

# Verifiable Streaming Computation and Step-by-Step Zero-Knowledge

Abtin Afshar, Rishab Goyal\*

## Abstract

We propose a new incrementally computable proof system, called Incrementally Verifiable *Streaming* Computation (IVsC). IVsC enables computing incremental proofs of correct execution for any RAM program  $\mathcal{M}$  on a *streaming* input  $x$ . Input  $x$  is called a *streaming* input if it is only available on-the-fly as part of an ongoing data generation/streaming process, and not available at once. We also propose a new notion of zero-knowledge features for IVsC that guarantees the proof can be incrementally verified w.r.t. an encrypted digest, where the proof and digest hide the full data stream. We design zero-knowledge IVsC from a wide variety of standard falsifiable assumptions (such as decision-linear/sub-exponential DDH/LWE).

We also introduce a new notion of non-interactive zero-knowledge proofs, that we call *step-by-step* zero-knowledge protocols. Such protocols have strong zero-knowledge guarantees, wherein the prover's entire internal memory is simulatable at any point during proof generation. That is, unlike standard zero-knowledge proofs, where only the final proof is simulatable, we can also simulate prover's state at every step of the computation. Such proof systems will be useful in settings where an adversary could corrupt an honest prover even before it generates the full proof. We show that a zero-knowledge IVsC system can be used (almost) as a black-box to design step-by-step zero-knowledge proof systems, therefore secure under standard assumptions.

---

\*Email: {abtin,rishab}@cs.wisc.edu. Research supported by Wisconsin Alumni Research Foundation.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Overview . . . . .	2
<b>2</b>	<b>Preliminaries</b>	<b>10</b>
2.1	RAM Programs . . . . .	10
2.2	Rate-1 Block Somewhere Extractable Hash (SEH) Families . . . . .	11
2.3	Rate-1 Somewhere Extractable Batch Arguments (seBARGs) . . . . .	13
2.4	2-to-1 Collision Resistant Hash Function . . . . .	14
2.5	Hash Tree . . . . .	14
2.6	Non-Interactive Zero-Knowledge Arguments (NIZK) . . . . .	15
2.7	(Customized) Public-Key Encryption System . . . . .	16
2.7.1	leakage resilient/RDM/randomness recoverable/rate-1 PKEs . . . . .	17
<b>3</b>	<b>Zero-Knowledge Incrementally Verifiable Streaming Computation (zk-IVsC)</b>	<b>18</b>
3.1	Definition. . . . .	18
3.2	Construction . . . . .	20
3.3	Analysis . . . . .	25
3.4	Further Optimizing zk-IVsC . . . . .	27
3.5	Predicated zk-IVsC . . . . .	27
3.5.1	Definition . . . . .	27
3.5.2	Construction . . . . .	29
3.5.3	Analysis . . . . .	31
<b>4</b>	<b>Incrementally Computable Zero-Knowledge Arguments (ICZK)</b>	<b>31</b>
4.1	Definition . . . . .	31
4.2	Construction . . . . .	32
4.3	Analysis . . . . .	34
<b>5</b>	<b>From IVsC to Deterministic Distributed Computation</b>	<b>35</b>
5.1	Binary Tree to Path Distributed Computation . . . . .	37
5.2	Class of Succinctly Transformable Distributed Computations to a Path . . . . .	38
<b>6</b>	<b>Benefits and Applications</b>	<b>39</b>
<b>A</b>	<b>zk-IVsC Complete Analysis.</b>	<b>45</b>
<b>B</b>	<b>Predicated zk-IVsC Complete Analysis.</b>	<b>62</b>
<b>C</b>	<b>ICZK Complete Analysis.</b>	<b>66</b>

# 1 Introduction

Incrementally verifiable computation (IVC) [Val08] and, its popular generalization, proof carrying data (PCD) [CT10] are powerful succinct argument protocols. They enable incremental proof generation for a distributed computation performed by a network of mutually distrustful parties. PCDs have major theoretical and practical significance [CTV13, CTV15, NT16, Min17, KB23, BMRS20, CCDW20]. They offer strong succinctness and soundness guarantees, which informally state that the proof size is very short and no cheating polytime attacker can create proofs for incorrect computation. In privacy-sensitive applications (e.g., zk-EVM [But22] and more [CTV15, BBB<sup>+</sup>18, XZC<sup>+</sup>22, ZGSX23]), it is extremely valuable to support zero-knowledge guarantees as well. That is, the proofs (both, intermediate and final) must hide the sensitive data.

The most common approaches to design PCDs rely on recursively composing succinct non-interactive arguments of knowledge (SNARKs), or related primitives with similar accumulation/folding features. Such PCDs are typically proven secure in oracle models (e.g., Random Oracles), or under non-falsifiable (knowledge) assumptions (e.g., knowledge-of-exponents) [CT10, BCCT13, BCTV14, BGH19, BCMS20, COS20, BDFG21, BCL<sup>+</sup>21, CY21, CCS22, KST22, KS22, HN23, CCG<sup>+</sup>23, BC23, CGSY24]. Despite numerous advantages (e.g., transparent setup via ROs, nearly-constant sized proofs, etc), these have a strong limitation. Their security relies on complex heuristics [BR94], or strong non-falsifiable assumptions [Nao03]. As is standard practice in cryptography, proving security under standard (falsifiable) cryptographic assumptions is more desirable.

Unfortunately, there are known strong barriers to designing PCDs from standard falsifiable assumptions [GW11]. In light of these, it is natural to ask whether we could design PCDs, for any subset of general distributed computation, secure under standard assumptions. Recently, two beautiful works by Paneth-Pass [PP22] and Devadas et al. [DGKV22] designed IVC for all *deterministic computation* assuming hardness of standard assumptions. IVCs are a sub-class of PCDs, where distributed computation occurs in a straight-line (i.e., each user hands off partial computation state to the next user). Although not yet practical, these works design the first fully-succinct IVC from standard falsifiable assumptions. In this work, we continue this investigation of designing PCDs (beyond IVC) under standard assumptions.

**This work.** We introduce the notion of incrementally verifiable *streaming* computation (IVsC), as an improvement over IVCs. We also propose a new notion of privacy for IVsC. Our goal is to capture computations over *‘streaming data’* while preserving data privacy. In words, IVsC can be used in applications, where data is not statically available, but comes in via an ongoing generation process. Similar to [PP22, DGKV22], we keep our focus on only *deterministic computation*.

We also introduce a new notion of non-interactive zero-knowledge (NIZK), that we call *step-by-step* zero-knowledge protocols. These protocols are standard NIZK systems, except they have two special features: (1) the NIZK proof can be incrementally computed, (2) the prover’s internal state is also simulatable at any step during proof generation. That is, unlike standard NIZKs where only the final proof can be simulated, we can simulate prover’s state after every computation step.<sup>1</sup>

We design IVsC (with privacy) from a wide variety of standard assumptions (decision linear, or sub-exponential DDH, or LWE). In addition, we design step-by-step ZK protocols starting from our IVsC protocol, obtaining them from same assumptions. We also briefly discuss another application

---

<sup>1</sup>E.g., consider a user proving correctness of PRF evaluation ( $y = F_K(x)$ ) computed via the GGM-tree [GGM86], then using a step-by-step ZK protocol we can guarantee that the prover’s state only leaks the partial PRF evaluation (i.e., after  $i$  PRG computations, only  $F_K(x[1 : i])$  could get leaked, instead of  $K$ ).

in developing a new approach to prove distributed computations beyond path-graphs.

**Organization.** In the next section, we provide a high level overview of our main results and techniques. We discuss some motivating examples highlighting the benefits of supporting streaming inputs and various practical use cases of our designed systems in Section 6. In Section 2 we provide standard cryptographic preliminaries. In Section 3 we formally introduce and design IVsC with privacy. In Section 4 we describe step-by-step ZK and design it via a useful extension of IVsC, called *predicated zero-knowledge* IVsC. Finally, in Section 5 we discuss some more applications.

## 1.1 Overview

Consider the computation of a RAM machine  $\mathcal{M}$  on an input  $x$ , where the computation of  $\mathcal{M}(x)$  is very long (i.e., cannot be finished by a single user). An IVC protocol enables distribution of this computation across multiple users. In IVC, the  $(u+1)$ -th user takes the entire intermediate machine configuration  $\text{cf}_t$  (i.e., computation state and memory) from  $u$ -th user along with a succinct proof  $\pi_t$ , and continues the computation for some  $\ell$  steps to generate updated configuration and proof,  $(\text{cf}_{t+\ell}, \pi_{t+\ell})$ . Here  $\text{cf}_0$  denotes the initial configuration (i.e., input  $x$ ) and  $\pi_t$  proves that  $\text{cf}_0 \rightarrow \text{cf}_t$  after  $t$  steps of  $\mathcal{M}$ . Importantly, the proof size,  $|\pi_t|$ , should never grow with  $t$ , and the running time for a single user should only grow with the number of computation steps it performs (i.e.,  $\ell$  for user  $u+1$  in the above). In the context of *deterministic computations*, the verifier gets both, the original and final, configuration (e.g.,  $\text{cf}_0, \text{cf}_t$ ) along with  $\pi_t$ .

**Streaming inputs and introducing IVsC.** As described above, IVC requires the full input  $x$  to be part of the original machine configuration,  $\text{cf}_0$ . Thus, it cannot support incrementally certifying long computations, if the entire input is *not* available (or cannot be given) to the first user when the computation begins. To handle dynamic/streaming inputs, we consider a new streaming variant called incrementally verifiable *streaming* computation (IVsC).

In IVsC, each user  $u$  gets a local input  $x_u$ , in addition to the (intermediate) machine configuration  $\text{cf}_t$  and proof  $\pi_t$ . Rather than viewing the input  $x$  to be an implicit part of the configuration  $\text{cf}$ , we consider a separate input tape for  $\mathcal{M}$  (in addition to its internal memory). Thus,  $\mathcal{M}$  can read bits off of the input tape as well as its memory. One significant difference is that  $\pi_t$  now certifies that  $\text{cf}_0 \rightarrow \text{cf}_t$  w.r.t. inputs  $x_1, x_2, \dots$ , where  $x_i$  is local data/input available to the  $u$ -th user and  $\text{cf}_0$  is some fixed initial configuration (e.g., all-zeros). Since we focus on deterministic computations (like [PP22, DGKV22]), thus the verifier receives the inputs  $x_1, x_2, \dots$  to verify the proof.

Prior works in IVC [PP22, DGKV22] (and RAM delegation [CJJ21], more broadly) also considered an explicit Digest algorithm to speed up verification. Digest takes as input a configuration and outputs a short digest. A verifier only needs the digest for the configurations, thus its running time could be independent of the input size. We consider a similar improvement in IVsC. Our Digest procedure produces a short digest for a *streaming* input  $x_1, x_2, \dots$ ; moreover, the digest is incrementally computable. That is, in addition to proof  $\pi_t$ , each user passes along the input digest (so far), and the next user updates the proof as well as the digest. Thus,  $\pi_t$  can be verified rather quickly (i.e., in time independent of the streaming input size) given the digest. An important aspect of the ‘input digest’, both in IVC and IVsC, is that the digest is independent of  $\mathcal{M}$ , thus it can be reused for verifying other computations on the same data. However, a soundness attacker still must output the full input, or otherwise, the computation becomes non-deterministic.

We formalize the IVsC framework in Section 3, and explain our design choices. Succinctness and soundness can be naturally extended from either IVC/PCD literature (for deterministic com-

putation). At a very high level, IVsC can be viewed as a PCD for straight-line graphs. (Note that since IVC does not support local data for intermediate users, thus IVC does not really correspond to PCDs for straight-line graphs.) Next, we sketch the main ideas behind our IVsC design.

**IVsC via recursive proof batching.** Our starting point is the elegant template developed by Paneth-Pass [PP22] and Devadas et al. [DGKV22]. Both works relied on a common master tool, called rate-1 somewhere extractable batch arguments (seBARGs)<sup>2</sup>. Rate-1 seBARGs are powerful batch argument systems that generate a succinct proof  $\pi$  proving the validity of a batch of  $k$  instances  $x_1, \dots, x_k$ , given corresponding NP witnesses  $\omega_1, \dots, \omega_k$ . These argument systems are very desirable for two reasons: (i) proof size,  $|\pi|$ , is equal to the size of a single witness,  $|\omega_i|$ , plus a fixed additive polynomial term and (ii) one can efficiently extract the NP witness, corresponding to a hidden *trapdoor* index, from any accepting proof.

While both [PP22, DGKV22] rely on rate-1 seBARGs, their IVC designs are somewhat different due to diverging intermediate abstractions developed in each work. Paneth-Pass introduced the notion of mergeable proofs of delegation, while Devadas et al. proposed the concept of hashed multi-hop seBARGs. We briefly summarize both of them, and refer the reader to [PP22, DGKV22] for a detailed overview.

*Prior intermediate abstractions.* Mergeable proofs [PP22] enable merging two delegation proofs into a single succinct proof. That is, given proof  $\pi_1, \pi_2$  that certify  $\text{cf}_0 \rightarrow \text{cf}_1$  and  $\text{cf}_1 \rightarrow \text{cf}_2$  (resp.), using mergeable proofs these can be accumulated into a single proof  $\pi'$  of fixed polynomial size, certifying  $\text{cf}_0 \rightarrow \text{cf}_2$ . On the other hand, hashed multi-hop seBARGs generalize seBARGs to allow unbounded batching of batch proofs, with a special feature that the verifier only needs the digest of instances for verification, where the prover also generates this digest after each batching. That is, given  $k$  proofs  $\pi_1, \dots, \pi_k$  for  $x_1, \dots, x_k$ , one generates a proof  $\pi$  and digest  $h$  such that verifier only needs  $(\pi, h)$ , and  $k$  such proof-digest pairs can be further batched polynomially many times. Both works design these two objects by relying on rate-1 seBARGs.

