{R} be a NP-relation and \mathbf{L} be the corresponding NP-language, where $\mathbf{L} = \{x : \exists w \text{ such that } (x, w) \in \mathcal{R}\}$. A *succinct argument of knowledge* consists of a pair of PPT algorithms $(\mathcal{P}, \mathcal{V})$. Given a public instance x , the prover \mathcal{P} , convinces the verifier \mathcal{V} , that $x \in \mathbf{L}$, where the prover additionally has as a witness w . We use the notation $b \leftarrow \langle \mathcal{P}(w), \mathcal{V} \rangle(x)$ to denote \mathcal{V} 's output in the interactive protocol involving \mathcal{P} and \mathcal{V} with w as \mathcal{P} 's input and x as the common input. The *knowledge-soundness* property says that if the verifier is convinced, then an efficient extractor algorithm given oracle access to the prover outputs a witness w such that $(x, w) \in \mathcal{R}$. An argument system is *succinct* if the communication complexity and the complexity of \mathcal{V} is polylogarithmic in the size of the witness. We provide formal definitions in Appendix A.2.

Fiat-Shamir. An interactive protocol is *public-coin* if the verifier's messages are uniformly random strings. Public-coin protocols can be transformed into non-interactive arguments in the Random Oracle Model (ROM) by using the Fiat-Shamir [FS87] heuristic to derive the verifier's messages as the output of a Random Oracle.

Modular Approach. A modular approach for designing efficient succinct arguments consists of two steps; constructing an information theoretic protocol in an idealized model, and then compiling the information-theoretic protocol via a cryptographic compiler to obtain an argument system. Informally, the prover and the verifier interact where the prover provides oracle access to a set of polynomials, and the verifier accepts or rejects by checking certain identities over the polynomials output by the prover and possibly public polynomials known to the verifier. Such a protocol can be compiled into a succinct argument of knowledge by realizing the polynomial oracles using a *polynomial commitment scheme*. A polynomial commitment scheme allows a prover to commit to polynomials, and later verifiably open evaluations at chosen points by giving evaluation proofs. This enables the verifier to probabilistically check polynomial identities at random points of \mathbb{F} . Many recent constructions of zk-SNARKs [BFS20, CHM⁺20, GWC19b] follow this approach where the information theoretic object is a polynomial interactive oracle proof (PIOP) (or a polynomial protocol), and the cryptoprimitive in the compiler is a polynomial commitment scheme.

3.2 Security Model

We describe public-coin interactive protocols in the structured reference string (SRS) model where both the parties have access to a SRS. The SRS in our protocols consists of encodings of monomials of the form $\{[x^i]_1\}_{a \leq i \leq b}, \{[x^i]_2\}_{c \leq i \leq d}$ for x chosen uniformly from \mathbb{F} and a, b, c, d are bounded by some polynomial in λ . It then follows from [BGM17] that such an SRS can be generated using a universal and updatable setup [GKM⁺18] requiring only one honest participant. In practice, this is a superior security model compared to requiring a fully trusted setup. We use $\text{srs} = (\text{srs}_1, \text{srs}_2)$ to denote the structured reference string of the above form. We say that the srs has degree Q if all the elements of srs_i , $i = 1, 2$ are of the form $[f(x)]_i$ for a polynomial $f \in \mathbb{F}_{<Q}[X]$.

Algebraic Group Model. We analyze security of our protocols in the Algebraic Group Model (AGM) introduced in [FKL18]. An adversary \mathcal{A} is called *algebraic* if every group element output by \mathcal{A} is accompanied by a representation of that group element in terms of all the group elements that \mathcal{A} has seen so far (input and output). In the AGM, an adversary \mathcal{A} is restricted to be *algebraic*, which in our SRS-based protocol means a PPT algorithm satisfying the following: for $i \in \{1, 2\}$, whenever \mathcal{A} outputs an element $A \in \mathbb{G}_i$, it is accompanied by its representation, \mathcal{A} also outputs a vector \mathbf{v} over \mathbb{F} such that $A = \langle \mathbf{v}, \text{srs}_i \rangle$.

Real and Ideal Pairing Checks. For an algebraic adversary \mathcal{A} interacting in a protocol with a degree Q SRS over a bilinear group, the verifier can use the pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ to perform "ideal check" of the form $(R_1 \cdot T_1) \cdot (R_2 \cdot T_2) \equiv 0$, where R_1, R_2 are vectors of polynomials over \mathbb{F} and T_1, T_2 are public matrices over \mathbb{F} . Under the Q -DLOG assumption stated below, the aforementioned ideal check is equivalent (except with a negligible probability) to a real pairing check $(a \cdot T_1) \cdot (T_2 \cdot b) = 0$ with a and b denoting vectors in \mathbb{F} encoding polynomials in R_1 and R_2 in groups \mathbb{G}_1 and \mathbb{G}_2 respectively (see [GWC19b, Lemma 2.2]).

Definition 1 (Q-DLOG Assumption [FKL18]). Fix an integer Q . The Q -DLOG assumption for $(\mathbb{G}_1, \mathbb{G}_2)$ states that given $[1]_1, [x]_1, \dots, [x^Q]_1, [1]_2, [x]_2, \dots, [x^Q]_2$ for uniformly chosen $x \leftarrow \mathbb{F}$, the probability of an efficient \mathcal{A} outputting x is $\text{negl}(\lambda)$.

3.3 KZG Commitment Scheme

A polynomial commitment scheme allows the prover to open evaluations of a committed polynomial succinctly (Appendix A.1). We use the KZG commitment scheme introduced in [KZG10] which satisfies succinctness, completeness and knowledge-soundness (extractability) in the algebraic group model, while additionally featuring a universal and updatable setup. We denote the KZG scheme by the tuple of PPT algorithms (KZG.Setup, KZG.Commit, KZG.Prove, KZG.Verify) as defined below.

Definition 2 (KZG Polynomial Commitment Scheme). Let $(\mathbb{F}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2, g_t)$ be output of bilinear group generator

$\text{BG}(1^\lambda)$ for security parameter λ . The KZG polynomial commitment scheme is defined as follows:

- KZG.Setup on input $(1^\lambda, d)$, where d is the degree bound, outputs $\text{srs} = (\{[\tau]_1, \dots, [\tau^d]_1\}, \{[\tau]_2, \dots, [\tau^d]_2\})$.
- KZG.Commit on input $(\text{srs}, p(X))$, where $p(X) \in \mathbb{F}_{\leq d}[X]$, outputs $C = [p(\tau)]_1$
- KZG.Prove on input $(\text{srs}, p(X), \alpha)$, where $p(X) \in \mathbb{F}_{\leq d}[X]$ and $\alpha \in \mathbb{F}$, outputs (v, π) such that $v = p(\alpha)$ and $\pi = [q(\tau)]_1$, for

$$q(X) = \frac{p(X) - p(\alpha)}{X - \alpha}$$

- KZG.Verify on input $(\text{srs}, C, v, \alpha, \pi)$, outputs 1 if the following equation holds, and 0 otherwise.

$$e(C - v[1]_1 + \alpha\pi, [1]_2) = e(\pi, [\tau]_2)$$

Note that both sides of the verification equation involve a fixed generator, and hence several proof verifications can be batched together to reduce the number of pairing computations. We also assume (w.l.o.g) analogues of KZG.Commit, KZG.Prove and

KZG.Verify defined over the group \mathbb{G}_2 . We shall use the (non-standard) notation $[p(X)]_i$ to denote $[p(\tau)]_i$ for $i \in \{1, 2\}$. This allows us a convenient shorthand for referring to “commitment of the polynomial $p(X)$ ” in group \mathbb{G}_i . Our protocols also use batched KZG proofs to show that polynomial $p(X)$ satisfies $p(\alpha_i) = v_i$ for $i \in [n]$. Let $\alpha = (\alpha_1, \dots, \alpha_n)$ denote the vector of evaluation points and $\mathbf{v} = (v_1, \dots, v_n)$ denote the vector of claimed evaluations. Then the batched version of KZG.Prove is described as follows:

- KZG.Prove on input $(\text{srs}, p(X), \alpha)$, where $p(X) \in \mathbb{F}_{\leq d}[X]$ and $\alpha \in \mathbb{F}^n$, outputs (\mathbf{v}, π) with $\mathbf{v} \in \mathbb{F}^n$ such that $v_i = p(\alpha_i)$ and $\pi = [q(\tau)]_1$ where

$$q(X) = \frac{p(X) - r(X)}{a(X)}$$

In the above equation, $a(X) = (X - \alpha_1) \cdots (X - \alpha_n)$, while $q(X)$ and $r(X)$ are the quotient and remainder polynomials when $p(X)$ is divided by $a(X)$.

- KZG.Verify on input $(\text{srs}, C, \mathbf{v}, \alpha, \pi)$, outputs 1 if the following equation is satisfied, and 0 otherwise.

$$e(C - [r(\tau)]_1, [1]_2) = e(\pi, [a(\tau)]_2)$$

Here, the verifier interpolates the polynomial $r(X) \in \mathbb{F}_{< n}[X]$ such that $r(\alpha_i) = v_i$.

KZG for Vectors. For $\mathbf{f} \in \mathbb{F}^N$, let $\text{Enc}_{\#}(\mathbf{f})$ denote the polynomial encoding of \mathbf{f} over \mathbb{H} given by $\sum_{i=1}^N f_i \mu_i(X)$. We use KZG to commit to *vectors* by committing to its polynomial encoding. In general a vector \mathbf{g} of size m is encoded by a polynomial $g(X) \in \mathbb{F}_{< m}[X]$ which interpolates \mathbf{g} over a subgroup \mathbb{V} consisting of m^{th} roots of unity in some canonical order. We will explicitly state the subgroups for all sizes of vectors that we consider.

3.4 Lookup Arguments

Prior works on lookup arguments [ZBK⁺22, PK22, ZGK⁺22, EFG22] consider proving sub-vector relation over committed vectors, i.e, given commitments c_t and c_v to vectors $\mathbf{t} \in \mathbb{F}^N$ and $\mathbf{v} \in \mathbb{F}^m$, one proves that for all $i \in [m]$, there exists $j \in [N]$ such that $v_i = t_j$. We will use $\mathbf{v} \preceq \mathbf{t}$ to denote that \mathbf{v} is a sub-vector of \mathbf{t} . The definition below summarizes the sub-vector relation as defined in prior works.

Component	Protocol	Prover Work	Verifier Work	Communication
Committed Sub-vector Lookup	CQ [EFG22]	$O(m \log m) \mathbb{F}, O(m) \mathbb{G}_1$	$5P$	$8\mathbb{G}_1, 3\mathbb{F}$
Committed Index Lookup	Figure 2	$O(m \log m) \mathbb{F}, O(m) \mathbb{G}_1$	$5P$	$8\mathbb{G}_1, 3\mathbb{F}$
Localized Update in RAM	Figure 3	$O(m \log^2 m) \mathbb{F}, O(m) \mathbb{G}_1$	$8P$	$19\mathbb{G}_1, 1\mathbb{G}_2, 10\mathbb{F}$
Table Specific Preprocessing	Fast KZG [FK23]	$O(N \log N) \mathbb{F}, \mathbb{G}$	-	-
Lookup from Approximate Setup	Section 7	$O((m + \delta) \log^2(m + \delta)) \mathbb{F}, O(m + \delta) \mathbb{G}_1$	-	-
Polynomial Protocol for RAM	Figure 11	$O(m \log m) \mathbb{F}, O(m) \mathbb{G}$	$7P$	$36\mathbb{G}_1, 30\mathbb{F}$
Batching-Efficient RAM	Figure 4	$\tilde{O}(\sqrt{mN}) \mathbb{F}, \mathbb{G}$	$9P$	$65\mathbb{G}_1, 1\mathbb{G}_2, 43\mathbb{F}$

Table 2: Asymptotic efficiency of the component protocols for our scheme. Here, N denotes the size of the RAM, m denotes the number of operations, and δ denotes Hamming distance of table for which pre-computed parameters are available from the current table. As before, we use $(\mathbb{F}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2, g_t)$ to denote a bilinear group, and P to denote a pairing evaluation. The performance figures reported here correspond to the CQ-based realization of our batching-efficient RAM scheme.

Definition 3. We define the committed sub-vector relation $\mathcal{R}_{\text{srs}, N, m}^{\text{subvec}}$ to consist of tuples $((c_t, c_v), (\mathbf{t}, \mathbf{v}))$ where $c_t, c_v \in \mathbb{G}_1$, $\mathbf{t} \in \mathbb{F}^N$, $\mathbf{v} \in \mathbb{F}^m$ such that $\mathbf{v} \preceq \mathbf{t}$ and $c_t = \text{KZG.Commit}(\text{srs}, \text{Enc}_{\mathbb{H}}(\mathbf{t}))$ and $c_v = \text{KZG.Commit}(\text{srs}, \text{Enc}_{\mathbb{V}}(\mathbf{v}))$.

A committed sub-vector argument is an argument of knowledge for the relation $\mathcal{R}_{\text{srs}, N, m}^{\text{subvec}}$. Next, we consider a slightly modified relation that we call *committed index lookup* (called indexed lookup in [STW23]) where there is a commitment to the indices where \mathbf{v} appears in \mathbf{t} . Formally, we define it as below:

Definition 4. We define the committed index lookup relation $\mathcal{R}_{\text{srs}, N, m}^{\text{lookup}}$ to consist of tuples of the form $((c_t, c_a, c_v), (\mathbf{t}, \mathbf{a}, \mathbf{v}))$ where $c_t, c_a, c_v \in \mathbb{G}_1$, $\mathbf{t} \in \mathbb{F}^N$, $\mathbf{a}, \mathbf{v} \in \mathbb{F}^m$ such that $v_i = \mathbf{t}[a_i] = t_{a_i}$ for all $i \in [m]$ and $c_t = \text{KZG.Commit}(\text{srs}, \text{Enc}_{\mathbb{H}}(\mathbf{t}))$, $c_a = \text{KZG.Commit}(\text{srs}, \text{Enc}_{\mathbb{A}}(\mathbf{a}))$ and $c_v = \text{KZG.Commit}(\text{srs}, \text{Enc}_{\mathbb{V}}(\mathbf{v}))$.

A committed lookup argument is a succinct argument of knowledge for the relation $\mathcal{R}_{\text{srs}, N, m}^{\text{lookup}}$.

4 Technical Overview

As we have alluded to earlier, existing memory-checking based techniques to model RAM computations incur a cost that is linear in the size of the RAM. We are interested in the setting where the number of operations whose execution is to be verified is much smaller than the size of the RAM. Thus, our goal is to achieve prover complexity which is *sublinear* in the size of the RAM. Before we proceed, we establish a working definition of RAM for the rest of the paper. Informally, a RAM maps indices (addresses) to values, where we assume that values come from a finite field \mathbb{F} , while indices come from a subset I of \mathbb{F} . For us, I will generally be the set $\{1, \dots, k\}$ for some integer k (which may be different from size of the RAM n). Finally, for an index, there should be at most one value in the RAM, i.e., the association is unambiguous. The formal definition of RAM is as follows:

Definition 5 (RAM). Given $n \in \mathbb{N}$, finite field \mathbb{F} and a set $\mathcal{I} \subseteq \mathbb{F}$, a RAM of size n over indices \mathcal{I} is a tuple $\mathbf{T} = (\mathbf{a}, \mathbf{v}) \in \mathcal{I}^n \times \mathbb{F}^n$ such that $\forall i, j \in [n] v_i = v_j$ whenever $a_i = a_j$. We think of \mathbf{T} as a table with vectors \mathbf{a} and \mathbf{v} denoting its columns. The set of all such tables will be denoted by $\text{RAM}_{\mathcal{I}, n}$.

For a table $\mathbf{T} = (\mathbf{a}, \mathbf{v}) \in \text{RAM}_{\mathcal{I}, n}$, we refer to tuples (a_i, v_i) , $i \in [n]$ as records of the table \mathbf{T} . We use the access notation $v = \mathbf{T}[a]$ to mean that (a, v) is a record of \mathbf{T} (note there can be multiple such records according to our definition). When we consider RAMs where the first column (of indices) is of the form $\mathbf{I}_n = (1, 2, \dots, n)$, we simply denote such RAMs by $\mathbf{T} \in \mathbb{F}^n$. For a RAM $\mathbf{T} \in \text{RAM}_{\mathcal{I}, n}$, a RAM operation is a three tuple (op, a, v) with $op \in \{0, 1\}$, $a \in \mathcal{I}$ and $v \in \mathbb{F}$. An operation with $op = 0$ is called a *load* operation which denotes reading a value v mapped to index a in the RAM. Similarly, an operation with $op = 1$ is called a *store* operation, which denotes associating the value v with index a in the RAM. We use $\mathcal{O}_{\mathcal{I}}$ to denote the set of all RAM operations with index set \mathcal{I} .

4.1 Batching-Efficient RAM: Blueprint

We will use a vectors in \mathbb{F}^N to denote the “large” RAMs, where index column is implicitly assumed to be $(1, \dots, N)$. Let $\mathbf{T}, \mathbf{T}' \in \mathbb{F}^N$ denote the initial and final RAM states, and let \mathbf{o} be a sequence of m

operations ($m < N$) which updates \mathbf{T} to \mathbf{T}' . Let $\mathbf{a} \in \mathbb{F}^m$ denote the vector of RAM indices referenced by the operations in \mathbf{o} , i.e. a_i is the index referenced by the i^{th} operation. To prove the transformation of \mathbf{T} to \mathbf{T}' via operation sequence \mathbf{o} , we proceed as follows:

- We isolate sub-tables $\mathbf{S} = (\mathbf{a}, \mathbf{v})$ and $\mathbf{S}' = (\mathbf{a}, \mathbf{v}')$ of \mathbf{T} and \mathbf{T}' consisting of rows corresponding to indices in \mathbf{a} . This requires proving $\mathbf{v} = \mathbf{T}[\mathbf{a}]$ and $\mathbf{v}' = \mathbf{T}'[\mathbf{a}]$, which we show using *committed index lookup* argument discussed in Section 6.1.
- On the isolated sub-tables \mathbf{S} and \mathbf{S}' of size m , we use the standard memory checking arguments (c.f. argument presented in Appendix B) to prove that sequence \mathbf{o} correctly updates \mathbf{S} to \mathbf{S}' with prover complexity of $\tilde{O}(m)$.
- Finally, we show that the RAMs \mathbf{T} and \mathbf{T}' are identical outside indices in \mathbf{a} . We describe the protocol for proving the same in Section 6.2.

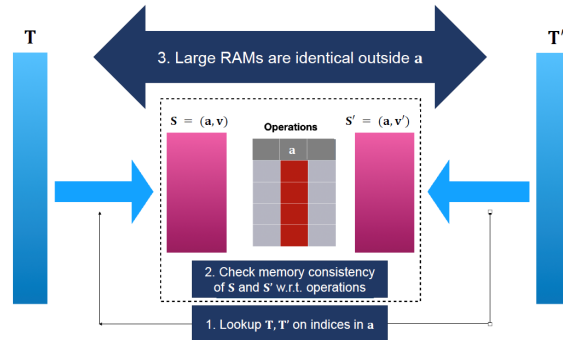


Fig. 1: Illustrating different steps of sublinear lookup protocol between large RAMs \mathbf{T} and \mathbf{T}' .

The blueprint for the above approach is illustrated in Figure 1.

4.2 Batching-Efficient RAM: Components

We now elaborate on the key technical components in realizing the above blueprint.

Committed Index Lookup. To limit the size of the RAM on which we use memory-checking techniques, our first step is to isolate sub-tables of RAMs \mathbf{T} and \mathbf{T}' corresponding to addresses which are involved in the operations. This is achieved by looking up RAMs \mathbf{T} and \mathbf{T}' at indices in the committed vector \mathbf{a} . We could leverage the recent work on efficient lookup arguments to verifiably extract m indices from a table of size N , in time dependent only on m . However, there are two technical challenges here. First, the aforementioned lookup arguments only prove the sub-vector relation, without linking the extracted vector to the indices in \mathbf{a} . This is easily solved, as there is an efficient realization of a *committed index lookup* from a *committed sub-vector* argument, where the commitment scheme is homomorphic. The details appear in Section 6.1, with the complete protocol presented in Figure 2. The second challenge is much more formidable: the efficiency of sub-vector arguments (and the committed index lookup argument derived from them) depends on *expensive* table-specific pre-processing. This is acceptable when the table in question is static, but is infeasible in our setting requiring updatable tables. This motivates our next technical component.

