

Incrementally Verifiable Computation via Rate-1 Batch Arguments

Omer Paneth^{*1} and Rafael Pass^{†1,2}

¹Tel Aviv University.

²Cornell Tech.

September 18, 2023

Abstract

Non-interactive delegation schemes enable producing succinct proofs (that can be efficiently verified) that a machine M transitions from c_1 to c_2 in a certain number of deterministic steps. We here consider the problem of efficiently *merging* such proofs: given a proof Π_1 that M transitions from c_1 to c_2 , and a proof Π_2 that M transitions from c_2 to c_3 , can these proofs be efficiently merged into a single short proof (of roughly the same size as the original proofs) that M transitions from c_1 to c_3 ? To date, the only known constructions of such a mergeable delegation scheme rely on strong non-falsifiable “knowledge extraction” assumptions. In this work, we present a provably secure construction based on the standard LWE assumption.

As an application of mergeable delegation, we obtain a construction of incrementally verifiable computation (IVC) (with polylogarithmic length proofs) for any (unbounded) polynomial number of steps based on LWE; as far as we know, this is the first such construction based on any falsifiable (as opposed to knowledge-extraction) assumption. The central building block that we rely on, and construct based on LWE, is a *rate-1 batch argument* (BARG): this is a non-interactive argument for NP that enables proving k NP statements x_1, \dots, x_k with communication/verifier complexity $m + o(m)$, where m is the length of one witness. Rate-1 BARGs are particularly useful as they can be recursively composed a super-constant number of times.

^{*}Supported by an Azrieli Faculty Fellowship, Len Blavatnik and the Blavatnik Foundation, the Blavatnik Interdisciplinary Cyber Research Center at Tel Aviv University, and ISF grant 1789/19.

[†]Work done primarily while being a distinguished scholar at the Institute for Advanced Study at Tel-Aviv University. Supported in part by NSF CNS-2149305, NSF Award SATC-1704788, NSF Award RI-1703846, CNS-2128519, AFOSR Award FA9550-18-1-0267, and a JP Morgan Faculty Award. This material is based upon work supported by DARPA under Agreement No. HR00110C0086. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government, the United States Airforce or DARPA.

Contents

1	Introduction	3
1.1	Our results	4
1.2	Technical Overview	5
1.2.1	Rate-1 batch arguments	6
1.2.2	Mergeable proofs and IVC	8
2	Preliminaries	9
2.1	Incrementally Verifiable Computation	9
2.2	RAM delegation	10
2.3	Tree Hash	12
2.4	Somewhere Extractable Hash	12
2.5	Batch Arguments	13
3	Rate-1 Batch Arguments	15
4	Mergeable Delegation	21
4.1	Incrementally Verifiable Computation	27

1 Introduction

Consider some very long computation spanning over multiple generations of humans. Is there a way for the current generation to ensure that the current state of the computation is correct? The notion of incrementally verifiable computation (IVC), proposed by Valiant [Val08], addresses exactly this problem.

An IVC scheme for a machine \mathcal{M} and time bound T consists of a key generation algorithm, and algorithms for updating and verifying proofs. The key-generation algorithm takes a security parameter λ and outputs a public key that is used to prove and verify statements of the form $(\text{cf}, \text{cf}', t)$ claiming that \mathcal{M} starting from configuration cf reaches configuration cf' after t steps. To prove a statement $(\text{cf}_0, \text{cf}_t, t)$ we start from the trivial statement $(\text{cf}_0, \text{cf}_0, 0)$ and the empty proof $\Pi_0 = \mathcal{E}$ which is always accepted. We then apply the proof updating algorithm t times. The update algorithm takes an accepting proof Π_i for the statement $(\text{cf}_0, \text{cf}_i, i)$ where $i < T$ and produces an accepting proof Π_{i+1} for the next statement $(\text{cf}_0, \text{cf}_{i+1}, i + 1)$.¹

The computational soundness requirement states that no polynomial-size attacker can find an accepting proof of a false statement with $t \leq T$. For the IVC to be useful, we additionally need the scheme to satisfy an efficiency requirement: the proof for a statement $x = (\text{cf}, \text{cf}', t)$ is of length $\text{poly}(\lambda, \log T)$ (that is, essentially independent of T) and the time to update and verify proofs is $|x| \cdot \text{poly}(\lambda, \log T)$. One may also consider *weakly-efficient* IVC where the proof length, as well as update and verification times are only required to be sublinear (as opposed to poly-logarithmic) in T .

The notion of IVC can be seen as a strengthening of non-interactive delegation of computation (also known as SNARGs for P). In delegation we require the same proof length and verification time requirements as IVC. A delegation proof for a statement $x = (\text{cf}, \text{cf}', t)$ can be generated in time $|x| \cdot \text{poly}(\lambda, T)$. However, in contrast to IVC, delegation proofs may not be updated.

IVC, on top of being a fascinating notion in its own right, has several applications to delegation:

- **Delegating computation with transferable intermediate proofs:** Consider a client that outsourcing a long computation to an untrusted server. Using delegation, the client can make sure that final result is correct. However, if the server does not perform the computation correctly, the client may only realize this once the computation is over. Instead, using IVC, the client can ask the server to provide the current state of the computation together with a proof of correctness at any point during the computation. If at some point the client detects cheating, or if the cloud is unable to continue the computation, the client can hire another server to continue the verifiable computation from its most recent verified state
- **Time and Space-preserving delegation:** While delegation schemes enable fast verification, the prover typically has significant computational overhead. In particular, in existing solutions based on standard assumption [KPY19, CJJ21, WW22] both the time and the space of the prover grow polynomially with running time of the original computation. This is the case even if the original computation requires small space. Ideally we would like a delegation scheme that preserves both the time and space complexity of the original computation up to quasi-linear overhead. While such efficient arguments have been the focus of recent work, existing solutions either have heuristic security [BC12, BCCT13, BHR⁺21], or are restricted to the designated-verifier setting [HR18].² An IVC directly enables proving the correctness of a time- T space- S computation in time $T \cdot \text{poly}(\lambda, \log T)$ and space $\text{poly}(\lambda, S, \log T)$.

As we shall discuss below, however, known construction of IVC are either based on heuristics, non-falsifiable assumptions (so-called knowledge-extraction assumptions) or only achieve weak efficiency.

¹We typically require the update algorithm to work for *any* accepting proof, even adversarially generated ones. This means that given any accepting proof for some partial computation, we are not only guaranteed that the current state is correct, but also that we can continue to update the state and its proof.

²In designated-verifier delegation, the public key is generated together with a secret key that is required to verify proofs.

IVC from proof merging. Valiant’s original work proposed an approach for constructing IVC based on *proof merging*: Assume that we are given a proof Π_1 for the statement (cf, cf', t_1) and another proof Π_2 for the statement (cf', cf'', t_2) . Can we efficiently (in time that is independent of t_1, t_2) “merge” them into a single proof for the statement $(cf, cf'', t_1 + t_2)$ that is roughly the same length as each of the original proofs? As observed by Valiant, proofs of length $\text{poly}(\lambda, \log T)$ that can be recursively merged $\log T$ times imply IVC. The idea is to incrementally construct a proof for each step of the computation as it is executed and merge the proofs into a single proof in the form of a binary tree. (Or, if the computation time t is not a power of two, into at most $\log t$ proofs.)

IVC under non-falsifiable assumptions. Valiant showed how to instantiate this approach and construct IVC and mergeable proofs, albeit, under a highly non-standard assumption: the existence succinct non-interactive arguments of knowledge for NP (SNARKs) such that for any attacker A of size $\leq \lambda^{\log T}$ there exists a knowledge extractor of *linear size* $|A| \cdot \text{poly}(\lambda)$. All known SNARK constructions are based on non-falsifiable so-called extractability/knowledge assumptions or heuristic assumptions such as Fiat-Shamir and are subject to strong limitations [BCPR14, BP15]. Currently, we do not know of any candidate SNARK construction with an explicit extractor, let alone a linear one.

To relax the knowledge extraction assumption, Bitansky et al. [BCCT13] showed that by merging proofs according to a tree of fan-out λ instead of a binary tree, we can reduce the depth of the “merge tree” to $\log_\lambda T$ instead of $\log T$. This gives IVC for any bound $T = \text{poly}(\lambda)$ based on SNARKs with a polynomial-size (as opposed to linear-size) knowledge extractor. Assuming SNARKs with sub-exponential security, their construction gives IVC for $T = \lambda^{o(\log \lambda)}$. A recent line of work [BCMS20, BCL⁺21, BDFG21] obtained IVC without relying directly on SNARKs, however, the assumptions underlying their constructions are known to imply SNARKS. Accordingly, they suffer from the same drawbacks and are subject to the same limitations as previous IVC constructions.

Weakly-efficient IVC under falsifiable assumptions. The work of Kalai, Paneth and Yang (KPY) [KPY19] provides the first non-trivial IVC scheme from a falsifiable assumptions on bilinear groups. Their scheme, however, only satisfies weak efficiency with proofs of length $2\sqrt{O(\log T \cdot \log \lambda)}$. Their IVC is based on a weak form a proof merging where many proofs are merged in one shot, and the merged proof is longer than the original proof by a factor of $\text{poly}(\lambda)$. To achieve sub-linear proof size, proofs must be merged in a tree of fan-out $> \lambda$.

Summarizing, to date, known IVC constructions require either non-falsifiable knowledge-extraction assumptions or heuristics, while weakly-efficient IVC is known under falsifiable assumptions. The state of the art thus leaves open the following question:

Is IVC possible under standard/falsifiable hardness assumptions?

In this work, we resolve this question by presenting an IVC scheme assuming polynomial hardness of the LWE assumption.

1.1 Our results

Our main result is a new delegation scheme that enables merging $\text{poly}(\lambda)$ many proofs into a single proof, by recursively merging pairs of proofs in a binary tree of depth $O(\log(\lambda))$. Our delegation scheme is for RAM computations (which implies delegation also for Turing machines).

Theorem 1.1 (Mergeable Delegation from LWE). *Assuming the LWE assumption holds, there exists a mergeable delegation scheme for RAM computations.*

Previously, proof systems that can be merged with such efficiently were only known based on non-falsifiable extractability/knowledge assumptions or heuristics. Following [Val08], any such a mergeable delegation scheme can be turned into an IVC scheme. This yields the first IVC scheme under any falsifiable assumption.

Theorem 1.2. *[IVC from LWE] Assuming the LWE assumption holds, there exists an IVC for any time bound $T = \text{poly}(\lambda)$.*

More generally, for any $T = T(\lambda)$, we get IVC for bound $T(\lambda)$ under the T -hardness of LWE. This gives the first IVC scheme with for quasi-polynomial (or larger) time bound T under any assumption other than Valiant’s SNARKs with linear-size knowledge extractor. In particular, by setting $\lambda = \text{polylog}(T)$ and relying on the sub-exponential hardness of LWE we can get IVC with proofs of length $\text{polylog } T$.

Finally, as observed above, any IVC scheme directly implies a time- and space-preserving delegation scheme, and as such we also get the first space-preserving delegation scheme under falsifiable assumptions.

Theorem 1.3 (Time and Space-Preserving Delegation from LWE). *For any $T = T(\lambda)$, as assuming the LWE problem is hard for algorithms of size $\text{poly}(T)$, there exists a non-interactive delegation scheme where proving a time- T space- S computation, the prover requires time $T \cdot \text{poly}(\lambda)$, and space $\text{poly}(\lambda, S, \log T)$.*

In fact, our theorem is stronger than stated. By directly using our the construction of Theorem 1.2 that works in the RAM model, we get time and space preserving delegation for RAM programs and, furthermore, the space complexity of the prover becomes linear, i.e. $S \cdot \text{poly}(\lambda, \log T)$.

Rate-1 batch arguments for NP. Our main tool in constructing mergeable delegation scheme is a new non-interactive argument for conjunctions of NP statements. Given k NP statements with witnesses of length m we can give a computationally sound proof of length $m + O(m/\lambda) + \text{poly}(\lambda, \log k)$ for their conjunction. This is a strengthening of the notion of batch arguments for NP [BHK17] where the proof is of length $\text{poly}(\lambda, m, \log k)$. Therefore, our arguments can be viewed as batch arguments with rate (the ratio between the witness length and the proof length) approaching 1. In contrast to plain batch arguments, where we must batch together many NP statements to achieve non-trivial efficiency, rate-1 batch arguments are useful even for two statements. Another important feature of rate-1 batch arguments is that they can be composed together recursively for a super-constant number of times without blowing up the proof size.

We note that in adaptive setting where the NP statements can be chosen by the adversary as a function of the public key, batch arguments (with any non-trivial rate) satisfying the standard notion of knowledge soundness are subject to strong limitations [BHK17]. Therefore, our arguments satisfy the relaxed notion of somewhere argument of knowledge for adaptively chosen statements introduced in [CJJ21]. This notion is sufficient for our application to mergeable arguments. (See the technical overview for more details.)

We show how to leverage the rate-1 FHE construction of Brakerski et al. [BDGM19] together with the recent construction of (plain) batch arguments of Choudhuri, Jain and Jin [CJJ21] (both based on LWE) to construct rate-1 batch arguments.

Theorem 1.4. *Assuming the LWE assumption holds, there exist rate-1 batch arguments for NP.*

On the Hardness Assumptions. Although above, we have stated our results based on LWE, as we explain in more detail below, all our results only rely on the existence of two primitives: (1) somewhere-extractable hash functions (SEH) [HW15, CJJ21], and (2) standard batch arguments (i.e., not necessarily rate-1).³ As noted above, these primitives are implied by LWE [HW15, CJJ21], which concludes the above theorems. We remark that future works [KLVW23] show how to also instantiate primitives (1) and (2) using other assumptions (e.g., DLIN or subexponential DDH); by relying on these results, we thus get all the above results also from those assumptions.

1.2 Technical Overview

We first elaborate on the construction of rate-1 batch arguments which is our main technical contribution. In Section 1.2.2 we describe the construction of mergeable proofs and IVC from from rate-1 batch batch arguments.

³We rely also on RAM delegation, however, as shown in [CJJ21], this primitive is implied by (1) and (2).

1.2.1 Rate-1 batch arguments

Our construction is based on standard batch arguments with arbitrary rate. Recall that a batch argument proves the conjunction of k NP statements x^1, \dots, x^k with a proof of length $\text{poly}(\lambda, m, \log k)$ where m is the size of one witness. In fact, we rely on a stronger notion known as batch arguments for the index language [CJJ21] where each statement x^i is simply the index i . In more detail, in a batch argument for the index language the prover and verifier are given the NP verification machine \mathcal{M} that takes the statement $i \in [k]$ and a witness w^i . The prover is also given witnesses w^1, \dots, w^k such that $\mathcal{M}(i, w^i)$ accepts for every $i \in [k]$. The batch argument verifier runs in time $|\mathcal{M}| \cdot \text{poly}(\lambda, m, \log k)$.⁴

The notion of soundness we require is somewhere argument of knowledge for adaptively chosen statements [CJJ21]. Roughly speaking, the requirement is that we can generate a programmed public key for the statement i together with a corresponding secret key such that the programmed public key is indistinguishable from an honestly generated public key (in particular, it hides i). Moreover, for every efficient adversary, if, given the programmed public key, the adversary produces an NP verification machine \mathcal{M} together with an accepting proof Π , then, given the secret key, we can extract a witness w from Π such that $\mathcal{M}(i, w) = 1$ with overwhelming probability. Batch arguments for the index language satisfying somewhere argument of knowledge for adaptively chosen statements were constructed in [CJJ21] based on the hardness of LWE.

Our approach. To construct rate-1 batch arguments, our high-level idea is as follows. Given an NP verification machine \mathcal{M} , and given witnesses $w^1, \dots, w^k \in \{0, 1\}^m$, we split each statement $i \in [k]$ into $n \leq m$ statements, each with a short witness. We then prove all $k \cdot n$ statements together with a single batch argument. In more detail, we define a new NP verification machine $\tilde{\mathcal{M}}$ for $k \cdot n$ statements indexed by pairs $(i, j) \in [k] \times [n]$ such that the witness \tilde{w}_j^i for the statement (i, j) is of length $\text{poly}(\lambda)$ (independent of m). Therefore, a batch argument for $\tilde{\mathcal{M}}$ has proof of length $\text{poly}(\lambda, \log k)$.