**Canonical IVC template.** Given these tools, [PP22, DGKV22] follow a somewhat similar template for designing IVC. Each IVC proof contains  $\lambda$  RAM delegation proofs  $\pi^1, \dots, \pi^\lambda$ , and  $\lambda$  configurations  $\text{cf}^1, \dots, \text{cf}^\lambda$ . Each  $\pi^i = \epsilon$  (i.e., empty), or it proves  $\text{cf}^i \rightarrow \text{cf}^{i+1}$ . The special feature of these is that if  $\pi^i \neq \epsilon$ , then it certifies correctness of exactly  $2^i$  computation steps. That is,  $\mathcal{M}$  starts from  $\text{cf}^i$  and reaches  $\text{cf}^{i+1}$  after  $2^i$  steps. Now  $\{\pi^i\}_i$  together certify  $\mathcal{M}$  that  $\text{cf}_0 \rightarrow \text{cf}_T$  (where  $\text{cf}_0$  and  $\text{cf}_T$  are starting/ending configs). And, the way this works is as: suppose  $T = \sum_i t_{i+1} 2^i$  (where  $t_i$  is the  $i$ -th least significant bit in binary representation of  $T$ ), then for each  $i$  s.t.  $t_i = 0$  we have that  $\pi^i = \epsilon$ , otherwise  $\pi^i$  certifies  $\text{cf}^i \rightarrow \text{cf}^{i+1}$  (proving correctness of  $2^i$  computation steps).

The advantage of designing IVC proofs as above is that it ensures proofs can be incrementally updated, yet their size does not grow with  $T$  (machine run-time). To understand a bit better, suppose a user is given an IVC proof certifying some  $T$  computation steps of  $\mathcal{M}$ , and it wants to run two more steps (i.e.,  $\text{cf}_T \rightarrow \text{cf}_{T+2}$ ). Given the above architecture, the user can first create a RAM delegation proof  $\pi'$  for certifying  $\text{cf}_T \rightarrow \text{cf}_{T+2}$ , and then do the following:

1. Start with  $i = 1$ .
2. Check if  $\pi^i$  is empty. If yes, it sets  $\pi^i = \pi'$ , appropriately updates the underlying configurations, and ends the routine. Else:
  - (a) “Accumulate”  $\pi^i$  and  $\pi'$ . (Note each of these prove a computation of size  $2^i$ .)

<sup>2</sup>Throughout, by rate-1 we mean almost-rate-1, where the proof to witness ratio is  $1 + o(1)$  and not exactly 1.

- (b) Update  $\pi'$  with the above accumulated proof. (Now  $\pi'$  proves a computation of size  $2^i + 2^i = 2^{i+1}$ .)
- (c) Update corresponding configurations appropriately, increase  $i$  to  $i + 1$ , and go back to step (2).

The above trick is colloquially referred to as the ‘powers-of-two’ trick and has been used in various contexts in cryptography [GKP<sup>+</sup>13, GHMR18, GV20]. In the context of IVC, it is extremely useful for incrementally updating RAM delegation proofs, while ensuring that the individual RAM delegation proofs (what we refer to as  $\{\pi^i\}_i$ ) do not get “merged” [PP22] or “batched” [DGKV22] too often. Clearly, at any point when we combine RAM proofs, the underlying configurations go from being  $2^i$  apart to being  $2^{i+1}$  apart. Thus, next time the resulting proof has to be merged will be only after  $2^{i+1}$  more computation steps have happened. This inherently guarantees desired succinctness for the proof size, and ensures updating a proof and verifying it can be done efficiently.

**Soundness of IVC.** Pass-Paneth execute the above template by using their mergeable proofs for *accumulation*, while Devadas et al. rely on hashed multi-hop seBARGs. The exact designs in both works slightly alter the above template to suit their intermediate object as well as to ensure the soundness proof goes through.

At an abstract level, we can view both their soundness proofs through a unified lens. We informally call it the *sliding window technique*, where the intuition is to view all  $T$  steps of computations (and their corresponding configurations) as a long linked list. Now the  $i$ -th element/node in the list corresponds to the computation that happens at the  $i$ -th step of  $\mathcal{M}$ ’s computation. Recall at any step,  $\mathcal{M}$  would read/write something from/on its memory, update the RAM machine state, and move the head along the memory. The core idea behind proving soundness is to fix a *constant-size* continuous sequence of nodes  $\text{cf}_i, \dots, \text{cf}_{i+w}$ , and extract the corresponding configurations by using extractability feature of underlying seBARGs. We refer to this sequence of nodes as a *window*, where  $w$  is the window size. Given the extracted information, the reduction algorithm can argue that the all these configurations must be consistent<sup>3</sup> w.r.t  $\mathcal{M}$  by relying on seBARG’s (/intermediate object’s) security. The next step is for the reduction to “slide” the window slightly in the forward direction. That is, say now it extracts  $\text{cf}_{i+1}, \dots, \text{cf}_{i+w+1}$  (i.e., moves it forward by one). Again, we can argue that these configurations must be consistent as before, and it appears that, by sliding the (extraction) window by one node at a time, we can argue that the starting and final configurations are consistent as per  $\mathcal{M}$ .

However, there are two caveats that require care. First, we need to guarantee that sliding the extraction window has only negligible drop in adversary’s success; and second, when we move across two overlapping windows, then the extracted configurations that appear in both are the same. This is where the works rely on the famous *no-signaling* proof techniques [KRR14, CJJ21], and show that by devising a very specific sequence of hybrids and case analysis, soundness goes through.

While the above is mostly accurate, there is one final technical issue. The configuration size is too big, thus if we want to extract even a single configuration entirely, then the underlying RAM delegation/batch proof size must grow with the (maximum) configuration size. At first, this appears to be a major bottleneck, but it turns out there is rather simple and elegant solution for addressing it. Simply denote each node with a ‘*psuedo*’-configuration  $\text{pcf}$ , instead of the full configuration  $\text{cf}$ . A *psuedo*-configuration can be viewed as a specialized digest of the full configuration. Its

<sup>3</sup>Throughout, when we say two (or more) configurations  $\text{cf}_1, \text{cf}_2$  ( $, \dots$ ) are consistent, we mean that if you run  $\mathcal{M}$  on  $\text{cf}_1$ , then you will reach  $\text{cf}_2$  (and so on).

size is a-priori fixed (thus, independent of  $|cf_i|$ ), and one can succinctly check if two (adjacent) pseudo-configurations are consistent. For security, they guarantee that an attacker cannot generate two pairs of pseudo-configurations  $(pcf_1, pcf_2)$  and  $(pcf'_1, pcf'_2)$  such that they both are consistent, yet  $pcf_1 = pcf'_1$  and  $pcf_2 \neq pcf'_2$ . To design such pseudo-configurations, prior works use special hash functions [HW15, OPWW15] implemented as Merkle trees. Combining all these ideas, one can avoid extracting the full configuration, and ensure the soundness proof goes through without blowing up the proof size. We refer the reader to [PP22, DGKV22] for a more detailed overview.

**Our IVsC template.** Our template builds upon the aforementioned IVC template to handle *streaming inputs*. Very briefly, the central reason the above IVC design cannot handle streaming inputs directly is that the input in IVC is programmed/hardwired inside the starting configuration. That is, since the input  $x$  is statically available, the first user simply treats it as part of the RAM memory. Thus, if some user (during the incremental computation) receives additional data/input, then it has to somehow merge it with its current configuration *while proving consistency*. The issue is that this significantly changes the pseudo-configuration, thus proving the pseudo-configurations are still consistent (succinctly) is challenging.

Our solution is to maintain a separate input tape, in addition to all other parts of the machine configuration. The advantage of separating this out would be: (1) we can add new data to the input tape dynamically, (2) it will not affect the current pseudo-configuration significantly<sup>4</sup>. However, just splitting (streaming) input and the full machine configuration is not enough. We need to ensure that an appropriate *succinct* digest for the input tape is added to the (new) pseudo-configuration, and this digest can be *incrementally* updated to support new incoming/streaming data. Finally, we also need to be careful that, despite these changes, the two necessary properties for pseudo-configurations still hold – (a) succinctly certifying two pseudo-configurations are adjacent, and (b) computational infeasibility of finding two *consistent* pseudo-configuration pairs with conflicting computation results.

At a high level, we achieve this by including  $\lambda$  input digests  $h_1, \dots, h_\lambda$  as an additional part of the incremental proof. Recall previously, they contained  $\lambda$  RAM delegation proofs and  $\lambda$  configurations. The purpose of maintaining these additional digests is to make sure that streaming inputs can be efficiently accumulated via the *powers-of-two* technique. Thus, the running time of update does not scale with the total size of the streaming input. Additionally, these digests can also be added to each pseudo-configuration  $pcf$ , and they will assist in succinctly proving adjacency of two pseudo-configurations. This is because each user (potentially) gets some new data/input, thus these digest values could significantly change when the user accumulates the latest input. To ensure this doesn't cause a problem, we prove that the input digests are also correctly updated by each prover. An important and desirable feature of the above input digests is that these can be reused for certifying different computations on the same data. Thus, if the same sequence of users run another machine  $\mathcal{M}'$  on the given data, then the input digests will be consistent. Therefore, a verifier does not need to recompute the digests for a given sequence of streaming inputs every time.

Overall, these ideas could put together to generalize the IVC template to handle streaming inputs. And, although the soundness proof has to be sufficiently expanded (due to the changes in pseudo-configurations and additional tape headers), the core technical ideas still rely on the sliding window technique coupled with no-signaling proof strategies. Next, we discuss how to formalize and achieve a non-trivial notion of *privacy* for IVsC.

---

<sup>4</sup>There would be some changes which include storing additional tape header for the input tape, etc. But these can be appropriately handled, thus we ignore them for this overview.



**Privacy (zero-knowledge) for IVsC.** As mentioned earlier, defining privacy (zero-knowledge) for IVsC is a rather delicate task. Recall that a zero-knowledge proof for any deterministic computation is simply  $\epsilon$ , an empty string. Because a verifier can verify its validity by performing the computation on its own, and an empty proof string is perfectly simulatable (vacuously). Since we study IVsC for deterministic computations, thus it is unclear if there exists a non-trivial notion of zero-knowledge. In words, *what can we hide from a malicious verifier that gets the (streaming) input for the (deterministic) computation?*

At first glance, it seems zero-knowledge for (incremental) deterministic computations is unlikely to be meaningful. However, it turns out that we indeed can formalize this non-trivially for IVsC (even IVC). Recall that a verifier gets a proof  $\pi$  and a short *digest*  $h$  as inputs, rather than the full input. This digest can either be statically computed for a fixed input (i.e., in IVC), or incrementally for a streaming input (i.e., in IVsC). Suppose we define privacy for the proof and digest *jointly* as follows:

$$\{(\pi, h, y = \mathcal{M}(x_1, x_2, \dots)) : (\pi, h) \leftarrow \text{IVsC.Prove}(\mathcal{M}, x_1, x_2, \dots)\} \approx_c \text{Sim}(\mathcal{M}, y) \quad (1)$$

The above states that an honestly computed proof-digest pair is indistinguishable from a simulated distribution, given the machine description and the output of the computation. Moreover, we can define this to hold *even for intermediate proofs, digests, and configurations*.

In other words, the guarantee is when a particular user passes off the IVsC proof, digests, and configurations after (say)  $T$  computation steps to the next user, then that user cannot learn anything about the previous inputs or intermediate configurations of  $\mathcal{M}$ , more than what is revealed by  $\text{cf}_T$ , i.e. the configuration of  $\mathcal{M}$  after running it for  $T$  steps on the (streaming) input so far.

While the above formulation may seem ideal and almost the best-possible privacy guarantee that could be feasibly defined for incremental deterministic computations, there is one technical nuance. If we juxtapose the soundness definition and Eq. (1), then we realize the threat models are not fully overlapping. Soundness states that an adversary cannot generate  $y, x_1, x_2, \dots$  and  $\pi$  such that  $y \neq \mathcal{M}(x_1, x_2, \dots)$  but  $(\text{Digest}(x_1, x_2, \dots), \pi)$  is accepted by the verifier. Notably, soundness attackers must output the *full* input. While Eq. (1) states that simulation only works for the proof and digest, and *not* the full input. This leaves a noticeable gap. Let us explain via an example.

*An example.* Consider a trusted authority that collects and signs large amounts of certain sensitive data. Due to resource constraints, it delegates the computation of a study/survey on this data to a cloud provider. Suppose it wants the cloud provider to publish the output publicly (e.g., census statistics and results) so that any individual user can verify the output. Now it might appear that we could readily rely on IVsC in this setting. The cloud runs IVsC to incrementally run the computation, generate a proof  $\pi$ , and digest  $h$ , and publishes these along with the output. Due to the zero-knowledge guarantee, we know that the proof-digest pair will not reveal too much about the data. But what about soundness!? We cannot rely on IVsC soundness to guarantee that a user will not accept the results, as the user does not have the entire data.

While it appears bypassing this would require all powerful succinct proofs for non-deterministic computations, we point out that (in many such applications) all a user might want to check is that the input digest  $h$  conforms to a publicly checkable predicate. That is,  $h$  is *well-formed*. E.g., the user just wants to check that every data was signed by the trusted (collection) authority. With this insight, we introduce a new notion that we call *predicate zero-knowledge* property for IVsC.

*Predicated zero-knowledge for IVsC.* Consider  $\phi$  to be the predicate against which a verifier wants



to check data conformity. That is, the verifier wants to check that

$$\exists x_1, x_2, \dots \text{ s.t. } y = \mathcal{M}(x_1, x_2, \dots) \wedge (\phi(x_i) = 1)_i$$

Syntactically, we do not make any significant adjustments to the IVsC abstraction, except with one notable difference. Since we require the proof-digest pair  $(\pi, h)$  to hide everything about the streaming inputs, beyond their predicate satisfiability, thus the digest  $h$  can no longer be deterministically computed from just the inputs  $(x_1, x_2, \dots)$ . Instead, we expect each digest  $h$  is accompanied with a predicate proof  $\pi_\phi$  such that: (1) a verifier can efficiently check the underlying data conforms as per  $\phi$ , and (2)  $\pi_\phi$  can also be incrementally updated as new streaming data comes in. Soundness and privacy (in the form of zero-knowledge) can be appropriately generalized to fit the abstraction. We refer the reader to the main body for detailed descriptions. We remark that we also keep track of the randomness used in creating/updating these digests. This is useful for defining stronger soundness notions as well as enabling auditability of the digest computation by the data owner.