Fast Lookup from Approximate Setup. We build upon the rich body of work on polynomial protocols enabling efficient lookups from static tables [ZBK⁺22,PK22,ZGK⁺22,EFG22], which rely on expensive table-dependent pre-computation to optimise online proving performance. We make the first attempt towards breaking this rigid dependence. Our key idea is to extend the utility of pre-computed parameters for a table \mathbf{T} , to proving lookups from tables $\mathbf{T}' \neq \mathbf{T}$. We show that for $\delta = \Delta(\mathbf{T}, \mathbf{T}')$, an argument for m lookups from \mathbf{T}' incurs an additional prover overhead of $(m + \delta) \log^2(m + \delta)$ over the lookup argument for static tables. We note that the overhead is *quasi-linear* in both m and δ . Our competitive overhead rests on several innovative applications of algebraic algorithms, which are summarised in Appendix D.1. We then leverage this ability to use “approximate” setup into a *base + cache* strategy; where at all times we maintain pre-computed parameters corresponding to a base table

\mathbf{T}_b , and use this setup to prove lookups from the current table \mathbf{T} . We achieve optimal prover effort on average by using parameters for \mathbf{T}_b till the current table is at a hamming distance at most \sqrt{mN} from \mathbf{T}_b , beyond which we recompute full parameters for the current table with $O(N \log N)$ prover effort. The cycle then repeats with current table as the base table.

Naive Approaches are Inadequate. We notice that the aforementioned constructions of lookup arguments require linear combination of encoded quotients of the form $[(T(X) - T(\xi^i))/(X - \xi^i)]_g$ for upto m values of i during the proof generation. While constructions [ZBK⁺22,PK22] consider quotients encoded in the group \mathbb{G}_2 , the protocol in [EFG22] encodes them in \mathbb{G}_1 . We use a generic $[\cdot]_g$ to account for protocol-specific choices. We also see that even a small change to the table requires one to update all the quotients (the polynomial $T(X)$ is common to all quotients). Updating all the quotients after each batch is clearly infeasible. One could consider delaying the updation of the quotients, till the time they are actually required in a proof, which happens when the corresponding index in the table is involved in lookup. However, each of the m quotients is now potentially “lagging” by δ updates, so we would need $\Omega(m\delta)$ group operations to refresh all of them. This gives us multiplicative degradation with δ , and is clearly unsustainable for reasonable values of δ . In Section 7, we present an efficient method to directly compute linear combination of upto $O(m)$ encoded quotients of the form $[(T(X) - T(\xi^i))/(X - \xi^i)]_g$.

Localizing changes in RAMs. While the above two components allow us to reliably extract sub-RAMs corresponding to indices in vector \mathbf{a} , we still need to prove that RAMs are identical outside indices in \mathbf{a} . Looking ahead, in terms of polynomials this requires proving that $T(\xi^i) = T^*(\xi^i)$ for $i \notin \{a_i : i \in [m]\}$. Assuming $Z_I(X)$ to be the vanishing polynomial of the set $\{\xi^{a_i} : i \in [m]\}$, this is equivalent to proving that $Z_I(X)(T(X) - T^*(X)) = D(X)Z_{\mathbb{H}}(X)$ for some polynomial D . However, naively this involves working with polynomials with degree $O(N)$, which is expensive. In Section 6.2 we show a polynomial protocol for the above relation which requires only $O(m \log^2 m)$ prover effort. The protocol appears in Figure 3.

Polynomial Protocol for Memory Checking. To complete the verification, we need to show that the smaller RAMs, $\mathbf{S} = (\mathbf{a}, \mathbf{v})$ and $\mathbf{S}' = (\mathbf{a}, \mathbf{v}')$ extracted from larger RAMs \mathbf{T}, \mathbf{T}' are consistent with respect to the operations. This can be accomplished using standard memory checking techniques based on address ordered transcripts, which we formalize in Section 5. Later in Appendix B and C, we assemble known techniques to present a polynomial protocol for memory consistency based on address ordered transcripts. This involves encoding several artefacts such as operations, transcripts etc., as polynomials and relations among them such as concatenation, permutation and monotonicity as polynomial identities. Our modelling is simple and implementation friendly, and helps in realizing a “circuit-free” overall construction. Complete polynomial protocol for memory checking appears in Figure 11, while constituent protocols appear in Figures 9, 8 and 10.

Efficiency. We conclude the overview with a discussion of efficiency achieved by our scheme, and how different components discussed in this section contribute to the overall efficiency. The asymptotic performance of our scheme using CQ [EFG22] is summarized in Table 2, with efficiency of the overall scheme highlighted in gray. The table also serves as a ready-reckoner for component protocols involved in the overall scheme. A more detailed discussion and break-up of the polynomial protocol for RAM appears in Table 6 in Appendix B. We note that the verification complexity of the overall solution is substantially less than the aggregate of component protocols; this is due to the fact that several pairing checks required for KZG verification proofs can be batched together. For concrete instantiation using BLS12-381 curve, RAM size of 1 million, the online cost of proving an update of 1000 operations on a table as a function of its Hamming distance from the “pre-processed” table is described in Figure 6. Other performance metrics are summarised in Tables 5 and 3. We refer to Section 8 for a more detailed performance evaluation and comparison with prior work. As is clear from the tables above, the (offline) parameter re-computation is the most expensive operation. We reiterate that all of our reported costs throughout the paper are for a single-threaded implementation on a consumer-grade laptop. We believe that parallel implementations can substantially speed up parameter re-computation as their cost is dominated by FFTs over group polynomials, which are highly parallelizable.

Continuity. To support applications such as rollups, we also consider it imperative to ensure that online proof generation does not halt during offline parameter re-generation. In other words, offline parameter re-generation should not hinder the operational continuity of the system. In our scheme, we can ensure this by carefully overlapping the offline computation with online proof generation such that the system can *instantly* switch to using the more recently generated parameters before the online

proving time becomes prohibitive. We present more concrete discussions around this scheduling at the end of Section 8. We note that prior works [OWWB20,CFH⁺22] have not addressed this issue of continuity in detail. To the best of our knowledge, we are the first to highlight the issue and present a discussion on a viable approach.

5 Memory Consistency for RAM

In this section, we briefly review and formalize existing memory-checking techniques to ensure correctness of RAM operations. The formal definitions for various relations involved in memory checking will be used to describe polynomial protocol for RAM in Appendix B.

5.1 Correctness of RAM Update

The versatility of the RAM primitive stems from its updatability. While a load operation leaves the RAM unchanged, the store operation updates the value in the RAM associated with the referenced index. We model the update via the function $\text{Upd}_{\mathcal{I}}$ which takes RAM $\mathbf{T} \in \text{RAM}_{\mathcal{I},n}$, operation $o = (op, a, v) \in \mathcal{O}_{\mathcal{I}}$ as inputs and returns an updated RAM $\mathbf{T}' \in \text{RAM}_{\mathcal{I},n}$. The updated RAM $\mathbf{T}' = \text{Upd}_{\mathcal{I}}(\mathbf{T}, o)$ satisfies $\mathbf{T}' = \mathbf{T}$ if $op = 0$ while for $op = 1$ it satisfies $\mathbf{T}'[a] = v$ and $\mathbf{T}'[x] = \mathbf{T}[x]$ for $x \neq a$. The central problem in verifiable RAM protocols is to establish that a sequence of operations $\mathbf{o} = (o_1, \dots, o_m)$ are correct with respect to the initial RAM state \mathbf{T} and the final RAM state \mathbf{T}' . This involves ensuring that all load operations read the value which is consistent with updates to the RAM as a result of preceding store operations, and that \mathbf{T}' is the final state. We say that an operation $o = (op, a, v)$ is *load-consistent* with respect to RAM \mathbf{T} if $v = \mathbf{T}[a]$ whenever o is a load operation (store operations are vacuously defined to be load-consistent). We formally define the notion of consistency below:

Definition 6 (Consistent Operations). *Let $n \in \mathbb{N}$ and $\mathbf{T}, \mathbf{T}' \in \text{RAM}_{\mathcal{I},n}$ for some index set I . We say that a sequence of operations $\mathbf{o} = (o_1, \dots, o_k) \in \mathcal{O}_{\mathcal{I}}^k$ over I is consistent with RAM states \mathbf{T}, \mathbf{T}' if for all $i \in [k]$, $\mathbf{T}_i = \text{Upd}_{\mathcal{I}}(\mathbf{T}_{i-1}, o_i)$ and operation o_i is load-consistent with respect to \mathbf{T}_{i-1} . Here we assume $\mathbf{T}_0 = \mathbf{T}$ and $\mathbf{T}_k = \mathbf{T}'$.*

For $m, n \in \mathbb{N}$, let $\text{LRAM}_{\mathcal{I},m,n}$ denote the language consisting of tuples $(\mathbf{T}, \mathbf{o}, \mathbf{T}')$ with $\mathbf{T}, \mathbf{T}' \in \text{RAM}_{\mathcal{I},n}$ and $\mathbf{o} \in (\mathcal{O}_{\mathcal{I}})^m$ such that \mathbf{o} is consistent with \mathbf{T}, \mathbf{T}' . Next, we formalize the folklore technique of checking correctness of RAM operations using *address-ordered transcripts*.

5.2 Consistency Check via Transcripts

A *transcript* is time-stamped sequence of operations executed on a RAM. More formally, given a RAM $\mathbf{T} = (\mathbf{a}, \mathbf{v}) \in \text{RAM}_{\mathcal{I},n}$, operation sequence $\mathbf{o} = (o_1, \dots, o_m)$ with $o_i = (op_i, \bar{a}_i, \bar{v}_i) \in \mathcal{O}_{\mathcal{I}}$ and RAM $\mathbf{T}' = (\mathbf{a}', \mathbf{v}') \in \text{RAM}_{\mathcal{I},n}$, the *time ordered transcript* for the tuple $(\mathbf{T}, \mathbf{o}, \mathbf{T}')$ is given by the table tr with $k = 2n + m$ rows and four columns $\text{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ defined as follows: (i) $\mathbf{t} = I_k = (1, \dots, k)$, (ii) $\mathbf{op} = 0^n \parallel (\bar{op}_1, \dots, \bar{op}_m) \parallel 0^n$, (iii) $\mathbf{A} = \mathbf{a} \parallel (\bar{a}_1, \dots, \bar{a}_m) \parallel \mathbf{a}'$ and (iv) $\mathbf{V} = \mathbf{v} \parallel (\bar{v}_1, \dots, \bar{v}_m) \parallel \mathbf{v}'$. The i^{th} row of the table tr is (t_i, op_i, A_i, V_i) for $i \in [k]$. The first n records in tr correspond to the contents of \mathbf{T} , the next m records correspond to the operations in \mathbf{o} and final n records correspond to contents of \mathbf{T}' . The timestamp column \mathbf{t} is added to order operations with the same index. Notationally, we write $\text{tr} = \text{TimeTr}(\mathbf{T}, \mathbf{o}, \mathbf{T}')$.

We call a transcript $\text{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ to be *address ordered* if $A_i \leq A_{i+1}$ for $i \in [k-1]$ and $t_i < t_{i+1}$ whenever $A_i = A_{i+1}$. For a transcript $\text{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ with k records and a permutation $\sigma : [k] \rightarrow [k]$, we use $\sigma(\text{tr})$ to denote the transcript $(\sigma(\mathbf{t}), \sigma(\mathbf{op}), \sigma(\mathbf{A}), \sigma(\mathbf{V}))$ obtained by permuting the records of tr according to the permutation σ . An address ordered transcript for tuple $(\mathbf{T}, \mathbf{o}, \mathbf{T}')$ is defined as $\text{tr}^* = \sigma(\text{tr})$ where $\text{tr} = \text{TimeTr}(\mathbf{T}, \mathbf{o}, \mathbf{T}')$ and σ is a permutation such that tr^* is address ordered. We denote it by $\text{tr}^* = \text{AddrTr}(\mathbf{T}, \mathbf{o}, \mathbf{T}')$. We say that an address ordered transcript $\text{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ satisfies *load-store correctness* if for all pairs of consecutive records (t_i, op_i, A_i, V_i) and $(t_{i+1}, op_{i+1}, A_{i+1}, V_{i+1})$ we have $V_{i+1} = V_i$ whenever $op_{i+1} = 0$ (load operation) and $A_i = A_{i+1}$, i.e., a load operation does not change the value at an index. We formally state the folklore technique for enforcing memory consistency in our setting.

Common Input: $\text{srs}, c_t, c_a, c_v, c_I = [I(X)]_1$ where $I(X) = \text{Enc}_{\mathbb{H}}(\mathbf{I})$ encodes the vector $\mathbf{I} = (1, \dots, N) \in \mathbb{F}^N$.

Prover's Input: Vectors $\mathbf{t} \in \mathbb{F}^N, \mathbf{a}, \mathbf{v} \in \mathbb{F}^m$.

1. \mathcal{V} sends $\gamma \leftarrow \mathbb{F}$.
2. \mathcal{P} and \mathcal{V} compute: $\tilde{c}_t = \gamma c_I + c_t, \tilde{c}_v = \gamma c_a + c_v$.
3. \mathcal{P} computes: $\tilde{\mathbf{t}} = \gamma \mathbf{I} + \mathbf{t}, \tilde{\mathbf{v}} = \gamma \mathbf{a} + \mathbf{v}$.
4. \mathcal{P} and \mathcal{V} run sub-vector argument $(\mathcal{P}_{\text{sv}}, \mathcal{V}_{\text{sv}})$ with $(\text{srs}, \tilde{c}_t, \tilde{c}_v)$ as the common input and $(\tilde{\mathbf{t}}, \tilde{\mathbf{v}})$ as \mathcal{P}_{sv} 's input.
5. \mathcal{V} outputs $b \leftarrow \langle \mathcal{P}_{\text{sv}}(\tilde{\mathbf{t}}, \tilde{\mathbf{v}}), \mathcal{V}_{\text{sv}} \rangle(\text{srs}, \tilde{c}_t, \tilde{c}_v)$.

Fig. 2: Committed Index Lookup Argument

Lemma 1. *Let \mathbb{F} be a finite field, $m, n \in \mathbb{N}$ be positive integers and $I \subseteq \mathbb{F}$. Then $(\mathbf{T}, \mathbf{o}, \mathbf{T}') \in \text{LRAM}_{\mathcal{I}, n, m}$ if and only if the address ordered transcript $\text{tr}^* = \text{AddrTr}(\mathbf{T}, \mathbf{o}, \mathbf{T}')$ satisfies load-store correctness.*

The consistency check in Lemma 1 can be encoded as an arithmetic circuit of size $\tilde{O}(m+n)$, thus yielding an argument of knowledge for the language $\text{LRAM}_{\mathcal{I}, n, m}$ with prover complexity quasi-linear in $m+n$. For completeness, we present a self-contained argument of knowledge for $\text{LRAM}_{\mathcal{I}, m, m}$ ($m=n$) based on the ‘‘polynomial protocol’’ framework defined in [GWC19b].

6 Improved Batching-Efficient RAM

We now detail the steps required to realize batching efficient RAM outlined in the technical overview.

6.1 Committed Index Lookup

In this section, we ‘‘lift’’ any committed sub-vector argument to a committed index lookup argument, where the latter makes a black-box use of the former. We use the trick of random linear combination of vectors to infer indexed lookup relation among them from sub-vector relation over the aggregated vectors. Similar use of random linear combinations has been made in the context of proving permutations in literature (e.g. [dSGOTV22]).

Lemma 2. *Let $\mathbf{t} \in \mathbb{F}^n$ and let $\mathbf{a}, \mathbf{v} \in \mathbb{F}^m$ for some positive integers m, n . Let \mathbf{I}_n denote the vector $(1, \dots, n)$. Then for $\gamma \leftarrow \mathbb{F}$, $(\mathbf{v} + \gamma \mathbf{a}) \preceq (\mathbf{t} + \gamma \mathbf{I}_n)$ implies $\mathbf{v} = \mathbf{t}[\mathbf{a}]$ except with probability $mn/|\mathbb{F}|$.*

Proof. We define vectors of linear polynomials $\mathbf{p} = (p_1, \dots, p_m)$ and $\mathbf{q} = (q_1, \dots, q_n)$ where $p_i(X) = v_i + a_i X, i \in [m]$ and $q_i(X) = t_i + iX, i \in [n]$. Now, we see that $\mathbf{v} = \mathbf{t}[\mathbf{a}]$ if and only if $\mathbf{p} \preceq \mathbf{q}$. For $\gamma \in \mathbb{F}$, let \mathbf{p}_γ and \mathbf{q}_γ denote the vectors $(p_1(\gamma), \dots, p_m(\gamma))$ and $(q_1(\gamma), \dots, q_n(\gamma))$ respectively. It is obvious that $\mathbf{p} \preceq \mathbf{q}$ implies $\mathbf{p}_\gamma \preceq \mathbf{q}_\gamma$ for all $\gamma \in \mathbb{F}$. Using Schwartz-Zippel Lemma, it can also be seen that $\Pr_{\gamma \leftarrow \mathbb{F}}[\mathbf{p} \not\preceq \mathbf{q} \mid \mathbf{p}_\gamma \preceq \mathbf{q}_\gamma] \leq mn/|\mathbb{F}|$. The bound follows from the observation that the event occurs only when γ is a common root of at least one pair of linear polynomials $\{(p_i(X), q_j(X)) : i \in [m], j \in [n]\}$.

In Figure 2, we invoke Lemma 2 to construct a committed index lookup argument using a committed sub-vector argument $(\mathcal{P}_{\text{sv}}, \mathcal{V}_{\text{sv}})$. We formally state the following lemma, whose proof essentially follows from Lemma 2.

Lemma 3. *Assuming that $(\mathcal{P}_{\text{sv}}, \mathcal{V}_{\text{sv}})$ is an argument of knowledge for the relation $\mathcal{R}_{\text{srs}, N, m}^{\text{subvec}}$ in the AGM, the interactive protocol in Figure 2 is an argument of knowledge for the relation $\mathcal{R}_{\text{srs}, N, m}^{\text{lookup}}$ in the AGM.*

6.2 Almost Identical RAM States

For a vector $\mathbf{a} \in [N]^m$, let $\text{uniq}(\mathbf{a}) = \{a_i : i \in [m]\}$ denote the subset of unique values in \mathbf{a} . We call two RAM states $\mathbf{T}, \mathbf{T}' \in \mathbb{F}^N$ to be \mathbf{a} -*identical* if $\mathbf{T}[i] = \mathbf{T}'[i]$ for all $i \notin \text{uniq}(\mathbf{a})$. As before, let $T(X), T^*(X)$ and $a(X)$ be polynomials encoding the vectors \mathbf{T}, \mathbf{T}' (over \mathbb{H}) and \mathbf{a} (over \mathbb{V}). Let c_T, c_T' and c_a be the

commitments to vectors \mathbf{T}, \mathbf{T}' and \mathbf{a} respectively in the group \mathbb{G}_1 . The polynomial protocol to prove that $\mathbf{T}, \mathbf{T}' \in \mathbb{F}^N$ and $\mathbf{a} \in \mathbb{F}^m$ are \mathbf{a} -identical requires proving the relation $Z_I(X)(T(X) - T^*(X)) = 0$ over the set $Z_{\mathbb{H}}$ where $I = \text{uniq}(\mathbf{a})$ and $Z_I(X) = \prod_{i \in I} (X - \xi^i)$ is the vanishing polynomial for the set $\mathbb{H}_I = \{\xi^i : i \in I\}$. To proceed, the honest prover commits to polynomial $Z_I(X)$ and proves (i) $Z_I(X) \cdot (T(X) - T^*(X)) = 0 \pmod{Z_{\mathbb{H}}}$ and (ii) the zeroes of $Z_I(X)$ form a subset of zeroes of $\mathbb{H}_I(X)$ as defined. Together, the two conditions imply that $T(\xi^i) = T^*(\xi^i)$ for $i \notin \text{uniq}(\mathbf{a})$. To prove the first relation, the prover computes the polynomial $D(X)$ as below:

$$\begin{aligned} D(X) &= \frac{(T(X) - T^*(X)) \cdot Z_I(X)}{Z_{\mathbb{H}}(X)} \\ &= \sum_{i \in I} \frac{(T(\xi^i) - T^*(\xi^i)) \mu_i(X)}{Z_{\mathbb{H}}(X)} Z_I(X) \end{aligned}$$

Substituting $\Delta_i = T(\xi^i) - T^*(\xi^i)$, $\mu_i(X) = Z_{\mathbb{H}}(X)/(Z'_{\mathbb{H}}(\xi^i)(X - \xi^i))$, we get

$$D(X) = \sum_{i \in I} \frac{\Delta_i}{Z'_{\mathbb{H}}(\xi^i)} \left(\frac{Z_I(X)}{X - \xi^i} \right) = \sum_{i \in I} \frac{\Delta_i Z'_I(\xi^i)}{Z'_{\mathbb{H}}(\xi^i)} \kappa_i(X) \quad (1)$$

In the above, the summation only runs over indices in I , as $\Delta_i = 0$ for $i \notin I$. In the final equality, we use $\kappa_i(X) = Z_I(X)/(Z'_I(\xi^i)(X - \xi^i))$ for $i \in I$ which we recognize as the lagrange basis polynomials for the set $\{\xi^i : i \in I\}$. Thus, Equation (1) implies that $D(X)$ is at most degree $|I|-1$ polynomial, with $D(\xi^i) = \Delta_i Z'_I(\xi^i)/Z'_{\mathbb{H}}(\xi^i)$ for $i \in I$. The prover can therefore interpolate $D(X)$ (in power basis) in $O(|I| \log^2 |I|)$ \mathbb{F} -operations and compute $[D(X)]_1$ in $O(|I|)$ \mathbb{G}_1 -operations. The prover sends the commitment $[D(X)]_1$ to the verifier. Finally, the verifier can check the identity $Z_I(X) \cdot (T(X) - T^*(X)) = D(X) \cdot Z_{\mathbb{H}}(X)$ by a pairing check. For this, since the tables are committed in \mathbb{G}_1 , prover will need to send $[Z_I(X)]_2$.