To split the statement $i \in [k]$, we first split the witness w^i into n blocks $w^i = (w_1^i, \dots, w_n^i)$ each of length $\ell = m/n$. Then, we emulate the execution of $\mathcal{M}(i, w^i)$ in a sequence of n intervals implemented by RAM machines $\mathcal{R}_1^i, \dots, \mathcal{R}_n^i$ executed sequentially, with each machine starting from the final configuration of the previously executed machine. For every $j \in [n]$, the machine \mathcal{R}_j^i has the witness block w_j^i hard-coded. For $j < n$ the machine \mathcal{R}_j^i writes w_j^i to memory and terminates. The final machine \mathcal{R}_n^i writes the final witness block w_n^i to memory and then, once the entire witness w^i is in memory, it emulates $\mathcal{M}(i, w^i)$ and accepts if and only if \mathcal{M} accepts. Let cf_0^i be the starting configuration of \mathcal{R}_1^i and for $j \in [n]$, let cf_j^i be the final configuration of \mathcal{R}_j^i .

Shrinking the sub-statements with RAM delegation. Note that we cannot simply include the configuration cf_j^i as part of the witness \tilde{w}_j^i since it is too long (recall that \tilde{w}_j^i should be of length $\text{poly}(\lambda)$ independent of m). Therefore, we instead verify the computation of each interval using a delegation scheme for RAM machines. Such delegation scheme allows us to verify that the machine \mathcal{R}_j^i transitions from configuration cf_{j-1}^i to cf_j^i without knowing the (potentially long) configurations. Instead, to verify the delegation proof we only need to know the short digests $\text{h}_{j-1}^i, \text{h}_j^i$ of $\text{cf}_{j-1}^i, \text{cf}_j^i$ respectively. The length of the delegation proof is $\text{poly}(\lambda)$ and the verification time is $|\mathcal{R}_j^i| \cdot \text{poly}(\lambda)$. The soundness of the delegation scheme states that an efficient adversary cannot produce configurations $\text{cf}_{j-1}^i, \text{cf}_j^i$ together with digests $\text{h}_{j-1}^i, \tilde{\text{h}}_j^i$ and an accepting proof Π_j^i such that \mathcal{R}_j^i transitions from cf_{j-1}^i to cf_j^i and h_{j-1}^i is the digest of cf_{j-1}^i , but $\tilde{\text{h}}_j^i$ is not the digest of cf_j^i .

We proceed to describe a simplified version of the NP verification machine $\tilde{\mathcal{M}}$ and witness \tilde{w}_j^i : For every $(i, j) \in [k] \times [0, n]$ we compute the digest h_j^i of the configuration cf_j^i as above, and generate a delegation proof Π_j^i that the machine \mathcal{R}_j^i transitions from cf_{j-1}^i to cf_j^i . We then compute a hash tree over the $k \cdot (n+1)$ pairs $\{y_j^i = (w_j^i, \text{h}_j^i)\}$ (we set $w_0^i = \perp$) and hard-code the root h of the tree into $\tilde{\mathcal{M}}$. The witness \tilde{w}_j^i contains

⁴Batch arguments for the index language imply standard batch arguments by considering a verification machine \mathcal{M} that has the instances x^1, \dots, x^k hard-coded in it.

the pairs y_{j-i}^i, y_j^i , their authentication paths to h , and the delegation proof Π_j^i . To verify \tilde{w}_j^i , the machine $\tilde{\mathcal{M}}$ checks that the authentication paths are valid and that the delegation proof for the machine \mathcal{R}_j^i (with w_j^i hard-coded) and the digests h_{j-1}^i, h_j^i is accepting. If $j = 1$ (resp. $j = n$), $\tilde{\mathcal{M}}$ also checks that h_0^i (resp. h_n^i) is the digest of the starting (resp. accepting) configuration.⁵

Towards somewhere argument of knowledge. If we simply use the underlying batch argument for $\tilde{\mathcal{M}}$ as a batch argument for \mathcal{M} we cannot prove that it satisfies the somewhere argument of knowledge property. Recall that to prove somewhere argument of knowledge we must be able to program the public key and extract a complete witness w^i for $\mathcal{M}(i)$ from any accepting proof. However, the underlying batch argument for $\tilde{\mathcal{M}}$ only lets us extract one short witness \tilde{w}_j^i . We may try to execute the adversary with different public keys and extract \tilde{w}_j^i for each $j \in [n]$. However, since the adversary may choose $\tilde{\mathcal{M}}$ adaptively, it may contain a different root h in each execution and, therefore, the extracted witness blocks may not be consistent.

To resolve this, we add to the proof a separate mechanism for extracting w^i based on the notion of somewhere extractable hashing [HW15, CJJ21]. A somewhere extractable hash is a collision resistant hash that can shrink k blocks y^1, \dots, y^k into a single root h and supports local opening. Additionally, in a somewhere extractable hash we can generate a programmed hash key for any index i (that is indistinguishable from an honestly generated key) together with a corresponding secret key that can be used to efficiently extract the block y^i from the root h . In fact, even if the root h is generated adversarially, as long as there exists a valid authentication path from y^i to h , we can extract y^i from h .

We modify the construction of $\tilde{\mathcal{M}}$ as follows: Instead of computing a hash tree directly over the $k \cdot n$ pairs $\{y_j^i = (w_j^i, h_j^i)\}$, we compute the hash in two stages. First, for every $j \in [n]$, we use a somewhere extractable hash to compute a hash tree over the k pairs y_j^1, \dots, y_j^k and obtain a root h_j . Then we compute a (standard) hash tree over the n roots h_1, \dots, h_n and hard-code the root h into $\tilde{\mathcal{M}}$. The witness \tilde{w}_j^i is unchanged except that we replace the direct authentication path from y_j^i to h with two paths: one from y_j^i to h_j and one from h_j to h (and similarly for y_{j-i}^i).

The final batch argument for \mathcal{M} consist of the underlying batch argument for $\tilde{\mathcal{M}}$ together with the n roots $\{h_j\}$. To verify, the batch argument for \mathcal{M} , we check that the underlying batch argument for $\tilde{\mathcal{M}}$ is accepting and that the root h hard-coded in $\tilde{\mathcal{M}}$ is indeed the hash of h_1, \dots, h_n .

To get a rate-1 proof, the somewhere extractable hash we use must also have good rate. We observe that instantiating the hash construction of [HW15] based on FHE, with the rate-1 FHE of [BDGM19] gives a somewhere extractable hash where each root h_j is of length $\ell + O(\ell/\lambda)$ (assuming that the block length ℓ is a sufficiently large polynomial in λ) and therefore the length of all n roots together is $m + O(m/\lambda)$.

Ensuring consistency. Our final step is to modify the NP verification machine $\tilde{\mathcal{M}}$ to check that the witness encoded in the machines $\mathcal{R}_1^i, \dots, \mathcal{R}_n^i$ is the same witness extracted from the roots h_1, \dots, h_n . Otherwise, there is no guarantee that the extracted witness is a valid. The idea is to change the way we compute the root h hard-coded in $\tilde{\mathcal{M}}$. Instead of computing a hash tree directly over the $k \cdot n$ pairs $\{y_j^i = (w_j^i, h_j^i)\}$, we compute a hash tree over the n roots h_1, \dots, h_n (recall that each h_j is itself the root of a hash tree over the k pairs $\{y_j^i\}$ using a somewhere extractable hash) and set h to be its root. Moreover, instead of a direct authentication path from y_j^i to h , we give two paths: one from y_j^i to h_j and one from h_j to h . Finally, to verify the batch argument for \mathcal{M} , we check that the underlying batch argument for $\tilde{\mathcal{M}}$ is accepting and that the root h hard-coded in $\tilde{\mathcal{M}}$ is indeed the hash of h_1, \dots, h_n .

The proof of security. We argue that our final construction satisfies the somewhere argument knowledge property: To extract a witness for the statement $i \in [k]$, we program the somewhere extractable hash key and for every $j \in [n]$, we extract $\bar{y}_j^i = (\bar{w}_j^i, \bar{h}_j^i)$ from the root h_j . We condition on the event that the batch argument proof for \mathcal{M} is accepting and we show that the extracted witness $\bar{w}^i = (\bar{w}_1^i, \dots, \bar{w}_n^i)$ must be valid with overwhelming probability. To this end, we emulate the machines $\mathcal{R}_1^i, \dots, \mathcal{R}_n^i$ with the hard-coded

⁵Assume WLOG that \mathcal{R}_n^i clears its memory before accepting and, therefore, it has a unique accepting configuration.

witness blocks $\bar{w}_1^i, \dots, \bar{w}_n^i$ and obtain the configurations $\bar{c}_0^i, \dots, \bar{c}_n^i$. We say that the digest \bar{h}_j^i extracted from h_j is “good” if it is indeed the digest of \bar{c}_j^i . Below we argue that for every $j \in [n]$, \bar{h}_j^i is good with overwhelming probability. However, before proving that, we show that assuming all digests \bar{h}_j^i are good, the witness \bar{w}^i must be valid.

To show that \bar{w}^i is valid, consider an experiment where we program the public key of the underlying batch argument and extract the witness \tilde{w}_n^i . Since the programmed public key is indistinguishable from the real key (even when extracting from the somewhere extractable hash) we still have that \bar{h}_n^i is good with overwhelming probability. Recall the witness \tilde{w}_n^i contains the pair $y_n^i = (w_n^i, h_n^i)$ and its authentication path to h . We have that with overwhelming probability $\tilde{\mathcal{M}}$ accepts \tilde{w}_j^i and, therefore, the authentication path is valid and h_n^i is the digest of the accepting configuration of \mathcal{R}_n^i . By collision resistance and the extraction guarantee of the somewhere extractable hash, we have that $\bar{h}_n^i = h_n^i$. Since \bar{h}_n^i is good and, again, using collision resistance, we have that \bar{c}_n^i must be the accepting configuration of \mathcal{R}_n^i and, thus, \bar{w}^i is valid.

It remains to argue that for every $j \in [n]$, \bar{h}_j^i is good with overwhelming probability. For $j = 0$, \bar{h}_0^i is good by definition. If for some $j > 1$, \bar{h}_{j-1}^i is good but \bar{h}_j^i is not with noticeable probability, then we break the soundness of the delegation scheme as follows. First, we program the public key of the underlying batch argument and extract the witness \tilde{w}_j^i . Since the programmed public key is indistinguishable from the real key (even when extracting from the somewhere extractable hash) we still have that \bar{h}_{j-1}^i is good but \bar{h}_j^i is not with noticeable probability. Recall the witness \tilde{w}_j^i contains the pairs $y_{j-i}^i = (w_{j-1}^i, h_{j-1}^i)$ and $y_j^i = (w_j^i, h_j^i)$, their authentication paths to h , and the delegation proof Π_j^i . We have that $\tilde{\mathcal{M}}$ accepts \tilde{w}_j^i with overwhelming probability and, therefore, the authentication paths and the delegation proof are all valid. By collision resistance and the extraction guarantee of the somewhere extractable hash we have that $\bar{y}_{j-1}^i = y_{j-1}^i$ and $\bar{y}_j^i = y_j^i$, thus to break the soundness of the delegation scheme for the machine \mathcal{R}_j^i with \bar{w}_j^i hard-coded, we produce the configurations $\bar{c}_{j-1}^i, \bar{c}_j^i$, the digests $\bar{h}_{j-1}^i, \bar{h}_j^i$ and the proof Π_j^i .

1.2.2 Mergeable proofs and IVC

We overview our construction of mergeable delegation for RAM machines based on rate-1 batch arguments. Recall that in delegation for a RAM machines, we can verify that a RAM machine \mathcal{R} transitions from configuration cf to cf' in t steps without knowing the (potentially long) configurations. Instead, the verifier is only given access to the short digests h, h' of cf, cf' respectively. Soundness states that an efficient adversary cannot produce (cf, cf', t) such that \mathcal{R} transitions from cf to cf' in t steps, together with a “hashed statement” (h, h', t) and an accepting proof Π such that h is indeed the digest of cf , but h' is not the digest of cf' . We note that delegation for RAM implies delegation for Turing machines since the verifier can compute the digests of the configurations in the statement itself.

The high-level idea behind our construction is as follows. Given two accepting proofs Π_1, Π_2 under some public key pk for the hashed statements $x_1 = (h, h', t_1)$ and $x_2 = (h', h'', t_2)$ respectively, we merge them into a single proof Π for $(h, h'', t_1 + t_2)$. The proof Π contains both statements x_1, x_2 as well as a batch argument for the following NP verification machine \mathcal{M} with $k = 2$. The machine \mathcal{M} has the hashed statements x_1, x_2 , the public key pk and the code of the RAM machine \mathcal{R} hard-coded. Given an index $i \in [2]$ and a witness Π_i , the machine \mathcal{M} checks that the original delegation verifier given the public key pk , the RAM machine \mathcal{R} and the hashed statement x_i accepts the proof Π_i . We verify a merged proof Π for a hashed statement (h, h', t) as follows. The proof Π contains a pair of instances x_1, x_2 where $x_i = (h_i, h'_i, t_i)$ and a batch argument proof. We first verify that the hashed statements are indeed of the correct form. That is, $(h_1, h'_2, t_1 + t_2) = (h, h', t)$ and $h'_1 = h_2$. Then, we accept if and only if the batch argument proof for the machine \mathcal{M} hard-coded with the hashed statements x_1, x_2 (as well as pk and \mathcal{R}) is accepting.

The proof of security. To argue the soundness of the merged proof, consider an adversary that cheats with noticeable probability ϵ . That is, the adversary outputs a true statement (cf, cf'', t) , a hashed statement (h, h'', t) and an accepting merged proof Π such that h is the digest of cf , but h'' is not the digest of cf'' . Let

x_1, x_2 be the hashed statements contained in the proof. If the proof is accepting these statements must be of the form $x_1 = (h, h', t_1)$ and $x_2 = (h', h'', t_2)$ where $t_1 + t_2 = t$. Let cf' be the configuration that follows t_1 steps after cf and let \tilde{h}' be its digest. Conditioned on the adversary cheating, we have that either $h' = \tilde{h}'$ or $h' \neq \tilde{h}'$ must hold with probability at least $1/2$. Say that the adversary cheats and $h' \neq \tilde{h}'$ with probability $\geq \epsilon/2$ (the proof in the other case is similar). To break the underlying delegation scheme we run the adversary with a programmed public key and extract the witness Π_1 for $i = 1$. Since the programmed public key is indistinguishable from an honestly generated one, the event that the adversary cheats and $h' \neq \tilde{h}'$ still occurs with roughly the same probability $\geq \epsilon/2$. By the somewhere argument of knowledge property, if the batch argument proof in Π is accepting, then $\mathcal{M}(1, \Pi_1)$ accepts with all but negligible probability and, therefore, Π_1 is an accepting proof for the hashed statement x_1 . To break the the underlying delegation scheme we output the true statement (cf, cf', t_1) , the hashed statement $x_1 = (h, h', t_1)$ where $h' \neq \tilde{h}'$ and the accepting proof Π_1 . Similarly, if $h' = \tilde{h}'$ with probability at least $\epsilon/2$, we program the batch argument public key, extract the witness Π_2 . To break the underlying delegation scheme we output the true statement (cf', cf'', t_2) , the hashed statement $x_2 = (h', h'', t_2)$ where $h' = \tilde{h}'$ and the accepting proof Π_2 .

Recursive merging. Since we are using rate-1 batch arguments, we can marge proofs recursively with the proof length the proof growing from m to $m + O(m/\lambda) + \text{poly}(\lambda)$ in each level. Therefore, we can recursively merge proofs $O(\lambda)$ times while keeping the length of the merged proofs bounded by $\text{poly}(\lambda)$. However, to prove soundness of recursively merged proofs we apply the argument above inductively, halving the adversary's advantage with each level. Therefore, to ensure that the advantage remains noticeable we must bound the depth of nested proofs to be logarithmic or, more generally, $\log T$ assuming T -hardness. This allows us to marge T proofs and, therefore, it is sufficient for constructing IVC for bound T .