We believe the above gives us the best possible guarantee for capturing privacy for incremental proofs for deterministic computations. Along the way, it answers an important question left open by Ananth et al. [ADKL19] asking if one could design incremental proofs that enjoy both succinctness and privacy guarantees. Next, we explain the main ideas behind our construction.

**Designing IVsC with predicated zero-knowledge.** We explain our final template in two steps. First, we design an IVsC protocol with *empty predicates*. That is, we do not check well-formedness of the digest in the first step. While this is not ideal, we fix this in the second step. In the second step, we take our base protocol and extend it for arbitrary non-empty predicates.

**STEP 1: IVsC for empty predicates.** In our previous IVsC template (without privacy), each proof contains  $\lambda$  RAM delegation proofs  $(\pi_1, \dots, \pi_\lambda)$ ,  $\lambda$  pseudo-configurations  $(\text{pcf}_1, \dots, \text{pcf}_\lambda)$ , and  $\lambda$  digests  $(h_1, \dots, h_\lambda)$ . Let us inspect why this does not already satisfy zero-knowledge as in Eq. (1). There are multiple reasons listed out below:

1. The digests and pseudo-configurations are (deterministic, Merkle-tree-style) hashes of the streaming input and configurations. Thus, they could reveal parts of the inputs and intermediate configurations.
2. The RAM delegation proofs themselves do not have any zero-knowledge guarantees.

In short, every component of the original proof leaks information about the inputs and computation so far. At a high level, our approach is two-fold: (a) commit to every input  $x_i$  and each configuration  $\text{cf}_j$ , and then hash down the corresponding commitments, and (b) equip RAM delegation proofs with necessary zero-knowledge properties by using NIZKs. The intuition is that by committing and then hashing the inputs/configurations, we can rely on hiding property of commitments. Similarly, by using NIZKs (inside RAM delegation proofs) we can prove correctness of each transition between two pseudo-configurations, while hiding the openings for the commitments.

The above serves as a good high level summary of the technical modifications needed. However, the exact composition of these have to done very carefully so that the soundness proof goes through. Notably, using NIZKs and commitments, we make all these digest and pseudo-configurations “lossy”. By lossy, we mean that they may no longer satisfy statistical binding properties that are necessary for executing the sliding-window-style proof strategy. Moreover, once we plug in NIZKs, then recursive composition of RAM delegation proofs can *not* be done succinctly. Because

even if the NIZK proofs are very short, then we need to compose them recursively where the instance size grows significantly with each composition. However, as we detail in the main body, these issues can be handled by relying on hiding commitments with somewhere-extractability guarantees. And, more importantly, we use NIZKs to only prove consistency of two pseudo-configurations that are 2 steps apart. In other words, we realize that it is sufficient to just have the first RAM delegation proof,  $\pi_1$ , satisfy zero-knowledge properties. And, rest of the RAM delegation proofs can still be composed as before, without using NIZKs. The proof of zero-knowledge is much more simple, since we can start by simulating NIZK proof that show consistency of all pseudo-configurations that are 2 steps apart. Once we change these, we can accumulate them as an honest prover does, and we can finish the proof by using hiding security of all (extractable) commitments.

More details are provided later in Section 3. We remark that we could feasibly introduce a new intermediate abstraction (similar to mergeable proofs/hashed multi-hop BARGs [PP22, DGKV22]), equip it with stronger zero-knowledge-style guarantees, and then use ideas from our IVsC template described earlier. Instead, we decided to give a more direct construction for IVsC by relying on rate-1 seBARGs (and related primitives) in the main body. This is mostly because, unlike prior works [DGKV22, PP22], we focus on privacy of incremental proofs. And, capturing a sufficient notion of zero-knowledge for those intermediate abstractions turned out problematic. Briefly, as we hinted above, we need two different RAM delegation proofs (one with zero-knowledge property and one without for ensuring optimal accumulation), thus any possible intermediate abstraction would be unfortunately quite complicated and make the soundness proof a highly non-black-box process.

**STEP 2: Extending to general predicates.** Finally, we show that we can upgrade the above IVsC protocol to generate predicate proofs along with the incremental proofs. Note that each predicate proof is almost like another incremental proof. That is, consider each prover computes two parallel IVsC proofs, where the first IVsC proof certifies computation of  $\mathcal{M}(x_1, \dots)$ , while the second proof certifies computation of  $((\phi(x_1) = 1) \wedge (\phi(x_2) = 1) \dots)$ . Note that if we generate both of these proofs using our above IVsC construction with empty predicates, then this offers a generic approach for designing IVsC for general predicates. We need some additional care in ensuring the soundness proof goes through. Because of the fact that if these two proofs are generated independently, then we cannot argue that their corresponding input digests are the same (since they are not deterministic). Thus, a black-box approach for extending to general predicates does fail. However, as we detail in Section 3.5, if we reuse the digests carefully, and rely on similar incremental computation techniques (as we do for our empty predicate IVsC), then the proofs of soundness, well-formedness of digest, and zero-knowledge work out. Moreover, we consider even more general predicates, where different portions of the streaming data could have separate predicates, and those could even be provided in a streaming fashion. Refer to Section 3.5 for more details.

The above concludes the overview of our IVsC protocols, and next we introduce the notion of step-by-step zero-knowledge protocols and discuss our main design ideas.

**Step-by-step zero knowledge.** The prover’s running time in any NIZK proof system (for NP) is always larger than the time it takes to compute  $\mathcal{R}(x, \omega)$ , where  $\mathcal{R}$  is the corresponding NP relation<sup>5</sup>. Consider a setting where  $\mathcal{R}$  is a really long computation, or  $x$  and  $\omega$  are being generated as part of a streaming computation. The question that we are interested in studying is: *can zero-knowledge proofs be incrementally computed?* Suppose the NIZK proof will be computed by more than one user (due to large runtime of  $\mathcal{R}$ ), then can we obtain any *intermediary* zero-knowledge guarantees

<sup>5</sup>Otherwise, the NIZK system will not satisfy either correctness, soundness, or zero-knowledge.

for the computation state that is passed to the next prover? Or, can we obtain any protections against attackers that can corrupt a prover in the middle of the computation?

To address these, we introduce a new notion of NIZK systems, that we call step-by-step zero-knowledge protocols. We design an *incrementally computable zero-knowledge* protocol (ICZK) that gives us desired step-by-step ZK guarantees. We define ICZK as a regular NIZK system with two special features: (a) the NIZK proof is computed as an incremental proof (i.e., the proof size does not scale with runtime of  $\mathcal{R}$ ), and (b) they satisfy *incremental zero-knowledge*. By incremental ZK, we mean that the following distributions (for any  $T$ ) are indistinguishable:

$$\{\text{state} : \text{state} \leftarrow \text{Prove}(x, \omega, T)\} \approx_c \{\text{Sim}(\text{cf}) : \text{cf} = \mathcal{R}(x, \omega; T)\}$$

In the above,  $\mathcal{R}(x, \omega; T)$  denotes running  $\mathcal{R}$  on  $(x, \omega)$  for  $T$  time steps, and **state** refers to the prover’s state after running  $T$  steps of the NP validation relation,  $\mathcal{R}$ . It is straightforward to see that the above tightly captures the intuition behind a step-by-step zero-knowledge protocol. Because we are quite literally defining simulatability at every step of the computation. As we show next, we can design ICZK quite naturally from our (predicated) zero-knowledge IVsC protocols.

**From zk-IVsC to ICZK.** A natural idea would be to directly use zk-IVsC to incrementally run  $\mathcal{R}$  on  $(x, \omega)$  as the joint input. While the resulting scheme would satisfy our incremental ZK property, we are unable to prove its soundness. The reason is that there is no obvious ‘meaningful’ predicate that we could check on  $(x, \omega)$  to ensure a cheating prover is really using a satisfying *witness* as part of the input, without revealing  $\omega$  as part of the proof.

At a high level, our plan is to include a redundant encoding of the witness  $\omega$  as an additional component of the proof, and then define the predicate to check that the witness portion of the input (in the zk-IVsC proof) is consistent with this encoding. To preserve zero-knowledge, we set the redundant encoding to be an encryption of  $\omega$ . That is, the prover starts by encrypting  $\omega$  to create a ciphertext  $\text{ct} = \text{Enc}(\omega; R)$  for some randomness  $R$ . Then, we run the zk-IVsC prover while setting the predicate  $\phi$  to check that  $\text{ct}$  is an encryption of the input witness  $\omega$ . However, there is a major hurdle in executing this strategy. The predicate  $\phi_{\text{ct}}$  is no longer a deterministic function. Encryption must be randomized for preserving zero-knowledge, thus we need to include the encryption randomness  $R$  as part of the zk-IVsC input.

Combining the above ideas, the ICZK prover starts by encrypting  $\omega$  as  $\text{ct} = \text{Enc}(\omega; R)$ , and then sets  $(x, \omega, R)$  as the input for zk-IVsC. With this, the ICZK prover runs the zk-IVsC protocol for computing  $\mathcal{R}$  on  $(x, \omega)$  portion of the input, and sets the predicate as  $\phi_{\text{ct}}(x, \omega, R) = (\text{ct} \stackrel{?}{=} \text{Enc}(\omega; R))$ . While this change makes the predicate deterministic, and we can prove soundness of the resulting system, we have made it difficult to prove (incremental) zero-knowledge. This is because the randomness  $R$  that we use to encrypt  $\omega$  is used as an input in the zk-IVsC protocol. Now to use semantic security of encryption, we have to argue that  $R$  is not leaked by the zk-IVsC portion of the proof. However, to argue that  $R$  is not leaked by the zk-IVsC proof, we need to make sure the predicate  $\phi_{\text{ct}}$  is independent of its input  $(x, \omega, R)$ .

This circularity prevents us from proving (incremental) zero-knowledge for the above design. Our final observation is that we can break this circularity by relying on two special PKE schemes, one that is leakage-resilient and satisfies randomness recovery<sup>6</sup> [PW08], and other that supports randomness-dependent-message security [BCPT13]. The first type of PKE schemes is used to encrypt  $\omega$ , while the second type is used to commit to the input  $(x, \omega, R)$  as part of zk-IVsC. Together

<sup>6</sup>We also need it to be rate-1, but as we show later, the rate can be improved quite easily.

these both guarantee that our zk-IVsC protocol satisfies a leakage-resilient-style zero-knowledge guarantee, and the semantic security of ct can be used despite the circularity. Fortunately, we observe that these PKE can be designed from lossy trapdoor functions [PW08], thus we are able to instantiate our template under a wide variety of standard assumption. Although there are some more technical subtleties that we need to handle along the way, the above nicely summarizes the key challenges and circular issues. We refer to Section 4 for a more in-depth discussion.

**New trade-offs in certifying general distributed computation via IVsC.** Finally, we also briefly discuss another application for IVsC in certifying distributed computations beyond path graphs. While we do not design general PCDs for deterministic computations, we show that IVsC can be deployed in a wide variety of distributed computation networks to enable a non-trivial trade-off. We show that if we start with any arbitrary distributed computation network with say  $N$  computation nodes, then using IVsC we can certify the computation with a proof size of  $\text{poly}(\lambda, \log N)$ . However, the total running time of the computation would grow linearly with  $N$ , rather than with the network diameter/depth (which could be as small as  $\log N$ ). Our idea is to modify the (more general) distributed network by inserting a few artificial communication edges between different nodes, and use those to convert it into an almost path graph by relying on an appropriate topological sorting. We believe this could be interesting model for future research, thus sketch a construction in Section 5.

**Related work.** While the scope of PCD constructions from standard falsifiable assumptions [Nao03] is limited to IVC protocols for *deterministic computation* [DGKV22, PP22], there are several PCD constructions such as [CT10, BDFG21, CY21, KS22, HN23] for general distributed computations from *non-falsifiable* assumptions. These works can be split to two different templates: (1) PCDs via recursive composition of succinct non-interactive arguments of knowledge (SNARKs) [BCCT13, BCTV14, BSCTV17, COS20], and (2) more efficient PCDs via accumulation schemes in the random oracle model [BCMS20, BCL<sup>+</sup>21, KST22, BC23, BC24, BMNW24]<sup>7</sup>. Moreover, there are a few works that construct PCDs with a transparent setup [BGH19, BCL<sup>+</sup>21, CCS22, CCG<sup>+</sup>23] without heuristically instantiating the random oracle in a non-black-box way.

## 2 Preliminaries

**Notation** We use overline for sets (e.g., for inputs to the RAM machine we let  $\bar{z} = (z_k)_{k \in [K]}$  and  $\bar{z}_i = (z_k)_{k \in [i]}$ ).

### 2.1 RAM Programs

A RAM program is typically defined as a deterministic machine  $\mathcal{M}$  represented using a fixed polynomial-sized set of states that has random access to a large memory (where the memory includes an explicit input and rest of it is initialized as zeros). In this work, we use the following representation for RAM machines as it enables a simpler exposition of our main ideas.

Any RAM machine  $\mathcal{M}$  that receives  $n$  bits of explicit input  $x = \{0, 1\}^n$ , is associated with a work tape of size  $S = S(n)$  and a fixed set of states  $Q$  that the machine can be in at any point during machine execution. Formally, a machine  $\mathcal{M}$  is associated with following components:

<sup>7</sup>All these works additionally rely on public-key assumption with a notable exception of [BMNW24] which is limited to bounded depth accumulation.