Next, the prover needs to show that zeroes of Z_I are indeed in the set $\mathbb{H}_I = \{\xi^i : i \in I\} = \{\xi^{a_i} : i \in [m]\}$. Clearly, it suffices to show that $Z_I(X)$ divides the polynomial $\prod_{i \in [m]} (X - \xi^{a_i})$. To obtain a polynomial protocol, the prover commits to an auxiliary polynomial $h(X) = \sum_{i=1}^m \xi^{a_i} \tau_i(X)$ which interpolates the vector $\mathbf{h} = (\xi^{a_1}, \dots, \xi^{a_m})$. The correctness of h polynomial can be established by showing that the interpolated vector \mathbf{h} satisfies committed index lookup relation $\mathbf{h} = \mathbf{T}_{\text{exp}}[\mathbf{a}]$ where $\mathbf{T}_{\text{exp}} = (\xi^1, \dots, \xi^N)$. Moreover, we notice that the polynomial interpolating the table \mathbf{T}_{exp} is particularly simple, i.e, $T_{\text{exp}}(X) = X$, and thus the setup need not be augmented with table-specific parameters for \mathbf{T}_{exp} . Finally, it remains to show that $Z_I(X)$ divides $K(X) = \prod_{i=1}^m (X - h(\nu^i))$. To do so, the prover commits to $K(X)$ and the quotient polynomial $q(X) = K(X)/Z_I(X)$. The verifier checks the polynomial identities at α , i.e $K(\alpha) = q(\alpha)Z_I(\alpha)$ and $K(\alpha) = \prod_{i=1}^m (\alpha - h(\nu^i))$. The former is easily accomplished using evaluation proofs for K, q and Z_I at α . For checking the latter, the prover commits to another polynomial $u(X)$ satisfying $u(\nu^i) = \prod_{j=1}^{i-1} ((\alpha - h(\nu^j))/(1 + \beta \tau_1(\nu^j)))$ for $i \in [m]$ where $\beta = K(\alpha) - 1$. The verifier ensures the correctness of $u(X)$ by checking:

$$\begin{aligned} \tau_1(X)(u(X) - 1) &= 0 \pmod{Z_{\mathbb{V}}} \\ u(\nu X)(1 + \beta \tau_1(X)) - u(X)(\alpha - h(X)) &= 0 \pmod{Z_{\mathbb{V}}}. \end{aligned} \quad (2)$$

We prove that the above constraints imply that $K(\alpha) = \prod_{i \in [m]} (\alpha - h(\nu^i))$ in Lemma 4. Note that in this protocol we require commitment to the polynomial Z_I in both \mathbb{G}_1 and \mathbb{G}_2 , and thus another pairing check is required to show that the $Z_I(X)$ committed in \mathbb{G}_1 is the same as the $Z_I(X)$ committed in \mathbb{G}_2 (used for the real pairing check). The complete protocol for checking that RAMs \mathbf{T} and \mathbf{T}' are identical outside indices in \mathbf{a} is given in Figure 3.

Lemma 4. *There exists a polynomial $u(X) \in \mathbb{F}[X]$ satisfying the identities in Equation (2) if and only if $K(\alpha) = 1 + \beta = \prod_{i \in [m]} (\alpha - h(\nu^i))$.*

Proof. Assume that the identities hold for some polynomial $u(X)$. The first identity implies $u(\nu) = 1$. From the second identity, we conclude that for all $i \in [m]$, we have $u(\nu^{i+1}) = u(\nu^i) \cdot ((\alpha - h(\nu^i))/(1 + \beta \tau_1(\nu^i)))$, and thus:

$$1 = u(\nu^{m+1})/u(\nu) = \prod_{i \in [m]} \left(\frac{\alpha - h(\nu^i)}{1 + \beta \tau_1(\nu^i)} \right).$$

We observe that the product of denominators in the above equation is simply $1 + \beta$ as $\tau_1(\nu^i) = 0$ for all $i \neq 1$, and thus $1 + \beta = \prod_{i=1}^m (\alpha - h(\nu^i))$. In the other direction, it is easy to check that $u(X)$ as defined for an honest prover, satisfies the identities in Equation 2.

Common Input: $\text{srs}, c_T, c'_T, c_a$.

Prover's Input: Vectors $\mathbf{T}, \mathbf{T}' \in \mathbb{F}^N$, $\mathbf{a} \in \mathbb{F}^m$. Polynomials $T(X), T^*(X)$ and $a(X)$ encoding \mathbf{T}, \mathbf{T}' and \mathbf{a} respectively.

Round 1: Prover commits to auxiliary polynomials

1. \mathcal{P} computes:
 - $I = \text{uniq}(\mathbf{a})$, $Z_I(X) = \prod_{i \in I} (X - \xi^i)$.
 - $D(X) = Z_I(X)(T(X) - T^*(X))/Z_{\mathbb{H}}(X)$.
 - $h(X)$ such that $h(\nu^i) = \xi^{a_i}$ for $i \in [m]$.
 - $K(X) = \prod_{i=1}^m (X - h(\nu^i))$, $q(X) = K(X)/Z_I(X)$.
2. \mathcal{P} sends $c_z = [Z_I(X)]_1$, $c'_z = [Z_I(X)]_2$, $c_d = [D(X)]_1$, $c_h = [h(X)]_1$, $c_k = [K(X)]_1$, $c_q = [q(X)]_1$.
3. \mathcal{V} sends $\alpha \leftarrow \mathbb{F}$.

Round 2: Prover commits to polynomial $u(X)$.

1. \mathcal{P} sets $\beta = K(\alpha) - 1$ and interpolates $u(X)$ on \mathbb{V} such that $u(\nu^i) = \prod_{j=1}^{i-1} ((\alpha - h(\nu^j))/(1 + \beta\tau_1(\nu^j)))$ for $i \in [m]$.
2. \mathcal{P} sends $c_u = [u(X)]_1$.
3. \mathcal{V} sends $r \leftarrow \mathbb{F}$.

Round 3: Prover batches checks in Eq (2).

1. \mathcal{P} computes $Q(X) = (u(\nu X)(1 + \beta\tau_1(X)) - u(X)(\alpha - h(X)) + r\tau_1(X)(u(X) - 1))/Z_{\mathbb{V}}(X)$
2. \mathcal{P} sends $c_Q = [Q(X)]_1$.
3. \mathcal{V} sends $s \leftarrow \mathbb{F}$.

Round 4: Prover sends evaluations.

1. \mathcal{P} sends $\langle z \rangle_\alpha = Z_I(\alpha)$, $\langle q \rangle_\alpha = q(\alpha)$, $\langle K \rangle_\alpha = K(\alpha)$, $\langle Q \rangle_s = Q(s)$, $\langle u \rangle_s = u(s)$, $\langle u \rangle_{\nu s} = u(\nu s)$, $\langle h \rangle_s = h(s)$.
2. \mathcal{V} sends $r_\alpha, r_s \leftarrow \mathbb{F}$.

Round 5: Prover batches evaluation proofs.

1. \mathcal{P} computes:
 - $p_\alpha(X) = Z_I(X) + r_\alpha q(X) + r_\alpha^2 K(X)$.
 - $p_s(X) = Q(X) + r_s u(X) + r_s^2 h(X)$.
 - $\Pi_\alpha = \text{KZG.Prove}(\text{srs}, p_\alpha, \alpha)$.
 - $\Pi_s = \text{KZG.Prove}(\text{srs}, p_s, s)$, $\Pi_{\nu s} = \text{KZG.Prove}(\text{srs}, u, \nu s)$.
2. \mathcal{P} sends $\Pi_\alpha, \Pi_s, \Pi_{\nu s}$.

Round 6: Verifier checks identities.

1. \mathcal{V} computes $[p_\alpha]_1 = c_z + r_\alpha c_q + r_\alpha^2 c_k$, $[p_z]_1 = c_Q + r_s c_u + r_s^2 c_h$.
2. \mathcal{V} checks:
 - $\langle z \rangle_\alpha \cdot \langle q \rangle_\alpha = \langle K \rangle_\alpha$.
 - $\langle u \rangle_{\nu s} (1 + \beta\tau_1(s)) - \langle u \rangle_s (\alpha - \langle h \rangle_s) + r\tau_1(s)(\langle u \rangle_s - 1) = \langle Q \rangle_s Z_{\mathbb{V}}(s)$.
 - $e(c_T - c'_T, c'_z) = e(c_d, [Z_{\mathbb{H}}(X)]_2)$.
 - $e([1]_1, c'_z) = e(c_z, [1]_2)$.
 - $\text{KZG.Verify}(\text{srs}, [p_\alpha]_1, \langle z \rangle_\alpha + r_\alpha \langle q \rangle_\alpha + r_\alpha^2 \langle K \rangle_\alpha, \alpha, \Pi_\alpha)$.
 - $\text{KZG.Verify}(\text{srs}, [p_z]_1, \langle Q \rangle_s + r_s \langle u \rangle_s + r_s^2 \langle K \rangle_s, s, \Pi_s)$.
 - $\text{KZG.Verify}(\text{srs}, c_u, \langle u \rangle_{\nu s}, \nu s, \Pi_{\nu s})$.

Round 7: Check correctness of polynomial h .

1. \mathcal{P} and \mathcal{V} execute committed index lookup argument (Fig 2) to check $([X]_1, c_a, c_h) \in \mathcal{R}_{\text{srs}, N, m}^{\text{lookup}}$.
2. \mathcal{V} accepts if the above argument accepts and all the preceding checks succeed.

Fig. 3: Argument for showing RAMs are identical outside small set of indices.

6.3 Batching-Efficient RAM: Combined Protocol

We put the entire protocol together now. Let I denote the set of indices $\{1, \dots, N\}$, and \mathbf{I}_N denote the vector $(1, \dots, N)$. We formally define the committed RAM relation for which we present an argument of knowledge in this section.

Definition 7. We define the committed ram relation $\mathcal{R}_{\text{srs},N,m}^{\text{ram}}$ to consist of tuples $((c_T, c'_T, c_{op}, c_a, c_w), (\mathbf{T}, \mathbf{T}', \mathbf{op}, \mathbf{a}, \mathbf{w}))$ such that:

- $(\mathbf{T}, \mathbf{o}, \mathbf{T}') \in \text{LRAM}_{\mathcal{I},N,m}$ for $\mathbf{o} = (o_1, \dots, o_m)$ where we have $o_i = (op_i, a_i, w_i) \in \mathcal{O}_{\mathcal{I}}$ for all $i \in [m]$ (here we implicitly view vectors \mathbf{T} and \mathbf{T}' as RAMs with index column \mathbf{I}_N).
- $c_T = \text{KZG.Commit}(\text{srs}, T(X))$, $c'_T = \text{KZG.Commit}(\text{srs}, T^*(X))$, $c_{op} = \text{KZG.Commit}(\text{srs}, op(X))$, $c_a = \text{KZG.Commit}(\text{srs}, a(X))$, $c_w = \text{KZG.Commit}(\text{srs}, w(X))$ where polynomials $T(X), T^*(X)$ encode vectors \mathbf{T}, \mathbf{T}' over \mathbb{H} , while $op(X), a(X)$ and $w(X)$ encode vectors $\mathbf{op} = (op_1, \dots, op_m)$, \mathbf{a} and \mathbf{w} over \mathbb{V} .

As outlined in the blueprint, the prover first commits to “smaller” RAMs $\mathbf{S} = (\mathbf{a}, \mathbf{v})$ and $\mathbf{S}' = (\mathbf{a}, \mathbf{v}')$ where $\mathbf{v} = \mathbf{T}[\mathbf{a}]$ and $\mathbf{v}' = \mathbf{T}'[\mathbf{a}]$. The prover commits to \mathbf{S} and \mathbf{S}' by sending commitments c_v and c'_v to \mathbf{v} and \mathbf{v}' . Then the prover and verifier execute the committed index lookup protocol to prove:

$$(c_T, c_a, c_v) \in \mathcal{R}_{\text{srs},N,m}^{\text{lookup}} \wedge (c'_T, c_a, c'_v) \in \mathcal{R}_{\text{srs},N,m}^{\text{lookup}} \quad (3)$$

The verifier uses a random challenge $\chi \leftarrow \mathbb{F}$ to reduce two instances of $\mathcal{R}_{\text{srs},N,m}^{\text{lookup}}$ to one instance $(c_T + \chi c'_T, c_a, c_v + \chi c'_v) \in \mathcal{R}_{\text{srs},N,m}^{\text{lookup}}$. Then, we show that RAMs \mathbf{T} and \mathbf{T}' are \mathbf{a} -identical using the protocol in Figure 3, described in Section 6.2. All that remains is to prove is that the operation sequence \mathbf{o} is consistent with small RAMs \mathbf{S} and \mathbf{S}' . We check this using the argument in Appendix B, which is obtained by compiling the polynomial protocol for RAM in Appendix C into an argument of knowledge in the AGM. Specifically, the prover and the verifier set $c_S = (c_a, c_v)$, $c'_S = (c_a, c'_v)$ and $c_o = (c_{op}, c_a, c_w)$, and execute the argument of knowledge for showing $(c_S, c_o, c'_S) \in \mathcal{R}_{\text{srs},m}^{\text{LRAM}}$ (see Definition 14). We provide the complete protocol listing in Figure 4. The protocol in Figure 4 assumes pre-computed parameters for the tables \mathbf{T} and \mathbf{T}' . The maintenance of these pre-computed parameters in the presence of updates is detailed in Section 7.

Theorem 3. The protocol in Figure 4 is a succinct argument of knowledge for the relation $\mathcal{R}_{\text{srs},N,m}^{\text{ram}}$ in the AGM, under the Q -DLOG assumption for the bilinear group $(\mathbb{F}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$.

7 Fast Lookups from Approximate Pre-Processing

In this section, we provide details of the algorithm to construct lookup argument for a table \mathbf{T} , using pre-computed parameters of a table which is a small hamming distance away. The dependence on pre-computed parameters in several recent lookup arguments such as [ZBK⁺22, PK22, ZGK⁺22, EFG22] stems from the need to compute an encoded quotient of the form:

$$[Q]_g = \sum_{i \in I} c_i \left[\frac{T(X) - T(\xi^i)}{X - \xi^i} \right]_g \quad (4)$$

for some $O(m)$ sized set I . The quotient in Equation (4) can be computed in $O(m)$ cost when the quotients $[(T(X) - T(\xi^i))/(X - \xi^i)]_g$ are available for all $i \in [N]$. In this section we exhibit an algorithm which computes the above with $O((m + \delta) \log^2(m + \delta))$ cost, given access to similar quotients for a table at hamming distance δ from \mathbf{T} . We now describe our approach.

7.1 Base + Cache approach

The key idea we employ is to express the current table $\mathbf{T} \in \mathbb{F}^N$ as $\mathbf{T}_b + \mathbf{T}_{ch}$, where \mathbf{T}_b is the table for which we assume that the encoded quotients are available (via the $O(N \log N)$ computation), and \mathbf{T}_{ch} captures the changes to the table since. We will periodically update (say after s batch updates) \mathbf{T}_b to current table state, and re-compute all the quotients (we call it the *offline* phase). We will revisit the question on choosing s optimally later. Let $I \subseteq [N]$ denote the set of indices in the current batch of m lookups. The *online* phase of our proof generation involves computing the sum in Equation (4) for the table \mathbf{T} . The following Theorem determines the efficiency of the online phase of our prover.

Setup $(1^\lambda, N, m, \mathbf{T}, \mathbf{T}')$:

- $\text{srs} = (\{\tau^i\}_{i=0}^N, \{\tau^i\}_{i=0}^N)$ for $\tau \leftarrow \mathbb{F}$
- Both \mathcal{P} and \mathcal{V} precomputes $[Z_{\mathbb{H}}(X)]_1, [Z_{\mathbb{H}}(X)]_2$
- \mathcal{P} precomputes the following:
 - $W_2^i = [Z_{\mathbb{H}}(X)/(X - \xi^i)]_2, i \in [N]$
- \mathcal{P} precomputes the following (with respect to $(\mathbf{T}, \mathbf{T}')$):
 - $W_1^i = [(T(X) - T(\xi^i))/(X - \xi^i)]_2, i \in [N],$
 - $W_1^{i'} = [(T^*(X) - T^*(\xi^i))/(X - \xi^i)]_2, i \in [N].$

Common Input: $\text{srs}, c_T, c'_T, c_{op}, c_a, c_w \in \mathbb{G}_1$.

Prover's Input: Vectors $\mathbf{T}, \mathbf{T}', \mathbf{op}, \mathbf{a}, \mathbf{w}$ and their encoding polynomials.

Round 1: Commit to sub RAMs.

1. \mathcal{P} computes $\mathbf{v} = \mathbf{T}[\mathbf{a}], \mathbf{v}' = \mathbf{T}'[\mathbf{a}]$ and the encoding polynomials $v(X)$ and $v^*(X)$.
2. \mathcal{P} sends $c_v = [v(X)]_1, c'_v = [v^*(X)]_1$.
3. \mathcal{V} sends $\chi \leftarrow \mathbb{F}$.

Round 2: Execute committed index lookup.

1. \mathcal{P} and \mathcal{V} compute $\hat{c}_T = c_T + \chi c'_T, \hat{c}_v = c_v + \chi c'_v$.
2. \mathcal{P} computes $\hat{\mathbf{T}} = \mathbf{T} + \chi \mathbf{T}', \hat{\mathbf{v}} = \mathbf{v} + \chi \mathbf{v}'$.
3. \mathcal{P} and \mathcal{V} execute committed index lookup argument in Fig 2, with $(\hat{c}_T, c_a, \hat{c}_v)$ as the common input and $(\hat{\mathbf{T}}, \mathbf{a}, \hat{\mathbf{v}})$ as prover's input.

Round 3: Prove RAMs are \mathbf{a} -identical.

1. \mathcal{P} and \mathcal{V} execute argument in Fig 3 with common input (c_T, c'_T, c_a) and prover's input as $(\mathbf{T}, \mathbf{T}', \mathbf{a})$.

Round 4: Prove sub RAMs are memory-consistent with update.

1. \mathcal{P} and \mathcal{V} execute argument in Fig 11 to check $(c_S, c_o, c'_S) \in \mathcal{R}_{\text{srs}, m}^{\text{LRAM}}$ with $c_S = (c_a, c_v), c'_S = (c_a, c'_v)$ and $c_o = (c_{op}, c_a, c_w)$.
2. \mathcal{V} accepts if all sub-protocols accept.