IVC from mergeable delegation. The construction of IVC from mergeable delegation follows the outline of [Val08]. The high-level idea is as follows. Given a statement (cf, cf', t) we split it into $d \leq \log t$ consecutive statements $\{(cf_i, cf_{i+1}, t_i = 2^{\ell_i}) \in \mathcal{T}_{\mathcal{M}}\}_{i \in [d-1]}$ such that $(cf_1, cf_d) = (cf, cf')$ and for every i , $\ell_i > \ell_{i+1}$. The proof Π for (cf, cf', t) consist of one mergeable delegation proof for each of the d segments, where the proof for the i -th segment is of level ℓ_i . To increment the proof Π by one step we create a new level-0 proof for the statement $(cf', cf'', 1)$ and append it to Π . If Π contains two proofs of the same level ℓ we merge them into a single proof of level $\ell + 1$ (note that such proofs will always be for two consecutive statements) . We repeat this merging operation until Π no longer contains two proof of the same level and it is a valid proof for the statement $(cf, cf'', t + 1)$.

2 Preliminaries

For a deterministic Turing machine \mathcal{M} , let $\mathcal{U}_{\mathcal{M}}$ be the language that contains (x, t) if and only if the machine \mathcal{M} accepts x in t steps. A deterministic RAM machine \mathcal{R} with a word size of λ has random access to memory of size 2^λ bits and a local state of size $O(\lambda)$. At every step, the machine reads or writes a single memory bit and updates its state. For a Turing or RAM machine \mathcal{M} , we refer to the machine's state and the content of its memory/tapes at a given timestep as its configuration. Let $\mathcal{T}_{\mathcal{M}}$ be the language that contains (cf, cf', t) if and only if the machine \mathcal{M} starting from configuration cf transitions to configuration cf' in t steps.

2.1 Incrementally Verifiable Computation

In this section we define incrementally verifiable computation scheme for deterministic Turing machines. The definition is adapted from [KPY19].

An incrementally verifiable computation scheme for deterministic Turing machines consists of algorithms (G, U, V) with the following syntax:

G: The randomized setup algorithm takes as input a security parameter $\lambda \in \mathbb{N}$. It outputs a public key pk .

U: The deterministic update algorithm takes as input the public key pk , a machine \mathcal{M} , a statement $(\text{cf}, \text{cf}', t)$ and a proof Π . It outputs a proof Π' .

V: The deterministic verifier algorithm takes as input the public key pk , a machine \mathcal{M} , a statement $(\text{cf}, \text{cf}', t)$ and a proof Π . It outputs an acceptance bit.

Definition 2.1. A $T(\cdot)$ -secure incrementally verifiable computation scheme for deterministic Turing machines satisfies the following requirements:

Incremental Completeness. For every $\lambda \in \mathbb{N}$ and machine \mathcal{M} :

- For every configuration cf :

$$\Pr \left[\text{pk} \leftarrow \text{G}(\lambda) \quad : \quad \text{V}(\text{pk}, \mathcal{M}, (\text{cf}, \text{cf}, 0), \mathcal{E}) = 1 \right] = 1 .$$

- For every $t < 2^\lambda$, pair of statements $x, x' \in \mathcal{T}_{\mathcal{M}}$ of the form $x = (\text{cf}, \text{cf}', t)$ and $x' = (\text{cf}, \text{cf}'', t+1)$ and a proof Π :

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow \text{G}(\lambda) \\ \Pi' \leftarrow \text{U}(\text{pk}, \mathcal{M}, x, \Pi) \end{array} \quad : \quad \begin{array}{l} \text{V}(\text{pk}, \mathcal{M}, x, \Pi) = 1 \\ \text{V}(\text{pk}, \mathcal{M}, x', \Pi') = 0 \end{array} \right] = 0 .$$

Efficiency. In the incremental completeness experiments above:

- The generation algorithm runs in time $\text{poly}(\lambda)$.
- The verifier algorithm runs in time $(|\mathcal{M}| + |x|) \cdot \text{poly}(\lambda)$.
- The update algorithm runs in time $\text{poly}(\lambda, |\mathcal{M}|, |x|)$.

$T(\cdot)$ -Soundness. For every $\text{poly}(T(\lambda))$ -size adversary Adv , there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow \text{G}(\lambda) \\ (\mathcal{M}, x = (\text{cf}, \text{cf}', t), \Pi) \leftarrow \text{Adv}(\text{pk}) \end{array} \quad : \quad \begin{array}{l} t \leq T(\lambda) \\ \text{V}(\text{pk}, \mathcal{M}, x, \Pi) = 1 \\ x \notin \mathcal{T}_{\mathcal{M}} \end{array} \right] \leq \mu(T(\lambda)) .$$

We say that the scheme is polynomially secure if it is $T(\cdot)$ -secure for every polynomial T .

2.2 RAM delegation

In this section, we define non-interactive delegation for RAM. The definition is a strengthening of the definition in [KPY19] (see discussion following Theorem 2.3). In a RAM delegation scheme, the prover convinces the verifier of a statement of the form $(\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}}$. Since the configurations might be long, the verifier only needs to know a hashed down version of the statement (h, h', t) where h, h' are short digests of cf, cf' respectively.

A non-interactive delegation scheme for RAM consists of algorithms $(\text{G}, \text{H}, \text{P}, \text{V})$ with the following syntax:

G: The randomized setup algorithm takes as input a security parameter $\lambda \in \mathbb{N}$. It outputs a public key pk .

H: The deterministic digest algorithm takes as input the public key pk and a configuration $\text{cf} \in \{0, 1\}^*$. It outputs a digest h .

P: The deterministic prover algorithm takes as input the public key pk , a RAM machine \mathcal{R} and a statement $(\text{cf}, \text{cf}', t)$. It outputs a proof Π .

V: The deterministic verifier algorithm takes as input the public key pk , a RAM machine \mathcal{R} , a hashed statement (h, h', t) and a proof Π . It outputs an acceptance bit.

Definition 2.2. A $T(\cdot)$ -secure non-interactive delegation scheme for RAM satisfies the following requirements.

Completeness. For every $\lambda \in \mathbb{N}$, RAM machine \mathcal{R} , $t \in [2^\lambda]$ and statement $(\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}}$ we have that:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow G(\lambda) \\ \text{h} \leftarrow H(\text{pk}, \text{cf}) \\ \text{h}' \leftarrow H(\text{pk}, \text{cf}') \\ \Pi \leftarrow P(\text{pk}, \mathcal{R}, (\text{cf}, \text{cf}', t)) \end{array} : \mathbf{V}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \right] = 1 .$$

Efficiency. In the completeness experiment above:

- The generation algorithm runs in time $\text{poly}(\lambda)$.
- The digest algorithm on cf runs in time $|\text{cf}| \cdot \text{poly}(\lambda)$ and outputs a digest of length $\text{poly}(\lambda)$.
- The prover algorithm runs in time $|\mathcal{R}| \cdot \text{poly}(\lambda, t)$ and outputs a proof of length $\text{poly}(\lambda)$.
- The verifier algorithm runs in time $|\mathcal{R}| \cdot \text{poly}(\lambda)$.⁶

$T(\cdot)$ -Collision Resistance. For every $\text{poly}(T(\lambda))$ -size adversary Adv , there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow G(\lambda) \\ (\text{cf}, \text{cf}') \leftarrow \text{Adv}(\text{pk}) \end{array} : \begin{array}{l} \text{cf} \neq \text{cf}' \\ H(\text{pk}, \text{cf}) = H(\text{pk}, \text{cf}') \end{array} \right] \leq \mu(T(\lambda)) .$$

$T(\cdot)$ -Soundness. For every $\text{poly}(T(\lambda))$ -size adversary Adv , there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow G(\lambda) \\ (\mathcal{R}, \text{cf}, \text{cf}', \text{h}, \text{h}', t, \Pi) \leftarrow \text{Adv}(\text{pk}) \end{array} : \begin{array}{l} t \leq T(\lambda) \\ \mathbf{V}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = H(\text{pk}, \text{cf}) \\ \text{h}' \neq H(\text{pk}, \text{cf}') \end{array} \right] \leq \mu(T(\lambda)) .$$

We say that the scheme is polynomially secure if it is $T(\cdot)$ -secure for every polynomial T .

Theorem 2.3 ([CJJ21]). Assuming the LWE problem is $T(\cdot)$ -hard, there exists a $T(\cdot)$ -secure non-interactive delegation scheme for RAM.

We remark that the notion of RAM delegation considered in [CJJ21] is slightly weaker than the notion in Definition 2.2. We explain how to modify the [CJJ21] construction to obtain our notion:

- In [CJJ21], the RAM machine \mathcal{R} is fixed while in our notion the prover and verifier are given the description of \mathcal{R} . We can realize this by using the scheme of [CJJ21] for the universal machine and encoding the description of \mathcal{R} as part of the machine's configuration.
- In [CJJ21], the setup algorithm is given the time t for the statement proven while in our notion the prover and verifier are given t as part of the statement. We can realize this by using λ copies the scheme of [CJJ21]: for every $i \in [\lambda]$ we generate a public key for time 2^i . To prove a statement with time t , we divide the t execution steps into at most λ segments, each of length that is a power of two. The proof contains each of the segments and its proof under one of the public keys.
- While [CJJ21] only argue polynomial security of their scheme, it is straightforward to extend their proof to show $T(\cdot)$ -security under the $T(\cdot)$ -hardness of LWE.

⁶The definition only bounds the verification time of honestly generated proofs. We can assume WLOG that the same bound holds also for maliciously generated proofs by capping the verifier execution time.

2.3 Tree Hash

A tree hash consist of algorithms (G, H, V) with the following syntax:

G: The randomized generation algorithm takes as input the security parameter λ . It outputs a public key pk .

H: The deterministic hashing algorithm takes as input the public key pk and messages $x_1, \dots, x_k \in \{0, 1\}^\ell$. It outputs a hash h and opening Π_1, \dots, Π_k .

V: The deterministic verification algorithm takes as input the public key pk , a hash h , a message $x \in \{0, 1\}^\ell$, an index $i \in [k]$ and a proof Π . It outputs an acceptance bit.

Definition 2.4. A $T(\cdot)$ -secure tree hash satisfies the following requirements:

Completeness. For every $\lambda \in \mathbb{N}$, $k \leq 2^\lambda$, messages $x_1, \dots, x_k \in \{0, 1\}^\ell$ and index $i \in [k]$:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow G(\lambda) \\ (h, (\Pi_1, \dots, \Pi_k)) \leftarrow H(\text{pk}, (x_1, \dots, x_k)) \end{array} : V(\text{pk}, h, x_i, i, \Pi_i) = 1 \right] = 1 .$$

Efficiency. In the completeness experiment above:

- The generation algorithm runs in time $\text{poly}(\lambda)$.
- The hashing algorithm runs in polynomial time in its input length and outputs a hash of length $\text{poly}(\lambda)$.
- The verification algorithm runs in time $\ell \cdot \text{poly}(\lambda)$.

$T(\cdot)$ -Collision Resistance. For every $\text{poly}(T(\lambda))$ -size adversary Adv there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow G(\lambda) \\ (h, x, x', i, \Pi, \Pi') \leftarrow \text{Adv}(\text{pk}) \end{array} : \begin{array}{l} x \neq x' \\ V(\text{pk}, h, x, i, \Pi) = 1 \\ V(\text{pk}, h, x', i, \Pi') = 1 \end{array} \right] \leq \mu(T(\lambda)) .$$

We say that the scheme is polynomially secure if it is $T(\cdot)$ -secure for every polynomial T .

Theorem 2.5 ([Mer87, Ajt96, GGH11]). Assuming the LWE problem is $T(\cdot)$ -hard, there exists a $T(\cdot)$ -secure tree hash.

2.4 Somewhere Extractable Hash

In this section, we define somewhere extractable hash. The definition is a strengthening of the definition in [CJJ21] restricted to the case when the set of “binding coordinates” is of size 1. We explicitly define the hash to operate on long strings (rather than on bits as in [CJJ21]) to enable rate-1 constructions. A somewhere extractable hash consist of algorithms (G, T, H, V, E) with the following syntax:

G: The randomized generation algorithm takes as input the security parameter λ . It outputs a public key pk .

T: The randomized trapdoor generation algorithm takes as input the security parameter λ and an index $i \in \mathbb{N}$. It outputs a public key pk and a secret key sk .

H: The deterministic hashing algorithm takes as input the public key pk and messages $x_1, \dots, x_k \in \{0, 1\}^\ell$. It outputs a hash h and openings Π_1, \dots, Π_k .

V: The deterministic verification algorithm takes as input the public key pk , a hash h , a message $x \in \{0, 1\}^\ell$, an index $i \in [k]$ and a proof Π . It outputs an acceptance bit.

E: The deterministic extraction algorithm takes as input the secret key sk and a hash h . It outputs a message $x \in \{0, 1\}^\ell$.

Definition 2.6. A $T(\cdot)$ -secure somewhere extractable hash satisfies the following requirements:

Completeness. For every $\lambda \in \mathbb{N}$, $k \leq 2^\lambda$, messages $x_1, \dots, x_k \in \{0, 1\}^\ell$ and index $i \in [k]$:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow \text{G}(\lambda) \\ (\text{h}, (\Pi_1, \dots, \Pi_k)) \leftarrow \text{H}(\text{pk}, (x_1, \dots, x_k)) \end{array} : \text{V}(\text{pk}, \text{h}, x_i, i, \Pi_i) = 1 \right] = 1 .$$

Efficiency. In the completeness experiment above:

- The generation, trapdoor generation and verification algorithms run in time $\ell \cdot \text{poly}(\lambda)$.⁷
- The hashing and extraction algorithms run in polynomial time in their input length.
- We defined the additive overhead, $\alpha(\lambda, \ell)$, of the scheme as $|\text{h}| - \ell$.

$T(\cdot)$ -Key Indistinguishability. For every $\text{poly}(T(\lambda))$ -size adversary Adv there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and index $i \leq 2^\lambda$:

$$\left| \begin{array}{l} \Pr [\text{pk} \leftarrow \text{G}(\lambda) : \text{Adv}(\text{pk}) = 1] \\ - \Pr [(\text{pk}, \text{sk}) \leftarrow \text{T}(\lambda, i) : \text{Adv}(\text{pk}) = 1] \end{array} \right| \leq \mu(T(\lambda)) .$$

Extraction. For every $\lambda \in \mathbb{N}$, hash h , message $x \in \{0, 1\}^\ell$, index $i \in [2^\lambda]$ and opening Π :

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{T}(\lambda, i) \\ x^* \leftarrow \text{E}(\text{sk}, \text{h}) \end{array} : \begin{array}{l} \text{V}(\text{pk}, \text{h}, x, i, \Pi) = 1 \\ x \neq x^* \end{array} \right] = 0 .$$

We say that the scheme is polynomially secure if it is $T(\cdot)$ -secure for every polynomial T .

Under the hardness of the LWE problem, we can construct a rate-1 somewhere extractable hash by instantiating the construction of [HW15] based on rate-1 FHE [BDGM19].

Theorem 2.7. Assuming the LWE problem is $T(\cdot)$ -hard, there exists a $T(\cdot)$ -secure somewhere-extractable hash with additive overhead $\alpha(\lambda, \ell) = \frac{\ell}{\lambda} + \text{poly}(\lambda)$.

Proof sketch. As shown in [BDGM19], assuming the LWE problem is $T(\cdot)$ -hard, there exists an FHE scheme that enables encrypting a message of length $\ell = \text{poly}(\lambda)$ (for sufficiently large ℓ) into a ciphertext of length $\ell + \frac{\ell}{\lambda}$. We next observe that the construction of somewhere extractable hash from [HW15, CJJ21] the hash value is simply an FHE ciphertext encrypting one of the message hashed. If employing the rate-1 FHE from [BDGM19], the construction will satisfy the desired additive overhead requirement. \square

2.5 Batch Arguments

In this section, we define non-interactive batch arguments for the index language. The definition is a strengthening of the definition in [CJJ21] (see discussion following Theorem 2.9).

A non-interactive batch argument for the index language consist of algorithms $(\text{G}, \text{T}, \text{P}, \text{V}, \text{E})$ with the following syntax:

G: The randomized generation algorithm takes as input the security parameter λ and the witness length m . It outputs a public key pk .