- A set of machine states  $\mathbb{Q}$ .
- A memory  $\mathbb{M}$  of size  $n + S$ . Without loss of generality, we assume that the explicit input  $x \in \{0, 1\}^n$  is written in the first  $n$  cells of the memory (i.e.,  $x = (\mathbb{M}_1, \dots, \mathbb{M}_n)$ ), and the rest of the memory (i.e.,  $(\mathbb{M}_{n+1}, \dots, \mathbb{M}_{n+S})$ ) corresponds to the work-tape of the machine. For ease of exposition, we also use  $\mathbb{W} = (\mathbb{W}[i])_{i \in [S]}$  to denote the work-tape. In our model, we view the work tape to be initialized as all zeros.
- A state transition function  $\delta$  with the following syntax:

$$\delta : \mathbb{Q} \times \{0, 1\} \rightarrow \mathbb{Q} \times \{0, 1\}^{\log(n+S)} \times \{0, 1\}^{\log S} \times \{0, 1\}$$

Here the transition function defines the execution of machine  $\mathcal{M}$  at each time step.

There is a circuit  $\mathbb{C}_{\mathcal{M}}$  that computes the next step as follows:

$$\mathbb{C}_{\mathcal{M}}(\mathbb{q}, \text{rbit}) = (\mathbb{q}', \text{ridx}', \text{widx}, \text{wbit})$$

We define the “configuration” of the RAM machine as the work part of memory, the current state, and the next reading index. Namely, we have  $\text{cf} = (\mathbb{W}, \mathbb{q}, \text{ridx})$ . Note that  $|\text{cf}| = S + \log |\mathbb{Q}| + \log(n + S)$ .

We let  $\text{out}$  to be the output of the RAM machine and require the machine to write out at the first  $n_{\text{out}} = |\text{out}|$  locations of the work tape  $\mathbb{W}$ . Namely, we have  $\text{out} = \mathbb{W}[1, \dots, n_{\text{out}}] = \text{cf}[1, \dots, n_{\text{out}}]$ .

In this work, our focus is on streaming inputs  $\bar{z} = (z_k)_{k \in [K]}$  for RAM machines. More specifically the machine is run by  $K$  users on their inputs where each user runs the machine for  $T$  steps. We use  $\text{cf}_{k,t}$  for the configuration of the machine at timestamp  $(k, t)$ , i.e., after running for  $t$  steps by user  $k$ . We let  $\text{cf}_{1,0}$  be the initial configuration of the machine and  $\text{cf}_{k,0} = \text{cf}_{k-1,T}$  for  $k > 1$ .

We denote by  $\mathcal{M}(\text{cf}_{k,t}, z_k; 1^{t'})$  running the machine  $\mathcal{M}$  starting from the intermediate configuration  $\text{cf}_{k,t}$ , on input  $z_k$  for  $t'$  steps. It follows by induction that

$$\text{cf}_{k,t} = \mathcal{M}(\mathcal{M}(\dots \mathcal{M}(\mathcal{M}(\text{cf}_{1,0}, z_1; 1^T), z_2; 1^T) \dots), z_k; 1^t).$$

We additionally use  $\mathcal{M}((z_k)_{k \in [K]}; 1^T) = \text{cf}_{K,T}$  for a complete computation of machine by  $K$  users.

The language of machine  $\mathcal{M}$  is defined as follows:

$$\mathcal{L}_{\mathcal{M},T} = \left\{ (\bar{z}, \text{out}) \quad : \quad \begin{array}{l} \text{cf}_{1,0} = \text{initial config, } \text{cf}_{K,T}[1, \dots, n_{\text{out}}] = \text{out,} \\ \forall k \in [K] : (|z_k| = n \wedge \mathcal{M}(\text{cf}_{k,0}, z_k; 1^T) = \text{cf}_{k,T}) \end{array} \right\} \quad (2)$$

**Remark 2.1.** We let w.l.o.g. the initial configuration  $\text{cf}_{1,0}$  to be all zeros.

## 2.2 Rate-1 Block Somewhere Extractable Hash (SEH) Families

In what follows we recall a *simplified* definition of a block somewhere extractable hash (SEH) family that suffices for our applications.

**Syntax.** A block somewhere extractable hash family SEH consists of the following polynomial time algorithms:

$\text{Gen}(1^\lambda, N, \Sigma, I) \rightarrow (\text{hk}, \text{td})$ . This is a probabilistic setup algorithm that takes as input a security parameter  $1^\lambda$  in unary, a number of blocks  $N$ , an alphabet  $\Sigma$ , and a subset  $I \subseteq [N]$ . It outputs a hash key  $\text{hk}$  along with trapdoor  $\text{td}$ .

$\text{Hash}(\text{hk}, x) \rightarrow v$ . This is a deterministic algorithm that takes as input a hash key  $\text{hk}$  and an input  $x \in \Sigma^N$ , and outputs a hash value  $v$ .

$\text{Extract}(\text{td}, v, j) \rightarrow y$ . This is a deterministic extraction algorithm that takes as input a trapdoor  $\text{td}$ , a hash value  $v$ , and an index  $j \in [I]$  and outputs a block  $y$ .

We use the notation  $\text{Extract}(\text{td}, v) = (\text{Extract}(\text{td}, v, j))_{j \in [I]}$ .

**Definition 2.2** (SEH). A somewhere extractable hash family  $\text{SEH} = (\text{Gen}, \text{Hash}, \text{Extract})$  is required to satisfy the following properties:

**Efficiency.** The size of the hash key  $\text{hk}$  and the hash value  $v$  is at most  $|I| \cdot |\Sigma| \cdot \text{poly}(\lambda)$ .

**Index hiding.** For any PPT adversary  $\mathcal{A}$ , any polynomial  $N = N(\lambda)$ , and any  $I_0, I_1 \subseteq [N]$  such that  $|I_0| = |I_1|$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \mathcal{A}(\text{hk}) = b : \begin{array}{l} b \leftarrow \{0, 1\} \\ (\text{hk}, \text{td}) \leftarrow \text{Gen}(1^\lambda, N, \Sigma, I_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

**Extraction correctness.** For any  $\lambda \in \mathbb{N}$ , any  $N \leq 2^\lambda$ , any subset  $I \subseteq [N]$ , and any  $x \in \Sigma^N$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ (x_i)_{i \in I} \neq \text{Extract}(\text{td}, v, I) : \begin{array}{l} (\text{hk}, \text{td}) \leftarrow \text{Gen}(1^\lambda, N, \Sigma, I) \\ v \leftarrow \text{Hash}(\text{hk}, x) \end{array} \right] \leq \text{negl}(\lambda).$$

**Remark 2.3.** Note that the extraction correctness of a SEH family, implies the following properties:

**Somewhere binding.** For any  $\lambda \in \mathbb{N}$ , any  $N \leq 2^\lambda$ , any subset  $I \subseteq [N]$ , any index  $i^* \in I$ , and any (all powerful) adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{Hash}(\text{hk}, x) = v \\ \wedge x_{i^*} \neq y_{i^*} \end{array} : \begin{array}{l} (\text{hk}, \text{td}) \leftarrow \text{Gen}(1^\lambda, N, \Sigma, I), \\ (v, x) \leftarrow \mathcal{A}(\text{hk}), \\ (y_i)_{i \in I} = \text{Extract}(\text{td}, v) \end{array} \right] \leq \text{negl}(\lambda).$$

**Somewhere binding w.r.t. path opening.** For any  $\lambda \in \mathbb{N}$ ,  $N = 2$ , any  $b \in [2]$ , and any  $x \in \Sigma^2$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$  and any  $\ell \in [\lambda]$ ,

$$\Pr \left[ \begin{array}{l} \text{VerifyAcc}(\overline{\text{hk}}, v, \rho, b) = 1 \\ \wedge v_{0,b} \neq v_0 \end{array} : \begin{array}{l} ((\text{hk}_d, \text{td}_d) \leftarrow \text{Gen}(1^\lambda, 2, \Sigma, b))_{d \in [\ell]}, \\ (v, \rho) \leftarrow \mathcal{A}(\overline{\text{hk}}), \\ v_0 = \text{Extract}'((\text{td}_d)_{d \in [\ell]}, v) \end{array} \right] \leq \text{negl}(\lambda).$$

where  $\overline{\text{hk}} = (\text{hk}_d)_{d \in [\ell]}$ ,  $\rho = ((v_{0,1}, v_{0,2}), \dots, (v_{\ell-1,1}, v_{\ell-1,2}))$ , and we define  $\text{VerifyAcc}(\overline{\text{hk}}, v, \rho, b) = 1$  if for  $d \in [\ell]$ ,  $\text{SEH.Hash}(\text{hk}_d, (v_{d-1,1}, v_{d-1,2})) = v_{d,b}$  where  $v_{\ell,b} = v$ . Additionally define  $\text{Extract}'((\text{td}_d)_{d \in [\ell]}, v)$  to recursively compute  $\text{Extract}(\text{td}_d, v_d) = v_{d-1}$  for  $d \in [\ell]$ , where  $v_\ell = v$ , and outputs  $v_0$ .



**Definition 2.4** (Almost Rate-1 SEH). A SEH scheme is said to be almost rate-1 if the hash value generated by  $\text{Hash}(\text{hk}, x)$  is of length  $(1 + c/\lambda) \cdot |I| \cdot |\Sigma| + \text{poly}(\lambda)$ .

**Remark 2.5** ([CG]). Assuming almost rate-1 seBARGs there exists almost rate-1 SEH.

### 2.3 Rate-1 Somewhere Extractable Batch Arguments (seBARGs)

**Syntax.** A (publicly verifiable and non-interactive) somewhere extractable batch argument scheme seBARG for an NP language  $\mathcal{L}$  (decided by relation  $\mathcal{R}$ ) consists of the following polynomial time algorithms:

**Setup** $(1^\lambda, k, n, i^*) \rightarrow \text{crs}$ . This is a probabilistic setup algorithm that takes as input a security parameter  $1^\lambda$ , number of instances  $k$ , input length  $n$ , and an index  $i^* \in [k]$ . It runs in time at most  $\text{poly}(\lambda, n, \log k)$  and outputs a common reference string  $\text{crs}$ .

**Prove** $(\text{crs}, x_1, \dots, x_k, w_1, \dots, w_k) \rightarrow \pi$ . This is a prover algorithm takes as input a  $\text{crs}$ ,  $k$  instances  $x_1, \dots, x_k$  and corresponding witnesses  $w_1, \dots, w_k$ , and outputs a proof  $\pi$ .

**Verify** $(\text{crs}, x_1, \dots, x_k, \pi) \rightarrow 0/1$ . The verification algorithm takes as input a common reference string  $\text{crs}$ ,  $k$  instances  $x_i$  for  $i \in [k]$ , and a proof  $\pi$ . It outputs 0 (reject) or 1 (accept).

**Definition 2.6** (seBARG). A somewhere-extractable batch argument scheme seBARG = (Setup, Prove, Verify) for  $\mathcal{L}$  is required to satisfy the following properties:

**Efficiency.** The size of the CRS and the proof is at most  $\text{poly}(\lambda, \log k, n, m)$ , where  $m$  is the witness length.

**Completeness.** For any  $\lambda \in \mathbb{N}$ , and any  $k = k(\lambda)$ ,  $n = n(\lambda)$  of size at most  $2^\lambda$ , any  $k$  instances  $x_1, \dots, x_k \in \mathcal{L}$ , and their corresponding witnesses  $w_1, \dots, w_k \in \{0, 1\}^m$ , and any index  $i^* \in [k]$ ,

$$\Pr \left[ \text{Verify}(\text{crs}, x_1, \dots, x_k, \pi) = 1 : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, k, n, i^*), \\ \pi \leftarrow \text{Prove}(\text{crs}, x_1, \dots, x_k, w_1, \dots, w_k) \end{array} \right] = 1.$$

**Index hiding.** For any PPT adversary  $\mathcal{A}$ , any polynomials  $k = k(\lambda)$  and  $n = n(\lambda)$ , and any indices  $i_0, i_1 \in [k]$  there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ b \leftarrow \mathcal{A}(\text{crs}) : \begin{array}{l} b \leftarrow \{0, 1\}, \\ \text{crs} \leftarrow \text{Setup}(1^\lambda, k, n, i_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

**Somewhere Extraction.** There exists a stateful PPT extractor  $\mathcal{E}$  such that for any PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for any polynomials  $k = k(\lambda)$  and  $n = n(\lambda)$ , and any index  $i^* \in [k]$ , for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{crs}, x_1, \dots, x_k, \pi) = 1 \\ \wedge \mathcal{R}(x_{i^*}, w^*) \neq 1 \end{array} : \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \mathcal{E}(1^\lambda, k, n, i^*) \\ (x_1, \dots, x_k, \pi) \leftarrow \mathcal{A}(\text{crs}) \\ w^* \leftarrow \mathcal{E}(\text{td}, \{x_i\}_{i \in [k]}, \pi) \end{array} \right] \leq \text{negl}(\lambda).$$

**Remark 2.7.** We note that the somewhere extraction property implies the following *semi-adaptive soundness* property which asserts that for any PPT adversary  $\mathcal{A}$ , any polynomials  $k = k(\lambda)$  and  $n = n(\lambda)$ , and any index  $i^* \in [k]$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{crs}, x_1, \dots, x_k, \pi) = 1 \\ \wedge x_{i^*} \notin \mathcal{L} \end{array} : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, k, n, i^*) \\ (x_1, \dots, x_k, \pi) \leftarrow \mathcal{A}(\text{crs}) \end{array} \right] \leq \text{negl}(\lambda).$$

**Definition 2.8** (Almost Rate-1 seBARG). An seBARG scheme ( $\text{Setup}, \text{Prove}, \text{Verify}$ ) is said to be almost rate-1 if the proof generated by  $\text{Prove}(\text{crs}, x_1, \dots, x_k, w_1, \dots, w_k)$  is of length  $(1 + c/\lambda)m + \text{poly}(\lambda)$  for some constant  $c$ .