Fig. 4: Our batching-efficient RAM protocol

Theorem 4. Let N, ξ be as defined previously. Given KZG proofs $\{W_i\}_{i=1}^N$ with

$$W_i = [T_b(X) - T_b(\xi^i)/(X - \xi^i)]_g,$$

where $T_b(X) = \text{Enc}_{\mathbb{H}}(\mathbf{T}_b)$ encodes a vector $\mathbf{T}_b \in \mathbb{F}^N$, for any $I \subseteq [N]$, there exists an algorithm to compute $[Q]_g$ as given in Equation (4) for polynomial $T(X) = \text{Enc}_{\mathbb{H}}(\mathbf{T})$ encoding the vector $\mathbf{T} \in \mathbb{F}^N$ using $O((\delta + |I|) \log^2(\delta + |I|))$ \mathbb{F} -operations and $O(\delta + |I|)$ \mathbb{G} -operations. Here, δ denotes the hamming distance between vectors \mathbf{T}_b and \mathbf{T} .

Proof. Let $\mathbf{T} = \mathbf{T}_b + \mathbf{T}_{\text{ch}}$ and thus $T(X) = T_b(X) + T_{\text{ch}}(X)$. Define $K = I \cup \{j \in [N] : \mathbf{T}_{\text{ch}}[j] \neq 0\}$ as a set which captures the indices where the current table \mathbf{T} differs from the base \mathbf{T}_b , where we explicitly also include the lookup indices I in K . For $j \in K$, let $\mathbf{T}_{\text{ch}}[j] = \Delta t_j$. Then $T_{\text{ch}}(X) = \sum_{j \in K} \Delta t_j \mu_j(X)$. We write the quotient $Q(X)$ as:

$$Q(X) = \sum_{i \in I} c_i \left(\frac{T_b(X) - T_b(\xi^i)}{X - \xi^i} \right) + \sum_{i \in I} c_i \left(\frac{T_{\text{ch}}(X) - T_{\text{ch}}(\xi^i)}{X - \xi^i} \right)$$

From above, we have $[Q(x)]_g = [Q_b(x)]_g + [Q_{\text{ch}}(x)]_g$ where

$$Q_b(X) = \sum_{i \in I} c_i (T_b(X) - T_b(\xi^i))/(X - \xi^i)$$

$$Q_{\text{ch}}(X) = \sum_{i \in I} c_i (T_{\text{ch}}(X) - T_{\text{ch}}(\xi^i))/(X - \xi^i)$$

We can compute $[Q_b(X)]_g$ from the pre-computed KZG openings of $T_b(X)$ at points $\xi^i, i \in I$ using $O(|I|)$ group operations and $O(|I| \log^2 |I|)$ field operations. Therefore, it suffices to compute $[Q_{\text{ch}}(X)]_g$ efficiently. Using $T_{\text{ch}}(X) = \sum_{j \in K} \Delta t_j \mu_j(X)$ we write $Q_{\text{ch}}(X)$ as linear combination of table-independent polynomials:

$$\begin{aligned} Q_{\text{ch}}(X) &= \sum_{i \in I} c_i \sum_{j \in K} \Delta t_j \frac{\mu_j(X) - \mu_j(\xi^i)}{X - \xi^i} \\ &= \sum_{i \in I} c_i \Delta t_i \frac{\mu_i(X) - 1}{X - \xi^i} + \sum_{i \in I} \sum_{j \in K \setminus \{i\}} c_i \Delta t_j \frac{\mu_j(X)}{X - \xi^i} \end{aligned}$$

Now, we can write $[Q_{\text{ch}}(X)]_g = [Q_{\text{ch}}^{(1)}(X)]_g + [Q_{\text{ch}}^{(2)}(X)]_g$ where:

$$Q_{\text{ch}}^{(1)}(X) = \sum_{i \in I} c_i \Delta t_i \frac{\mu_i(X) - 1}{X - \xi^i}, \quad Q_{\text{ch}}^{(2)}(X) = \sum_{i \in I} \sum_{j \in K \setminus \{i\}} c_i \Delta t_j \frac{\mu_j(X)}{X - \xi^i}$$

The term $[Q_{\text{ch}}^{(1)}(X)]_g$ can be computed using $O(|I|)$ group operations by augmenting the setup with pre-computed KZG opening proofs of polynomials $\mu_i(X)$ at ξ^i for $i \in [N]$. This adds $O(N)$ to the setup parameters, while the computation can be done in $O(N \log N)$ time with methods similar to existing pre-computed parameters. This eventually leaves us with $[Q_{\text{ch}}^{(2)}(X)]_g$. Next, we synthesize the polynomial $Q_{\text{ch}}^{(2)}(X)$ in a form that reduces group operations required to compute its encoding.

$$\begin{aligned} Q_{\text{ch}}^{(2)}(X) &= \sum_{i \in I} c_i \sum_{j \in K \setminus \{i\}} \Delta t_j \mu_j(X) / (X - \xi^i) \\ &= \sum_{i \in I} c_i \sum_{j \in K \setminus \{i\}} \frac{\Delta t_j}{Z'_{\mathbb{H}}(\xi^j)} \frac{Z_{\mathbb{H}}(X)}{(X - \xi^i)(X - \xi^j)} \\ &= N^{-1} \sum_{i \in I} c_i \sum_{j \in K \setminus \{i\}} \frac{\xi^j \Delta t_j}{\xi^i - \xi^j} \left(\frac{Z_{\mathbb{H}}(X)}{X - \xi^i} - \frac{Z_{\mathbb{H}}(X)}{X - \xi^j} \right) \\ &= N^{-1} \sum_{i \in I} \left(c_i \cdot \sum_{j \in K \setminus \{i\}} \frac{\xi^j \Delta t_j}{\xi^i - \xi^j} \right) \frac{Z_{\mathbb{H}}(X)}{X - \xi^i} \\ &\quad + N^{-1} \sum_{j \in K} \left(\xi^j \Delta t_j \cdot \sum_{i \in I \setminus \{j\}} \frac{c_i}{\xi^j - \xi^i} \right) \frac{Z_{\mathbb{H}}(X)}{X - \xi^j} \end{aligned} \tag{5}$$

In the first step, we substituted $\mu_j(X)$, while in the final step we re-arranged the summation to accumulate the scalar factor for each distinct polynomial of the form $Z_{\mathbb{H}}(X)/(X - \xi^i)$. Define scalars $a_i, i \in I$ and $b_j, j \in K$ as below:

$$a_i = \sum_{j \in K \setminus \{i\}} \frac{\xi^j \Delta t_j}{\xi^i - \xi^j}, \quad i \in I \quad b_j = \sum_{i \in I \setminus \{j\}} \frac{c_i}{\xi^j - \xi^i}, \quad j \in K \tag{6}$$

Now, define $W_3^j := [Z_{\mathbb{H}}(X)/(X - \xi^j)]_g$. We see that W_3^j is just the KZG opening proof of the polynomial $Z_{\mathbb{H}}(X)$ evaluated at ξ^j for $j \in [N]$. These can be precomputed one time and it adds $O(N)$ to the setup parameters and the computation can be done in $O(N \log N)$ time.

Now, we see that $[Q_{\text{ch}}^{(2)}(X)]_g$ can be written as linear combination of $O(|K| + |I|)$ group elements.

$$[Q_{\text{ch}}^{(2)}(X)]_g = N^{-1} \left(\sum_{i \in I} (c_i a_i) \cdot W_3^i + \sum_{j \in K} (\xi^j \Delta t_j b_j) \cdot W_3^j \right) \tag{7}$$

Now, c_i are known constants depending on the specific lookup scheme. So, given $\{a_i\}_{i \in I}, \{b_j\}_{j \in K}$, $[Q_{\text{ch}}^{(2)}(X)]_g$ can be computed in $O(|I| + |K|)$ group operations. While we have diligently reduced the

group operations, we still seem to need $O(|I||K|) = O(m\delta)$ field operations. We clearly need better than naive way of computing the scalars in (6) to obtain additive overhead in δ . This is what we consider next. Let $d_j := \xi^j \Delta t_j$. Then we have from Eq (6):

$$a_i = \sum_{j \in K \setminus \{i\}} \frac{d_j}{\xi^i - \xi^j}, i \in I \quad b_j = \sum_{i \in I \setminus \{j\}} \frac{c_i}{\xi^j - \xi^i}, j \in K \quad (8)$$

So, to compute a_i and b_j , it suffices to compute *reciprocal sums* efficiently. Our next lemma claims that such reciprocal sums can be computed efficiently. We defer the full proof of Lemma 5 to the Appendix, but illustrate the key ideas in the proof.

Lemma 5. *Let $I \subset K \subset [N]$ and let a_i for all $i \in I$ and b_j for all $j \in K$ be as described above. Then, a_i for all $i \in I$ and b_j for all $j \in K$ can be computed in $O(|K| \log^2 |K|) \mathbb{F}$ operations.*

Proof (Proof-Sketch). We sketch the proof here for a_i . First, we mention that the special case of the lemma when $d_j = 1$ for all $j \in K$ admits an efficient computation due to the following identity proved in Lemma 9.

$$\frac{Z_K''(\xi^i)}{Z_K'(\xi^i)} = 2 \sum_{j \in K \setminus \{i\}} \frac{1}{\xi^i - \xi^j}$$

for $Z_K(X) = \prod_{i \in K} (X - \xi^i)$. The polynomial Z_K can be computed in $O(|K| \log^2 |K|)$ and subsequent evaluations of its first two derivatives can also be evaluated on the set $\{\xi^i : i \in I\}$ with the same complexity. However, to deal with arbitrary values of d_j we need more ingenuity. We will *imagine* d_j to be $p(\xi^j)$ for some polynomial $p(X)$. Moreover, we demand that $p(\xi^j) = 0$ for $j \notin K$. We will not compute such a polynomial p , as it has degree $O(N)$, but view it as an “oracle” which we can hopefully query at the points we need. Then it can be seen that $a_i = g_i(\xi^i) - r_i(\xi^i)$ for rational functions $g_i(X)$ and $r_i(X)$ defined by:

$$g_i(X) = \sum_{j \in [N] \setminus i} \frac{p(X)}{X - \xi^j}, \quad r_i(X) = \sum_{j \in [N] \setminus i} \frac{p(X) - p(\xi^j)}{X - \xi^j} \quad (9)$$

Now, $g_i(\xi^i)$ for $i \in I$ turns out to be (using the special case above):

$$p(\xi^i) \sum_{j \in K \setminus \{i\}} 1/(\xi^i - \xi^j) = d_i (Z_K''(\xi^i)/Z_K'(\xi^i))/2$$

Defining $u(X, Y) = (p(X) - p(Y))/(X - Y)$, we can write $r_i(\xi^i)$ as:

$$r_i(X) = \sum_{j \in [N]} u(X, \xi^j) - u(X, \xi^i) \quad (10)$$

Observe that $u(X, X) = p'(X)$ and so $u(X, X)$ gives the formal derivative of polynomial $p(X)$. We get $r_i(\xi^i) = r(\xi^i) - p'(\xi^i)$ for all $i \in I$, where $r(X) = \sum_{j \in [N]} u(X, \xi^j)$. Fortunately, $r(X)$ is simply $Nu(X, 0) = N(p(X) - p(0))/X$, a fact that follows from uni-variate sum-check. The problem thus reduces to being able to compute derivatives $p'(\xi^i)$ for $i \in I$ and the value $p(0)$. Before concluding the proof-sketch, we briefly highlight the structure of the polynomial $p(X)$. Since $p(X)$ vanishes for $p(\xi^i)$ for $i \notin K$, it can be written as the product $\widehat{Z}_K(X)q(X)$ where \widehat{Z}_K is the vanishing polynomial of “complementary” roots of unity $\{\xi^i : i \notin K\}$ and q is a low-degree ($< K$) polynomial. Assuming we can interpolate $q(X)$, we can write:

$$p'(\xi^i) = \widehat{Z}_K(\xi^i)q'(\xi^i) + \widehat{Z}'_K(\xi^i)q(\xi^i)$$

In the above expression, we require evaluations of high-degree polynomials $\widehat{Z}_K(X)$ and $\widehat{Z}'_K(X)$ at ξ^i , $i \in I$. This is discussed in Lemma 11 and other related lemmas in Appendix D.2, and motivates the at times tedious algebra there. We conclude the proof-sketch, deferring the missing details to the full proof in Appendix D.3.

From the Lemma 5, we conclude that the scalars $a_i, i \in I$ and $b_j, j \in K$ can be computed in time $O(|K| \log^2 |K|)$, which proves the bound in Theorem 4.




Fig. 5: Comparison of R1CS constraints incurred by existing approaches for batching efficient RAMs. MT20 refers to Merkle-tree with depth 20, instantiated using Poseidon hash.

7.2 Amortized Sublinear Batching

We now return to the question of how frequently should we run the offline phase to compute full parameters. For concrete analysis, let s be the period after which the rebasing takes place; i.e., after s batches of m operations each, we set the base table \mathbf{T}_b to the current table, setting $\mathbf{T}_{ch} = \mathbf{0}$. At this point we also compute all encoded quotients for \mathbf{T}_b using the $O(N \log N)$ algorithm of [FK23]. Consider $\delta \leq ms$ as the upper-bound on δ , and distributing the cost of re-basing, the amortized overhead for the batch of m operations is: $O(ms \log^2(ms) + \frac{N \log N}{s})$ \mathbb{F} -operations and $O(ms + \frac{N \log N}{s})$ \mathbb{G} -operations. Ignoring the logarithmic factors, the cost is minimized by setting $s \approx \sqrt{N/m}$, resulting in amortized prover overhead of $\tilde{O}(\sqrt{mN})$. We note that the above analysis considers the worst case scenario, where each update affects a distinct position in the table. In settings, where frequency of updates is non-uniform across positions in the table (e.g, in the blockchain example, if bulk of transactions come from small number of clients), we may be able to defer the offline phase even longer. Same is also true for settings where updates to the table are infrequent.

8 Experiments

In this section we present a concrete evaluation of our batching efficient RAM and compare it to the prior works on batching efficient RAM [OWWB20,CFH⁺22]. We also separately benchmark the effectiveness of our approach of computing encoded quotients presented in Section 7 over the naïve approach. Our implementation is built on top of existing implementation [cau] of lookup argument Caulk [ZBK⁺22]. We make our implementation available at <https://github.com/nitsatiisc/caulk/tree/updateable-ram>.

Experimental Test-bed. All the benchmarks were run single-threaded on a commodity configuration featuring a 2.1GHz Intel-I5 processor, 16 GB memory running Ubuntu Linux 22.04. The implementation was compiled using `--release` flag in Rust. Our protocol was instantiated for BLS12-381 curve, using the scheme CQ [EFG22] as the underlying sub-vector argument.

Online Proof Generation. In Figure 6, we benchmark the time to prove a batch of $m = 2^{10}$ updates on a table of size $N = 2^{20}$. Here, the values on the x -axis denote the hamming distance between the table on which the updates are being proved and the table whose pre-processed parameters are used for the proof generation. Naturally, we expect the proving time to increase as the table becomes more distant from the one used to generate pre-processed parameters. Our proof generation time stays under a minute till the two table differ in almost 2×10^5 positions. In conjunction with offline pre-processing time from table 3, the graph in Figure 6 determines how the offline parameter generation should be scheduled to achieve optimal performance on average.

Offline Pre-Processing. In Table 3 we provide the time to compute table-specific parameters as a function of table size. This is the most computationally intensive step in our scheme involving FFTs




Fig. 6: Online proving time of updates on table \mathbf{T} given the pre-processing parameters for table \mathbf{T}_b , plotted against the Hamming distance δ between \mathbf{T}_b and \mathbf{T} . Here m denotes the batch size of updates, while the RAM size is $N = 2^{20}$. The blue plot corresponds to our scheme instantiated with CQ as the sub-vector argument (in a black-box manner). The red plot corresponds to the non-black-box adaptation of Caulk+ described in Appendix E.

over group polynomials as the main bottle-neck. We believe offline pre-processing can be made an order of magnitude more efficient by leveraging parallel implementation for the FFTs.

Table-Size (N)	2^{10}	2^{12}	2^{14}	2^{16}	2^{18}	2^{20}
Time (s)	7	29	135	620	2766	12000

Table 3: Pre-processing time for tables of different sizes.

Proof Size and Verification Costs. Our proof sizes and verification costs are independent of the size of the table and the number of updates in a batch. For the instantiation of our scheme with BLS12-381 curve, we incur a proof size of $4.4KB$ while the verification takes around $15ms$.

Fast Update Benchmarks. We also benchmark the efficiency of the algorithm to compute scalar coefficients in Lemma 5, required to assemble the encoded quotients from pre-computed quotients. This algorithm is implemented and tested in the `fastupdate` module of the referenced repository. In Table 4, we compare it against the naive computation of the quotients. In the table, we vary the sizes of set I in Lemma 5 in the set $\{2^i : 7 \leq i \leq 10\}$ and set $K = 2^7|I|$. We clearly see $5 \times -20 \times$ advantage over the naive computation.

Continuity. To maintain *continuity* of the system (this is particularly important for applications such as rollups), we must carefully align the online prover performance curve with the cost of offline computation. As an example, suppose \mathbf{T}_0 is the initial pre-processed table at time $t = 0$. We generate proofs using pre-computed parameters for \mathbf{T}_0 till the time $t = t_1$, when the table state \mathbf{T}_1 is at hamming distance 2^{17} from \mathbf{T}_0 . At this point, from Figure 6, online proof generation takes around $40s$ for batch of 1000 updates. At $t = t_1$, we also start an offline parameter computation for the table \mathbf{T}_1 , while continuing to generate online proofs using parameters for \mathbf{T}_0 . We can generate the next 2^7 batches of updates at an average of approximately $12000/128 \approx 94s$ each, thus finishing with a table state \mathbf{T}_2 at hamming distance at most 2^{18} from \mathbf{T}_0 at $t_2 = t_1 + 12000$. At this point, we should have the pre-computed parameters for \mathbf{T}_1 , which is at update distance of 2^{17} from \mathbf{T}_2 , and thus online proof generation can switch to parameters for the table \mathbf{T}_1 . This alignment gives us a proving time of $94s$ per batch of 1000, while ensuring system is live at all times. Clearly, a faster offline pre-computation using parallel implementation would allow us to stay at the cheaper end of online proving performance.

$ I $	$ K $	Lemma 5	Naive
2^7	2^{14}	3.3s	12s
2^8	2^{15}	7.7s	48s
2^9	2^{16}	16.8s	198s
2^{10}	2^{17}	39.2s	839s

Table 4: Comparison of Lemma 5 and naive computation for calculating scalar coefficients for encoded quotients.

Comparison with Prior Works. The proof generation in the prior batching efficient RAM constructions using RSA accumulators [OWWB20,CFH⁺22] involves two key steps (i) generation of a SNARK proof showing knowledge of witness for a relation modeled as arithmetic circuit/R1CS and (ii) computing the witness for the proof generation. A platform agnostic metric to express the cost of the first step is the number of R1CS constraints needed to encode the relation, which is also the metric used in [CFH⁺22] for comparison. Using the R1CS constraints reported in [OWWB20,CFH⁺22] (see Figure 5), we benchmark single-threaded proof generation using the **Groth16** protocol (used in prior works) for R1CS of equivalent size on our test-bed. We use publicly available benchmarking suite in [www] for **Groth16** benchmarks. Since the prior works only report performance of parallel implementation of the second step which is common to both (with degree of parallelism not explicitly mentioned), we will use the reported parallel performance to estimate the overall proving time with this caveat.

Even without the benefit of parallelization, our average proof generation time of $\approx 90s$ for batch size of 2^{10} and RAM size of 2^{20} is $3 - 5\times$ faster than the prior works for the same setting. The proof sizes and verification complexity are constant for prior work and our work. Concrete proof sizes are smaller in [OWWB20,CFH⁺22] owing to their usage of **Groth16** proving backend, while our verification times are competitive with [OWWB20] and substantially less than [CFH⁺22]. One way to reduce proof size in our scheme would be to use a SNARK to prove memory-checking on smaller sub-RAMs, instead of explicit polynomial protocol that we employ for the same. For completeness, we also include the batching efficient RAM using Merkle tree (with Poseidon hash) in the performance comparison in Table 5 which was considered in prior works. Note that the Merkle tree-based approach is faster than that of [OWWB20] for batch size of 2^{10} (the break-even point reported to be batch size of ≈ 1200).