T: The randomized trapdoor generation algorithm takes as input the security parameter λ , the witness length m and an index $i \in \mathbb{N}$. It outputs a public key pk and a secret key sk .

⁷The definition only bounds the verification time of honestly generated hash and opening. We can assume WLOG that the same bound holds also for maliciously generated hash and opening by capping the verifier execution time.

- P:** The deterministic prover algorithm takes as input the public key pk , a machine \mathcal{M} and witnesses $w_1, \dots, w_k \in \{0, 1\}^m$. It outputs a proof Π .
- V:** The deterministic verifier algorithm takes as input the public key pk , a machine \mathcal{M} , and a proof Π . It outputs an acceptance bit.
- E:** The deterministic extraction algorithm takes as input the secret key sk and a proof Π . It outputs a witness $w \in \{0, 1\}^m$.

If the generation and trapdoor generation algorithms do not require m , we say that the batch argument has *unbounded witness length* and omit m from the algorithms input.

Definition 2.8. A $T(\cdot)$ -secure non-interactive batch argument for the index language satisfies the following requirements:

Completeness. For every $\lambda \in \mathbb{N}$, $k, m, t \leq 2^\lambda$, machine \mathcal{M} and witnesses $w_1, \dots, w_k \in \{0, 1\}^m$ such that for every $i \in [k]$: $((i, w_i), t) \in \mathcal{U}_{\mathcal{M}}$ we have that:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow \text{G}(\lambda, m) \\ \Pi \leftarrow \text{P}(\text{pk}, \mathcal{M}, (w_1, \dots, w_k)) \end{array} : \text{V}(\text{pk}, \mathcal{M}, \Pi) = 1 \right] = 1 .$$

Efficiency. In the completeness experiment above:

- The generation and trapdoor generation algorithms run in time $\text{poly}(\lambda, m)$, or $\text{poly}(\lambda)$ if the argument has unbounded witness length.
- The prover algorithm runs in time $|\mathcal{M}| \cdot \text{poly}(\lambda, m, k, t)$ and outputs a proof of length $\text{poly}(\lambda, m)$.
- The verification algorithm runs in time $|\mathcal{M}| \cdot \text{poly}(\lambda, m)$.⁸
- The extraction algorithm runs in polynomial time in its input length.
- We defined the additive overhead, $\alpha(\lambda, \ell)$, of the scheme as $|\Pi| - m$.

$T(\cdot)$ -Key Indistinguishability. For every $\text{poly}(T(\lambda))$ -size adversary Adv and polynomial $m(\lambda)$ there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and index $i \leq 2^\lambda$:

$$\left| \begin{array}{l} \Pr [\text{pk} \leftarrow \text{G}(\lambda, m) : \text{Adv}(\text{pk}) = 1] \\ - \Pr [(\text{pk}, \text{sk}) \leftarrow \text{T}(\lambda, m, i) : \text{Adv}(\text{pk}) = 1] \end{array} \right| \leq \mu(T(\lambda)) .$$

$T(\cdot)$ -Somewhere Argument of Knowledge. For every $\text{poly}(T(\lambda))$ -size adversary Adv and polynomials $k(\lambda), m(\lambda), t(\lambda)$ there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and index $i \in [k]$:

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{T}(\lambda, m, i) \\ (\mathcal{M}, \Pi) \leftarrow \text{Adv}(\text{pk}) \\ w \leftarrow \text{E}(\text{sk}, \Pi) \end{array} : \begin{array}{l} \text{V}(\text{pk}, \mathcal{M}, \Pi) = 1 \\ ((i, w), t) \notin \mathcal{U}_{\mathcal{M}} \end{array} \right] \leq \mu(T(\lambda)) .$$

We say that the scheme is polynomially secure if it is $T(\cdot)$ -secure for every polynomial T .

Theorem 2.9 ([CJJ21]). Assuming the LWE problem is $T(\cdot)$ -hard, there exist a $T(\cdot)$ -secure non-interactive batch argument for the index language.

We remark that the notion of batch arguments defined in [CJJ21] is slightly different than the notion in Definition 2.8. We explain how to modify the [CJJ21] construction to obtain our notion:

⁸The definition only bounds the verification time of honestly generated proofs. We can assume WLOG that the same bound holds also for maliciously generated proofs by capping the verifier execution time.

- In [CJJ21], the prover and verifier are given a circuit implementing the NP verification procedure and the setup algorithm is given the size of this circuit. In our notion, the NP verification procedure is given by a Turing machine and the setup algorithm is not given a bound on the size or running time of the machine but only the witness length m . As discussed in [CJJ21, Section 6], this can be realized by combining the notion of [CJJ21] with a RAM delegation scheme satisfying Definition 2.2. Such a RAM delegation scheme is also known under the LWE assumption (see Theorem 2.3).
- In [CJJ21], the setup algorithm is given the number of statements k while in our notion the the number of statements is not fixed ahead of time. We can realize this by using λ copies the scheme of [CJJ21]: for every $i \in [\lambda]$, we generate a public key for 2^i statements. To prove k statements, divide the statements into at most λ groups, each of size that is a power of two. The final proof contains, for each group, a proof under one of the public keys. Here we use the fact that in the batch arguments of [CJJ21], the setup time is poly-logarithmic in the number of statements.
- While [CJJ21] only argue polynomial security of their scheme, it is straightforward to extend their proof to show $T(\cdot)$ -security under the $T(\cdot)$ -hardness of LWE.

3 Rate-1 Batch Arguments

In this section, we prove Theorem 3.1 giving batch arguments with unbounded witness length and small additive overhead.

Theorem 3.1. *Assume the existence of:*

- A $T(\cdot)$ -secure non-interactive batch argument for the index language.
- A $T(\cdot)$ -secure somewhere extractable hash with additive overhead $\alpha(\lambda, \ell) = \frac{\ell}{\lambda} + \text{poly}(\lambda)$.
- A $T(\cdot)$ -secure non-interactive delegation scheme for RAM.

Then there exist a $T(\cdot)$ -secure non-interactive batch argument for the index language with unbounded witness length and with additive overhead $\alpha(\lambda, m) = \frac{3m}{\lambda} + \text{poly}(\lambda)$.

Construction. We construct a $T(\cdot)$ -secure, non-interactive batch argument $(G'_{\text{BA}}, T'_{\text{BA}}, P'_{\text{BA}}, V'_{\text{BA}}, E'_{\text{BA}})$ with additive overhead $\frac{3m}{\lambda} + \text{poly}(\lambda)$ using the following building blocks:

- A $T(\cdot)$ -secure non-interactive batch argument for the index language $(G_{\text{BA}}, T_{\text{BA}}, P_{\text{BA}}, V_{\text{BA}}, E_{\text{BA}})$.
- A $T(\cdot)$ -secure somewhere extractable hash $(G_{\text{EH}}, T_{\text{EH}}, H_{\text{EH}}, V_{\text{EH}}, E_{\text{EH}})$ with additive overhead $\frac{\ell}{\lambda} + Q_1(\lambda)$ for some polynomial Q_1 .
- A $T(\cdot)$ -secure non-interactive delegation scheme for RAM $(G_{\text{DL}}, H_{\text{DL}}, P_{\text{DL}}, V_{\text{DL}})$.
- A $T(\cdot)$ -secure tree hash $(G_{\text{TH}}, H_{\text{TH}}, V_{\text{TH}})$ with hash of length $Q_2(\lambda)$ for some polynomial Q_2 .

We are given an NP verification machine \mathcal{M} and a sequence of witnesses (w^1, \dots, w^k) , each of length m . Let t be the running time of \mathcal{M} on each witness. We first split each witness w^i into n blocks $w^i = (w_1^i, \dots, w_n^i)$ where each block is of length $\ell = m/n$. For every $i \in [k]$, we define a computation checking that $((i, w^i), t) \in \mathcal{U}_{\mathcal{M}}$. The computation is made of a sequence of n intervals implemented by RAM machines $\mathcal{R}_1^i, \dots, \mathcal{R}_n^i$. We run the machines sequentially on the same memory, that is, each machine starts from the final configuration of the previous one. For every $j \in [n]$, the machine \mathcal{R}_j^i has the witness block w_j^i hard-coded. For $j < n$ the machine \mathcal{R}_j^i writes w_j^i to memory and terminates. The final machine \mathcal{R}_n^i writes the final witness block w_n^i to memory and then, once the entire witness w^i is in memory, it emulates $\mathcal{M}(i, w^i)$ for t steps and accepts if and only if \mathcal{M} accepts. Let cf_0^i be the starting configuration of \mathcal{R}_1^i and for $j \in [n]$, let cf_j^i be the final configuration of \mathcal{R}_j^i . Using a RAM delegation scheme we compute the digest h_j^i of the configuration cf_j^i and generate a proof Π_j^i that the machine \mathcal{R}_j^i indeed transitions from cf_{j-1}^i to cf_j^i . Therefore, given the witness

w^i and the digests $\{h_j^i\}_j$ we can verify that $((i, w^i), t) \in \mathcal{U}_{\mathcal{M}}$ by checking that h_0^i and h_n^i indeed are the digests of the correct starting and accepting configurations respectively, and that all the proofs $\{\Pi_j^i\}_{i,j}$ are accepting.

The next step is to prove that all proofs are accepting using a single batch argument. To this end, we use the witness blocks and the digests to define a collection of NP statements and hash them down. In more details, we first use a somewhere extractable hash to compute for every $j \in [n]$, the hash h_j of the k pairs $\{w_j^i, h_j^i\}_i$. Then, we further hash the n hash values $\{h_j\}_j$ to a single hash h using a tree hash. We can now define the NP verification machine \mathcal{M}' for which we provide the batch argument. The machine \mathcal{M}' given index (i, j) accepts a witness that contains (a) valid openings of h to h_{j-1} and h_j , (b) valid openings of h_{j-1} and h_j to (w_j^i, h_j^i) and (w_j^i, h_j^i) respectively, and (c) an accepting proof Π_j^i for the RAM machine \mathcal{R}_j^i (recall that the machine \mathcal{R}_j^i can be computed from w_j^i) and the digests (h_{j-1}^i, h_j^i) . The final proof consist of the hash values $\{h_j\}_j$ and the batch argument proof for \mathcal{M}' .

We proceed with a formal description of the construction.

The machine \mathcal{R}_j^i . For $i \in [k]$, $j \in [n]$, a string $x \in \{0, 1\}^\ell$ and a machine \mathcal{M} , let $\mathcal{R}_j^i[x, \mathcal{M}]$ be the following RAM machine:

- Writes x to memory starting at location $(j-1) \cdot \ell + 1$.
- If $j = n$:
 - Emulate \mathcal{M} on input (i, w^i) where w^i is the given in the first $m = n \cdot \ell$ memory locations.
 - Empty the memory and accept if and only if \mathcal{M} accepts.

Let $t_j^i = (\ell + |\mathcal{M}| + t) \cdot \text{poly}(\lambda)$ be the running time of the machine $\mathcal{R}_j^i[x, \mathcal{M}]$. Let $\text{cf}_{\text{start}}^i$ be the starting configuration of $\mathcal{R}_1^i[x, \mathcal{M}]$ and let $\text{cf}_{\text{accept}}^i$ be the accepting configuration of \mathcal{R}_n^i with empty memory.

The machine \mathcal{M}' . Given a public key $\text{pk} = (\text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}})$, a machine \mathcal{M} and hash value h , let $\mathcal{M}'[\text{pk}, \mathcal{M}, h]$ be the machine on input an index i' and a witness w' :

$$w' = ((x^b, h_{\text{DL}}^b, \Pi_{\text{EH}}^b, h_{\text{EH}}^b, \Pi_{\text{TH}}^b)_{b \in \{0,1\}}, \Pi_{\text{DL}}) ,$$

proceeds as follows:

- Let $i \in [k]$, $j \in [n]$ be such that $i' = (i-1) \cdot n + j$.
- For $b \in \{0, 1\}$, check that $V_{\text{TH}}(\text{pk}_{\text{TH}}, h, h_{\text{EH}}^b, j-b, \Pi_{\text{TH}}^b) = 1$.
- For $b \in \{0, 1\}$, check that $V_{\text{EH}}(\text{pk}_{\text{EH}}, h_{\text{EH}}^b, (x^b, h_{\text{DL}}^b), i, \Pi_{\text{EH}}^b) = 1$.
- Check that $V_{\text{DL}}(\text{pk}_{\text{DL}}, \mathcal{R}_j^i[x_0, \mathcal{M}], h_{\text{DL}}^1, h_{\text{DL}}^0, t_j^i, \Pi_{\text{DL}}) = 1$.
- If $j = 1$ check that $h_{\text{DL}}^1 = H_{\text{DL}}(\text{pk}_{\text{DL}}, \text{cf}_{\text{start}}^i)$.
- If $j = n$ check that $h_{\text{DL}}^0 = H_{\text{DL}}(\text{pk}_{\text{DL}}, \text{cf}_{\text{accept}}^i)$.
- Accept if and only if all checks pass.

Let $m' = \ell \cdot \text{poly}(\lambda)$ be the length of the witness w' .

The batch argument algorithms.

G'_{BA} : Given as input the security parameter λ the generation algorithm is as follows:

- Set $\ell = \lambda \cdot (Q_1(\lambda) + Q_2(\lambda))$.
- Set $\text{pk}_{\text{BA}} \leftarrow G_{\text{BA}}(\lambda, m')$.
- Set $\text{pk}_{\text{TH}} \leftarrow G_{\text{TH}}(\lambda)$.
- Set $\text{pk}_{\text{EH}} \leftarrow G_{\text{EH}}(\lambda)$.
- Set $\text{pk}_{\text{DL}} \leftarrow G_{\text{DL}}(\lambda)$.
- Output $\text{pk} = (\text{pk}_{\text{BA}}, \text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}})$.

T'_{BA} : Given as input the security parameter λ and an index i , the trapdoor generation algorithm is as follows:

- Set $\ell = \lambda \cdot (Q_1(\lambda) + Q_2(\lambda))$.
- Set $\text{pk}_{\text{BA}} \leftarrow G_{\text{BA}}(\lambda, m')$.
- Set $\text{pk}_{\text{TH}} \leftarrow G_{\text{TH}}(\lambda)$.
- Set $(\text{pk}_{\text{EH}}, \text{sk}) \leftarrow T_{\text{EH}}(\lambda, i)$.
- Set $\text{pk}_{\text{DL}} \leftarrow G_{\text{DL}}(\lambda)$.
- Output $(\text{pk} = (\text{pk}_{\text{BA}}, \text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}}), \text{sk})$.

P'_{BA} : Given as input the public key pk , a machine \mathcal{M} and witnesses $w_1, \dots, w_k \in \{0, 1\}^m$, the prover algorithm is as follows:

- Let $\text{pk} = (\text{pk}_{\text{BA}}, \text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}})$.
- Set $n = \frac{m}{\ell}$.
- For $i \in [k]$ let $w^i = (w_1^i, \dots, w_n^i)$ where each w_j^i is of length ℓ .
- For $i \in [k]$ set $w_0^i = 0^\ell$.
- For $i \in [k]$ set $\text{cf}_0^i = \text{cf}_{\text{start}}^i$.
- For $i \in [k]$ and $j \in [n]$ compute the final configuration cf_j^i of $\mathcal{R}_j^i[w_j^i, \mathcal{M}]$ when starting from configuration cf_{j-1}^i .
- For $i \in [k]$ and $j \in [n]$ set $\tilde{\Pi}_j^i \leftarrow P(\text{pk}_{\text{DL}}, \mathcal{R}, \text{cf}_{j-1}^i, \text{cf}_j^i, t_j^i)$.
- For $i \in [k]$ and $j \in [0, n]$ set $\mathbf{h}_j^i \leftarrow H_{\text{DL}}(\text{pk}_{\text{DL}}, \text{cf}_j^i)$.
- For $j \in [0, n]$ set $(\mathbf{h}_j, (\Pi_j^1, \dots, \Pi_j^k)) \leftarrow H_{\text{EH}}(\text{pk}_{\text{EH}}, ((w_j^1, \mathbf{h}_j^1), \dots, (w_j^k, \mathbf{h}_j^k)))$.
- Set $(\mathbf{h}, (\Pi_0, \dots, \Pi_n)) \leftarrow H_{\text{TH}}(\text{pk}_{\text{TH}}, (\mathbf{h}_0, \dots, \mathbf{h}_n))$.
- For $i \in [k]$ and $j \in [n]$ set $w'_{(i-1) \cdot n + j} = ((w_{j-b}^i, \mathbf{h}_{j-b}^i, \Pi_{j-b}^i, \mathbf{h}_{j-b}, \Pi_{j-b})_{b \in \{0, 1\}}, \tilde{\Pi}_j^i)$.
- Set $\Pi_{\text{BA}} \leftarrow P_{\text{BA}}(\text{pk}_{\text{BA}}, \mathcal{M}'[(\text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}}), \mathcal{M}, \mathbf{h}], (w'_1, \dots, w'_{k \cdot n}))$.
- Output $\Pi = ((\mathbf{h}_0, \dots, \mathbf{h}_n), \Pi_{\text{BA}})$.