**Remark 2.9** ([DGKV22, PP22, KLVW23]). Assuming either  $\text{LWE}/\text{DLIN}/\text{sub-exponential DDH}$  (and  $\text{QR}$ ), there exists almost rate-1 seBARGs.

## 2.4 2-to-1 Collision Resistant Hash Function

**Syntax.** A 2-to-1 collision resistant hash function (CRHF) consists of the following polynomial time algorithms:

$\text{Gen}(1^\lambda) \rightarrow \text{hk}$ . This is a probabilistic key generation algorithm that takes as input a security parameter  $1^\lambda$ , and outputs a hash key  $\text{hk}$ .

$\text{Hash}(\text{hk}, x) \rightarrow \text{h}$ . This is the hashing algorithm that takes as input the hash key  $\text{hk}$  and an input  $x \in \{0, 1\}^{2\lambda}$ , and outputs a hash value  $\text{h} \in \{0, 1\}^\lambda$ .

**Definition 2.10** (CRHF). A CRHF  $\text{H} = (\text{Gen}, \text{Hash})$  is required to satisfy the following property:

**Collision Resistance.** For any PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{Hash}(x_1) = \text{Hash}(x_2) \\ x_1 \neq x_2 \end{array} \wedge : \begin{array}{l} (\text{hk}) \leftarrow \text{Gen}(1^\lambda), \\ (x_1, x_2) \leftarrow \mathcal{A}(\text{hk}) \end{array} \right] \leq \text{negl}(\lambda).$$

## 2.5 Hash Tree

**Syntax.** A hash tree consists of the following polynomial time algorithms:

$\text{Gen}(1^\lambda) \rightarrow \text{hk}$ . This is a probabilistic key generation algorithm that takes as input a security parameter  $1^\lambda$ , and outputs a hash key  $\text{hk}$ .

$\text{Hash}(\text{hk}, x) \rightarrow (\text{rt}, \text{tree})$ . This is the hashing algorithm that takes as input the hash key  $\text{hk}$  and an input  $x \in \{0, 1\}^*$ , and outputs a hash root  $\text{rt}$  and a hash tree  $\text{tree}$ .

$\text{Read}(\text{hk}, \text{tree}, i) \rightarrow (b, \text{rop})$ . This is the hash opening algorithm that takes as input the hash key  $\text{hk}$ , a hash tree  $\text{tree}$ , and an index  $i$ , and outputs a bit  $b$  and a reading opening  $\text{rop}$ .

$\text{Write}(\text{hk}, \text{tree}, i, b) \rightarrow (\text{rt}', \text{tree}', \text{wop})$ . This is the writing algorithm that takes as input the hash key  $\text{hk}$ , a hash tree  $\text{tree}$ , an index  $i$  and a bit  $b$ , and outputs a hash root  $\text{rt}'$ , a hash tree  $\text{tree}'$ , a writing opening  $\text{wop}$ .

$\text{VfyRd}(\text{hk}, \text{rt}, i, b, \text{rop}) \rightarrow 0/1$ . This is a read-verification algorithm that takes as input the hash key  $\text{hk}$ , a hash root  $\text{rt}$ , an index  $i$ , a bit  $b$  and a reading opening  $\text{rop}$ , and outputs 0 (reject) or 1 (accept).

$\text{VfyWt}(\text{hk}, \text{rt}, i, b, \text{rt}', \text{wop}) \rightarrow 0/1$ . This is a write-verification algorithm that takes as input the hash key  $\text{hk}$ , a hash root  $\text{rt}$ , an index  $i$ , a bit  $b$ , a new hash root  $\text{rt}'$  and a writing opening  $\text{wop}$ , and outputs 0 (reject) or 1 (accept).

**Definition 2.11** (Hash Tree). A hash tree  $\text{HT} = (\text{Gen}, \text{Hash}, \text{Read}, \text{Write}, \text{VfyRd}, \text{VfyWt})$  is required to satisfy the following properties:

**Efficiency.** The size of hash key  $\text{hk}$  is at most  $\text{poly}(\lambda)$ , the size of hash root  $\text{rt}$  is  $\lambda$ , and the size of openings is at most  $\text{poly}(\lambda, \log n)$  where  $n$  is the input size.

**Reading Soundness.** For any PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{VfyRd}(\text{hk}, \text{rt}, i, b^{(1)}, \text{rop}^{(1)}) = 1 \wedge \\ \text{VfyRd}(\text{hk}, \text{rt}, i, b^{(2)}, \text{rop}^{(2)}) = 1 \wedge \\ b^{(1)} \neq b^{(2)} \end{array} : \begin{array}{l} \text{hk} \leftarrow \text{Gen}(1^\lambda), \\ (\text{rt}, i, b^{(1)}, \text{rop}^{(1)}, b^{(2)}, \text{rop}^{(2)}) \leftarrow \mathcal{A}(\text{hk}) \end{array} \right] \leq \text{negl}(\lambda).$$

**Writing Soundness.** For any PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{VfyWt}(\text{hk}, \text{rt}, i, b, \text{rt}^{(1)}, \text{wop}^{(1)}) = 1 \wedge \\ \text{VfyWt}(\text{hk}, \text{rt}, i, b, \text{rt}^{(2)}, \text{wop}^{(2)}) = 1 \wedge \\ \text{rt}^{(1)} \neq \text{rt}^{(2)} \end{array} : \begin{array}{l} (\text{hk}) \leftarrow \text{Gen}(1^\lambda), \\ (\text{rt}, i, b, \text{rt}^{(1)}, \text{wop}^{(1)}, \text{rt}^{(2)}, \text{wop}^{(2)}) \leftarrow \mathcal{A}(\text{hk}) \end{array} \right] \leq \text{negl}(\lambda).$$

**Remark 2.12.** For all the algorithms and definitions above, instead of a single bit  $b$  at location  $i$  we can consider a set  $(b_{i_j})_{i_j \in I}$  at locations in the set  $I = (i_1, \dots, i_t)$ . Then the size of opening will grow linearly with  $t$ .

**Remark 2.13.** Reading soundness implies collision resistance of the hash tree.

**Remark 2.14** ([Mer87]). Hash trees can be constructed using any family  $\mathcal{F} : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$  of collision-resistance hash functions.

## 2.6 Non-Interactive Zero-Knowledge Arguments (NIZK)

Consider an NP language  $\mathcal{L} = \{x \mid \exists w : \mathcal{R}(x, w) = 1\}$  defined w.r.t. a relation  $\mathcal{R}$ .

**Syntax.** A non-interactive zero-knowledge (NIZK) argument consists of the following polynomial time algorithms:

$\text{Setup}(1^\lambda, 1^{n_x}, 1^{n_w}) \rightarrow \text{crs}$ . The probabilistic setup algorithm takes as input a security parameter  $\lambda$ , an instance length  $n_x$ , an witness length  $n_w$ , and outputs a common reference string  $\text{crs}$ .

$\text{Prove}(\text{crs}, x, w) \rightarrow \pi$ . The prover algorithm takes as input a common reference string  $\text{crs}$ , an instance  $x$ , and a witness  $w$  and outputs a proof  $\pi$ .

$\text{Verify}(\text{crs}, x, \pi) \rightarrow 0/1$ . The verifier algorithm takes as input a common reference string  $\text{crs}$ , an instance  $x$ , and a proof  $\pi$ . It outputs 0 (reject) or 1 (accept).

**Definition 2.15** (NIZK). A non-interactive zero-knowledge proof ( $\text{Setup}, \text{Prove}, \text{Verify}$ ) for  $\mathcal{L}$  is required to satisfy the following properties:

**Completeness.** For all  $\lambda, n_x, n_w \in \mathbb{N}$  and  $(x, w) \in \mathcal{R}$  where  $|x| = n_x$  and  $|w| = n_w$  we have:

$$\Pr[\text{Verify}(\text{crs}, x, \pi) = 1 : \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^{n_x}, 1^{n_w}), \pi \leftarrow \text{Prove}(\text{crs}, x, w)] = 1.$$

**Adaptive Soundness.** For any PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda, n_x \in \mathbb{N}$ :

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{crs}, x, \pi) = 1 \wedge \\ x \notin \mathcal{L} \end{array} : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^{n_x}, 1^{n_w}), \\ (x, \pi) \leftarrow \mathcal{A}(\text{crs}), |x| = n_x \end{array} \right] \leq \text{negl}(\lambda).$$

**Zero-Knowledge.** There exists a stateful PPT simulator  $\mathcal{S}$  such that for any PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda, n_x, n_w \in \mathbb{N}$ :

$$\left| \Pr[\mathcal{A}^{\text{Prove}(\text{crs}, \cdot, \cdot)}(\text{crs}) = 1 : \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^{n_x}, 1^{n_w})] - \Pr[\mathcal{A}^{\mathcal{O}^{\mathcal{S}}(\cdot, \cdot)}(\text{crs}) = 1 : \text{crs} \leftarrow \mathcal{S}(1^\lambda, 1^{n_x}, 1^{n_w})] \right| \leq \text{negl}(\lambda)$$

where  $\mathcal{O}^{\mathcal{S}}(x, w)$  outputs  $\mathcal{S}(x)$  if  $x \in \mathcal{L}$  and  $\perp$  otherwise.

**Knowledge Extractor.** There exists a stateful PPT extractor  $\mathcal{E}$  such that for any non-uniform PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda, n_x, n_w \in \mathbb{N}$ :

$$\Pr \left[ \begin{array}{l} \text{Verify}(\widetilde{\text{crs}}, x, \pi) = 1, \\ \wedge (\mathcal{R}(x, w) = 0 \vee |w| > n_w) \end{array} : \begin{array}{l} (\widetilde{\text{crs}}, \text{td}) \leftarrow \mathcal{E}(1^\lambda, 1^{n_x}, 1^{n_w}), \\ (x, \pi) \leftarrow \mathcal{A}(\widetilde{\text{crs}}), \\ |x| = n_x, \\ w \leftarrow \mathcal{E}(\text{td}, x, \pi) \end{array} \right] \leq \text{negl}(\lambda),$$

and  $\widetilde{\text{crs}}$  and  $\text{crs} \leftarrow \text{Setup}(1^\lambda, 1^{n_x}, 1^{n_w})$  are computationally indistinguishable.

**Remark 2.16** ([CW23, BW23, BKP+23]). Assuming PKE and seBARGs there exists NIZKs.

## 2.7 (Customized) Public-Key Encryption System

**Syntax.** A public key encryption (PKE) scheme for the message space  $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$  consists of the following polynomial time algorithms.

$\text{Setup}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$ . The probabilistic setup algorithm takes as input a security parameter  $1^\lambda$  and outputs the public and secret key pair  $(\text{pk}, \text{sk})$ .

$\text{Enc}(\text{pk}, m) \rightarrow \text{ct}$ . The probabilistic encryption algorithm takes as input the public key  $\text{pk}$ , a message  $m \in \mathcal{M}_\lambda$ , and outputs the ciphertext  $\text{ct}$ .

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m'$ . The decryption algorithm takes as input secret key  $\text{sk}$ , ciphertext  $\text{ct}$ , and outputs  $m'$ .

**Definition 2.17** (PKE). A public-key encryption system  $(\text{Setup}, \text{Enc}, \text{Dec})$  for  $m \in \mathcal{M}_\lambda$  is required to satisfy the following properties:

**Correctness.** For any  $\lambda \in \mathbb{N}$ ,  $m \in \mathcal{M}_\lambda$ , we have that  $\text{Dec}(\text{sk}, \text{ct}) = m$  where  $\text{ct} \leftarrow \text{Enc}(\text{pk}, m)$  and  $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ .

**Security.** For any stateful PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda \in \mathbb{N}$ :

$$\left| \Pr \left[ 1 \leftarrow \mathcal{A}^{\text{Enc}(\text{pk}, \cdot)}(1^\lambda, \text{pk}) \right] - \Pr \left[ 1 \leftarrow \mathcal{A}^{\text{Enc}(\text{pk}, 0^{|m|})}(1^\lambda, \text{pk}) \right] \right| \leq \text{negl}(\lambda)$$

where  $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ .

### 2.7.1 leakage resilient/RDM/randomness recoverable/rate-1 PKEs

**Definition 2.18** (PKE with Randomness Recovery). A public-key encryption system  $(\text{Setup}, \text{Enc}, \text{Dec})$  for  $m \in \mathcal{M}_\lambda$  is said to have randomness recovery if the decryption algorithm additionally outputs the randomness, i.e.,  $\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m; r)) \rightarrow (m, r)$ .

**Definition 2.19** (Leakage Resilient PKE). A public-key encryption system  $(\text{Setup}, \text{Enc}, \text{Dec})$  for  $m \in \mathcal{M}_\lambda$  is said to be leakage resilient if for any stateful PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda \in \mathbb{N}$ :

$$\left| \Pr \left[ 1 \leftarrow \mathcal{A}^{\text{Enc}'(\text{pk}, \cdot)}(1^\lambda, \text{pk}) \right] - \Pr \left[ 1 \leftarrow \mathcal{A}^{\text{Enc}'(\text{pk}, 0^{|m|}, \cdot)}(1^\lambda, \text{pk}) \right] \right| \leq \text{negl}(\lambda)$$

where  $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$  and  $\text{Enc}'(\text{pk}, m, f)$  outputs  $(\text{Enc}(\text{pk}, m; r), f(r))$ .

**Remark 2.20** ([PW08], §4.1). Assuming lossy trapdoor functions (LTDFs) there exists a leakage resilient PKE with randomness recovery. The construction given in section 4.1 of [PW08] immediately implies randomness recovery. Moreover, in their security analysis, in addition to  $c_1$  we can consider more leakage on the randomness, thus the min entropy of the randomness goes down by the length of the additional leakage. Therefore, we only need to increase the randomness size by the leakage size  $n_\ell$  ( $\lambda$  bits in our applications) and use  $(n + n_\ell, k + n_\ell)$ -LTDFs. Note that LTDFs with better parameter can be constructed by combining two LTDFs.