Scheme	\mathcal{P} (s)	\mathcal{V} (ms)	$ \pi $ (KB)
MT20	450	7	0.26
OWWB20 [OWWB20]	$550 + 43$	7	0.26
CFHKKO22 [CFH ⁺ 22]	$226 + 43$	120	1.3
Our Work	94	15	4.4

Table 5: Comparing performance of our batching-efficient RAM with prior works. \mathcal{P} denotes proof generation time, \mathcal{V} denotes verification time while $|\pi|$ denotes argument size. We mention proof generation time as $a + b$ for [OWWB20,CFH⁺22] where a denote proving time of **Groth16** and b denotes witness generation time. The latter is reported in respective works for a parallel implementation.

Acknowledgments

The research of the second author is supported in part by IBM Global University Program Academic Award and Rising Star Award, Intel Corporation.

References

- AST24. Arasu Arun, Srinath Setty, and Justin Thaler. Jolt: Snarks for virtual machines via lookups. In *Advances in Cryptology – EUROCRYPT 2024: 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26–30, 2024, Proceedings, Part VI*, page 3–33, Berlin, Heidelberg, 2024. Springer-Verlag.

- bar. barryWhiteHat. rollup. https://github.com/barryWhiteHat/roll_up.
- BBF19. Dan Boneh, Benedikt Bünz, and Ben Fisch. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 561–586. Springer, Heidelberg, August 2019.
- BCG⁺13. Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, August 2013.
- BCTV14a. Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 276–294. Springer, Heidelberg, August 2014.
- BCTV14b. Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 781–796. USENIX Association, August 2014.
- BFR⁺13a. Benjamin Braun, Ariel J. Feldman, Zuocheng Ren, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. Verifying computations with state. SOSP '13, New York, NY, USA, 2013.
- BFR⁺13b. Benjamin Braun, Ariel J. Feldman, Zuocheng Ren, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. Verifying computations with state (extended version). Cryptology ePrint Archive, Report 2013/356, 2013. <https://eprint.iacr.org/2013/356>.
- BFS20. Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 677–706. Springer, Heidelberg, May 2020.
- BGH19. Sean Bowe, Jack Grigg, and Daira Hopwood. Halo: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Report 2019/1021, 2019. <https://eprint.iacr.org/2019/1021>.
- BGH20. Sean Bowe, Jack Grigg, and Daira Hopwood. Halo2, 2020. <https://github.com/zcash/halo2>.
- BGM17. Sean Bowe, Ariel Gabizon, and Ian Miers. Scalable multi-party computation for zk-SNARK parameters in the random beacon model. Cryptology ePrint Archive, Report 2017/1050, 2017. <https://eprint.iacr.org/2017/1050>.
- cau. Caulk. <https://github.com/caulk-crypto/caulk>.
- CFF⁺24. Matteo Campanelli, Antonio Faonio, Dario Fiore, Tianyu Li, and Helger Lipmaa. Lookup arguments: Improvements, extensions and applications to zero-knowledge decision trees. In *PKC 2024*, volume 14602, pages 337–369, 2024.
- CFH⁺22. Matteo Campanelli, Dario Fiore, Semin Han, Jihye Kim, Dimitris Kolonelos, and Hyunok Oh. Succinct zero-knowledge batch proofs for set accumulators. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 455–469. ACM Press, November 2022.
- 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.
- CL02. Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 61–76. Springer, Heidelberg, August 2002.
- dSGOTV22. Cyprien Delpech de Saint Guilhem, Emmanuela Orsini, Titouan Tanguy, and Michiel Verbauwhede. Efficient proof of RAM programs from any public-coin zero-knowledge system. In *SCN 2022*, pages 615–638, 2022.
- 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>.
- FK23. Dankrad Feist and Dmitry Khovratovich. Fast amortized KZG proofs. Cryptology ePrint Archive, Report 2023/033, 2023. <https://eprint.iacr.org/2023/033>.
- 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.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.
- 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>.
- 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.

- 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>.
- GWC19a. 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>.
- GWC19b. Ariel Gabizon, Zachary J. Williamson, and Oana-Madalina Ciobotaru. Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. *IACR Cryptol. ePrint Arch.*, 2019:953, 2019.
- Hab22. Ulrich Haböck. Multivariate lookups based on logarithmic derivatives. Cryptology ePrint Archive, Report 2022/1530, 2022. <https://eprint.iacr.org/2022/1530>.
- 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.
- Mer88. Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *CRYPTO’87*, volume 293 of *LNCS*, pages 369–378. Springer, Heidelberg, August 1988.
- OWWB20. Alex Ozdemir, Riad S. Wahby, Barry Whitehat, and Dan Boneh. Scaling verifiable computation using efficient set accumulators. In Srdjan Capkun and Franziska Roesner, editors, *USENIX Security 2020*, pages 2075–2092. USENIX Association, August 2020.
- PH23. Shahar Papini and Ulrich Haböck. Improving logarithmic derivative lookups using GKR. *IACR Cryptol. ePrint Arch.*, page 1284, 2023.
- 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>.
- STW23. Srinath Setty, Justin Thaler, and Riad Wahby. Unlocking the lookup singularity with lasso. Cryptology ePrint Archive, Paper 2023/1216, 2023. <https://eprint.iacr.org/2023/1216>.
- WB15. Michael Walfish and Andrew J. Blumberg. Verifying computations without reexecuting them. *Commun. ACM*, page 74–84, 2015.
- WSR⁺15. Riad S. Wahby, Srinath T. V. Setty, Zuo Cheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *NDSS 2015*. The Internet Society, February 2015.
- www. www.arkworks.com. ark-groth16.
- ZBK⁺22. 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.
- ZGK⁺18. Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Pappamanthou. 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.
- 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>.
- ZSG24. Yuncong Zhang, Shifeng Sun, and Dawu Gu. Efficient kzg-based univariate sum-check and lookup argument. In *PKC 2024*, volume 14602, pages 400–425, 2024.

A More Preliminaries

A.1 Polynomial Commitment Scheme

The notion of a polynomial commitment scheme (PCS) that allows the prover to open evaluations of the committed polynomial succinctly was introduced in [KZG10] who gave a construction under the trusted setup assumption. A polynomial commitment scheme over \mathbb{F} is a tuple $\text{PC} = (\text{Setup}, \text{Commit}, \text{open}, \text{eval})$ where:

- $\text{pp} \leftarrow \text{Setup}(1^\lambda, D)$. On input security parameter λ , and an upper bound $D \in \mathbb{N}$ on the degree, Setup generates public parameters pp .
- $(C, \tilde{c}) \leftarrow \text{Commit}(\text{pp}, f(X), d)$. On input the public parameters pp , and a univariate polynomial $f(X) \in \mathbb{F}[X]$ with degree at most $d \leq D$, Commit outputs a commitment to the polynomial C , and additionally an opening hint \tilde{c} .
- $b \leftarrow \text{open}(\text{pp}, f(X), d, C, \tilde{c})$. On input the public parameters pp , the commitment C and the opening hint \tilde{c} , a polynomial $f(X)$ of degree $d \leq D$, open outputs a bit indicating accept or reject.

- $b \leftarrow \text{eval}(\text{pp}, C, d, x, v; f(X))$. A public coin interactive protocol $\langle P_{\text{eval}}(f(X)), V_{\text{eval}} \rangle(\text{pp}, C, d, z, v)$ between a PPT prover and a PPT verifier. The parties have as common input public parameters pp , commitment C , degree d , evaluation point x , and claimed evaluation v . The prover has, in addition, the opening $f(X)$ of C , with $\deg(f) \leq d$. At the end of the protocol, the verifier outputs 1 indicating accepting the proof that $f(x) = v$, or outputs 0 indicating rejecting the proof.

A polynomial commitment scheme must satisfy completeness, binding and extractability.

Definition 8 (Completeness). For all polynomials $f(X) \in \mathbb{F}[X]$ of degree $d \leq D$, for all $x \in \mathbb{F}$,

$$\Pr \left(\begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda, D) \\ b = 1 : (C, \tilde{c}) \leftarrow \text{Commit}(\text{pp}, f(X), d) \\ \quad \quad \quad v \leftarrow f(x) \\ b \leftarrow \text{eval}(\text{pp}, C, d, x, v; f(X)) \end{array} \right) = 1.$$

Definition 9 (Binding). A polynomial commitment scheme PC is binding if for all PPT \mathcal{A} , the following probability is negligible in λ :

$$\Pr \left(\begin{array}{l} \text{open}(\text{pp}, f_0, d, C, \tilde{c}_0) = 1 \wedge \\ \text{open}(\text{pp}, f_1, d, C, \tilde{c}_1) = 1 \wedge \\ f_0 \neq f_1 \end{array} : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda, D) \\ (C, f_0, f_1, \tilde{c}_0, \tilde{c}_1, d) \leftarrow \mathcal{A}(\text{pp}) \end{array} \right).$$

Definition 10 (Knowledge Soundness). For any PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a PPT algorithm \mathcal{E} such that the following probability is negligible in λ :

$$\Pr \left(\begin{array}{l} b = 1 \wedge \\ \mathcal{R}_{\text{eval}}(\text{pp}, C, x, v; \tilde{f}, \tilde{c}) = 0 \end{array} : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda, D) \\ (C, d, x, v, \text{st}) \leftarrow \mathcal{A}_1(\text{pp}) \\ (\tilde{f}, \tilde{c}) \leftarrow \mathcal{E}^{\mathcal{A}_2}(\text{pp}) \\ b \leftarrow \langle \mathcal{A}_2(\text{st}), V_{\text{eval}} \rangle(\text{pp}, C, d, x, v) \end{array} \right).$$

where the relation $\mathcal{R}_{\text{eval}}$ is defined as follows:

$$\mathcal{R}_{\text{eval}} = \{((\text{pp}, C \in \mathbb{G}, x \in \mathbb{F}, v \in \mathbb{F}); (f(X), \tilde{c})) : (\text{open}(\text{pp}, f, d, C, \tilde{c}) = 1) \wedge v = f(x)\}$$

We denote by $\text{Prove}, \text{Verify}$, the non-interactive prover and verifier algorithms obtained by applying FS to the eval public-coin interactive protocol, giving a non-interactive PCS scheme $(\text{Setup}, \text{Commit}, \text{Prove}, \text{Verify})$.

Definition 11 (Succinctness). We require the commitments and the evaluation proofs to be of size independent of the degree of the polynomial, that is the scheme is proof succinct if $|C|$ is $\text{poly}(\lambda)$, $|\pi|$ is $\text{poly}(\lambda)$ where π is the transcript obtained by applying FS to eval . Additionally, the scheme is verifier succinct if eval runs in time $\text{poly}(\lambda) \cdot \log(d)$ for the verifier.

A.2 Succinct Argument of Knowledge

Let \mathcal{R} be a NP-relation and \mathbb{L} be the corresponding NP-language, where $\mathbb{L} = \{x : \exists w \text{ such that } (x, w) \in \mathcal{R}\}$. Here, a prover \mathcal{P} aims to convince a verifier \mathcal{V} that $x \in \mathbb{L}$ by proving that it knows a witness w for a public statement x such that $(x, w) \in \mathcal{R}$. An interactive argument of knowledge for a relation \mathcal{R} consists of a PPT algorithm Setup , that takes an input the security parameter λ , and outputs the public parameter pp , and a pair of interactive PPT algorithms $\langle \mathcal{P}, \mathcal{V} \rangle$, where \mathcal{P} takes as input (pp, x, w) and \mathcal{V} takes as input (pp, x) . An interactive argument of knowledge $\langle \mathcal{P}, \mathcal{V} \rangle$, must satisfy completeness and knowledge soundness.

Definition 12 (Completeness). For all security parameter $\lambda \in \mathbb{N}$ and statement x and witness w such that $(x, w) \in \mathcal{R}$, we have

$$\Pr \left(b = 1 : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda) \\ b \leftarrow \langle \mathcal{P}(w), \mathcal{V} \rangle(\text{pp}, x) \end{array} \right) = 1.$$

Definition 13 (Knowledge Soundness). For any PPT malicious prover $\mathcal{P}^* = (\mathcal{P}_1^*, \mathcal{P}_2^*)$, there exists a PPT algorithm \mathcal{E} such that the following probability is negligible:

$$\Pr \left(\begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda) \\ b = 1 \wedge (x, \text{st}) \leftarrow \mathcal{P}_1^*(1^\lambda, \text{pp}) \\ (x, w) \notin \mathcal{R} : b \leftarrow \langle \mathcal{P}_2^*(\text{st}), \mathcal{V} \rangle(\text{pp}, x) \\ w \leftarrow \mathcal{E}^{\mathcal{P}_2^*}(\text{pp}, x) \end{array} \right).$$

A succinct argument of knowledge $\langle \mathcal{P}, \mathcal{V} \rangle$ for a relation \mathcal{R} , must satisfy completeness and knowledge soundness and additionally ensure that the communication complexity between prover and verifier, as well as the verification complexity is bounded by $\text{poly}(\lambda, \log |w|)$, where w is the witness for the relation.

B Argument for RAM From Polynomial Protocols

In this section, we give a self-contained argument of knowledge for membership in the language $\text{LRAM}_{\mathcal{I}, m, m}$ introduced in Section 5. We first consider the polynomial encoding of different RAM artefacts.

B.1 Polynomial Encoding

Let $k = 3m$ and let ω be a primitive k^{th} root of unity in \mathbb{F} . Let $\nu = \omega^3$, and thus ν is a primitive m^{th} root of unity in \mathbb{F} (We assume, these roots exist in \mathbb{F}). We recall \mathbb{V} as the subgroup consisting of m^{th} roots of unity with associated Lagrange basis polynomials $\{\tau_i(X)\}_{i \in [m]}$, while we additionally introduce the set \mathbb{K} of k^{th} roots of unity with $\{\lambda_i(X)\}_{i \in [k]}$ as the associated Lagrange polynomials.

$$\mathbb{K} = \{\omega, \dots, \omega^k\}, \quad \mathbb{V} = \{\nu, \dots, \nu^m\} \quad (11)$$

As before, we define the encoding of vectors in $\mathbf{f} \in \mathbb{F}^k$ as $\text{Enc}_{\mathbb{K}}(\mathbf{f}) = \sum_{i \in [k]} f_i \lambda_i(X)$. We canonically extend the encoding of vectors to encode RAM, operations and transcripts by encoding their component vectors. Thus, for a RAM $\mathbf{T} = (\mathbf{a}, \mathbf{v}) \in \text{RAM}_{\mathcal{I}, m}$, we define its encoding $\tilde{T} = (a(X), v(X))$ where $a(X), v(X) \in \mathbb{F}_{< m}[X]$ encode vectors \mathbf{a}, \mathbf{v} respectively. Given an operation sequence $\mathbf{o} = (o_1, \dots, o_m)$ with $o_i = (\bar{o}p_i, \bar{a}_i, \bar{v}_i)$ we encode \mathbf{o} as $\tilde{O} = (\bar{o}p(X), \bar{a}(X), \bar{v}(X))$ where $\bar{o}p(X)$ encodes the vector $\mathbf{op} = (\bar{o}p_1, \dots, \bar{o}p_m)$, $\bar{a}(X)$ encodes the vector $(\bar{a}_1, \dots, \bar{a}_m)$ and $\bar{v}(X)$ encodes the vector $(\bar{v}_1, \dots, \bar{v}_m)$. Finally, a transcript $\text{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ for tuples $(\mathbf{T}, \mathbf{o}, \mathbf{T}')$ where \mathbf{T}, \mathbf{T}' are RAMs of size m , and \mathbf{o} is an operation sequence of size m is encoded as $\tilde{\text{tr}} = (t(X), op(X), A(X), V(X))$ where the polynomials $t(X), op(X), V(X)$ and $A(X)$ encode the respective vectors in \mathbb{F}^k (See Section 5).

B.2 Relations over Polynomial Encodings

In this section, we describe polynomial checks for two important relations we need in subsequent sections, viz, (i) checking concatenation of vectors and (ii) checking monotonicity and load-store consistency of a transcript. The lemma below specifies the polynomial identities for verifying that vector $\mathbf{v} \in \mathbb{F}^k$ is concatenation of vectors $\mathbf{a}, \mathbf{b}, \mathbf{c}$ in \mathbb{F}^m .

Lemma 6. Let $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{F}^m$ and $v \in \mathbb{F}^k$ be vectors encoded by polynomials $a(X), b(X), c(X)$ and $v(X)$ respectively. Then,

$$a(X^3) - v(X) = 0 \quad \text{mod } Z(X) \quad (\text{A1})$$

$$b(X^3) - v(\omega^m X) = 0 \quad \text{mod } Z(X) \quad (\text{A2})$$

$$c(X^3) - v(\omega^{2m} X) = 0 \quad \text{mod } Z(X) \quad (\text{A3})$$

for $Z(X) = \prod_{i=1}^m (X - \omega^i)$ if and only if $\mathbf{v} = \mathbf{a} \parallel \mathbf{b} \parallel \mathbf{c}$.

Proof. Assume that the polynomial identities hold. Substituting $X = \omega^i$ for $i \in [m]$ in above equations implies for $i \in [m]$: $a_i = v_i$ (Eq (A1)), $b_i = v_{m+i}$ (Eq (A2)) and $c_i = v_{2m+i}$ (Eq (A3)), which together imply $\mathbf{v} = \mathbf{a} \parallel \mathbf{b} \parallel \mathbf{c}$. Converse follows by observing that $\mathbf{v} = \mathbf{a} \parallel \mathbf{b} \parallel \mathbf{c}$ implies that $v(X) = a(X^3)$, $v(\omega^m X) = b(X^3)$ and $v(\omega^{2m} X) = c(X^3)$ holds for all $X = \omega^i, i \in [m]$. Thus, the equalities hold modulo the polynomial $Z(X)$ as defined above.

Next, we specify polynomial checks on the encoding of a transcript to ensure it satisfies address-ordering and load-store consistency. Let N be an upper bound on the values of \mathbf{A} , i.e, the index set $I \subseteq [N]$. Let $\mathbf{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ be a transcript encoded as $\tilde{\mathbf{tr}} = (t(X), op(X), A(X), V(X))$. Recall that we need to check two conditions on \mathbf{tr} , viz, (i) *monotonicity*: the transcript is sorted by address and timestamp respectively, i.e, $A_i \leq A_{i+1}$ for all $i < k$ and $t_i < t_{i+1}$ whenever $A_i = A_{i+1}$, (ii) *load-store consistency*: whenever $op_{i+1} = 0$ and $A_i = A_{i+1}$, we have $V_i = V_{i+1}$. To do so, we exhibit disjoint sets I_1, I_2 with $I_1 \uplus I_2 = [k-1]$ such that: (i) for all $i \in I_1$, $A_i < A_{i+1}$, (ii) for all $i \in I_2$, $(A_i = A_{i+1}) \wedge (t_i < t_{i+1})$ and (iii) for all $i \in I_2$, $(op_i = 1) \vee (V_i = V_{i+1})$. Note that the conditions on the sets I_1 and I_2 ensures monotonicity. Moreover, it can be seen that load-store consistency requirements are satisfied for all $i \in I_1$ (as $A_i \neq A_{i+1}$). Similarly, load-store consistency also holds for all $i \in I_2$. It remains to exhibit the sets and show that they satisfy the above invariants using polynomials, as in the following lemma:

Lemma 7. *Let $\tilde{\mathbf{tr}}$ be a polynomial encoding of transcript \mathbf{tr} of size k , given by polynomials $t(X)$, $op(X)$, $A(X)$ and $V(X)$, with index set $[N]$. Then assuming $kN < |\mathbb{F}|$, \mathbf{tr} is address ordered and satisfies load-store consistency if and only if there exist polynomials $Z_1, Z_2, \delta_T, \delta_A$ such that the following hold:*

$$A(\omega X) - A(X) - \delta_A(X) = 0 \text{ mod } Z_1(X) \quad (\text{C1})$$

$$A(\omega X) - A(X) = 0 \text{ mod } Z_2(X) \quad (\text{C2})$$

$$t(\omega X) - t(X) - \delta_T(X) = 0 \text{ mod } Z_2(X) \quad (\text{C3})$$

$$(op(X) - 1)(V(\omega X) - V(X)) = 0 \text{ mod } Z_2(X) \quad (\text{C4})$$

$$Z_1(X) \cdot Z_2(X) \cdot (X - 1) = Z_{\mathbb{K}}(X) \quad (\text{C5})$$

$$1 \leq A(\omega^i) \leq N \quad (\text{C6})$$

$$1 \leq t(\omega^i) \leq N, 1 \leq \delta_A(\omega^i) \leq N, 1 \leq \delta_T(\omega^i) \leq N \text{ for } i \in [k] \quad (\text{C7})$$

Proof. Suppose there exist polynomials $Z_1(X), Z_2(X), \delta_T(X)$ and $\delta_A(X)$ satisfying above identities. From Equation (C5), we conclude that their exist sets I_1, I_2 with $I_1 \uplus I_2 = [k-1]$ such that $Z_b(X)$, $b \in \{1, 2\}$ is the vanishing polynomial of the set $\{\omega^i : i \in I_b\}$. We now note that the following are true for $i \in I_1$:

– $A(\omega^{i+1}) - A(\omega^i) = \delta_A(\omega^i)$. Since $1 \leq \delta_A(\omega^i) \leq N$, this ensures $A_i < A_{i+1}$ for the vector \mathbf{A} encoded by $A(X)$. We note that $kN < |\mathbb{F}|$ implies there is no overflow modulo the field characteristic.