V'_{BA} : Given as input the public key pk , a machine \mathcal{M} and a proof Π , the verifier algorithm is as follows:

- Let $\text{pk} = (\text{pk}_{\text{BA}}, \text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}})$.
- Let $\Pi = ((\mathbf{h}_0, \dots, \mathbf{h}_n), \Pi_{\text{BA}})$.
- Set $(\mathbf{h}, (\Pi_0, \dots, \Pi_n)) \leftarrow H_{\text{EH}}(\text{pk}_{\text{EH}}, (\mathbf{h}_0, \dots, \mathbf{h}_n))$.
- Output the same as $V_{\text{BA}}(\text{pk}_{\text{BA}}, \mathcal{M}'[(\text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}}), \mathcal{M}, \mathbf{h}], \Pi_{\text{BA}})$.

E'_{BA} : Given as input the secret key sk and a proof Π , the extraction algorithm is as follows:

- Let $\Pi = ((\mathbf{h}_0, \dots, \mathbf{h}_n), \Pi_{\text{BA}})$.
- For $j \in [n]$, set $(w_j, \mathbf{h}'_j) \leftarrow \text{E}_{\text{EH}}(\text{sk}, \mathbf{h}_j)$
- Output $w = (w_1, \dots, w_n)$.

The completeness and efficiency requirement follow directly from the completeness and efficiency of the building blocks. The key indistinguishability requirement follows directly from the key indistinguishability of the somewhere extractable hash. We focus on proving that the construction has low additive overhead and that satisfies the somewhere argument of knowledge requirement.

Additive overhead. Next we analyze the additive overhead of the batch argument. The proof Π consists of $n + 1$ hash values $\mathbf{h}_0, \dots, \mathbf{h}_n$ and a batch argument Π_{BA} . Using the additive overhead of the somewhere extractable hash and since $\ell = \lambda \cdot (Q_1(\lambda) + Q_2(\lambda))$, we get that each hash \mathbf{h}_j of string of length $\ell + Q_2(\lambda)$ is of length:

$$(\ell + Q_2(\lambda))(1 + \frac{1}{\lambda}) + Q_1(\lambda) \leq \ell + 3(Q_1(\lambda) + Q_2(\lambda)) .$$

The length of the batch argument Π_{BA} is $\text{poly}(\lambda, m') = \text{poly}(\lambda, \ell)$. Overall we have that:

$$|\Pi| = (n + 1) \cdot (\ell + 3(Q_1(\lambda) + Q_2(\lambda))) + \text{poly}(\lambda, \ell) = m + \frac{3m}{\lambda} + \text{poly}(\lambda, \ell) .$$

Somewhere argument of knowledge. Fix a $\text{poly}(T(\lambda))$ -size adversary Adv and polynomials k, m, t . For $\lambda \in \mathbb{N}, i \in [k]$, let EXP be the experiment of the somewhere argument of knowledge requirement:

$$\left\{ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{T}'_{\text{BA}}(\lambda, i) \\ (\mathcal{M}, \Pi) \leftarrow \text{Adv}(\text{pk}) \\ w \leftarrow \text{E}'_{\text{BA}}(\text{sk}, \Pi) \end{array} \right\}$$

Assume towards contradiction that there exists a polynomial P such that for infinitely many $\lambda \in \mathbb{N}$, there exists $i \in [k]$ such that:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} \text{V}'_{\text{BA}}(\text{pk}, \mathcal{M}, \Pi) = 1 \\ ((i, w), t) \notin \mathcal{U}_{\mathcal{M}} \end{array} \right] \geq \frac{1}{P(T(\lambda))} .$$

Fix any sufficiently large λ and i such that the above holds. Recall that in the experiment EXP :

- $\text{pk} = (\text{pk}_{\text{BA}}, \text{pk}_{\text{TH}}, \text{pk}_{\text{EH}}, \text{pk}_{\text{DL}})$.
- $\Pi = ((\mathbf{h}_0, \dots, \mathbf{h}_n), \Pi_{\text{BA}})$.
- $w = (w_1^i, \dots, w_n^i)$ where $(w_j, \mathbf{h}_j^i) \leftarrow \text{E}_{\text{EH}}(\text{sk}, \mathbf{h}_j)$.

On a high-level, our proof strategy is as follows. We first use w to recompute the digests $\hat{\mathbf{h}}_0^i, \dots, \hat{\mathbf{h}}_n^i$ following the honest prover algorithm. Then, we compare these honest digests to the digests $\mathbf{h}_0^i, \dots, \mathbf{h}_n^i$ extracted. If all the extracted digests are honest, but w is not a valid witness, then the final extracted digest \mathbf{h}_n^i corresponds to a rejecting configuration. Therefore, \mathbf{h}_n^i must be different from $\mathbf{h}_{\text{accept}}^i$, the digest of the accepting configuration $\text{cf}_{\text{accept}}^i$, or we find a digest collision. If this event occurs with noticeable probability over the the choice of the public key then, by the key indistinguishability requirement of the batch argument, it must also occur with noticeable probability if we program the batch argument public key and extract the witness for the final block of w . The extracted witness contains the opening of the digest $\bar{\mathbf{h}}_n^i$ and if the witness is valid, this opening is accepting and we can use the collision resistance of the tree hash and somewhere extractable hash to show that $\bar{\mathbf{h}}_n^i = \mathbf{h}_n^i$. However, since \mathcal{M}' explicitly checks that $\bar{\mathbf{h}}_n^i = \text{cf}_{\text{accept}}^i$, the witness cannot be valid, contradicting the somewhere argument of knowledge property of the batch argument. Similarly, if the first extracted digest \mathbf{h}_0^i is not honest, we can reach a contradiction since \mathcal{M}' explicitly checks that $\bar{\mathbf{h}}_0^i$ is the digest of the starting configuration.

It is left to deal with the case that h_0^i is honest but h_n^i is not. Therefore, there must be an index j such that, with noticeable probability over the choice of the public key, the digest h_{j-1}^i is honest, but the subsequent digest h_j^i is not. Therefore, by the key indistinguishability requirement of the batch argument, this event must also occur with noticeable probability if we program the batch argument public key and extract the witness for the j -th block of w . The extracted witness contains openings of the witness block \bar{w}_j^i , and the two digests $\bar{h}_{j-1}^i, \bar{h}_j^i$. If the witness is valid, all opening are accepting, and we can use the collision resistance of the tree hash and somewhere extractable hash to show that $(\bar{w}_j^i, \bar{h}_{j-1}^i, \bar{h}_j^i) = (w_j^i, h_{j-1}^i, h_j^i)$. The extracted witness also contains a RAM delegation proof that \mathcal{R}_j^i hard-coded with w_j^i indeed transitions from h_{j-1}^i to h_j^i . Since h_{j-1}^i is honest but h_j^i is not, we reach a contradiction to the soundness of the delegation scheme.

We proceed with the formal proof. We first define the digests $\hat{h}_1^i, \dots, \hat{h}_n^i$ that result from the execution of the machines $\mathcal{R}_1^i, \dots, \mathcal{R}_n^i$ on w , just as the honest prover does:

- Let $\hat{cf}_0^i = cf_{\text{start}}^i$ be the starting configuration of $\mathcal{R}_1^i[w_1^i, \mathcal{M}]$.
- For $j \in [n]$, compute the final configuration \hat{cf}_j^i of $\mathcal{R}_j^i[w_j^i, \mathcal{M}]$ when starting from configuration \hat{cf}_{j-1}^i .
- For $j \in [0, n]$, set $\hat{h}_j^i \leftarrow H_{\text{DL}}(\text{pk}_{\text{DL}}, \hat{cf}_j^i)$.

By the definition of the machines $\mathcal{R}_1^i, \dots, \mathcal{R}_n^i$, we have that if $((i, w), t) \notin \mathcal{U}_{\mathcal{M}}$ then $\hat{cf}_n^i \neq cf_{\text{accept}}^i$ since cf_{accept}^i is an accepting configuration of \mathcal{R}_n^i . Therefore:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} V'_{\text{BA}}(\text{pk}, \mathcal{M}, \Pi) = 1 \\ \hat{cf}_n^i \neq cf_{\text{accept}}^i \end{array} \right] \geq \frac{1}{P(T(\lambda))} .$$

Let $h_{\text{accept}}^i = H_{\text{DL}}(\text{pk}_{\text{DL}}, cf_{\text{accept}}^i)$. By the collision resistance requirement of the delegation scheme we have that:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} V'_{\text{BA}}(\text{pk}, \mathcal{M}, \Pi) = 1 \\ \hat{h}_n^i \neq h_{\text{accept}}^i \end{array} \right] \geq \frac{1}{2P(T(\lambda))} .$$

Let A_1, \dots, A_n be the following events:

- A_1 is the event that $h_1^i \neq \hat{h}_1^i$.
- For $j \in [2, n-1]$, A_j is the event that $h_{j-1}^i = \hat{h}_{j-1}^i$ and $h_j^i \neq \hat{h}_j^i$.
- A_n is the event that $h_{n-1}^i = \hat{h}_{n-1}^i$.

Since one of these events must occur, there exists $j \in [n]$ such that:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} V'_{\text{BA}}(\text{pk}, \mathcal{M}, \Pi) = 1 \\ A_j \wedge \hat{h}_n^i \neq h_{\text{accept}}^i \end{array} \right] \geq \frac{1}{2n \cdot P(T(\lambda))} .$$

Fix such j . Let EXP' be the experiment that is defined like EXP except that we use the trapdoor generation algorithm of the batch argument and extract a witness w' for block (i, j) :

- Set $(\text{pk}_{\text{BA}}, \text{sk}_{\text{BA}}) \leftarrow T_{\text{BA}}(\lambda, m', (i-1) \cdot n + j)$.
- Set $w' \leftarrow E_{\text{BA}}(\text{sk}_{\text{BA}}, \Pi_{\text{BA}})$.
- Let $w' = ((\bar{w}_{j-b}^i, \bar{h}_{j-b}^i, \bar{\Pi}_{j-b}^i, \bar{h}_{j-b}^i, \bar{\Pi}_{j-b}^i)_{b \in \{0,1\}}, \bar{\Pi}_j^i)$.

Since the event A_j can be tested in polynomial time, without knowing sk_{BA} , it follows from the key indistinguishability requirement of the underlying batch argument that:

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} V'_{\text{BA}}(\text{pk}, \mathcal{M}, \Pi) = 1 \\ A_j \wedge \hat{h}_n^i \neq h_{\text{accept}}^i \end{array} \right] \geq \frac{1}{3n \cdot P(T(\lambda))} .$$

Let $(\mathbf{h}, (\Pi_0, \dots, \Pi_n)) \leftarrow \mathbf{H}_{\text{EH}}(\text{pk}'_{\text{EH}}, (\mathbf{h}_0, \dots, \mathbf{h}_n))$ be the hash computed by \mathbf{V}'_{BA} . Since \mathbf{V}'_{BA} only accepts if \mathbf{V}_{BA} accepts Π_{BA} , it follows from the somewhere argument of knowledge requirement of the underlying batch argument that:

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} ((i-1) \cdot n + j, w') \in \mathcal{U}_{\mathcal{M}'[(\text{pk}_{\text{TH}}, \text{pk}'_{\text{EH}}, \text{pk}_{\text{DL}}), \mathcal{M}, \mathbf{h}]} \\ \mathbf{A}_j \wedge \hat{\mathbf{h}}_n^i \neq \mathbf{h}_{\text{accept}}^i \end{array} \right] \geq \frac{1}{4n \cdot P(T(\lambda))} .$$

Let $\mathbf{h}_{\text{start}}^i = \mathbf{H}_{\text{DL}}(\text{pk}_{\text{DL}}, \text{cf}_{\text{start}}^i)$. By the definition of \mathcal{M}' :

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} \forall b \in \{0, 1\} : \mathbf{V}_{\text{TH}}(\text{pk}_{\text{TH}}, \mathbf{h}, \bar{\mathbf{h}}_{j-b}, j-b, \bar{\Pi}_{j-b}) = 1 \\ \forall b \in \{0, 1\} : \mathbf{V}_{\text{EH}}(\text{pk}'_{\text{EH}}, \bar{\mathbf{h}}_{j-b}, (\bar{w}_{j-b}^i, \bar{\mathbf{h}}_{j-b}^i), i, \bar{\Pi}_{j-b}^i) = 1 \\ \mathbf{V}_{\text{DL}}(\text{pk}_{\text{DL}}, \mathcal{R}_j^i[\bar{w}_j^i, \mathcal{M}], \bar{\mathbf{h}}_{j-1}^i, \bar{\mathbf{h}}_j^i, t_j^i, \bar{\Pi}_j^i) = 1 \\ j=1 \rightarrow \bar{\mathbf{h}}_{j-1}^i = \mathbf{h}_{\text{start}}^i \\ j=n \rightarrow \bar{\mathbf{h}}_j^i = \mathbf{h}_{\text{accept}}^i \\ \mathbf{A}_j \wedge \hat{\mathbf{h}}_n^i \neq \mathbf{h}_{\text{accept}}^i \end{array} \right] \geq \frac{1}{4n \cdot P(T(\lambda))} .$$

Since $\mathbf{h}_{\text{start}}^i = \hat{\mathbf{h}}_0^i$ we have that:

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} \forall b \in \{0, 1\} : \mathbf{V}_{\text{TH}}(\text{pk}_{\text{TH}}, \mathbf{h}, \bar{\mathbf{h}}_{j-b}, j-b, \bar{\Pi}_{j-b}) = 1 \\ \forall b \in \{0, 1\} : \mathbf{V}_{\text{EH}}(\text{pk}'_{\text{EH}}, \bar{\mathbf{h}}_{j-b}, (\bar{w}_{j-b}^i, \bar{\mathbf{h}}_{j-b}^i), i, \bar{\Pi}_{j-b}^i) = 1 \\ \mathbf{V}_{\text{DL}}(\text{pk}_{\text{DL}}, \mathcal{R}_j^i[\bar{w}_j^i, \mathcal{M}], \bar{\mathbf{h}}_{j-1}^i, \bar{\mathbf{h}}_j^i, t_j^i, \bar{\Pi}_j^i) = 1 \\ j=1 \rightarrow \bar{\mathbf{h}}_{j-1}^i = \hat{\mathbf{h}}_0^i \\ j=n \rightarrow \bar{\mathbf{h}}_j^i \neq \hat{\mathbf{h}}_n^i \\ \mathbf{A}_j \end{array} \right] \geq \frac{1}{4n \cdot P(T(\lambda))} .$$

By the collision resistance requirement of the tree hash we have that there exists a negligible function μ such that for every $b \in \{0, 1\}$:

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} \mathbf{V}_{\text{TH}}(\text{pk}_{\text{TH}}, \mathbf{h}, \bar{\mathbf{h}}_{j-b}, j-b, \bar{\Pi}_{j-b}) = 1 \\ \bar{\mathbf{h}}_{j-b} \neq \mathbf{h}_{j-b} \end{array} \right] \leq \mu(T(\lambda)) .$$

Therefore:

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} \forall b \in \{0, 1\} : \mathbf{V}_{\text{EH}}(\text{pk}'_{\text{EH}}, \bar{\mathbf{h}}_{j-b}, (\bar{w}_{j-b}^i, \bar{\mathbf{h}}_{j-b}^i), i, \bar{\Pi}_{j-b}^i) = 1 \\ \mathbf{V}_{\text{DL}}(\text{pk}_{\text{DL}}, \mathcal{R}_j^i[\bar{w}_j^i, \mathcal{M}], \bar{\mathbf{h}}_{j-1}^i, \bar{\mathbf{h}}_j^i, t_j^i, \bar{\Pi}_j^i) = 1 \\ j=1 \rightarrow \bar{\mathbf{h}}_{j-1}^i = \hat{\mathbf{h}}_0^i \\ j=n \rightarrow \bar{\mathbf{h}}_j^i \neq \hat{\mathbf{h}}_n^i \\ \mathbf{A}_j \end{array} \right] \geq \frac{1}{5n \cdot P(T(\lambda))} .$$

By the extraction requirement of the somewhere extractable hash, we have that there exists a negligible function μ such that for every $b \in \{0, 1\}$:

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} \mathbf{V}_{\text{EH}}(\text{pk}'_{\text{EH}}, \bar{\mathbf{h}}_{j-b}, (\bar{w}_{j-b}^i, \bar{\mathbf{h}}_{j-b}^i), i, \bar{\Pi}_{j-b}^i) = 1 \\ (\bar{w}_{j-b}^i, \bar{\mathbf{h}}_{j-b}^i) \neq \mathbf{E}_{\text{EH}}(\text{sk}, \mathbf{h}_{j-b}) \end{array} \right] \leq \mu(T(\lambda)) .$$

Therefore, since we set $(w_{j-b}, \mathbf{h}_{j-b}^i) \leftarrow \mathbf{E}_{\text{EH}}(\text{sk}, \mathbf{h}_{j-b})$ we have that:

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} \mathbf{V}_{\text{DL}}(\text{pk}_{\text{DL}}, \mathcal{R}_j^i[w_j^i, \mathcal{M}], \mathbf{h}_{j-1}^i, \mathbf{h}_j^i, t_j^i, \bar{\Pi}_j^i) = 1 \\ j=1 \rightarrow \mathbf{h}_{j-1}^i = \hat{\mathbf{h}}_0^i \\ j=n \rightarrow \mathbf{h}_j^i \neq \hat{\mathbf{h}}_n^i \\ \mathbf{A}_j \end{array} \right] \geq \frac{1}{6n \cdot P(T(\lambda))} .$$

By the definition of the event A_j :

$$\Pr_{\text{EXP}'} \left[\begin{array}{l} V_{\text{DL}}(\text{pk}_{\text{DL}}, \mathcal{R}_j^i[w_j^i, \mathcal{M}], \mathbf{h}_{j-1}^i, \mathbf{h}_j^i, t_j^i, \tilde{\Pi}_j^i) = 1 \\ \mathbf{h}_{j-1}^i = \hat{\mathbf{h}}_{j-1}^i \wedge \mathbf{h}_j^i \neq \hat{\mathbf{h}}_j^i \end{array} \right] \geq \frac{1}{6n \cdot P(T(\lambda))} .$$

Let Adv' be the adversary that given as input the delegation public key pk_{DL} emulates the experiment EXP' using pk_{DL} and outputs:

$$(\mathcal{R}_j^i[w_j^i, \mathcal{M}], \hat{\mathbf{c}}\mathbf{f}_{j-1}^i, \hat{\mathbf{c}}\mathbf{f}_j^i, \mathbf{h}_{j-1}^i, \mathbf{h}_j^i, \tilde{\Pi}_j^i) .$$

We have that:

$$\Pr_{\substack{\text{pk}_{\text{DL}} \leftarrow \text{G}_{\text{DL}}(\lambda) \\ \text{Adv}'(\text{pk}_{\text{DL}})}} \left[\begin{array}{l} V(\text{pk}_{\text{DL}}, \mathcal{R}_j^i[w_j^i, \mathcal{M}], \mathbf{h}_{j-1}^i, \mathbf{h}_j^i, t_j^i, \tilde{\Pi}_j^i) = 1 \\ (\hat{\mathbf{c}}\mathbf{f}_{j-1}^i, \hat{\mathbf{c}}\mathbf{f}_j^i, t_j^i) \in \mathcal{T}_{\mathcal{R}_j^i[w_j^i, \mathcal{M}]} \\ \mathbf{h}_{j-1}^i = \hat{\mathbf{h}}_{j-1}^i = \text{H}_{\text{DL}}(\text{pk}_{\text{DL}}, \hat{\mathbf{c}}\mathbf{f}_{j-1}^i) \\ \mathbf{h}_j^i \neq \hat{\mathbf{h}}_j^i = \text{H}_{\text{DL}}(\text{pk}_{\text{DL}}, \hat{\mathbf{c}}\mathbf{f}_j^i) \end{array} \right] \geq \frac{1}{6n \cdot P(T(\lambda))} ,$$

contradicting the soundness requirement of the RAM delegation scheme.

Upgrading batch arguments without using rate-1 SHE. We note that we are crucially relying on the rate-1 property of the somewhere extractable hash. We, however, observe (as a result that will not be relevant in the sequel but may be of independent interest) that even if we just have a “plain” somewhere extractable hash, essentially the same construction but with a different parametrization can still be used to upgrade a batch argument to satisfy a stronger efficiency requirement (albeit not the rate-1 requirement). In particular, we now give a variant of Theorem 3.1 that uses any somewhere extractable hash (regardless of its additive overhead) to transform any batch argument into one with unbounded witness length and where the proof length is $m \cdot \text{poly}(\lambda)$ where m is the witness length (yet, with no guarantee on its additive overhead as in given by Theorem 3.1). This is in contrast to previous constructions of batch arguments [CJJ21] where the length of the public key grows with m and the proof length is $\text{poly}(\lambda, m)$.

Theorem 3.2. *Assume the existence of:*

- A $T(\cdot)$ -secure non-interactive batch argument for the index language.
- A $T(\cdot)$ -secure somewhere extractable hash.
- A $T(\cdot)$ -secure non-interactive delegation scheme for RAM.

Then there exists a $T(\cdot)$ -secure non-interactive batch argument for the index language with unbounded witness length and proof of length $m \cdot \text{poly}(\lambda)$.

The proof of Theorem 3.2 uses the same construction as in the proof of Theorem 3.1 except that we divide the witness into $n = m$ blocks, each of length $\ell = 1$.

4 Mergeable Delegation

A mergeable delegation scheme for RAM is a non-interactive delegation scheme for RAM that is equipped with an additional merging algorithm M . Given the proofs for two statements $(\mathbf{c}\mathbf{f}_1, \mathbf{c}\mathbf{f}_2, t_1), (\mathbf{c}\mathbf{f}_2, \mathbf{c}\mathbf{f}_3, t_2) \in \mathcal{T}_{\mathcal{R}}$ the merging algorithm efficiently generates a proof for the combined statement $(\mathbf{c}\mathbf{f}_1, \mathbf{c}\mathbf{f}_3, t_1 + t_2) \in \mathcal{T}_{\mathcal{R}}$. What makes this notion non-trivial is the *compactness* requirement which states that the merged proof is not much longer than (the longer one of) the original proofs. In essence, merging compresses two proofs into a single proof of roughly the same size.

Our notion of mergeable delegation allows proofs to be merged recursively. In more detail, for each proof we maintain a level, where newly generated proofs are of level 0 and merging two proofs of level i results in

a proof of level $i + 1$. Similarly to the notion of levelled-FHE, We restrict attention to schemes that support an a-priori bounded number of levels, say λ .

Formally, a mergeable non-interactive delegation scheme for RAM is a non-interactive delegation scheme for RAM (G, H, P, V) augmented with algorithms (M, L) with the following syntax:

M: The deterministic merging algorithm takes as input a public key pk , a RAM machine \mathcal{R} and a pair of hashed statements and proofs $((h_i, h'_i, t_i), \Pi_i)_{i \in [2]}$. It outputs a new proof Π .

L: The deterministic level algorithm takes as input a proof Π , it outputs a level $\ell \in \mathbb{N}$.

Definition 4.1. A $T(\cdot)$ -secure non-interactive delegation scheme for RAM is mergeable if it satisfies the following requirements:

Completeness with Level 0. For every $\lambda \in \mathbb{N}$, $t \leq 2^\lambda$, RAM machine \mathcal{R} and statement $(\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}}$ we have that:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow G(\lambda) \\ \text{h} \leftarrow H(\text{pk}, \text{cf}) \\ \text{h}' \leftarrow H(\text{pk}, \text{cf}') \\ \Pi \leftarrow P(\text{pk}, \mathcal{R}, (\text{cf}, \text{cf}', t)) \end{array} : \begin{array}{l} V(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ L(\Pi) = 0 \end{array} \right] = 1 .$$

Completeness of Merge. For every $\lambda \in \mathbb{N}$, public key pk , $\ell < \lambda$, RAM machine \mathcal{R} and a pair of hashed statements and proofs $(x_i = (h_i, h'_i, t_i), \Pi_i)_{i \in [2]}$ such that $h'_1 = h_2$ and for every $i \in [2]$, $V(\text{pk}, \mathcal{R}, x_i, \Pi_i) = 1$ and $L(\Pi_i) = \ell$ we have that:

$$\Pr \left[\Pi \leftarrow M(\text{pk}, \mathcal{R}, (x_i, \Pi_i)_{i \in [2]}) : \begin{array}{l} V(\text{pk}, \mathcal{R}, (\text{h}_1, \text{h}'_2, t_1 + t_2), \Pi) = 1 \\ L(\Pi) = \ell + 1 \end{array} \right] = 1 .$$

Efficiency. In the completeness experiments above:

- The merging algorithm runs in time $\text{poly}(\lambda, |\mathcal{R}|)$ and outputs a proof of length $\text{poly}(\lambda)$.
- The level algorithm runs in time $\text{poly}(\lambda)$.

$T(\cdot)$ -Soundness for Bounded Level Proofs. For every $\text{poly}(T(\lambda))$ -size adversary Adv there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$:

$$\Pr \left[\begin{array}{l} \text{pk} \leftarrow G(\lambda) \\ (\mathcal{R}, \text{cf}, \text{cf}', \text{h}, \text{h}', t, \Pi) \leftarrow \text{Adv}(\text{pk}) \end{array} : \begin{array}{l} t \leq T(\lambda) \\ L(\Pi) \leq \log T(\lambda) \\ V(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = H(\text{pk}, \text{cf}) \\ \text{h}' \neq H(\text{pk}, \text{cf}') \end{array} \right] \leq \mu(T(\lambda)) .$$

We say that the scheme is polynomially secure if it is $T(\cdot)$ -secure for every polynomial T .

We make some remarks on Definition 4.1:

- The completeness with level 0 requirement is exactly the same as the plain completeness requirement of delegation except that we require that freshly created proofs are of level 0.
- The soundness for bounded levels requirement is exactly the same as the plain soundness requirement except that we restrict the level of the proof to be $O(\log T(\lambda))$.
- We only require completeness of merge for a pair of proofs at the same level and up to level λ . Therefore, constructing a proof of level ℓ requires merging 2^ℓ proofs in a full binary tree (some of these proof may be for 0 computation steps).

Theorem 4.2. *Assume the existence of:*

- A $T(\cdot)$ -secure non-interactive batch arguments for the index language with unbounded witness length and additive overhead $\alpha(\lambda, m) = \frac{O(m)}{\lambda} + \text{poly}(\lambda)$.
- A $T(\cdot)$ -secure non-interactive delegation scheme for RAM.

Then there exists a $T(\cdot)$ -secure mergeable non-interactive delegation scheme for RAM.

Construction. We construct a mergeable non-interactive delegation scheme $(G'_{\text{DL}}, H'_{\text{DL}}, P'_{\text{DL}}, V'_{\text{DL}}, M'_{\text{DL}}, L'_{\text{DL}})$ using the following building blocks:

- A $T(\cdot)$ -secure non-interactive batch argument for the index language $(G_{\text{BA}}, T_{\text{BA}}, P_{\text{BA}}, V_{\text{BA}}, E_{\text{BA}})$ with unbounded witness length and additive overhead $\alpha(\lambda, m) = \frac{O(m)}{\lambda} + \text{poly}(\lambda)$.
- A $T(\cdot)$ -secure non-interactive delegation scheme for RAM $(G_{\text{DL}}, H_{\text{DL}}, P_{\text{DL}}, V_{\text{DL}})$.

We proceed with a formal description of the construction. We start by defining a verification procedure for proofs of each level. The level-0 verification algorithm V^0 is simply the verifier V_{DL} of the underlying delegation scheme. For $\ell \in [\lambda]$, the level- ℓ verification algorithm V^ℓ is given as input the public key $\text{pk} = (\text{pk}_0, \dots, \text{pk}_\ell)$, a RAM machine \mathcal{R} , a hashed statement $x = (\mathbf{h}, \mathbf{h}', t)$ and a proof Π . It proceeds as follows:

- Let $\Pi = (x_1 = (\mathbf{h}_1, \mathbf{h}'_1, t_1), x_2 = (\mathbf{h}_2, \mathbf{h}'_2, t_2), \Pi')$.
- If $\mathbf{h}'_1 \neq \mathbf{h}_2$ or $(\mathbf{h}_1, \mathbf{h}'_2, t_1 + t_2) \neq (\mathbf{h}, \mathbf{h}', t)$ then reject.
- Let $\mathcal{M}^\ell = \mathcal{M}^\ell[\text{pk}, \mathcal{R}, (x_1, x_2)]$ be the machine that on $(i, \tilde{\Pi})$ outputs $V^{\ell-1}((\text{pk}_0, \dots, \text{pk}_{\ell-1}), \mathcal{R}, x_i, \tilde{\Pi})$.
- Output $V_{\text{BA}}(\text{pk}_\ell, \mathcal{M}^\ell, \Pi')$

The delegation scheme algorithms.

G'_{DL} : Given as input a security parameter $\lambda \in \mathbb{N}$, the setup algorithm is as follows:

- Set $\text{pk}_0 \leftarrow G_{\text{DL}}(\lambda)$.
- For $\ell \in [\lambda]$, set $\text{pk}_\ell \leftarrow G_{\text{BA}}(\lambda)$.
- Output $\text{pk} = (\text{pk}_0, \dots, \text{pk}_\lambda)$.

H'_{DL} : Given as input the public key $\text{pk} = (\text{pk}_0, \dots, \text{pk}_\lambda)$ and a configuration $\text{cf} \in \{0, 1\}^*$, the digest algorithm outputs $H_{\text{DL}}(\text{pk}_0, \text{cf})$.

P'_{DL} : Given as input the public key pk , a RAM machine \mathcal{R} , and a statement $(\text{cf}, \text{cf}', t)$ the prover algorithm is as follows:

- Let $\text{pk} = (\text{pk}_0, \dots, \text{pk}_\lambda)$.
- Set $\Pi \leftarrow P_{\text{DL}}(\text{pk}_0, \mathcal{R}, (\text{cf}, \text{cf}', t))$.
- Output $(0, \Pi)$

V'_{DL} : Given as input the public key $\text{pk} = (\text{pk}_0, \dots, \text{pk}_\lambda)$, a RAM machine \mathcal{R} , a hashed statement x and a proof $\Pi = (\ell, \Pi')$ the verifier algorithm outputs $V^\ell((\text{pk}_0, \dots, \text{pk}_\ell), \mathcal{R}, x, \Pi')$.

M'_{DL} : Given as input a public key pk , a RAM machine \mathcal{R} and a pair of hashed statements and proofs $(x_i = (\mathbf{h}_i, \mathbf{h}'_i, t_i), \Pi_i)_{i \in [2]}$, the merging algorithm is as follows:

- Let $\text{pk} = (\text{pk}_0, \dots, \text{pk}_\lambda)$.
- For $i \in [2]$, let $\Pi_i = (\ell_i, \Pi'_i)$.
- If $\ell_1 \neq \ell_2$ then reject.

- Set $\ell = \ell_1 + 1$
- Let $\mathcal{M}^\ell = \mathcal{M}^\ell[\text{pk}, \mathcal{R}, (x_1, x_2)]$ be the machine that on $(i, \tilde{\Pi})$ outputs $\mathbf{V}^{\ell-1}((\text{pk}_0, \dots, \text{pk}_{\ell-1}), \mathcal{R}, x_i, \tilde{\Pi})$.
- Set $\Pi \leftarrow \text{P}_{\text{BA}}(\text{pk}_\ell, \mathcal{M}^\ell, (\Pi'_1, \Pi'_2))$
- Output $(\ell, (x_1, x_2, \Pi))$

L'_{DL} : Given as input a proof $\Pi = (\ell, \Pi')$, the level algorithm outputs ℓ .