**Remark 2.21** ([PW08, FGK<sup>+</sup>10]). Assuming LWE, DLIN, and DDH there exists lossy trapdoor functions. Thus under the same set of assumptions there exists a leakage resilient PKE with randomness recovery.

**Definition 2.22** (Rate-1 PKE). A PKE scheme  $(\text{Setup}, \text{Enc}, \text{Dec})$  is said to be rate-1 if the ciphertext generated by  $\text{Enc}(\text{pk}, m)$  is of length  $|m| + \text{poly}(\lambda)$

**Remark 2.23** (Rate-1 leakage resilient PKE with randomness recovery.). Assuming PRGs and PKEs, there exists rate-1 PKEs. Furthermore, if we start with a leakage resilient PKE with randomness recovery, we get a rate-1 leakage resilient PKE with randomness recovery. The construction works as follows: (1) **Setup** compute  $\text{pk}' = (\text{pk}, G)$  where  $\text{pk} \leftarrow \text{PKE.Setup}(\lambda)$  and  $G$  is a PRG with  $3\lambda$ -bits seeds (that is secure if  $s$  has high entropy), (2) the encryption algorithm samples a PRG seed  $s$ , and computes  $\text{ct}_1 \leftarrow \text{PKE.Enc}(\text{pk}, s)$ ,  $\text{ct}_2 = G(s) \oplus m$ , and (3) the decryption algorithm first decrypts  $\text{ct}_1$  to find  $s$  and then computes  $m = \text{ct}_2 \oplus G(s)$ .

The security follows from the security of the underlying PKE and PRG. Additionally, if the underlying PKE is randomness recoverable, since the decryption algorithm also finds  $s$ , the resulting algorithm is also randomness recoverable. For leakage resilience we rely on leakage resilience of the underlying PKE, and high entropy of  $s$  given the leakage. Moreover, since the encrypted message is  $3\lambda$ -bits long, the required randomness for PKE is  $\text{poly}(\lambda)$ , so the resulting scheme is rate-1.

**Definition 2.24** (Randomness-Dependent Message (RDM) PKE). A public-key encryption system (Setup, Enc, Dec) for  $m \in \mathcal{M}_\lambda$  is said to be randomness-dependent message if for any stateful PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda \in \mathbb{N}$ :

$$\left| \Pr \left[ 1 \leftarrow \text{exp}_0^{\mathcal{A}}(1^\lambda, \text{pk}) \right] - \Pr \left[ 1 \leftarrow \text{exp}_1^{\mathcal{A}}(1^\lambda, \text{pk}) \right] \right| \leq \text{negl}(\lambda)$$

where definitions of  $\text{exp}_0^{\mathcal{A}}$  and  $\text{exp}_1^{\mathcal{A}}$  are provided in Figure 1.

$\text{exp}_0^{\mathcal{A}}(1^\lambda, \text{pk})$ . This is the real experiment where  $\mathcal{A}$  first receives  $\text{pk} \leftarrow \text{Setup}(1^\lambda)$  from the challenger. Then  $\mathcal{A}$  iteratively sends  $f_i$  to the challenger, receives  $f_i(r)$ , then sends  $m_i$  and receives  $\text{Enc}(\text{pk}, m_i; r)$ . After this,  $\mathcal{A}$  outputs guess  $b'$ . Output  $b'$ .

$\text{exp}_1^{\mathcal{A}}(1^\lambda, \text{pk})$ . This is the ideal experiment where  $\mathcal{A}$  first receives  $\text{pk} \leftarrow \text{Setup}(1^\lambda)$  from the challenger. Then  $\mathcal{A}$  iteratively sends  $f_i$  to the challenger, receives  $f_i(r)$ , then sends  $m_i$  and receives  $\text{Enc}(\text{pk}, 0^{m_i}; r)$ . After this,  $\mathcal{A}$  outputs guess  $b'$ . Output  $b'$ .

Figure 1: Real and ideal experiments for RDM-PKE security.

**Remark 2.25** ([BCPT13]). Assuming PKE there exists an RDM-PKE.

### 3 Zero-Knowledge Incrementally Verifiable Streaming Computation (zk-IVsC)

In this section we first formally define zk-IVsC, and construct and analyze our zk-IVsC. We finally discuss how to further extend and optimize our zk-IVsC.

#### 3.1 Definition.

**Syntax.** A zero-knowledge incrementally verifiable streaming computation (zk-IVsC) scheme for a RAM machine  $\mathcal{M}$  consists of the following PPT algorithms:

$\text{Gen}(1^\lambda, n, S, T) \rightarrow \text{crs}$ . The probabilistic setup algorithm takes as input a security parameter  $1^\lambda$ , the input length  $n$ , the maximum configuration size  $S$ , and the running time of machine  $T$ , and outputs a common reference string  $\text{crs}$ .

$\text{Update}(\text{crs}, z_k, \text{st}_{k-1}, \text{hz}_{k-1}, \Pi_{k-1}; \text{ps}_k) \rightarrow (\text{st}_k, \text{hz}_k, \Pi_k)$ . The update algorithm takes as input a CRS  $\text{crs}$ , an input  $z_k$ , an intermediate state  $\text{st}_{k-1}$ , an intermediate digest  $\text{hz}_{k-1}$ , a proof  $\Pi_{k-1}$ , and some private state/randomness  $\text{ps}_k$ . It outputs an updated state  $\text{st}_k$ , an updated digest  $\text{hz}_k$ , and an updated proof  $\Pi_k$ . Note that this algorithm has RAM access to the memory.

NOTE. The update algorithm takes  $\text{ps}_k$  as explicit randomness, and it could use more randomness. However, we are not making the entire randomness explicit for simplicity.



$\text{Digest}(\text{crs}, \bar{z}, \bar{\text{ps}}) \rightarrow \text{hz}$ . The digest algorithm takes as input a CRS  $\text{crs}$ , a set of inputs  $\bar{z} = (z_k)_{k \in [K]}$ , and a set of private states  $\bar{\text{ps}} = (\text{ps}_k)_{k \in [K]}$ , and outputs a hash digest  $\text{hz}$ .

$\text{Verify}(\text{crs}, (\text{hz}, \text{out}), \Pi) \rightarrow 0/1$ . The verifier algorithm takes as input a CRS  $\text{crs}$ , a hash digest  $\text{hz}$ , an output  $\text{out}$ , and a proof  $\Pi$ , and outputs a bit to signal whether the proof is valid or not.

**Remark 3.1** (Prover). We will define the prover  $\text{Prove}(\text{crs}, (z_k)_{k \in [K]}; (\text{ps}_k)_{k \in [K]})$  for a zk-IVsC inductively as follows (where  $\text{st}_0$  is the initial state and  $\Pi_0 = \epsilon$ ):

If  $K = 1$  then output  $(\text{st}_1, \text{hz}_1, \Pi_1) \leftarrow \text{Update}(\text{crs}, z_1, \text{st}_0, \Pi_0; \text{ps}_1)$ , otherwise inductively compute  $(\text{st}_{K-1}, \text{hz}_{K-1}, \Pi_{K-1}) = \text{Prove}(\text{crs}, (z_k)_{k \in [K-1]}; \bar{\text{ps}}_{K-1})$  and output  $(\text{st}_K, \text{hz}_K, \Pi_K)$  that is computed as  $(\text{st}_K, \text{hz}_K, \Pi_K) \leftarrow \text{Update}(\text{crs}, z_K, \text{st}_{K-1}, \text{hz}_{K-1}, \Pi_{K-1}; \text{ps}_K)$ .

**Informal definition.** Informally speaking, zk-IVsC requires that for a RAM computation  $\mathcal{M}(\bar{z}) = \text{cf}$  where  $\text{out} = \text{cf}[1, \dots, n_{\text{out}}]$ , there exists an efficient prover who generates a proof whose size is independent of computation time and space. Moreover, no PPT adversary can find (1) a collision: i.e. two different inputs whose digests are the same, (2) a valid proof w.r.t. an *honest* digest of  $\bar{z}$  for some incorrect output  $\text{out}$ , (3) valid proofs  $\Pi^{(0)}$  and  $\Pi^{(1)}$  w.r.t. a unique digest  $\text{hz}$  for  $\text{out}^{(0)} \neq \text{out}^{(1)}$ , (4) any information about  $\bar{z}$  from  $(\text{hz}, \Pi)$  more than what is revealed by  $\text{cf}$ .

**Definition 3.2** (zk-IVsC). A zero-knowledge incrementally verifiable streaming computation scheme  $(\text{Gen}, \text{Update}, \text{Digest}, \text{Verify})$  for  $\mathcal{M}$  has to satisfy the following properties:

**Completeness.** For every  $\lambda, n, S, T \in \mathbb{N}$ , any  $\bar{z} = (z_k)_{k \in [K]}$  s.t.  $z_k \in \{0, 1\}^n$ ,

$$\Pr \left[ \begin{array}{l} (\bar{z}, \text{out}) \in \mathcal{L}_{\mathcal{M}, T} \wedge \\ \text{hz} = \text{Digest}(\text{crs}, \bar{z}, \bar{\text{ps}}) \wedge \\ \text{Verify}(\text{crs}, (\text{hz}, \text{out}), \Pi) = 1 \end{array} : \begin{array}{l} \text{crs} \leftarrow \text{Gen}(1^\lambda, n, S, T) \\ (\text{st}, \text{hz}, \Pi) \leftarrow \text{Prove}(\text{crs}, \bar{z}; \bar{\text{ps}}) \\ \text{find } \text{cf} \in \text{st}, \text{out} = \text{cf}[1, \dots, n_{\text{out}}] \end{array} \right] = 1.$$

**Efficiency.** For the completeness experiment above,

- Size of  $\text{crs}$ , and setup running time are  $\text{poly}(\lambda, \log S, \log n)$ .
- Size of  $\Pi$ , and its verification time are  $\text{poly}(\lambda, \log S, \log n, \log T, \log K)$ .
- Update's running time is  $\text{poly}(\lambda, n) + \text{poly}(\lambda, S) + T \cdot \text{poly}(\lambda, \log S, \log n, \log T, \log K)$ . (We also consider - and later achieve - a more efficient update that runs in  $T \cdot \text{poly}(\lambda, \log S, \log n, \log T, \log K)$ , when it is allowed to preprocess the  $n$ -bit input.)

**Collision Resistance.** For any stateful PPT adversary  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{hz}^{(0)} = \text{hz}^{(1)} \wedge \\ \bar{z}^{(0)} \neq \bar{z}^{(1)} \end{array} : \begin{array}{l} (n, S, T) \leftarrow \mathcal{A}(1^\lambda), \text{crs} \leftarrow \text{Gen}(1^\lambda, n, S, T) \\ (\bar{z}^{(b)}, \bar{\text{ps}}^{(b)})_{b \in \{0,1\}} \leftarrow \mathcal{A}(\text{crs}) \\ (\text{hz}^{(b)} = \text{Digest}(\text{crs}, \bar{z}^{(b)}, \bar{\text{ps}}^{(b)}))_{b \in \{0,1\}} \end{array} \right] \leq \text{negl}(\lambda).$$

**Soundness.** For any admissible stateful PPT adversary  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} (\bar{z}, \text{out}) \notin \mathcal{L}_{\mathcal{M}, T} \wedge \\ \text{Verify}(\text{crs}, (\text{hz}, \text{out}), \pi) = 1 \end{array} : \begin{array}{l} (n, S, T) \leftarrow \mathcal{A}(1^\lambda), \text{crs} \leftarrow \text{Gen}(1^\lambda, n, S, T) \\ (\bar{z}, \text{out}, \bar{\text{ps}}, \Pi, 1^{K \cdot T}) \leftarrow \mathcal{A}(\text{crs}), \text{hz} = \text{Digest}(\text{crs}, \bar{z}, \bar{\text{ps}}) \end{array} \right] \leq \text{negl}(\lambda).$$

where an adversary is admissible if  $\forall k \in [K], z_k \in \{0, 1\}^n$ .

**Strong Soundness.** For any stateful PPT adversary  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{crs}, (\text{hz}, \text{out}^{(0)}), \Pi^{(0)}) = 1 \wedge \\ \text{Verify}(\text{crs}, (\text{hz}, \text{out}^{(1)}), \Pi^{(1)}) = 1 \wedge \\ \text{out}^{(0)} \neq \text{out}^{(1)} \end{array} : \begin{array}{l} (n, S, T) \leftarrow \mathcal{A}(1^\lambda), \\ \text{crs} \leftarrow \text{Gen}(1^\lambda, n, S, T), \\ (\text{hz}, (\text{out}^{(b)}, \Pi^{(b)})_{b \in \{0,1\}}, 1^{K \cdot T}) \leftarrow \mathcal{A}(\text{crs}) \end{array} \right] \leq \text{negl}(\lambda).$$

**Remark 3.3.** Strong soundness implies the normal soundness by having the challenger run  $\mathcal{M}$  on  $((z_k)_{k \in [K]})$  (received from the adversary), and generate  $(\text{out}^*, \Pi^*)$  corresponding to the correct computation. Additionally, strong soundness implies collision resistance property.

**Zero-Knowledge.** There exists a stateful PPT simulator  $\mathcal{S}$  such that for any PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} (n, S, T) \leftarrow \mathcal{A}(1^\lambda), b \leftarrow \{0, 1\} \\ \text{crs}_0 \leftarrow \text{Gen}(1^\lambda, n, S, T), \text{crs}_1 \leftarrow \mathcal{S}(1^\lambda, n, S, T), \\ b = b' : \bar{z} \leftarrow \mathcal{A}(\text{crs}_b), \text{cf} = \mathcal{M}(\bar{z}; 1^T) \\ (\text{st}_0, \text{hz}_0, \Pi_0) \leftarrow \text{Prove}(\text{crs}_0, \bar{z}), (\text{st}_1, \text{hz}_1, \Pi_1) \leftarrow \mathcal{S}(\text{crs}_1, \text{cf}), \\ b' \leftarrow \mathcal{A}(\text{st}_b, \text{hz}_b, \Pi_b) \end{array} \right] \leq \text{negl}(\lambda).$$

## 3.2 Construction

In this section we first define the notation, then present an overview of our construction, followed by the building blocks and the construction itself.