Similarly, it can be seen that for $i \in I_2$, we must have (i) $A_i = A_{i+1} \wedge t_i < t_{i+1}$ and (ii) $op_i = 1 \vee V_i = V_{i+1}$. Together these imply that the encoded transcript is address-ordered.

Protocols facilitating the checks mentioned in Lemma 6 and Lemma 7 are presented in Figure 7 and 8 respectively.

C Succinct Argument for Verifiable RAM

The polynomial encodings in the previous section can be used to polynomial protocol for checking the membership in the language $\text{LRAM}_{\mathcal{L}, m, m}$ for $m \in \mathbb{N}$. The polynomial protocol can be subsequently be compiled into a succinct argument using an extractable polynomial commitment scheme. In this section, we use KZG polynomial commitment scheme to obtain a succinct argument for checking membership in $\text{LRAM}_{\mathcal{L}, m, m}$ in the Algebraic Group Model (AGM). At a high level, to prove $(\mathbf{T}, \mathbf{o}, \mathbf{T}') \in \text{LRAM}_{\mathcal{L}, m, m}$, the prover constructs time ordered transcript \mathbf{tr} and then permutes it to obtain the address sorted transcript \mathbf{tr}^* . It then sends the polynomial encodings of $\mathbf{T}, \mathbf{o}, \mathbf{T}', \mathbf{tr}$ and \mathbf{tr}^* to the verifier, who verifies that:

1. The time ordered transcript is correctly constructed, i.e, $\mathbf{tr} = \text{TimeTr}(\mathbf{T}, \mathbf{o}, \mathbf{T}')$. This is achieved using the protocol in Figure 9.
2. The transcript \mathbf{tr}^* is a permutation of the transcript \mathbf{tr} , i.e, $\mathbf{tr}^* = \sigma(\mathbf{tr})$ for some permutation σ of $[k]$. The protocol for this check appears in Figure 10.
3. The transcript \mathbf{tr}^* is address ordered and satisfies load-store consistency. We describe the protocol to check this property of transcripts in Figure 8.

Common Input: Commitments c_a, c_b, c_c, c_v , and $[Z]_1$ (to the polynomial $Z(X) = \prod_{i=1}^m (X - \omega^i)$).
Prover's Input: Vectors $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{F}^m$ and $\mathbf{v} \in \mathbb{F}^k$.

1. \mathcal{V} sends $\gamma \leftarrow \mathbb{F}$.
2. \mathcal{P} computes the following:
 - $h(X) = a(X) + \gamma b(X) + \gamma^2 c(X)$.
 - $Q(X) = (h(X^3) - v(X) - \gamma v(\omega^m X) - \gamma^2 v(\omega^{2m} X))/Z(X)$.
3. \mathcal{P} sends commitment $[Q]_1 = [Q(X)]_1$.
4. \mathcal{V} sends $s \leftarrow \mathbb{F}$.
5. \mathcal{P} sends evaluations $\langle v \rangle_s = v(s)$, $\langle v \rangle_{\omega^m s} = v(\omega^m s)$, $\langle v \rangle_{\omega^{2m} s} = v(\omega^{2m} s)$, $\langle h \rangle_{s^3} = h(s^3)$, $\langle Q \rangle_s = Q(s)$ and $\langle Z \rangle_s = Z(s)$.
6. \mathcal{V} sends $r \leftarrow \mathbb{F}$.
7. \mathcal{P} computes KZG proofs:
 - $\Pi_v = \text{KZG.Prove}(\text{srs}, v, (s, \omega^m s, \omega^{2m} s))$.
 - $\Pi_h = \text{KZG.Prove}(\text{srs}, h, s^3)$.
 - $\Pi_f = \text{KZG.Prove}(\text{srs}, f, s)$ where $f(X) = Z(X) + rQ(X)$.
8. \mathcal{P} sends Π_v, Π_h and Π_f .
9. \mathcal{V} computes commitments $[h]_1$ and $[f]_1$.^a
10. \mathcal{V} checks:
 - $\text{KZG.Verify}(\text{srs}, [v]_1, \mathbf{e}_v, \mathbf{p}_v, \Pi_v)$ where $\mathbf{p}_v = (s, \omega^m s, \omega^{2m} s)$ and $\mathbf{e}_v = (\langle v \rangle_s, \langle v \rangle_{\omega^m s}, \langle v \rangle_{\omega^{2m} s})$.
 - $\text{KZG.Verify}(\text{srs}, [h]_1, \langle h \rangle_{s^3}, s^3, \Pi_h)$.
 - $\text{KZG.Verify}(\text{srs}, [f]_1, \langle Z \rangle_s + r \langle Q \rangle_s, s, \Pi_f)$.
 - $\langle Q \rangle_s \cdot \langle Z \rangle_s = \langle h \rangle_{s^3} - \langle v \rangle_s - \gamma \langle v \rangle_{\omega^m s} - \gamma^2 \langle v \rangle_{\omega^{2m} s}$.
11. \mathcal{V} outputs accept if all the above checks succeed, else it outputs reject.

^a This can be done locally by leveraging the linearity of the operation

Fig. 7: Check concatenation over committed vectors.

We check above conditions over commitments. Let srs denote a KZG setup over a bi-linear group, with prime order groups $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T . We canonically commit to RAM, operation sequences and transcripts by committing to their polynomial encodings. Commitment of an encoding represented as tuple of polynomials is simply the tuple consisting of commitments of the component polynomials. We now define the relation $\mathcal{R}_{\text{srs}, m}^{\text{LRAM}}$ below, and present a succinct argument for the same.

Definition 14. Let $\mathcal{R}_{\text{srs}, m}^{\text{LRAM}}$ consist of tuples $((c_T, c_o, c'_T), (\mathbf{T}, \mathbf{o}, \mathbf{T}'))$ where $c_T = \text{KZG.Commit}(\text{srs}, \tilde{T})$, $c'_T = \text{KZG.Commit}(\text{srs}, \tilde{T}')$, $c_o = \text{KZG.Commit}(\text{srs}, \tilde{O})$ commit to \mathbf{T}, \mathbf{T}' and \mathbf{o} with $(\mathbf{T}, \mathbf{o}, \mathbf{T}') \in \text{LRAM}_{\mathcal{I}, m, m}$.

In the above definition we have $c_T = (c_a, c_v)$ where c_a and c_v are KZG commitments to polynomials $a(X)$ and $v(X)$ in the encoding $\tilde{T} = (a(X), v(X))$. Similarly we parse $c'_T = (c'_a, c'_v)$ and $c_o = (\tilde{c}_{op}, \tilde{c}_a, \tilde{c}_v)$ (see Appendix B.1 for polynomial encodings). For proving relation 14, prover's input consists of initial RAM state $\mathbf{T} = (\mathbf{a}, \mathbf{v})$, final RAM state $\mathbf{T}' = (\mathbf{a}', \mathbf{v}')$, operation sequence $\mathbf{o} = (o_1, \dots, o_m)$ with $o_i = (\tilde{op}_i, \tilde{a}_i, \tilde{v}_i)$, time-ordered transcript $\text{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ and address-ordered transcript $\text{tr}^* = (\mathbf{t}^*, \mathbf{op}^*, \mathbf{A}^*, \mathbf{V}^*)$ obtained from tr using a permutation $\sigma : [k] \rightarrow [k]$. Verifier's input consists of the commitments c_T, c_o and c'_T as described above.

The prover starts the protocol by sending commitments c_{tr} and c_{tr}^* to the transcripts tr and tr^* respectively. To show that tr is correctly formed, the prover needs to prove the concatenations: (i) $\mathbf{op} = 0^m || (\tilde{op}_1, \dots, \tilde{op}_m) || 0^m$, (ii) $\mathbf{A} = \mathbf{a} || (\tilde{a}_1, \dots, \tilde{a}_m) || \mathbf{a}'$ and (iii) $\mathbf{V} = \mathbf{v} || (\tilde{v}_1, \dots, \tilde{v}_m) || \mathbf{v}'$. Note that the time-stamp column \mathbf{t} is implicitly assumed to be $(1, \dots, k)$. The verifier checks the concatenations using Lemma 6. It uses a random challenge β to reduce the three concatenations to one concatenation, and uses another challenge γ to reduce the three polynomial checks in Lemma 6 to a single check. The complete polynomial protocol is detailed in Figure 9.

Next, we show a polynomial protocol for proving that the transcript tr^* is a permutation of the transcript tr . We first recall the permutation argument for vectors from [GWC19a].

Lemma 8 (Permutation Check [GWC19a]). Let $f(X), g(X)$ be polynomials in $\mathbb{F}[X]$. Then, the vectors $\mathbf{f}, \mathbf{g} \in \mathbb{F}^k$ encoded by the polynomials are permutations of each other if and only if with overwhelming probability over the choice of $\alpha \leftarrow \mathbb{F}$, there exists a polynomial $z(X)$ satisfying the polynomial constraints:

- Common Input:** Commitments c_t, c_{op}, c_A and c_V to $\mathbf{t}, \mathbf{op}, \mathbf{A}$ and \mathbf{V} constituting the transcript \mathbf{tr} .
Prover's Input: $\mathbf{tr} = (\mathbf{t}, \mathbf{op}, \mathbf{A}, \mathbf{V})$ and its polynomial encoding $\tilde{\mathbf{tr}} = (t(X), op(X), A(X), V(X))$.
- 1 Prover determines sets I_1, I_2 as described in Appendix A.2.
 - 2 Prover computes polynomials $Z_1(X), Z_2(X), \delta_T(X), \delta_A(X)$.
 - 3 \mathcal{P} sends $[Z_1(X)]_1, [Z_2(X)]_1, [\delta_T(X)]_1, [\delta_A(X)]_1$.
 - 4 \mathcal{V} sends $\gamma \leftarrow \mathbb{F}$.
 - 5 \mathcal{P} computes the following:
 - $Q_1(X) = (A(\omega X) - A(X) - \delta_A(X))/Z_1(X)$.
 - $Q_2(X) = [(A(\omega X) - A(X)) + \gamma(t(\omega X) - t(X) - \delta_T(X)) + \gamma^2(op(X) - 1)(V(\omega X) - V(X))]/Z_2(X)$
 - 6 \mathcal{P} sends commitments $[Q_1(X)]_1, [Q_2(X)]_1$.
 - 7 \mathcal{V} sends $s \leftarrow \mathbb{F}$.
 - 8 \mathcal{P} sends evaluations $\langle A \rangle_s = A(s), \langle A \rangle_{\omega s} = A(\omega s), \langle \delta_A \rangle_s = \delta_A(s), \langle t \rangle_s = t(s), \langle t \rangle_{\omega s} = t(\omega s), \langle \delta_T \rangle_s = \delta_T(s), \langle op \rangle_s = op(s), \langle V \rangle_s = V(s), \langle V \rangle_{\omega s} = V(\omega s), \langle Q_1 \rangle_s = Q_1(s), \langle Q_2 \rangle_s = Q_2(s), \langle Z_1 \rangle_s = Z_1(s), \langle Z_2 \rangle_s = Z_2(s)$.
 - 9 \mathcal{V} checks:
 - $\langle Q_1 \rangle_s \cdot \langle Z_1 \rangle_s = (\langle A \rangle_{\omega s} - \langle A \rangle_s - \langle \delta_A \rangle_s)$.
 - $\langle Q_2 \rangle_s \cdot \langle Z_2 \rangle_s = (\langle A \rangle_{\omega s} - \langle A \rangle_s) + \gamma(\langle t \rangle_{\omega s} - \langle t \rangle_s - \langle \delta_T \rangle_s) + \gamma^2(\langle op \rangle_s - 1)(\langle V \rangle_{\omega s} - \langle V \rangle_s)$.
 - $\langle Z_1 \rangle_s \cdot \langle Z_2 \rangle_s = s^{3m} - 1$.
 - 10 \mathcal{V} sends $r_1, r_2 \leftarrow \mathbb{F}$.
 - 11 \mathcal{P} computes the following:
 - $\Phi_{\omega s}(X) = A(X) + r_1 t(X) + r_2^2 V(X)$.
 - $\Phi_s(X) = A(X) + r_2 \delta_A(X) + r_2^2 t(X) + r_2^3 \delta_T(X) + r_2^4 op(X) + r_2^5 V(X) + r_2^6 Q_1(X) + r_2^7 Q_2(X) + r_2^8 Z_1(X) + r_2^9 Z_2(X)$.
 - $\Pi_{\omega s} = \text{KZG.Prove}(\text{srs}, \Phi_{\omega s}(X), \omega s)$.
 - $\Pi_s = \text{KZG.Prove}(\text{srs}, \Phi_s(X), s)$.
 - 12 \mathcal{P} sends $\Pi_{\omega s}, \Pi_s$.
 - 13 \mathcal{V} computes:
 - $[\Phi_{\omega s}(X)]_1 = c_A + r_1 c_t + r_2^2 c_V$.
 - $[\Phi_s(X)]_1 = c_A + r_2 [\delta_A(X)]_1 + r_2^2 c_t + r_2^3 [\delta_T(X)]_1 + r_2^4 c_{op} + r_2^5 c_V + r_2^6 [Q_1(X)]_1 + r_2^7 [Q_2(X)]_1 + r_2^8 [Z_1(X)]_1 + r_2^9 [Z_2(X)]_1$.
 - $V_{\omega s} = \langle A \rangle_{\omega s} + r_1 \langle t \rangle_{\omega s} + r_2^2 \langle V \rangle_{\omega s}$.
 - $V_s = \langle A \rangle_s + r_2 \langle \delta_A \rangle_s + r_2^2 \langle t \rangle_s + r_2^3 \langle \delta_T \rangle_s + r_2^4 \langle op \rangle_s + r_2^5 \langle V \rangle_s + r_2^6 \langle Q_1 \rangle_s + r_2^7 \langle Q_2 \rangle_s + r_2^8 \langle Z_1 \rangle_s + r_2^9 \langle Z_2 \rangle_s$.
 - 14 \mathcal{V} checks:
 - $\text{KZG.Verify}(\text{srs}, [\Phi_{\omega s}]_1, V_{\omega s}, \omega s, \Pi_{\omega s})$.
 - $\text{KZG.Verify}(\text{srs}, [\Phi_s]_1, V_s, s, \Pi_s)$.
 - 15 \mathcal{P} and \mathcal{V} invoke sub-vector arguments $(\mathcal{P}_{\text{sv}}, \mathcal{V}_{\text{sv}})$ (eg. [EFG22]) to prove that $(\text{srs}, c_A, c_t), (\text{srs}, c_t, c_I), (\text{srs}, [\delta_A(X)]_1, c_I)$ and $(\text{srs}, [\delta_T(X)]_1, c_I)$ are in $\mathcal{R}_{\text{srs}, N, m}^{\text{subvec}}$.
 - 16 \mathcal{V} outputs accept if all checks succeed and the sub-vector arguments outputs accept. Otherwise it outputs reject.

Fig. 8: Check that transcript is address ordered and load-store consistent.

$$\lambda_1(X)(z(X) - 1) = 0 \text{ mod } Z_{\mathbb{K}}(X) \quad (\text{B1})$$

$$(\alpha - g(X))z(\omega X) = (\alpha - f(X))z(X) \text{ mod } Z_{\mathbb{K}}(X) \quad (\text{B2})$$

The polynomial protocol in Figure 10 essentially invokes the above argument on the random linear combination of the columns of the respective transcripts.

Finally, we see that Lemma 7 implies a polynomial protocol to check that the transcript \mathbf{tr}^* is address ordered and satisfies load-store consistency, which essentially involves the prover identifying sets I_1, I_2 as described in Appendix B.2 and sending auxiliary polynomials $Z_1(X), Z_2(X), \delta_A^*(X)$ and $\delta_T^*(X)$ to the verifier. The verifier then checks the identities (C1)-(C6) in Lemma 7. The range checks in (C7) can be checked using polynomial protocols in sub-vector lookup arguments such as [PK22, EFG22, CFF⁺24, ZSG24]. The protocol (compiled using KZG commitments in AGM) can be found in Figure 8. The overall protocol for $\mathcal{R}_{\text{srs}, m}^{\text{LRAM}}$ which combines invokes protocols in Figures 9, 10 and 8 as sub-protocols is presented in Figure 11.

Efficiency. We provide a break-up of costs incurred by different components involved in construction of RAM based on memory-checking techniques in Table 6. To reduce pairing checks we use standard technique of batching pairing checks involving common generators. In addition, to reduce communi-

Component	Protocol	Prover Work	Verifier Work	Communication
Concatenation of transcripts	Figure 9	$O(m \log m) \mathbb{F}, O(m) \mathbb{G}_1$	$2P$	$4\mathbb{G}_1, 6\mathbb{F}$
Permutation of transcripts	Figure 10	$O(m \log m) \mathbb{F}, O(m) \mathbb{G}_1$	$2P$	$4\mathbb{G}_1, 5\mathbb{F}$
Memory consistency & address ordering of transcript	Figure 8	$O(m \log m) \mathbb{F}, O(m) \mathbb{G}_1$	$6P$	$20\mathbb{G}_1, 19\mathbb{F}$
Polynomial Protocol for RAM	Figure 11	$O(m \log m) \mathbb{F}, O(m) \mathbb{G}_1$	$7P$	$36\mathbb{G}_1, 30\mathbb{F}$

Table 6: Efficiency parameters for components of polynomial protocol for RAM. Here m denotes both the size of the RAM and number of operations (the special case we consider). P denotes a pairing evaluation, while \mathbb{G}_1 , \mathbb{G}_2 and \mathbb{F} denote the groups and the scalar field of the bilinear group used for instantiating the protocol.