Analysis. The completeness requirements follow directly from the completeness requirement of the building blocks. The collision resistance requirement follows directly from the collision resistance requirement of the underlying delegation scheme. We focus on proving that the efficiency and soundness requirements.

Efficiency. The efficiency of the level algorithm follows directly from the construction. The efficiency of the generation, digest and prover algorithms follow directly from the efficiency requirement of the underlying delegation scheme. It remain to bound the running time of the verifier and merging algorithms and length of the merged proofs.

Let m_ℓ denote the length of level- ℓ proofs. By the efficiency requirement of the underlying delegation scheme, we have that $m_0 = \text{poly}(\lambda)$. Recall that to create level- ℓ proofs, the merging algorithms creates a batch argument proof with witness of length $m_{\ell-1}$ and appends to it the hashed statements that are of length $\text{poly}(\lambda)$. Since the batch argument has additive overhead $\alpha(\lambda, m) = \frac{O(m)}{\lambda} + \text{poly}(\lambda)$, we have that $m_\ell \leq m_{\ell-1} + \frac{O(m_{\ell-1})}{\lambda} + \text{poly}(\lambda)$. Overall, we have that $m_\ell \leq \ell \cdot \text{poly}(\lambda) \cdot (1 + \frac{1}{\lambda})^{O(\ell)}$. Therefore, for every $\ell \leq \lambda$, we have that $m_\ell = \text{poly}(\lambda)$.

The verifier algorithm runs the algorithm \mathbf{V}^ℓ for some level $\ell \leq \lambda$, and \mathbf{V}^ℓ runs the batch argument verifier algorithm V_{BA} with the machine \mathcal{M}^ℓ . The machine \mathcal{M}^ℓ is of size $|\mathcal{R}| + \text{poly}(\lambda)$ and it is taking as witnesses level- $(\ell - 1)$ merged proofs of length $m_{\ell-1} = \text{poly}(\lambda)$. Therefore, the running time of verifier algorithm is $|\mathcal{R}| \cdot \text{poly}(\lambda)$.

The merging algorithm runs the batch argument prover algorithm P_{BA} with the machine \mathcal{M}^ℓ . The machine \mathcal{M}^ℓ is of size $|\mathcal{R}| + \text{poly}(\lambda)$ and it is taking as witnesses level- $(\ell - 1)$ merged proofs of length $m_{\ell-1} = \text{poly}(\lambda)$. The machine \mathcal{M}^ℓ runs the algorithm $\mathbf{V}^{\ell-1}$ and therefore it runs in time $|\mathcal{R}| \cdot \text{poly}(\lambda)$. Therefore, the running time of merging algorithm is $\text{poly}(\lambda, |\mathcal{R}|)$.

Soundness for bounded level proofs. For an adversary Adv , Let EXP_{Adv} be the experiment of the soundness requirement:

$$\left\{ \begin{array}{l} \text{pk} = (\text{pk}_0, \dots, \text{pk}_\lambda) \leftarrow \text{G}'_{\text{DL}}(\lambda) \\ (\mathcal{R}, \text{cf}, \text{cf}', \text{h}, \text{h}', t, \Pi) \leftarrow \text{Adv}(\text{pk}) \end{array} \right\}$$

The main step of the proof is given by the following lemma that transforms a cheating prover for level $\ell > 0$ into a cheating prover for level $\ell - 1$.

Claim 4.3. *For every $\text{poly}(T(\lambda))$ -size adversary Adv , there exists an adversary Adv' of size $|\text{Adv}| + \text{poly}(T(\lambda))$ and a negligible function μ such that for every $\lambda \in \mathbb{N}$ and $\ell \in [\lambda]$:*

$$\Pr_{\text{EXP}_{\text{Adv}'}} \left[\begin{array}{l} t \leq T(\lambda) \\ \text{L}'_{\text{DL}}(\Pi) = \ell - 1 \\ \text{V}'_{\text{DL}}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = \text{H}'_{\text{DL}}(\text{pk}, \text{cf}) \\ \text{h}' \neq \text{H}'_{\text{DL}}(\text{pk}, \text{cf}') \end{array} \right] \geq \frac{1}{2} \cdot \Pr_{\text{EXP}_{\text{Adv}}} \left[\begin{array}{l} t \leq T(\lambda) \\ \text{L}'_{\text{DL}}(\Pi) = \ell \\ \text{V}'_{\text{DL}}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = \text{H}'_{\text{DL}}(\text{pk}, \text{cf}) \\ \text{h}' \neq \text{H}'_{\text{DL}}(\text{pk}, \text{cf}') \end{array} \right] - \mu(T(\lambda)) .$$

Before proving Claim 4.3, we use it to complete the soundness proof. Fix an $\text{poly}(T(\lambda))$ -size adversary Adv and let $d = \log T(\lambda)$. In what follows we construct a sequence of cheating prover for decreasing levels from d down to 0. Let $\text{Adv}_0 = \text{Adv}$ and for every $i \in [d]$ let Adv_i be the adversary given by Claim 4.3 for

Adv_{i-1} . Let $\bar{T}(\lambda) = \text{poly}(T(\lambda))$ denote the size of Adv_d . By Claim 4.3, for every adversary Adv^* of size at most \bar{T} there exists a negligible function μ^* for which the claim holds. By considering the adversary with maximal advantage, there exists a single negligible function μ such that Claim 4.3 holds for every adversary of size at most \bar{T} and, in particular, for every Adv_i .

Assume toward contradiction that there exists a polynomial P such that for infinitely many $\lambda \in \mathbb{N}$:

$$\Pr_{\text{EXP}_{\text{Adv}}} \left[\begin{array}{l} t \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) \leq d \\ V'_{\text{DL}}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = H'_{\text{DL}}(\text{pk}, \text{cf}) \\ \text{h}' \neq H'_{\text{DL}}(\text{pk}, \text{cf}') \end{array} \right] \geq \frac{1}{P(T(\lambda))} .$$

Therefore, for infinitely many $\lambda \in \mathbb{N}$ there exists a level $\ell \in [0, d]$ such that:

$$\Pr_{\text{EXP}_{\text{Adv}}} \left[\begin{array}{l} t \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) = \ell \\ V'_{\text{DL}}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = H'_{\text{DL}}(\text{pk}, \text{cf}) \\ \text{h}' \neq H'_{\text{DL}}(\text{pk}, \text{cf}') \end{array} \right] \geq \frac{1}{d \cdot P(T(\lambda))} .$$

By ℓ applications of Claim 4.3, for infinitely many $\lambda \in \mathbb{N}$:

$$\Pr_{\text{EXP}_{\text{Adv}_\ell}} \left[\begin{array}{l} t \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) = 0 \\ V'_{\text{DL}}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = H'_{\text{DL}}(\text{pk}, \text{cf}) \\ \text{h}' \neq H'_{\text{DL}}(\text{pk}, \text{cf}') \end{array} \right] \geq \frac{1}{T(\lambda) \cdot d \cdot P(T(\lambda))} - d \cdot \mu(T(\lambda)) .$$

Recall that a proof Π of level 0 is simply a proof of the underlying delegation scheme under pk_0 . Since given pk_0 , we can sample the rest of the public key $\text{pk}_1, \dots, \text{pk}_\lambda$ ourselves, we reach a contradiction to the soundness of the underlying delegation scheme.

Proof of Claim 4.3. Fix a $\text{poly}(T(\lambda))$ -size adversary Adv . For every λ and $\ell \in [\lambda]$ let ϵ denote the probability:

$$\epsilon = \Pr_{\text{EXP}_{\text{Adv}}} \left[\begin{array}{l} t \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) = \ell \\ V'_{\text{DL}}(\text{pk}, \mathcal{R}, (\text{h}, \text{h}', t), \Pi) = 1 \\ (\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{R}} \\ \text{h} = H'_{\text{DL}}(\text{pk}, \text{cf}) \\ \text{h}' \neq H'_{\text{DL}}(\text{pk}, \text{cf}') \end{array} \right] .$$

If $L'_{\text{DL}}(\Pi) = \ell > 0$ and V'_{DL} accepts then by the construction of V^ℓ :

- Π is of the form $(\ell, (x_1, x_2, \Pi'))$ where $x_i = (\text{h}_i, \text{h}'_i, t_i)$.
- $\text{h}'_1 = \text{h}_2$ and $(\text{h}_1, \text{h}'_2, t_1 + t_2) = (\text{h}, \text{h}', t)$
- $V_{\text{BA}}(\text{pk}_\ell, \mathcal{M}^\ell[\text{pk}, \mathcal{R}, (x_1, x_2)], \Pi') = 1$

Let cf'' be the unique configuration such that $(\text{cf}, \text{cf}'', t_1) \in \mathcal{T}_{\mathcal{R}}$ and $(\text{cf}'', \text{cf}', t_2) \in \mathcal{T}_{\mathcal{R}}$. Note that cf'' can be computed from cf in time $\text{poly}(T(\lambda))$. Let $\text{h}'' = H'_{\text{DL}}(\text{pk}, \text{cf}'')$ and let $(\tilde{\text{cf}}_1, \tilde{\text{cf}}'_1) = (\text{cf}, \text{cf}'')$ and

$(\tilde{cf}_2, \tilde{cf}'_2) = (cf'', cf')$. Whether $h'' = h'_1$ or $h'' \neq h'_1$, we have that:

$$\Pr_{\text{EXP}_{\text{Adv}}} \left[\begin{array}{l} t \leq T(\cdot) \\ L'_{\text{DL}}(\Pi) = \ell \\ \exists i \in [2] : \\ (\tilde{cf}_i, \tilde{cf}'_i, t_i) \in \mathcal{T}_{\mathcal{R}} \\ h_i = H'_{\text{DL}}(\text{pk}, \tilde{cf}_i) \\ h'_i \neq H'_{\text{DL}}(\text{pk}, \tilde{cf}'_i) \\ V_{\text{BA}}(\text{pk}_\ell, \mathcal{M}^\ell[\text{pk}, \mathcal{R}, (x_1, x_2)], \Pi') = 1 \end{array} \right] \geq \epsilon .$$

Fix $i \in [2]$ such that:

$$\Pr_{\text{EXP}_{\text{Adv}}} \left[\begin{array}{l} t_i \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) = \ell \\ (\tilde{cf}_i, \tilde{cf}'_i, t_i) \in \mathcal{T}_{\mathcal{R}} \\ h_i = H'_{\text{DL}}(\text{pk}, \tilde{cf}_i) \\ h'_i \neq H'_{\text{DL}}(\text{pk}, \tilde{cf}'_i) \\ V_{\text{BA}}(\text{pk}_\ell, \mathcal{M}^\ell[\text{pk}, \mathcal{R}, (x_1, x_2)], \Pi') = 1 \end{array} \right] \geq \frac{\epsilon}{2} .$$

Let EXP'_{Adv} be the experiment that is defined like EXP_{Adv} except that we use the trapdoor generation algorithm of the batch argument to sample pk_ℓ and extract a witness Π_i from the proof:

- $(\text{pk}_\ell, \text{sk}_\ell) \leftarrow G_{\text{BA}}(\lambda, i)$
- Set $\Pi_i \leftarrow E_{\text{BA}}(\text{sk}_\ell, \Pi')$.

By the key indistinguishability requirement of the batch argument, we have that:

$$\Pr_{\text{EXP}'_{\text{Adv}}} \left[\begin{array}{l} t_i \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) = \ell \\ (\tilde{cf}_i, \tilde{cf}'_i, t_i) \in \mathcal{T}_{\mathcal{R}} \\ h_i = H'_{\text{DL}}(\text{pk}, \tilde{cf}_i) \\ h'_i \neq H'_{\text{DL}}(\text{pk}, \tilde{cf}'_i) \\ V_{\text{BA}}(\text{pk}_\ell, \mathcal{M}^\ell[\text{pk}, \mathcal{R}, (x_1, x_2)], \Pi') = 1 \end{array} \right] \geq \frac{\epsilon}{2} - \text{negl}(T(\lambda)) .$$

By the somewhere argument of knowledge requirement of the batch argument we have that:

$$\Pr_{\text{EXP}'_{\text{Adv}}} \left[\begin{array}{l} t_i \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) = \ell \\ (\tilde{cf}_i, \tilde{cf}'_i, t_i) \in \mathcal{T}_{\mathcal{R}} \\ h_i = H'_{\text{DL}}(\text{pk}, \tilde{cf}_i) \\ h'_i \neq H'_{\text{DL}}(\text{pk}, \tilde{cf}'_i) \\ \mathcal{M}^\ell[\text{pk}, \mathcal{R}, (x_1, x_2)](i, \Pi_i) = 1 \end{array} \right] \geq \frac{\epsilon}{2} - \text{negl}(T(\lambda)) .$$

By the definition of \mathcal{M}^ℓ and V'_{DL} :

$$\Pr_{\text{EXP}'_{\text{Adv}}} \left[\begin{array}{l} t_i \leq T(\lambda) \\ L'_{\text{DL}}(\Pi) = \ell \\ (\tilde{cf}_i, \tilde{cf}'_i, t_i) \in \mathcal{T}_{\mathcal{R}} \\ h_i = H'_{\text{DL}}(\text{pk}, \tilde{cf}_i) \\ h'_i \neq H'_{\text{DL}}(\text{pk}, \tilde{cf}'_i) \\ V'_{\text{DL}}(\text{pk}, \mathcal{R}, x_i, (\ell - 1, \Pi_i)) = V^{\ell-1}(\text{pk}, \mathcal{R}, x_i, \Pi_i) = 1 \end{array} \right] \geq \frac{\epsilon}{2} - \text{negl}(T(\lambda)) .$$

Let Adv' be the adversary that, given $\text{pk} = (\text{pk}_0, \dots, \text{pk}_\lambda)$, emulates the experiment EXP' using the public key:

$$\text{pk}' = (\text{pk}_0, \dots, \text{pk}_{\ell-1}, \text{pk}'_{\ell}, \dots, \text{pk}'_\lambda)$$

where $\text{pk}'_{\ell}, \dots, \text{pk}'_\lambda$ are sampled as in EXP' . Adv' outputs $(\mathcal{R}, \tilde{\text{cf}}_i, \tilde{\text{cf}}'_i, h_i, h'_i, t_i, \Pi = (\ell - 1, \Pi_i))$. We have that:

$$\Pr_{\text{EXP}'_{\text{Adv}'}} \left[\begin{array}{l} t_i \leq T(\lambda) \\ \text{L}'_{\text{DL}}(\Pi) = \ell - 1 \\ \text{V}'_{\text{DL}}(\text{pk}', \mathcal{R}, x_i = (h_i, h'_i, t_i), \Pi) = 1 \\ (\tilde{\text{cf}}_i, \tilde{\text{cf}}'_i, t_i) \in \mathcal{T}_{\mathcal{R}} \\ h_i = \text{H}'_{\text{DL}}(\text{pk}, \tilde{\text{cf}}_i) \\ h'_i \neq \text{H}'_{\text{DL}}(\text{pk}, \tilde{\text{cf}}'_i) \end{array} \right] \geq \frac{\epsilon}{2} - \text{negl}(T(\lambda)) .$$

Since pk and pk' agree on the first ℓ public keys $(\text{pk}'_0, \dots, \text{pk}'_{\ell-1})$ and since $\text{V}^{\ell-1}$ only takes the first ℓ public keys, we have that:

$$\Pr_{\text{EXP}'_{\text{Adv}'}} \left[\begin{array}{l} t_i \leq T(\lambda) \\ \text{L}'_{\text{DL}}(\Pi) = \ell - 1 \\ \text{V}'_{\text{DL}}(\text{pk}, \mathcal{R}, (h_i, h'_i, t_i), \Pi) = 1 \\ (\tilde{\text{cf}}_i, \tilde{\text{cf}}'_i, t_i) \in \mathcal{T}_{\mathcal{R}} \\ h_i = \text{H}'_{\text{DL}}(\text{pk}', \tilde{\text{cf}}_i) \\ h'_i \neq \text{H}'_{\text{DL}}(\text{pk}', \tilde{\text{cf}}'_i) \end{array} \right] \geq \frac{\epsilon}{2} - \text{negl}(T(\lambda)) ,$$

as required. \square

4.1 Incrementally Verifiable Computation

In this section, we use mergeable delegation to construct incrementally verifiable computation. The construction follows the outline in [Val08].