**Notations.** We implement our algorithms in the RAM model and to indicate a RAM access, we use *underline in the construction*. For ease of exposition (and w.l.o.g.) we assume  $T$  is a power of 2, every user has an input of size  $n$ , and the work tape is of size  $S$ .

For a RAM machine  $\mathcal{M}$  we use  $\text{cf} = (\text{W}, \text{q}, \text{ridx})$  to denote a configuration and  $\text{pcf} = (\text{rt}, \text{q}, \text{ridx})$  to denote a pseudo-configuration, where  $\text{W}$  is the work tape,  $\text{rt}$  is a short digest of the work tape  $\text{W}$  (root of the hash tree computed on  $\text{W}$ ),  $\text{q}$  is the state of the RAM machine, and  $\text{ridx}$  (reading index) is the pointer on the memory. For a configuration (resp. pseudo-configuration) that corresponds to user  $k$  after  $t$  steps of the computation, we use  $\text{cf}_{k,t} = (\text{W}_{k,t}, \text{q}_{k,t}, \text{ridx}_{k,t})$  (resp.  $\text{pcf}_{k,t} = (\text{rt}_{k,t}, \text{q}_{k,t}, \text{ridx}_{k,t})$ ). Note that we are splitting the memory into two parts, input tape and work tape, and often we denote the digest of the input tape by  $\text{rt}_{\text{in}}$ . Moreover, if  $\text{ridx} \leq n$  then  $\mathcal{M}$  is reading the  $\text{ridx}^{\text{th}}$  index from the input tape, otherwise,  $\mathcal{M}$  is reading the  $(\text{ridx} - n)^{\text{th}}$  index of the work tape.

**Construction overview.** We break our zk-IVsC proof to small steps, and show how each step builds upon the previous step. Our starting point is the fact the hash trees allow one to succinctly prove a single step of RAM computation. Namely, consider a RAM transition  $\text{cf} \xrightarrow{z} \text{cf}'$  which consists of a single-bit reading from input/work tape  $z_i/\text{W}_i$  and a single-bit writing in the work tape. One can succinctly prove such a transition by generating a reading proof  $\text{rop}$  and a writing proof  $\text{wop}$  corresponding to the digest of the input/work tapes. Hence,  $(\text{rop}, \text{wop})$  is a proof for the RAM transition  $\text{pcf} \xrightarrow{\text{rt}_{\text{in}}} \text{pcf}'$ . Namely, we use  $\text{is-nxt-cnfg}_{\mathcal{M}}(\text{hk}_{\text{ht}}, \text{rt}_{\text{in}}, \text{pcf}, \text{pcf}', \text{rbit}, \text{rop}, \text{wop}) \rightarrow 0/1$  (see Fig. 2) (where  $\text{hk}_{\text{ht}}$  is a hash-key for a hash tree) as a function that efficiently checks the adjacency of two [consecutive] pseudo-configurations w.r.t. their corresponding set of reading and writing proofs  $(\text{rop}, \text{wop})$ .

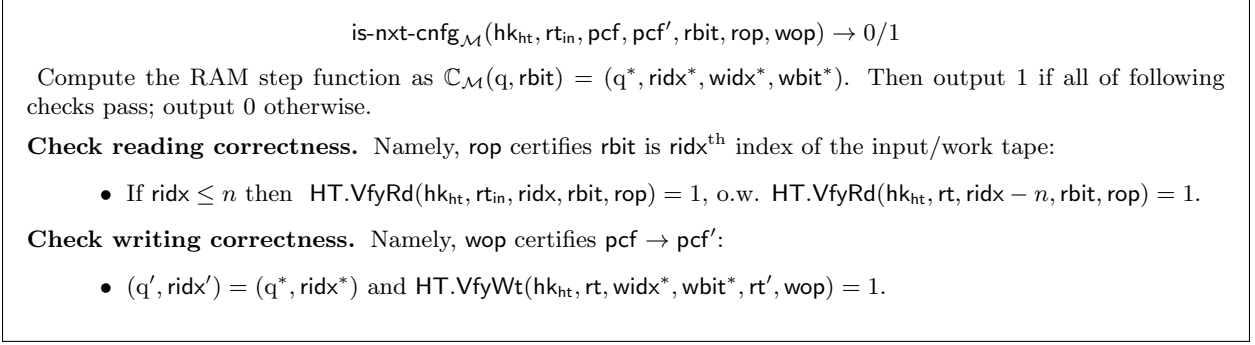


Figure 2: Description of  $\text{is-nxt-cnfg}_{\mathcal{M}}$ .

Next, we show how to make such a proof zero-knowledge by encrypting the pseudo-configurations and input digest and generating a NIZK proof  $\pi_{\text{zk}}$  for such a transition. Namely, we let  $(\text{ct}_{\text{in}}, \text{epcf}, \text{epcf}')$  be encryptions of  $(\text{rt}_{\text{in}}, \text{pcf}, \text{pcf}')$ , and by  $\text{epcf} \xrightarrow{\text{ct}_{\text{in}}} \text{epcf}'$  we denote there exist  $(\text{rt}_{\text{in}}, \text{pcf}, \text{pcf}')$ , s.t.  $\text{pcf} \xrightarrow{\text{rt}_{\text{in}}} \text{pcf}'$ . Thus we define  $\mathcal{L}_{\text{zk}}$  (see Fig. 3) to be the NIZK language for such a transition which checks the correctness of all the encryptions (w.r.t. an encryption public-key  $\text{pk}$  and the corresponding randomnesses), in addition to checking whether  $\text{pcf} \xrightarrow{\text{rt}_{\text{in}}} \text{pcf}'$  holds.

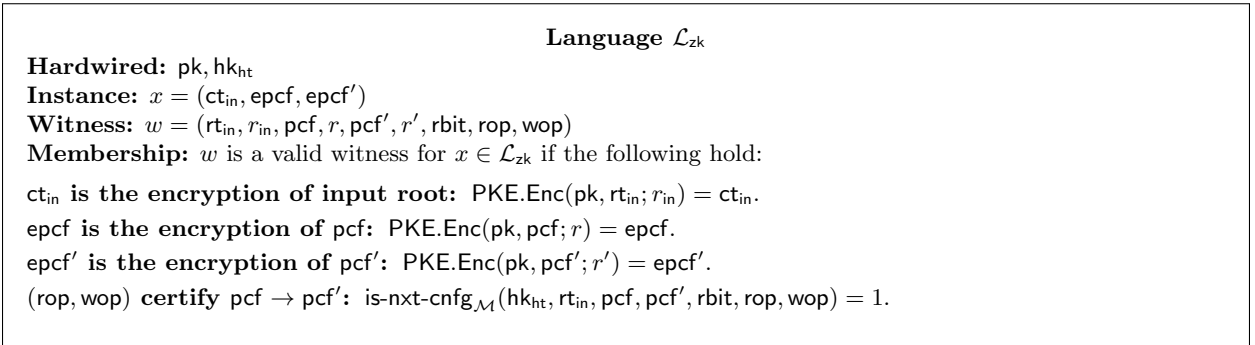


Figure 3: Description of language  $\mathcal{L}_{\text{zk}}$ .

We let  $(\mathbf{v} = (\text{epcf}, \text{epcf}'), \pi_{\text{zk}})$  to be a level 0 proof (w.r.t.  $\text{ct}_{\text{in}}$ ), i.e.  $\pi_0 \in \Pi$  in our zk-IVsC construction. Next we show how to iteratively merge any two proofs  $\pi_{\ell-1,1}$  and  $\pi_{\ell-1,2}$  (each for  $2^\ell$  steps of the computation) to get a proof  $\pi_\ell$  (for  $2^{\ell+1}$  steps of the computation). First note that by our choice of  $T$  (being a power of 2), a proof for  $T$  steps of the computation (performed by any user  $k$ ) would only be a proof  $\pi_{\log T-1}$ . Thus until level  $\ell-2$  we use the same  $\text{ct}_{\text{in}}$  to generate/check the proofs, and only after that the  $\text{ct}_{\text{in}}$  values are being merges. Therefore, in our construction, we only see  $\mathbf{h}$  in the proofs  $\pi_\ell$  for  $\ell \geq \log T - 1$  (where  $\mathbf{h} = \text{ct}_{\text{in}}$  at level  $\log T - 1$ ). Moreover, every proof has two path opening from  $\mathbf{v}$  to its leftmost and rightmost childs/epcfs. Since the leftmost and rightmost childs of  $\mathbf{v}$  at level 0 are part of  $\mathbf{v}$  itself, we do not consider path openings for  $\pi_0$ . Next we show how to merge two level  $\ell-1$  proofs for  $\ell \geq \log T$ . Note that this is the more general case and includes all the checks we need to consider for merging.

First note that our construction samples separate SEH keys and seBARG CRSs for every level of merging, hence we use  $\text{crs}_{\text{bp}, \ell}$  and  $\text{hk}_{\text{cf}, \ell}$  for merging level  $\ell-1$  proofs. Now consider two proofs  $\pi_{\ell-1,1} = (\mathbf{v}_1, \rho_{1,1}, \rho_{1,2}, \mathbf{h}_1, \hat{\pi}_1)$  and  $\pi_{\ell-1,2} = (\mathbf{v}_2, \rho_{2,1}, \rho_{2,2}, \mathbf{h}_2, \hat{\pi}_2)$ . To merge such proofs, (1) we use a CRHF to hash  $(\mathbf{h}_1, \mathbf{h}_2)$  to get  $\mathbf{t-h}$ , (2) we use SEH to hash  $(\mathbf{v}_1, \mathbf{v}_2)$  to get  $\mathbf{t-v}$ , (3) we compute an

seBARG proof  $\mathbf{t}\text{-}\pi^*$  for  $\mathcal{L}_{\text{bp},\ell}$  (see Fig. 4) that proves (a) the correctness of  $\mathbf{t}\text{-h}$  w.r.t.  $(\mathbf{h}_1, \mathbf{h}_2)$ , (b) the correctness of  $\mathbf{t}\text{-v}$  w.r.t.  $(\mathbf{v}_1, \mathbf{v}_2)$ , (c) the validity of  $(\mathbf{v}_i, \mathbf{h}_i)$  w.r.t.  $\hat{\pi}_i$ , (d) the validity of  $\rho_{i,1}$  and  $\rho_{i,2}$  w.r.t.  $\mathbf{v}$ , and (e) the equality of  $\mathbf{v}_1$  and  $\mathbf{v}_2$  at their intersection, i.e. the rightmost child of  $\mathbf{v}_1$  is the same as the leftmost child of  $\mathbf{v}_2$ , and (4) we extend  $\rho_{1,1}$  (resp.  $\rho_{2,2}$ ) to get  $\mathbf{t}\text{-}\rho_1$  (resp.  $\mathbf{t}\text{-}\rho_2$ ), path openings to the leftmost (resp. rightmost) child of  $\mathbf{t}\text{-v}$ .

<b>Language <math>\mathcal{L}_{\text{bp},\ell}</math></b>
<b>Hardwired:</b> $\text{pk}, \text{hk}_{\text{ht}}, \text{hk}_{\text{ch}}, \text{crs}_{\text{zk}}, (\text{crs}_{\text{bp},j})_{j \in [\ell-1]}, (\text{hk}_{\text{cf},j})_{j \in [\ell]}$ .
<b>Instance:</b> $x = (\mathbf{h}, \mathbf{v}, i)$ .
<b>Witness:</b> $w = (\mathbf{v}_1, \mathbf{v}_2, \hat{w})$ .
<b>Membership:</b> $w$ is a valid witness for $x \in \mathcal{L}_{\text{bp},\ell}$ if the following hold:
$\mathbf{v}$ is the hash of $(\mathbf{v}_1, \mathbf{v}_2)$ . Namely, $\text{SEH.Hash}(\text{hk}_{\text{cf},\ell}, (\mathbf{v}_1, \mathbf{v}_2)) = \mathbf{v}$ .
$\hat{w}$ certifies the validity of $\mathbf{v}_i$ . Namely:
<ul style="list-style-type: none"> <li>• If <math>\ell = 1</math>, then parse <math>\mathbf{v}_i = (\text{epcf}, \text{epcf}')</math> and <math>\hat{w} = \pi</math>, and check <math>\text{NIZK.Verify}(\text{crs}_{\text{zk}}, (\mathbf{h}, \text{epcf}, \text{epcf}'), \pi) = 1</math>.</li> <li>• If <math>1 &lt; \ell \leq \log T - 1</math>, then parse <math>\hat{w} = (\rho_1, \rho_2, \pi)</math>, and check <math>\text{BARG.Verify}(\text{crs}_{\text{bp},\ell-1}, (\mathbf{h}, \mathbf{v}_i, j)_{j \in [2]}, \pi) = 1</math>.</li> <li>• If <math>\log T \leq \ell</math>, then parse <math>\hat{w} = (\rho_1, \rho_2, \mathbf{h}_1, \mathbf{h}_2, \pi)</math>, and check <math>\text{BARG.Verify}(\text{crs}_{\text{bp},\ell-1}, (\mathbf{h}_i, \mathbf{v}_i, j)_{j \in [2]}, \pi) = 1</math> and <math>\text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\mathbf{h}_1, \mathbf{h}_2)) = \mathbf{h}</math>.</li> </ul>
$\rho_i \in \hat{w}$ is a valid path from $\mathbf{v}_i$ . Namely, if $\ell > 1$ check $\text{SEH.VerifyAcc}((\text{hk}_{\text{cf},d})_{d \in [\ell-1]}, \mathbf{v}_i, \rho_i, 3 - i)$ .
<b>Intersection of <math>\mathbf{v}_1</math> and <math>\mathbf{v}_2</math> is equal.</b> Namely, for $i \in [2]$ , if $\ell = 1$ parse $\mathbf{v}_i = (\mathbf{v}'_{i,1}, \mathbf{v}'_{i,2})$ , otherwise, parse $\rho_i = ((\mathbf{v}_{0,i,1}, \mathbf{v}_{0,i,2}), \dots, (\mathbf{v}_{\ell-2,i,1}, \mathbf{v}_{\ell-2,i,2}))$ and $\mathbf{v}_{0,i,3-i} = (\mathbf{v}'_{i,1}, \mathbf{v}'_{i,2})$ . Then check $\mathbf{v}'_{1,2} = \mathbf{v}'_{2,1}$ .