<p>Common Input: Commitments $c_T = (c_a, c_v)$, $c_o = (\bar{c}_{op}, \bar{c}_a, \bar{c}_v)$, $c'_T = (c'_a, c'_v)$ and $c_{tr} = (c_t, c_{op}, c_A, c_V)$ to $\mathbf{T}, \mathbf{o}, \mathbf{T}'$ and tr (which is supposed to be the time ordered transcript) respectively. Commitment $[Z(X)]_1$ to the polynomial $Z(X) = \prod_{i=1}^m (X - \omega^i)$.</p> <p>Prover's Input: $\text{tr}, \mathbf{T}, \mathbf{T}', \mathbf{o}$ and their polynomial encodings, $Z(X)$.</p> <ol style="list-style-type: none"> \mathcal{V} sends $\beta, \gamma \leftarrow \mathbb{F}$. \mathcal{P} computes the following: <ul style="list-style-type: none"> $G_1(X) = a(X) + \beta v(X)$, $G_2(X) = \bar{a}(X) + \beta \bar{v}(X) + \beta^2 \bar{op}(X)$ $G_3(X) = a^*(X) + \beta v^*(X)$, $G(X) = A(X) + \beta V(X) + \beta^2 op(X)$ $H(X) = G_1(X) + \gamma G_2(X) + \gamma^2 G_3(X)$ $Q(X) = [(H(X^3) - G(X) - \gamma G(\omega^m X) - \gamma^2 G(\omega^{2m} X))] / Z(X)$ \mathcal{P} sends commitment $[Q]_1$ to $Q(X)$. \mathcal{V} sends $s \leftarrow \mathbb{F}$. \mathcal{P} sends evaluations $\langle G \rangle_s = G(s)$, $\langle G \rangle_{\omega^m s} = G(\omega^m s)$, $\langle G \rangle_{\omega^{2m} s} = G(\omega^{2m} s)$, $\langle H \rangle_{s^3} = H(s^3)$, $\langle Q \rangle_s = Q(s)$ and $\langle Z \rangle_s = Z(s)$. \mathcal{V} sends $r \leftarrow \mathbb{F}$. \mathcal{P} sends the following KZG proofs: <ul style="list-style-type: none"> $\Pi_G = \text{KZG.Prove}(\text{srs}, G(X), (s, \omega^m s, \omega^{2m} s))$. $\Pi_H = \text{KZG.Prove}(\text{srs}, H(X), s^3)$. $\Pi_F = \text{KZG.Prove}(\text{srs}, F(X), s)$ where $F(X) = Z(X) + rQ(X)$. \mathcal{V} computes $[G(X)]_1, [H(X)]_1$ and $[F(X)]_1$.^a \mathcal{V} checks: <ul style="list-style-type: none"> $\text{KZG.Verify}(\text{srs}, [G]_1, (\langle G \rangle_s, \langle G \rangle_{\omega^m s}, \langle G \rangle_{\omega^{2m} s}), (s, \omega^m s, \omega^{2m} s), \Pi_G)$. $\text{KZG.Verify}(\text{srs}, [H]_1, \langle H \rangle_{s^3}, s^3, \Pi_H)$. $\text{KZG.Verify}(\text{srs}, [F]_1, \langle Z \rangle_s + r \langle Q \rangle_s, s, \Pi_F)$. $\langle Q \rangle_s \cdot \langle Z \rangle_s = \langle H \rangle_{s^3} - \langle G \rangle_s - \gamma \langle G \rangle_{\omega^m s} - \gamma^2 \langle G \rangle_{\omega^{2m} s}$. \mathcal{V} outputs accept if all the above checks succeeds, otherwise it outputs reject. <p>^a This can be done locally by leveraging the linearity of the operation</p>
--

Fig. 9: Check the correctness of time-ordered transcript.

cation, instead of naively invoking four instances of sub-vector argument in Step 15 of the protocol in Figure 8, we concatenate the four vectors using a variant of protocol for concatenation of vectors in Figure 7, and then use the sub-vector argument to show that the concatenated vector is a sub-vector of the vector $(1, \dots, N)$. For CQ [EFG22] based instantiation, this reduces the total communication of this check from $4 \times (8\mathbb{G}_1 + 3\mathbb{F})$ to $(4\mathbb{G}_1 + 6\mathbb{F}) + (8\mathbb{G}_1 + 3\mathbb{F})$, a saving of $\approx 20\mathbb{G}_1$. The reported overheads in Table 6 take into account such optimizations.

D Proof of Lemma 5

Before starting the proof, we collect some preliminaries which will be useful in the proof.

Common Input: Commitments $c_{\text{tr}} = (c_t, c_{op}, c_A, c_V)$ and $c_{\text{tr}}^* = (c_t^*, c_{op}^*, c_A^*, c_V^*)$ of transcripts tr and tr^* respectively.

Prover's Input: Transcripts tr, tr^* and their polynomial encodings, permutation σ such that $\text{tr}^* = \sigma(\text{tr})$.

- 1 \mathcal{V} sends $\alpha, \beta, \chi \leftarrow \mathbb{F}$.
- 2 \mathcal{P} computes the following:
 - $f(X) = t(X) + \beta op(X) + \beta^2 A(X) + \beta^3 V(X)$.
 - $g(X) = t^*(X) + \beta op^*(X) + \beta^2 A^*(X) + \beta^3 V^*(X)$.
- 3 \mathcal{P} then computes polynomials $z(X), q(X)$ as:
 - Interpolate polynomial $z(X)$ of degree $k-1$ such that $z(\omega) = 1$ and $z(\omega^{i+1}) = \prod_{j=1}^i (\alpha - f(\omega^j)) / (\alpha - g(\omega^j))$ for $1 \leq i \leq k-1$.
 - $q(X) = ((\alpha - g(X))z(\omega X) - (\alpha - f(X))z(X) + \chi \lambda_1(X)(z(X) - 1)) / \mathbb{Z}_{\mathbb{K}}(X)$.
- 4 \mathcal{P} sends commitments $[z(X)]_1$ and $[q(X)]_1$ to polynomials $z(X)$ and $q(X)$ respectively.
- 5 \mathcal{V} computes commitments $[f]_1, [g]_1$.^a
- 6 \mathcal{V} checks that $q(X)Z_{\mathbb{K}}(X) = (\alpha - g(X))z(\omega X) - (\alpha - f(X))z(X) + \chi \lambda_1(X)(z(X) - 1)$ by requesting evaluations and KZG proofs of polynomials f, g, q, z at a random point, say s and evaluation and KZG proof of z at ωs .
- 7 \mathcal{V} outputs accept if all the checks succeed, else it outputs reject.

^a This can be done locally by leveraging the linearity of the operation

Fig. 10: Check that transcripts are permutations of each other.

Common Input: Commitments $c_T = (c_a, c_v)$, $c_o = (\bar{c}_{op}, \bar{c}_a, \bar{c}_v)$, $c'_T = (c'_a, c'_v)$.

Prover's Input: $\mathbf{T}, \mathbf{T}', \mathbf{o}$ and their polynomial encodings.

1. \mathcal{P} computes the following:
 - tr (time ordered transcript corresponding to $\mathbf{T}, \mathbf{o}, \mathbf{T}'$), its polynomial encoding, and its commitment $c_{\text{tr}} = (c_t, c_{op}, c_A, c_V)$.
 - $Z(X) = \prod_{i=1}^m (X - \omega^i)$ and its commitment $[Z(X)]_1$.
2. \mathcal{P} sends $c_{\text{tr}} = (c_t, c_{op}, c_A, c_V)$ and $[Z(X)]_1$.
3. \mathcal{P} and \mathcal{V} run the protocol for checking correctness of time ordered transcript (Figure 9).
4. \mathcal{P} computes the address ordered transcript tr^* (along with its polynomial encoding) and the permutation σ from the time ordered transcript tr , such that $\text{tr}^* = \sigma(\text{tr})$.
5. \mathcal{P} computes the commitment $c_{\text{tr}}^* = (c_t^*, c_{op}^*, c_A^*, c_V^*)$ of tr^* and sends c_{tr}^* .
6. \mathcal{P} and \mathcal{V} run the protocol for checking that the two transcripts are permutations of each other (Figure 10).
7. \mathcal{P} and \mathcal{V} run the protocol for checking the constraints given in Lemma 7 (Figure 8.)
8. \mathcal{V} outputs accept if all the three sub-protocols lead to accept, else it outputs reject.

Fig. 11: Overall protocol for the relation $\mathcal{R}_{\text{srs}, m}^{\text{LRAM}}$

D.1 Computational Algebra Preliminaries

Let \mathbb{F} be a finite field of prime order p and \mathbb{G} be a cyclic additive group of order p with generator g . For $s \in \mathbb{F}$, we use the notation $[s]$ to denote the group element $s \cdot g$. Assume that \mathbb{F} contains n^{th} root of unity ξ satisfying $\xi^n = 1$ for large n (All polynomial degrees are assumed less than n).

Fact D1 (Fast Evaluation) Let $f \in \mathbb{F}[X]$ be a polynomial of degree $< d$ and $(\xi_1, \dots, \xi_r) \in \mathbb{F}^r$ be distinct points in \mathbb{F} . Then the vector $(f(\xi_1), \dots, f(\xi_r))$ can be computed in $O((d+r) \log(d+r))$ \mathbb{F} operations if ξ_1, \dots, ξ_r form roots of unity, and in $O((d+r) \log^2(d+r))$ \mathbb{F} operations otherwise.

Fact D2 (Fast Interpolation) Let ξ_1, \dots, ξ_d be distinct points in \mathbb{F} and $(v_1, \dots, v_d) \in \mathbb{F}^d$. Then $(f_0, \dots, f_{d-1}) \in \mathbb{F}^d$ can be computed in $O(d \log^2 d)$ operations in \mathbb{F} such that $f(\xi_i) = v_i$ for all $i \in [d]$ where $f(X) = \sum_{i=0}^{d-1} f_i X^i$.

Fact D3 (Fast Multiplication) Let ξ_1, \dots, ξ_r be distinct points in \mathbb{F} . Then coefficients of $f(X) = \prod_{i=1}^r (X - \xi_i)$ can be computed in $O(r \log^2 r)$ operations in \mathbb{F} .

Fact D4 (Multi KZG proofs [FK23]) Let $\{[x^i]\}_{i=1}^d$ be given for some $x \in \mathbb{F}$. Then for set of r distinct points ξ_1, \dots, ξ_r , and a polynomial $f(X) \in \mathbb{F}[X]$ of degree $< d$, the vector $([h_1(x)], \dots, [h_r(x)])$, where $h_i(X) = (f(X) - f(\xi_i))/(X - \xi_i)$ can be computed in $O((r+d) \log(r+d))$ group and field operations when ξ_1, \dots, ξ_r are roots of unity, and in $O(r \log^2 r + d \log d)$ group and field operations otherwise.

Fact D5 (Lagrange Polynomials) Let $\mathbb{S} = \{\xi_1, \dots, \xi_r\}$ be a set of r distinct points and let $\tau_1(X), \dots, \tau_r(X)$ be the corresponding lagrange polynomials of degree $r - 1$ each. Let $Z_{\mathbb{S}}(X) = \prod_{i=1}^r (X - \xi_i)$ denote the vanishing polynomial for \mathbb{S} . Then we have:

$$\sum_{i=1}^r \tau_i(X) = 1$$

$$\tau_i(X) = \frac{Z_{\mathbb{S}}(X)}{Z'_{\mathbb{S}}(\xi_i)(X - \xi_i)} \text{ for all } i \in [r]$$

Formal Derivative: For a polynomial $p(X) \in \mathbb{F}[X]$, we define the formal derivative of $p(X)$ as the polynomial $u(X, X)$ where $u(X, Y) = \frac{p(X) - p(Y)}{X - Y}$. It can be seen that $u(X, X)$ equals the polynomial $p'(X)$ obtained by differentiating $p(X)$ according to regular rules of calculus. Thus, this definition agrees with the one given earlier in the preliminaries.

D.2 Some Useful Results

We state and prove some facts which are used later throughout the proof.

Lemma 9. For $K \subset [N]$, define H_K to be $\{\xi^i : i \in K\}$. Let $p(X)$ be the vanishing polynomial of H_K . Let $p'(X)$ and $p''(X)$ denote the formal first derivative and second derivative of $p(X)$ respectively. Then, $p''(\xi^i)/p'(\xi^i) = 2 \cdot \sum_{j \in K \setminus \{i\}} 1/(\xi^i - \xi^j)$ for all $i \in K$

Proof. Observe that $p'(X) = \sum_{i \in K} \prod_{j \in K \setminus \{i\}} (X - \xi^j)$ and $p''(X) = \sum_{i \in K} \sum_{j \in K \setminus \{i\}} \prod_{k \in K \setminus \{i, j\}} (X - \xi^k)$. Thus for $r \in K$, we have:

$$p'(\xi^r) = \prod_{j \in K \setminus \{r\}} (\xi^r - \xi^j),$$

$$p''(\xi^r) = \sum_{j \in K \setminus \{r\}} \prod_{k \in K \setminus \{r, j\}} (\xi^r - \xi^k) + \sum_{i \in K \setminus \{r\}} \prod_{k \in K \setminus \{r, i\}} (\xi^r - \xi^k)$$

Note that only non-zero products in the expansion of $p''(\xi^r)$ occur when $i = r$ or $j = r$, resulting in the two summands for the same in the above equation. Moreover, we notice that both summands are the same, giving us $p''(\xi^r) = 2 \sum_{i \in K \setminus \{r\}} \prod_{k \in \{r, i\}} (\xi^r - \xi^k)$. One may now verify that $p'(\xi^r)/p''(\xi^r)$ gives the desired result claimed in the lemma.

Lemma 10 (Sumcheck). Let $u(X, Y)$ be a bi-variate polynomial over a finite field \mathbb{F} with degree less than N in each of the variables and \mathbb{H} be defined as the group of N^{th} roots of unity ($N \ll |\mathbb{F}|$) with generator $\xi \in \mathbb{F}$. Then $\sum_{i \in [N]} u(X, \xi^i) = Nu(X, 0)$

Proof. For some $d < N$, we write $u(X, Y) = a_0 + a_1 Y + a_2 Y^2 + \dots + a_d Y^d$ where each a_i is a polynomial in X of degree less than N . Now we write the sum:

$$\sum_{i \in [N]} u(X, \xi^i) = Na_0 + a_1(\xi + \xi^2 + \dots + \xi^N) + a_2(\xi^2 + \xi^4 + \dots + \xi^{2N}) + \dots + a_d(\xi^d + \dots + \xi^{Nd})$$

But for any $\alpha = \xi^k$ for $k < N$, $\alpha + \alpha^2 + \dots + \alpha^N = 0$. Thus, all terms vanish except the first term and hence $\sum_{i \in [N]} u(X, \xi^i) = Na_0$. The lemma follows by observing that $a_0 = u(X, 0)$.

We use the following standard observation for our next lemma:

Fact D6 If polynomials f, g of degree $< N$ agree on N points, then they are equal as polynomials, that is, $f(X) = g(X)$

Lemma 11. Let $Z_{\mathbb{H}}(X)$ be the vanishing polynomial for \mathbb{H} , let $\widehat{Z}_K(X)$ and $Z_K(X)$ be the vanishing polynomials for $H_{[N]\setminus K}$ and H_K respectively. Let $\mu_1(X), \dots, \mu_N(X)$ be Lagrange polynomials for the set $\mathbb{H} = \{\xi, \dots, \xi^N\}$. Then:

$$\widehat{Z}_K(X) = \sum_{j \in K} \frac{Z'_{\mathbb{H}}(\xi^j)}{Z'_K(\xi^j)} \mu_j(X), \quad (12)$$

$$\widehat{Z}'_K(X) = \sum_{j \in K} \frac{Z'_{\mathbb{H}}(\xi^j)}{Z'_K(\xi^j)} \mu'_j(X) \quad (13)$$

Proof. Note that the second equation follows from the first by linearity of derivatives, so it suffices to prove the first equation. Both the sides of the identity are polynomials of degree $< N$, so it suffices to show their evaluations are identical over N distinct points. In particular we show their evaluations are identical over \mathbb{H} . Consider evaluating LHS and RHS at ξ^i for $i \in [N] \setminus K$. The left side is 0 by definition of $\widehat{Z}_K(X)$, while the right hand side is zero by the properties of Lagrange polynomials. Now consider evaluations LHS and RHS at ξ^i for $i \in K$. The RHS is $\frac{Z'_{\mathbb{H}}(\xi^i)}{Z'_K(\xi^i)}$ by properties of Lagrange polynomials, while the LHS is $\prod_{j \in [N] \setminus K} (\xi^i - \xi^j)$. Multiplying dividing by $\prod_{j \in K \setminus \{i\}} (\xi^i - \xi^j)$ gives:

$$LHS = \frac{\prod_{j \in [N] \setminus \{i\}} (\xi^i - \xi^j)}{\prod_{j \in K \setminus \{i\}} (\xi^i - \xi^j)}$$

Which is $\frac{Z'_{\mathbb{H}}(\xi^i)}{Z'_K(\xi^i)}$, the same as the right hand side. This proves the claim.

Lemma 12. Let μ_1, \dots, μ_N be the lagrange polynomials for the set $\mathbb{H} = \{\xi^i : i \in [N]\}$ of the N^{th} roots of unity. Then we have:

$$\mu'_i(\xi^j) = \begin{cases} \frac{(N-1)}{2\xi^i} & \text{if } j = i \\ \frac{\xi^i}{\xi^j(\xi^j - \xi^i)} & \text{otherwise} \end{cases}$$

Proof. Let us first consider the case where $i \neq j$. We know that $\mu_i(X) = \frac{Z_H(X)}{Z'_H(\xi^i)(X - \xi^i)}$. Thus, by applying quotient rule (note that μ_i is defined at ξ^j as $j \neq i$):

$$\mu'_i(X) \cdot Z'_H(\xi^i) = \frac{(X - \xi^i)(N \cdot X^{N-1}) - (X^N - 1)}{(X - \xi^i)^2}$$

Substituting X by ξ^j , we get:

$$\mu'_i(\xi^j) \cdot \frac{N}{\xi^i} = \frac{N(\xi^j - \xi^i)}{\xi^j(\xi^j - \xi^i)^2}$$

Thus, we get:

$$\mu'_i(\xi^j) = \frac{\xi^i}{\xi^j(\xi^j - \xi^i)}$$

Now, for the second case where $i = j$, we have:

$$\mu_i(X) = \frac{\prod_{j \in [N] \setminus \{i\}} (X - \xi^j)}{Z'_H(\xi^i)}$$

$$\text{or, } \mu_i(X) \cdot Z'_H(\xi^i) = \prod_{j \in [N] \setminus \{i\}} (X - \xi^j)$$

Differentiating the above equation on both sides, we get:

$$\mu'_i(X) \cdot \frac{N}{\xi^i} = \sum_{j \in [N] \setminus \{i\}} \prod_{k \in [N] \setminus \{i, j\}} (X - \xi^k)$$

Substituting $X = \xi^i$ in the above equation yields:

$$\begin{aligned}
\mu'_i(\xi^i) \cdot \frac{N}{\xi^i} &= \sum_{j \in [N] \setminus \{i\}} \prod_{k \in [N] \setminus \{i, j\}} (\xi^i - \xi^k) \\
&= \sum_{j \in [N] \setminus \{i\}} \frac{\prod_{k \in [N] \setminus \{i\}} (\xi^i - \xi^k)}{\xi^i - \xi^j} \\
&= \prod_{k \in [N] \setminus \{i\}} (\xi^i - \xi^k) \sum_{j \in [N] \setminus \{i\}} \frac{1}{\xi^i - \xi^j} \\
&= Z'_H(\xi^i) \sum_{j \in [N] \setminus \{i\}} \frac{1}{\xi^i - \xi^j} \\
&= N/\xi^i \sum_{j \in [N] \setminus \{i\}} \frac{1}{\xi^i - \xi^j}
\end{aligned}$$

We divide on both sides by N/ξ^i in the above, and use Lemma 9 to obtain:

$$\mu'_i(\xi^i) = \sum_{j \in [N] \setminus \{i\}} \frac{1}{\xi^i - \xi^j} = \frac{Z''_H(\xi^i)}{2Z'_H(\xi^i)} = \frac{N-1}{2\xi^i}$$

Lemma 13. Let $K \subseteq \mathbb{N}$ be a set of cardinality k and $\mathcal{X} = \{x_j : j \in K\}$ be a set where x_j for $j \in K$ are distinct elements of \mathbb{F} . Let $Z_{\mathcal{X}}(X) = z_k X^k + \dots + z_0$ denote the vanishing polynomial of \mathcal{X} and $\{\tau_j(X)\}_{j \in K}$ denote the Lagrange polynomials such that $\tau_i(x_j) = \delta_{ij}$ for $i, j \in K$. Then for all $j \in K$, we have $\tau'_j(x_j) = F_K(x_j)/Z'_{\mathcal{X}}(x_j)$ where the polynomial $F_K(X)$ is defined as

$$F_K(X) = \binom{k}{2} z_k X^{k-2} + \dots + \binom{2}{2} z_2 = \sum_{j=2}^k z_j \binom{j}{2} X^{j-2}$$

Proof. For $j \in K$ we have:

$$\tau_j(X) = \frac{Z_{\mathcal{X}}(X)}{(X - x_j)Z'_{\mathcal{X}}(x_j)} = \frac{1}{Z'_{\mathcal{X}}(x_j)} \frac{Z_{\mathcal{X}}(X)}{X - x_j}$$

by definition of Lagrange polynomials. By long division of $Z_{\mathcal{X}}(X)$ by $(X - x_j)$, we have:

$$\begin{aligned}
\tau_j(X) &= \frac{1}{Z'_{\mathcal{X}}(x_j)} (z_k X^{k-1} + (x_j z_k + z_{k-1}) X^{k-2} + \dots + (x_j^{k-1} z_k + \dots + z_1)) \\
&= \frac{1}{Z'_{\mathcal{X}}(x_j)} \sum_{p=0}^{k-1} \left(\sum_{q=p+1}^k z_q x_j^{q-p-1} \right) X^p
\end{aligned}$$

Differentiating both sides, we have:

$$\begin{aligned}
\tau'_j(X) &= \frac{1}{Z'_{\mathcal{X}}(x_j)} \sum_{p=0}^{k-1} \left(\sum_{q=p+1}^k z_q x_j^{q-p-1} \right) p X^{p-1} \\
&= \frac{1}{Z'_{\mathcal{X}}(x_j)} \sum_{p=1}^{k-1} p \sum_{q=p+1}^k z_q x_j^{q-p-1} X^{p-1}
\end{aligned}$$

Substituting $X = x_j$, we get:

$$\begin{aligned}
\tau_j'(x_j) &= \frac{1}{Z'_{\mathcal{X}}(x_j)} \sum_{p=1}^{k-1} p \sum_{q=p+1}^k z_q x_j^{q-2} \\
&= \frac{1}{Z'_{\mathcal{X}}(x_j)} \sum_{q=2}^k z_q x_j^{q-2} \sum_{p=1}^{q-1} p \\
&= \frac{1}{Z'_{\mathcal{X}}(x_j)} \sum_{q=2}^k z_q \binom{q}{2} x_j^{q-2} \\
&= \frac{F_K(x_j)}{Z'_{\mathcal{X}}(x_j)}
\end{aligned}$$

This completes the proof.