Theorem 4.4. *If there exists a $T(\cdot)$ -secure mergeable non-interactive delegation scheme for RAM then there exists a $T(\cdot)$ -secure incrementally verifiable computation scheme for deterministic Turing machines.*

Construction. We construct an incrementally verifiable computation scheme $(\text{G}_{\text{IVC}}, \text{P}_{\text{IVC}}, \text{V}_{\text{IVC}}, \text{U}_{\text{IVC}})$ using a mergeable non-interactive delegation scheme for RAM $(\text{G}_{\text{DL}}, \text{H}_{\text{DL}}, \text{P}_{\text{DL}}, \text{V}_{\text{DL}}, \text{M}_{\text{DL}}, \text{L}_{\text{DL}})$. The outline of the construction is as follows. Given a statement $(\text{cf}, \text{cf}', t) \in \mathcal{T}_{\mathcal{M}}$, we split it into $d \leq \log t$ consecutive statements $\{(\text{cf}_i, \text{cf}_{i+1}, t_i = 2^{\ell_i}) \in \mathcal{T}_{\mathcal{M}}\}_{i \in [d-1]}$ such that $(\text{cf}_1, \text{cf}_d) = (\text{cf}, \text{cf}')$ and for every i , $\ell_i > \ell_{i+1}$. The proof for $(\text{cf}, \text{cf}', t)$ consist of one mergeable delegation proof for each of the d segments, where the proof for the i -th segment is of level ℓ_i . To increment such proofs by one step, we create a level-0 proof for the statement $(\text{cf}', \text{cf}'', 1) \in \mathcal{T}_{\mathcal{M}}$ and append it to the proof. We then repeatedly merge pairs of proofs of the same level until we get a valid proof.

We proceed with a formal description of the construction. For a Turing machine \mathcal{M} , let $\mathcal{R}[\mathcal{M}]$ denote the RAM machine that emulates \mathcal{M} . That is, $\mathcal{T}_{\mathcal{M}} = \mathcal{T}_{\mathcal{R}[\mathcal{M}]}$ (we assume for simplicity and without loss of generality that each step of \mathcal{M} is emulated by a single step of $\mathcal{R}[\mathcal{M}]$).

G_{IVC} : Given as input a security parameter $\lambda \in \mathbb{N}$, the setup algorithm samples $\text{pk} \leftarrow \text{G}_{\text{DL}}(\lambda)$ and outputs pk .

P_{IVC} : Given as input the public key pk , a machine \mathcal{M} , and a statement $(\text{cf}, \text{cf}', t)$, the prover algorithm is as follows:

- Set $\Pi_0 = \emptyset$.
- Set $\text{cf}_0 = \text{cf}$.
- For $i = 1$ to t :
 - Set $\Pi_i \leftarrow \text{U}_{\text{IVC}}(\text{pk}, \mathcal{M}, (\text{cf}, \text{cf}_{i-1}, i - 1), \Pi_{i-1})$.

- Let cf_i be the configuration such that $(\text{cf}_{i-1}, \text{cf}_i, 1) \in \mathcal{T}_{\mathcal{M}}$.
- Output Π_t .

V_{IVC} : Given as input the public key pk , a machine \mathcal{M} , a statement $(\text{cf}, \text{cf}', t)$ and a proof Π , the verifier algorithm is as follows:

- If $\text{cf} = \text{cf}'$, $t = 0$ and $\Pi = \emptyset$, then accept.
- Let $\ell_1, \dots, \ell_d \in [0, \lambda]$ be the distinct indexes such that $t = \sum_{i \in [d]} 2^{\ell_i}$ in decreasing order (i.e., the indexes of the 1s in the λ -ary decomposition of t).
- Let $\Pi = \{(x_i = (\mathbf{h}_i, \mathbf{h}'_i, t_i), \Pi_i)\}_{i \in [d]}$ and reject if Π is not of this form.
- Check that $(\text{H}_{\text{DL}}(\text{pk}, \text{cf}), \text{H}_{\text{DL}}(\text{pk}, \text{cf}')) = (\mathbf{h}_1, \mathbf{h}'_d)$.
- For every $i \in [d-1]$, check that $\mathbf{h}'_i = \mathbf{h}_{i+1}$.
- For every $i \in [d]$, check that $t_i = 2^{\ell_i}$.
- For every $i \in [d]$, check that $\text{L}_{\text{DL}}(\Pi_i) = \ell_i$.
- For every $i \in [d]$, check that $\text{V}(\text{pk}, \mathcal{R}[\mathcal{M}], (\mathbf{h}_i, \mathbf{h}'_i, t_i), \Pi_i)$ accepts.
- Accept if and only if all checks pass.

U_{IVC} : Given as input a public key pk , a machine \mathcal{M} , a statement $(\text{cf}, \text{cf}', t)$, and a proof Π , the update algorithm is as follows:

- Let cf'' be the configuration such that $(\text{cf}', \text{cf}'', 1) \in \mathcal{T}_{\mathcal{M}}$.
- Set $x^* \leftarrow (\text{H}_{\text{DL}}(\text{pk}, \text{cf}'), \text{H}_{\text{DL}}(\text{pk}, \text{cf}''), 1)$.
- Set $\Pi^* \leftarrow \text{P}(\text{pk}, \mathcal{R}, x^*)$.
- Add (x^*, Π^*) to Π .
- While there exist $\{(x_i = (\mathbf{h}_i, \mathbf{h}'_i, t_i), \Pi_i)\}_{i \in [2]} \subseteq \Pi$ such that $\mathbf{h}'_1 = \mathbf{h}_2$ and $t_1 = t_2$:
 - Remove (x_1, Π_1) and (x_2, Π_2) from Π .
 - Set $x^* \leftarrow (\mathbf{h}_1, \mathbf{h}'_2, t_1 + t_2)$.
 - Set $\Pi^* \leftarrow \text{M}(\text{pk}, \mathcal{R}, (x_i, \Pi_i)_{i \in [2]})$.
 - Add (x^*, Π^*) to Π .
- Output Π .

Analysis. The incremental completeness and efficiency requirements follow by construction and by the completeness and efficiency requirements of the underlying mergeable delegation scheme. We focus on proving that the soundness requirement.

Fix a $\text{poly}(T(\lambda))$ -size adversary Adv . Let EXP_{Adv} be the experiment of the soundness requirement: For $\lambda \in \mathbb{N}$, let EXP be the experiment of the soundness requirement:

$$\left\{ \begin{array}{l} \text{pk} \leftarrow \text{G}_{\text{IVC}}(\lambda) \\ (\mathcal{M}, x = (\text{cf}, \text{cf}', t), \Pi) \leftarrow \text{Adv}(\text{pk}) \end{array} \right\}$$

Assume towards contradiction that there exists a polynomial P such that for infinitely many $\lambda \in \mathbb{N}$:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} t \leq T(\lambda) \\ \text{V}_{\text{IVC}}(\text{pk}, \mathcal{M}, x, \Pi) = 1 \\ x \notin \mathcal{T}_{\mathcal{M}} \end{array} \right] \geq \frac{1}{P(T(\lambda))} .$$

Fix any sufficiently large λ such that the above holds. In the above experiment, let $\ell_1, \dots, \ell_d \in [0, \lambda]$ denote the distinct indexes such that $t = \sum_{i \in [d]} 2^{\ell_i}$ in decreasing order. Let $\text{cf}_0 = \text{cf}$ and for $i \in [d]$, let cf_i be

the configuration such that $(\bar{c}_{i-1}, \bar{c}_i, \ell_i) \in \mathcal{T}_{\mathcal{M}}$. Note that $\bar{c}_0, \dots, \bar{c}_d$ can be computed in time $\text{poly}(T(\lambda))$ given \mathcal{M} and cf . Since $(\bar{c}_0, \bar{c}_d, t) \in \mathcal{T}_{\mathcal{M}}$, it follows that if $x \notin \mathcal{T}_{\mathcal{M}}$ then $\bar{c}_d \neq \text{cf}'$. That is:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} t \leq T(\lambda) \\ \text{V}_{\text{IVC}}(\text{pk}, \mathcal{M}, x, \Pi) = 1 \\ \bar{c}_d \neq \text{cf}' \end{array} \right] \geq \frac{1}{P(T(\lambda))} .$$

For $i \in [0, d]$, let $\bar{h}_i = \text{H}_{\text{DL}}(\text{pk}, \text{cf}_i)$. By the collision resistance requirement of the delegation scheme, we have that, for sufficiently large λ ,

$$\Pr_{\text{EXP}} \left[\begin{array}{l} t \leq T(\lambda) \\ \text{V}_{\text{IVC}}(\text{pk}, \mathcal{M}, x, \Pi) = 1 \\ \bar{h}_d \neq \text{H}_{\text{DL}}(\text{pk}, \text{cf}') \end{array} \right] \geq \frac{1}{2 \cdot P(T(\lambda))} .$$

If $\text{V}_{\text{IVC}}(\text{pk}, \mathcal{M}, x, \Pi)$ accepts, then Π is of the form $\{(x_i = (h_i, h'_i, 2^{\ell_i}), \Pi_i)\}_{i \in [d]}$ and:

- $(\text{H}_{\text{DL}}(\text{pk}, \text{cf}), \text{H}_{\text{DL}}(\text{pk}, \text{cf}')) = (h_1, h'_d)$.
- $\forall i \in [d-1] : h'_i = h_{i+1}$.
- $\forall i \in [d] : \text{L}_{\text{DL}}(\Pi_i) = \ell_i$.
- $\forall i \in [d] : \text{V}_{\text{DL}}(\text{pk}, \mathcal{R}[\mathcal{M}], (h_i, h'_i, 2^{\ell_i}), \Pi_i) = 1$.

$$\Pr_{\text{EXP}} \left[\begin{array}{l} t \leq T(\lambda) \\ \bar{h}_0 = h_1 \\ \bar{h}_d \neq h'_d \\ \forall i \in [d-1] : h'_i = h_{i+1} \\ \forall i \in [d] : \text{L}_{\text{DL}}(\Pi_i) = \ell_i \\ \forall i \in [d] : \text{V}_{\text{DL}}(\text{pk}, \mathcal{R}[\mathcal{M}], (h_i, h'_i, 2^{\ell_i}), \Pi_i) = 1 \end{array} \right] \geq \frac{1}{2 \cdot P(T(\lambda))} .$$

Denoting $h'_0 = h_1$, we have that:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} t \leq T(\lambda) \\ \bar{h}_0 = h'_0 \\ \bar{h}_d \neq h'_d \\ \forall i \in [d] : h'_{i-1} = h_i \\ \forall i \in [d] : \text{L}_{\text{DL}}(\Pi_i) = \ell_i \\ \forall i \in [d] : \text{V}_{\text{DL}}(\text{pk}, \mathcal{R}[\mathcal{M}], (h_i, h'_i, 2^{\ell_i}), \Pi_i) = 1 \end{array} \right] \geq \frac{1}{2 \cdot P(T(\lambda))} .$$

If $\bar{h}_0 = h'_0$ but $\bar{h}_d \neq h'_d$, then there must exist some $i \in [d]$ such that $\bar{h}_{i-1} = h'_{i-1}$ but $\bar{h}_i \neq h'_i$. Therefore, there exists $i \in [d]$ such that:

$$\Pr_{\text{EXP}} \left[\begin{array}{l} t \leq T(\lambda) \\ \bar{h}_{i-1} = h'_{i-1} = h_i \\ \bar{h}_i \neq h'_i \\ \text{L}_{\text{DL}}(\Pi_i) = \ell_i \\ \text{V}_{\text{DL}}(\text{pk}, \mathcal{R}[\mathcal{M}], (h_i, h'_i, 2^{\ell_i}), \Pi_i) = 1 \end{array} \right] \geq \frac{1}{2 \cdot \lambda \cdot P(T(\lambda))} .$$

Let Adv' be an adversary that given pk emulates EXP using pk and outputs:

$$(\mathcal{R}[\mathcal{M}], \bar{c}_{i-1}, \bar{c}_i, h_i, h'_i, 2^{\ell_i}, \Pi_i) .$$

It follows that:

$$\Pr_{\substack{\text{pk} \leftarrow \text{G}_{\text{DL}}(\lambda) \\ \text{Adv}'(\text{pk})}} \left[\begin{array}{l} 2^{\ell_i} \leq T(\lambda) \\ \text{L}_{\text{DL}}(\Pi_i) \leq \log T(\lambda) \\ \text{V}_{\text{DL}}(\text{pk}, \mathcal{R}[\mathcal{M}], (h_i, h'_i, 2^{\ell_i}), \Pi_i) = 1 \\ (\bar{c}_{i-1}, \bar{c}_i, \ell_i) \in \mathcal{T}_{\mathcal{M}} \\ h_i = \text{H}_{\text{DL}}(\text{pk}, \bar{c}_{i-1}) \\ h'_i \neq \text{H}_{\text{DL}}(\text{pk}, \bar{c}_i) \end{array} \right] \geq \frac{1}{2 \cdot \lambda \cdot P(T(\lambda))} .$$

contradicting the soundness for bounded level proofs requirement of the delegation scheme.

References

- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108. ACM, 1996.
- [BC12] Nir Bitansky and Alessandro Chiesa. Succinct arguments from multi-prover interactive proofs and their efficiency benefits. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 255–272. Springer, 2012.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKS and proof-carrying data. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 111–120. ACM, 2013.
- [BCL⁺21] Benedikt Bünz, Alessandro Chiesa, William Lin, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data without succinct arguments. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 681–710. Springer, 2021.
- [BCMS20] Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Recursive proof composition from accumulation schemes. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part II*, volume 12551 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2020.
- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 505–514. ACM, 2014.
- [BDFG21] Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Halo infinite: Proof-carrying data from additive polynomial commitments. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 649–680. Springer, 2021.
- [BDGM19] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part II*, volume 11892 of *Lecture Notes in Computer Science*, pages 407–437. Springer, 2019.
- [BHK17] Zvika Brakerski, Justin Holmgren, and Yael Tauman Kalai. Non-interactive delegation and batch NP verification from standard computational assumptions. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 474–482. ACM, 2017.
- [BHR⁺21] Alexander R. Block, Justin Holmgren, Alon Rosen, Ron D. Rothblum, and Pratik Soni. Time- and space-efficient arguments from groups of unknown order. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part IV*, volume 12828 of *Lecture Notes in Computer Science*, pages 123–152. Springer, 2021.

- [BP15] Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II*, volume 9453 of *Lecture Notes in Computer Science*, pages 236–261. Springer, 2015.
- [CJJ21] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Snargs for P from LWE. *IACR Cryptol. ePrint Arch.*, page 808, 2021.
- [GGH11] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation - In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*, volume 6650 of *Lecture Notes in Computer Science*, pages 30–39. Springer, 2011.
- [HR18] Justin Holmgren and Ron Rothblum. Delegating computations with (almost) minimal time and space overhead. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 124–135. IEEE Computer Society, 2018.
- [HW15] Pavel Hubáček and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In Tim Roughgarden, editor, *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 163–172. ACM, 2015.
- [KLVW23] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Boosting batch arguments and RAM delegation. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1545–1552. ACM, 2023.
- [KPY19] Yael Tauman Kalai, Omer Paneth, and Lisa Yang. How to delegate computations publicly. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 1115–1124. ACM, 2019.
- [Mer87] Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *Advances in Cryptology - CRYPTO '87, A Conference on the Theory and Applications of Cryptographic Techniques, Santa Barbara, California, USA, August 16-20, 1987, Proceedings*, volume 293 of *Lecture Notes in Computer Science*, pages 369–378. Springer, 1987.
- [Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *TCC*, pages 1–18, 2008.
- [WW22] Brent Waters and David J. Wu. Batch arguments for NP and more from standard bilinear group assumptions. *IACR Cryptol. ePrint Arch.*, page 336, 2022.