D.3 Proof of Lemma 5

Proof. We only present the detailed proof for the computation of a_i for all $i \in I$. We then briefly mention the modifications needed to compute b_j for all $j \in K$.

Computing a_i . Recall that for each $i \in I$, we have:

$$a_i = \sum_{j \in K \setminus \{i\}} \frac{d_j}{\xi^i - \xi^j} \quad (14)$$

Also recall that $I \subset K$ in this case. To compute a_i , we first define a polynomial $p(X)$ of degree at most $N - 1$ such that $p(\xi^j) = d_j$ for $j \in K$ and $p(\xi^j) = 0$ for $j \in [N] \setminus K$. Then, the vanishing polynomial of $H_{[N] \setminus K}$ divides $p(X)$ and there exists a polynomial $q(X)$ of degree at most $|K| - 1$ such that:

$$p(X) = \hat{Z}_K(X) \cdot q(X) \quad (15)$$

where $\hat{Z}_K(X) = \prod_{i \in [N] \setminus K} (X - \xi^i)$ is the vanishing polynomial of $H_{[N] \setminus K}$. Now, we introduce the rational functions:

$$f_i(X) = \sum_{j \in [N] \setminus \{i\}} \frac{p(\xi^j)}{X - \xi^j}, i \in I \quad (16)$$

$$g_i(X) = \sum_{j \in [N] \setminus \{i\}} \frac{p(X)}{X - \xi^j}, i \in I \quad (17)$$

$$r_i(X) = \sum_{j \in [N] \setminus \{i\}} \frac{p(X) - p(\xi^j)}{X - \xi^j}, i \in I \quad (18)$$

Note that, by the definition of $p(X)$, $f_i(\xi^i) = a_i \forall i$. Thus, it suffices to compute $f_i(\xi^i)$ for all $i \in I$. Since $f_i(X) = g_i(X) - r_i(X)$ for $i \in I$, we have that $a_i = g_i(\xi^i) - r_i(\xi^i)$. Thus, we need to compute $g_i(\xi^i)$ and $r_i(\xi^i)$ for all $i \in I \subset K$.

$$\begin{aligned}
g_i(\xi^i) &= p(\xi^i) \sum_{j \in [N] \setminus \{i\}} \frac{1}{\xi^i - \xi^j} \\
&= \frac{p(\xi^i) Z''_{\mathbb{H}}(\xi^i)}{2Z'_{\mathbb{H}}(\xi^i)} \quad (\text{from Lemma 9}) \\
&= \frac{(N-1)d_i}{2\xi^i}
\end{aligned}$$

In the above, we used $Z_H(X) = X^N - 1$ and that $p(\xi^i) = d_i$. In other words, $g_i(\xi^i)$ for all i can be obtained in $O(|I|)$ operations. Therefore, it suffices to compute $r_i(\xi^i)$ for all $i \in I$ efficiently. To this end, we write $r_i(X)$ as:

$$r_i(X) = \sum_{j \in [N]} \frac{p(X) - p(\xi^j)}{X - \xi^j} - \frac{p(X) - p(\xi^i)}{X - \xi^i}$$

By defining the bi-variate polynomial

$$u(X, Y) = (p(X) - p(Y))/(X - Y)$$

we get

$$r_i(X) = \sum_{j \in [N]} u(X, \xi^j) - u(X, \xi^i)$$

Defining $r(X) = \sum_{j \in [N]} u(X, \xi^j)$, we have:

$$r_i(X) = r(X) - u(X, \xi^i)$$

Substituting $X = \xi^i$ in the above, we have:

$$r_i(\xi^i) = r(\xi^i) - u(\xi^i, \xi^i) = r(\xi^i) - p'(\xi^i)$$

where $p'(\xi^i) = u(\xi^i, \xi^i)$ by the definition of formal derivative. Now, using $r(X) = Nu(X, 0)$ (Lemma 10), we have:

$$r(X) = N \frac{(p(X) - p(0))}{X}$$

Finally, substituting $X = \xi^i$ above, we have:

$$r(\xi^i) = N \frac{(d_i - p(0))}{\xi^i}$$

Thus, it remains to compute $p(0)$ and $p'(\xi^i)$ efficiently for each $i \in I$.

Computing the polynomial $q(X)$. Recall from Equation (15) that

$$q(\xi^j) = \frac{p(\xi^j)}{\hat{Z}_K(\xi^j)}$$

for all $j \in K$. Furthermore, by Lemma 11, we have:

$$\hat{Z}_K(\xi^j) = \frac{Z_{\mathbb{H}}'(\xi^j)}{Z_K'(\xi^j)} = \frac{N/\xi^j}{Z_K'(\xi^j)}$$

for each $j \in K$. Observe that, given the set K , we can compute the polynomial $Z_K(X)$ in $O(|K| \log^2 |K|)$ operations using the fast multiplication, and we can then obtain $Z_K'(X)$ in additional $O(|K|)$ operations. Finally, $Z_K'(\xi^j)$ can be evaluated for $j \in K$ in additional $O(|K| \log^2 |K|)$ operations. Thus we can efficiently compute $q(\xi^j)$ for all $j \in K$ in $O(|K| \log^2 |K|)$ operations. Since degree of $q(X)$ is strictly less than $|K|$, we can further interpolate to obtain the polynomial $q(X)$ in $O(|K| \log^2 |K|)$ field operations.

Computing $p(0)$. From Equation (15), we have

$$p(0) = \hat{Z}_K(0) \cdot q(0)$$

Additionally, since we have

$$\hat{Z}_K(0) = \frac{Z_{\mathbb{H}}(0)}{Z_K(0)} = \frac{-1}{Z_K(0)}$$

this enables us to compute $p(0)$ since $q(0)$ and $Z_K(0)$ are just the constant terms of the known polynomials $q(X)$ and $Z_K(X)$.

Computing $p'(\xi^i)$. We now show how to compute $p'(\xi^i)$ for each $i \in I$. Using the product rule for derivatives, we have:

$$p'(X) = q(X) \hat{Z}_K'(X) + q'(X) \hat{Z}_K(X)$$

We have shown how to compute $q(\xi^i)$ and $\hat{Z}_K(\xi^i)$ in $O(|K| \log^2 |K|)$ field operations. By differentiating the polynomial $q(X)$ from earlier, we obtain $q'(X)$. Then, by fast evaluation, we get evaluations of $q'(X)$ at ξ^i for all $i \in I$, again in $O(|K| \log^2 |K|)$ field operations. So it only remains to evaluate $\hat{Z}_K'(\xi^i)$ for each $i \in I$, which we show next. From the second equation of Lemma 11, we have:

$$\hat{Z}_K'(\xi^i) = \sum_{j \in K \setminus \{i\}} \frac{Z_{\mathbb{H}}'(\xi^j)}{Z_K'(\xi^j)} \mu_j'(\xi^i) + \frac{Z_{\mathbb{H}}'(\xi^i)}{Z_K'(\xi^i)} \mu_i'(\xi^i)$$

Using Lemma 12, this becomes:

$$\widehat{Z}'_K(\xi^i) = N\xi^{-i} \sum_{j \in K \setminus \{i\}} \frac{1}{Z'_K(\xi^j)(\xi^i - \xi^j)} + \frac{N(N-1)}{2\xi^{2i}Z'_K(\xi^i)}$$

In other words, it suffices to efficiently compute φ_i for all $i \in I$, where

$$\varphi_i = \sum_{j \in K \setminus \{i\}} \frac{1}{Z'_K(\xi^j)(\xi^i - \xi^j)}$$

. To this end, we define the following polynomial:

$$\Phi_i(X) = \sum_{j \in K \setminus \{i\}} \frac{1}{Z'_K(\xi^j)(X - \xi^j)}$$

Let $\{\eta_i(X)\}_{i \in K}$ be the set of Lagrange polynomials for the set $H_K = \{\xi^i : i \in K\}$. Then, since $\frac{\eta_j(X)}{Z_K(X)} = \frac{1}{Z'_K(\xi^j)(X - \xi^j)}$, $\Phi_i(X)$ can be rewritten as:

$$\begin{aligned} \Phi_i(X) &= \sum_{j \in K \setminus \{i\}} \frac{\eta_j(X)}{Z_K(X)} \\ &= \sum_{j \in K \setminus \{i\}} \frac{\eta_j(X)/(X - \xi^i)}{Z_K(X)/(X - \xi^i)} \end{aligned}$$

Substituting $X = \xi^i$ in the above, we have:

$$\begin{aligned} \varphi_i = \Phi_i(\xi^i) &= \left(\sum_{j \in K \setminus \{i\}} \frac{\eta_j(X)/(X - \xi^i)}{\prod_{k \in K \setminus \{i\}} (X - \xi^k)} \right) (\xi^i) \\ &= \sum_{j \in K \setminus \{i\}} \left(\frac{\eta_j(X)/(X - \xi^i)}{\prod_{k \in K \setminus \{i\}} (X - \xi^k)} \right) (\xi^i) \\ &= \sum_{j \in K \setminus \{i\}} \frac{(\eta_j(X)/(X - \xi^i))(\xi^i)}{\left(\prod_{k \in K \setminus \{i\}} (X - \xi^k) \right) (\xi^i)} \\ &= \frac{1}{Z'_K(\xi^i)} \sum_{j \in K \setminus \{i\}} (\eta_j(X)/(X - \xi^i))(\xi^i) \end{aligned}$$

Now, note that for all $j \neq i$, $(\eta_j(X)/(X - \xi^i))(\xi^i)$ is just the evaluation of the polynomial $\frac{\eta_j(X) - \eta_j(\xi^i)}{X - \xi^i}$ at the point ξ^i . This is just $\eta'_j(\xi^i)$ by definition of formal derivative of the polynomial $\eta_j(X)$. Thus, we get:

$$\varphi_i = \Phi_i(\xi^i) = \frac{1}{Z'_K(\xi^i)} \sum_{j \in K \setminus \{i\}} \eta'_j(\xi^i)$$

Using that fact that $\sum_{j \in K} \eta_j(X) = 1$ (and hence, $\sum_{j \in K} \eta'_j(X) = 0$), we have

$$\sum_{j \in K \setminus \{i\}} \eta'_j(\xi^i) = \sum_{j \in K} \eta'_j(\xi^i) - \eta'_i(\xi^i) = -\eta'_i(\xi^i)$$

Thus, we get:

$$\varphi_i = \frac{-\eta'_i(\xi^i)}{Z'_K(\xi^i)}$$

At this point, it suffices to efficiently compute $\eta'_i(\xi^i)$ for $i \in I$. For this, we can use Lemma 13, with $\mathcal{X} = H_K = \{\xi^j : j \in K\}$ and $Z_{\mathcal{X}}(X) = Z_K(X)$ as the vanishing polynomial of \mathcal{X} , to obtain:

$$\eta'_i(\xi^i) = \frac{F_K(\xi^i)}{Z'_K(\xi^i)}$$

where $F_K(X) = \sum_{j=2}^k z_j \binom{j}{2} X^{j-2}$ as defined in Lemma 13. Hence, it suffices to compute $F_K(\xi^i)$ for all $i \in I$, where z_0, \dots, z_k are the coefficients of the polynomial $Z_K(X)$ computed earlier. This concludes the proof of computation of a_i for $i \in I$.

Modifications for Computing b_j for $j \in K$. For computing b_j , we proceed as in the case of a_i , with the roles of sets I and K swapped (all of the corresponding lemmas can be modified accordingly). The only additional technical subtlety arises when we need to compute $\varphi_j = \Phi_j(\xi^j)$ for all $j \in K$, where the polynomial $\Phi_j(X)$ is defined as:

$$\Phi_j(X) = \sum_{i \in I \setminus \{j\}} \frac{\eta_i(X)}{Z_I(X)}$$

Now, we consider two cases: $j \in I$ and $j \in K \setminus I$. We handle the second case first. For each $j \in K \setminus I$, we can very easily compute $\varphi_j = \Phi_j(\xi^j)$ as

$$\begin{aligned} \Phi_j(\xi^j) &= \sum_{i \in I \setminus \{j\}} \frac{\eta_i(\xi^j)}{Z_I(\xi^j)} \\ &= \frac{1}{Z_I(\xi^j)} \sum_{i \in I} \eta_i(\xi^j) \\ &= \frac{1}{Z_I(\xi^j)} \end{aligned}$$

This is efficiently computed by evaluating $Z_I(\xi^j)$ for each $j \in K$ in $O(|K| \log^2 |K|)$ operations. Next, we consider the case where $j \in I$. For this, we can again proceed as in the analysis for computing a_i (with the roles of sets I and K swapped) till we need to compute

$$\varphi_j = \Phi_j(\xi^j) = \frac{-\eta'_j(\xi^j)}{Z'_I(\xi^j)}$$

for all $j \in I$. First of all, note that during the prior computation to reach this stage, we would have already computed $Z'_I(\xi^j)$ for all $j \in K$, and thus, for all $j \in I \subset K$. Next, observe that we also computed $\eta'_i(\xi^i)$ for $i \in I$ during the computation of a_i . This completes the computation of b_j for all $j \in K$, and finishes the proof of lemma 5.

E Committed Index Lookup from Caulk+

In this section, we present an explicit (non-black-box) adaptation of [PK22] to obtain a committed index lookup, which again incurs costs comparable to a single instance of the underlying sub-vector protocol. Let $m, N \in \mathbb{N}$ be fixed parameters with $m < N$ and let srs denote a KZG setup of degree $d \geq N$ over bi-linear group $(\mathbb{F}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2, [1]_t)$. Recall that the committed index lookup relation in Definition 4 involves the prover showing knowledge of vectors $\mathbf{T} \in \mathbb{F}^N$, $\mathbf{a} \in \mathbb{F}^m$ and $\mathbf{v} \in \mathbb{F}^m$ corresponding to public commitments c_T, c_a and c_v such that they satisfy $v_i = \mathbf{T}[a_i] = T_{a_i}$. We present a polynomial protocol for the same, which is an adaptation of the lookup protocol from Caulk+ [PK22]. However, here we do not aim for zero-knowledge. Let $T(X) = \text{Enc}_{\mathbb{H}}(\mathbf{t})$, $a(X) = \text{Enc}_{\mathbb{V}}(\mathbf{a})$ and $v(X) = \text{Enc}_{\mathbb{V}}(\mathbf{v})$ denote the polynomials encoding the vectors \mathbf{t}, \mathbf{a} and \mathbf{v} respectively. The verifier knows commitments to these polynomials at the start of the protocol. Now $v_i = \mathbf{t}[a_i]$ for $i \in [m]$ is equivalent to $v(\nu^i) = T(\xi^{a(\nu^i)})$ for $i \in [m]$. To obtain a polynomial protocol, the prover interpolates a polynomial $h(X) = \sum_{i=1}^m \xi^{a_i} \tau_i(X)$, which satisfies $h(\nu^i) = \xi^{a(\nu^i)}$. To show that polynomial h correctly “exponentiates” evaluations of $a(X)$, we consider the inverting polynomial $\ell(X) = \sum_{i=1}^N i \mu_i(X)$ which behaves like “log” over \mathbb{H} by evaluating to i on ξ^i . Now, we see that all constraints are encoded as polynomial identities below:

$$\begin{aligned} \ell(h(X)) &= a(X) \pmod{Z_{\mathbb{V}}} \\ T(h(X)) &= v(X) \pmod{Z_{\mathbb{V}}} \\ Z_{\mathbb{H}}(h(X)) &= 0 \pmod{Z_{\mathbb{V}}} \end{aligned} \tag{19}$$

The last polynomial identity ensures that evaluations of h on \mathbb{V} lie in \mathbb{H} (the set of roots of $Z_{\mathbb{H}}$). Since the polynomial ℓ is one-one over \mathbb{H} , the first equation implies $h(\nu^i) = \xi^{a_i}$ for all $i \in [m]$.

The desired relation $v_i = T_{a_i}$ now follows from the second identity. The above formulation involves composition with polynomials ℓ, T and $Z_{\mathbb{H}}$ of degree $O(N)$, which is inefficient. We use the trick from [PK22], where we work with low-degree restrictions of $O(N)$ -degree polynomials such as T, ℓ over the set $\mathbb{H}_I = \{h(\nu^i) : i \in [m]\} = \{\xi^{a_i} : i \in I\} \subseteq \mathbb{H}$, where $I = \{a_i : i \in [m]\}$. The prover commits to the polynomials $Z_I(X) = \prod_{i \in I} (X - \xi^i)$, $h(X)$ and low degree ($< m$) restrictions T_I, ℓ_I of T and ℓ on the \mathbb{H}_I respectively. The polynomial protocol then checks the following:

$$\begin{aligned} T(X) - T_I(X) &= 0 \pmod{Z_I}, & T_I(h(X)) &= v(X) \pmod{Z_{\mathbb{V}}} \\ \ell(X) - \ell_I(X) &= 0 \pmod{Z_I}, & \ell_I(h(X)) &= a(X) \pmod{Z_{\mathbb{V}}} \\ Z_{\mathbb{H}}(X) &= 0 \pmod{Z_I}, & Z_I(h(X)) &= 0 \pmod{Z_{\mathbb{V}}} \end{aligned} \tag{20}$$

It must be noted that the above identities imply the earlier polynomial identities in (19). This is so because evaluations of h on \mathbb{V} are roots of Z_I , which implies $T_I(h(\nu^i)) = T(h(\nu^i))$, $\ell_I(h(\nu^i)) = \ell(h(\nu^i))$ and $Z_{\mathbb{H}}(h(\nu^i)) = 0$ over \mathbb{V} . While the identities on the left still involve a degree N polynomial, we can use the srs to check the polynomial identity at the point τ encoded in the srs. For example, we can evaluate the encoded quotient $[Q(X)]_2 = \left[\frac{T(X) - T_I(X)}{Z_I(X)} \right]_2$ using the relation:

$$\left[\frac{T(X) - T_I(X)}{Z_I(X)} \right]_2 = \sum_{i \in I} \frac{1}{Z_I'(\xi^i)} \left[\frac{T(X) - t_i}{X - \xi^i} \right]_2$$

By pre-computing the KZG proofs $W_1^i = \left[\frac{T(X) - t_i}{X - \xi^i} \right]_2$ for all $i \in [N]$, the encoded quotient can be evaluated using $O(m)$ \mathbb{G}_2 -operations and $O(m \log^2 m)$ \mathbb{F} -operations. The identity is then checked using a real pairing check

$$e([T(X)]_1 - [T_I(X)]_1, [1]_2) = e([Z_I(X)]_1, [Q(X)]_2).$$

Similarly, we also pre-compute the encoded quotients $W_2^i = \left[\frac{\ell(X) - i}{X - \xi^i} \right]_2$ and $W_3^i = \left[\frac{Z_{\mathbb{H}}(X)}{X - \xi^i} \right]_2$ for all $i \in [N]$. The quotients can be computed in time $O(N \log N)$ using the techniques in [FK23]. Using KZG commitment scheme the polynomial relations over $Z_{\mathbb{V}}$ can be checked in a standard manner by having the prover send evaluation proofs for the committed polynomials at a random point chosen by the verifier. The total prover effort incurred is $O(m^2)$ group and field operations. Thus, we have:

Lemma 14. *Assuming KZG is extractable polynomial commitment scheme, there exists a succinct argument of knowledge for the relation $\mathcal{R}_{\text{srs}, N, m}^{\text{lookup}}$ with prover complexity of $O(m^2)$, given access to pre-computed parameters of size $O(N)$.*