Figure 4: Description of language  $\mathcal{L}_{\text{bp},\ell}$ .

Now we put everything together to get  $\mathbf{t}\text{-}\pi = (\mathbf{t}\text{-v}, \mathbf{t}\text{-}\rho_1, \mathbf{t}\text{-}\rho_2, \mathbf{t}\text{-h}, \mathbf{t}\text{-}\pi^*)$ , which in the next iteration, (if  $\pi_\ell \in \Pi$  is not empty,) will be parsed as  $\pi_{\ell,2}$  and get further merged with the already existing level  $\ell$  proof in  $\Pi$ . This iterative process ends when  $\pi_\ell \in \Pi$  is empty and then we let  $(\epsilon, \dots, \epsilon, \mathbf{t}\text{-}\pi, \pi_{\ell+1}, \dots, \pi_\lambda)$  to be the new proof. The last piece of the puzzle is to generate a proof  $\pi_{\text{out}}$  for the actual output out of  $\mathcal{M}$ . We do this by generating a hash tree reading proof w.r.t. the final pcf in the computation/proof.

The verification checks the validity of all the proofs, and verifies every consecutive non-empty hash values  $\mathbf{v} \in \pi_\ell$  and  $\mathbf{v}' \in \pi_{\ell'}$  are equal in their intersecting epcf (w.r.t. their corresponding path openings), and the starting epcf in the proof is consistent with the initial configuration and the output is consistent with the last epcf in the proof (w.r.t.  $\pi_{\text{out}}$ ). We summarize our notation in Table 1.

**Building blocks and parameters.** Let HT be a hash tree with output root of size  $\lambda$  on any input, CRHF be a collision resistant  $\text{Hash} : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ , seBARG be an almost rate-1 seBARG, SEH be an almost rate-1 SEH, NIZK be a NIZK, and PKE be a public-key encryption. We use the following parameters to instantiate such primitives. Let  $|\text{pcf}| = \lambda + |\mathbf{q}| + \log(n + S)$ ,  $|\text{epcf}| = |\text{PKE.Enc}(\text{pk}, 0^{|\text{pcf}|})|$ , and  $n_{\text{cf},0} = 2 \cdot |\text{epcf}|$ . For  $\ell \in [\lambda]$ , let  $n_{\text{cf},\ell} = |\text{SEH.Hash}(\text{hk}_{\text{cf},\ell}, 0^{2 \cdot n_{\text{cf},\ell-1}})|$ ,  $n_{\text{bp},\ell} = \lambda + n_{\text{cf},\ell} + 1$ , and  $\Sigma_{\text{cf}}^\ell = \{0, 1\}^{n_{\text{cf},\ell}}$  be the alphabet for block SEH. Moreover, let  $|\text{ct}_{\text{in}}| = |\text{PKE.Enc}(\text{pk}, 0^\lambda)|$ ,  $|r|$  be the size of randomness used by PKE, and  $n_{\text{zk},x} = |\text{ct}_{\text{in}}| + 2|\text{epcf}|$  and  $n_{\text{zk},w} = \lambda + 2|\text{pcf}| + 3|r| + 1 + |\text{rop}| + |\text{wop}|$ .

Table 1: zk-IVsC Noatation Summary

pk/sk	public-key/secret-key for encryption
hk <sub>ht</sub> /hk <sub>ch</sub>	hash-key for hash tree/collision-resistant hash function
crs <sub>zk</sub> /td <sub>zk</sub>	CRS/trapdoor for the NIZK scheme
crs <sub>bp,ℓ</sub> /td <sub>bp,ℓ</sub>	CRS/trapdoor for seBARG used to generate proofs at level ℓ
hk <sub>cf,ℓ</sub> /td <sub>cf,ℓ</sub>	hash-key/trapdoor for SEH used to hash level ℓ − 1 hashes of epcf <sub>s</sub>
z/W/q	RAM machine's input tape/work tape/state
rbit/ridx	RAM reading bit/reading index
widx/wbit	RAM writing bit/writing index
cf	configuration = (W, q, ridx)
pcf	pseudo-configuration = (rt, q, ridx) where rt is hash root of W
epcf	encrypted pseudo-configuration ← PKE.Enc(pk, pcf)
pcf <sub>1,0</sub>	the universal starting pseudo-configuration
rop/wop	reading/writing proof w.r.t. hash roots
rt <sub>in</sub> /tree <sub>in</sub>	hash root/hash tree of RAM input z
r <sub>in</sub> /ct <sub>in</sub>	randomness/ciphertext of rt <sub>in</sub> , i.e. ct <sub>in</sub> = PKE.Enc(pk, rt <sub>in</sub> ; r <sub>in</sub> )
h	powers-of-2 hashes of ct <sub>in</sub> values, i.e. at level log T − 1, h = ct <sub>in</sub> and at any level ℓ ≥ log T, h = CRHF.Hash(hk <sub>ch</sub> , (h <sub>1</sub> , h <sub>2</sub> )) for level ℓ − 1 hashes h <sub>1</sub> and h <sub>2</sub>
v	powers-of-2 hashes of epcf values, i.e. at level 0, v = (epcf, epcf <sup>ℓ</sup> ) and at any level ℓ ≥ 1, v = SEH.Hash(hk <sub>cf,ℓ</sub> , (v <sub>1</sub> , v <sub>2</sub> )) for level ℓ − 1 hashes v <sub>1</sub> and v <sub>2</sub>
ρ	a path opening from hash value v to its rightmost or leftmost child
π <sub>ℓ</sub> ∈ Π	= (v, ρ <sub>1</sub> , ρ <sub>2</sub> , h, π'), a level ℓ proof in the powers-of-2 structure where h is empty if ℓ < log T − 1 and if ℓ = 0, π' is a NIZK proof, o.w. it is a seBARG proof

**Construction 3.4.** In this construction, we always parse  $\mathbf{hz} = (h_{\log T-1}, \dots, h_\lambda)$ ,  $\mathbf{crs} = (\mathbf{pk}, \mathbf{hk}_{\text{ht}}, \mathbf{hk}_{\text{ch}}, \mathbf{crs}_{\text{zk}}, (\mathbf{crs}_{\text{bp},\ell}, \mathbf{hk}_{\text{cf},\ell})_{\ell \in [\lambda]}, \mathbf{pcf}_{1,0})$ ,  $\Pi = (\pi_{\text{out}}, \text{epcf}, \pi_0, \dots, \pi_\lambda)$ ,  $\mathbf{pcf} = (\mathbf{rt}, \mathbf{q}, \mathbf{ridx})$ , and  $\mathbf{st}_i = (\mathbf{cf}_i, \mathbf{pcf}_i, r_i)$  wherever needed. Moreover, we let  $\mathbf{st}_0 = (0^{|\mathbf{cf}|}, \mathbf{pcf}_{1,0}, 0^{|\mathbf{r}|})$ , and  $\Pi_0 = (\epsilon, \text{epcf}_{1,0}, \epsilon, \dots, \epsilon)$ ,  $\mathbf{hz}_0 = (\epsilon, \dots, \epsilon)$  where  $\mathbf{pcf}_{1,0}$  is as defined in Gen and  $\text{epcf}_{1,0} = \text{PKE.Enc}(\mathbf{pk}, \mathbf{pcf}_{1,0}; 0^{|\mathbf{r}|})$ .

$\text{Gen}(1^\lambda, n, S, T) \rightarrow \mathbf{crs}$ . This algorithm lets  $\alpha, \beta = 1^{\lambda^8}$  and does the following:

1. Sample parameters:  $\mathbf{hk}_{\text{ht}} \leftarrow \text{HT.Gen}(1^\lambda)$ ,  $\mathbf{hk}_{\text{ch}} \leftarrow \text{CRHF.Gen}(1^\lambda)$ ,  $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{PKE.Gen}(1^\lambda)$ ,  $\mathbf{crs}_{\text{zk}} \leftarrow \text{NIZK.Setup}(1^\lambda, 1^{n_{\text{zk},x}}, 1^{n_{\text{zk},w}})$  for  $\mathcal{L}_{\text{zk}}$ ,  $(\mathbf{hk}_{\text{cf},\ell}, \mathbf{td}_{\text{cf},\ell}) \leftarrow \text{SEH.Gen}(1^\lambda, 2, \Sigma_{\text{cf}}^{\ell-1}, \alpha[\ell])$  for  $\ell \in [\lambda]$ , and  $(\mathbf{crs}_{\text{bp},\ell}, \mathbf{td}_{\text{bp},\ell}) \leftarrow \text{seBARG.Setup}(1^\lambda, 2, n_{\text{bp},\ell}, \beta[\ell])$  for  $\ell \in [\lambda]$ .
2. Get the honest starting pseudo-configuration:  $\mathbf{pcf}_{1,0} = (\text{HT.Hash}(\mathbf{hk}_{\text{ht}}, 0^S), 0^{|\mathbf{q}|}, 0^{|\mathbf{ridx}|})$ .
3. Output  $\mathbf{crs} = (\mathbf{pk}, \mathbf{hk}_{\text{ht}}, \mathbf{hk}_{\text{ch}}, \mathbf{crs}_{\text{zk}}, (\mathbf{crs}_{\text{bp},\ell}, \mathbf{hk}_{\text{cf},\ell})_{\ell \in [\lambda]}, \mathbf{pcf}_{1,0})$ .

$\text{Update}(\mathbf{crs}, z_k, \mathbf{st}_{k-1}, \mathbf{hz}_{k-1}, \Pi_{k-1}; \mathbf{ps}_k) \rightarrow (\mathbf{st}_k, \mathbf{hz}_k, \Pi_k)$ . This algorithm does the following:

1. Define the starting state for user  $k = \text{last state for user } k - 1$ :  $(\mathbf{st}_{k,0}, \Pi_{k,0}) = (\mathbf{st}_{k-1}, \Pi_{k-1})$ .

<sup>8</sup>While both  $\alpha$  and  $\beta$  have fixed values in the construction, we specifically define them in the construction as we will change these values in the security proofs.

2. Compute the work/input tape roots: let  $W_{k,0} \in \text{st}_{k,0}$ ,  $(\text{rt}_{k,0}, \text{tree}_{k,0}) = \text{HT.Hash}(\text{hk}_{\text{ht}}, W_{k,0})$  and  $(\text{rt}_{\text{in}}, \text{tree}_{\text{in}}) = \text{HT.Hash}(\text{hk}_{\text{ht}}, z_k)$ .
3. Encrypt input root: let  $r_{\text{in}} = \text{ps}_k$  and compute  $\text{ct}_{\text{in}} = \text{PKE.Enc}(\text{pk}, \text{rt}_{\text{in}}; r_{\text{in}})$ .
4. Define input auxiliary information required for proof generation:  $\text{aux}_{\text{in}} = (\text{rt}_{\text{in}}, \text{tree}_{\text{in}}, r_{\text{in}}, \text{ct}_{\text{in}})$ .
5. Run  $\mathcal{M}$  and update its corresponding proof for  $T$  steps: for  $t \in [T]$  run  $(\text{st}_{k,t}, \text{tree}_{k,t}, \Pi_{k,t}) = \text{UpdateSingleStep}(\text{crs}, \text{aux}_{\text{in}}, \text{st}_{k,t-1}, \text{tree}_{k,t-1}, \Pi_{k,t-1})$ .<sup>9</sup>
6. Update the input digest:  $\text{hz}_k = \text{nxt-hz}(\text{hk}_{\text{ch}}, \text{hz}_{k-1}, \text{ct}_{\text{in}})$ .
7. output  $(\text{st}_{k,T}, \text{hz}_k, \Pi_{k,T})$ .

$\text{Digest}(\text{crs}, (z_k)_{k \in [K]}, (\text{ps}_k)_{k \in [K]}) \rightarrow \text{hz}$ . This algorithm does the following:

1. Let  $\text{hz}_0 = (\text{h}_{\log T-1}, \dots, \text{h}_\lambda)$  where  $\text{h}_\ell$  is empty for  $\ell \in [\log T - 1, \lambda]$ .
2. For  $k \in [K]$  compute root of  $z_k$ , then encrypt the root and update  $\text{hz}$  by adding the encrypted root to it; namely, compute the following:
  - (a)  $(\text{rt}_k, \text{tree}_k) = \text{HT.Hash}(\text{hk}_{\text{ht}}, z_k)$ , and  $\text{ct}_{\text{in}} = \text{PKE.Enc}(\text{pk}, \text{rt}_k; \text{ps}_k)$ .
  - (b)  $\text{hz}_k = \text{nxt-hz}(\text{hk}_{\text{ch}}, \text{hz}_{k-1}, \text{ct}_{\text{in}})$ .
3. Output  $\text{hz} = (\text{h}_{\log T-1}, \dots, \text{h}_\lambda)$ .




Figure 5: Description of  $\text{nxt-hz}$ .

$\text{Verify}(\text{crs}, (\text{hz}, \text{out}), \Pi) \rightarrow 0/1$ . This algorithm does the following:

1. Parse  $\Pi = (\pi_{\text{out}}, \text{epcf}, \pi_0, \dots, \pi_\lambda)$ ,  $\text{hz} = (\text{h}_{\log T-1}, \dots, \text{h}_\lambda)$ .
2. Verify that  $\pi_{\text{out}}$  proves  $\text{out}$  is the corresponding output of  $\text{epcf}$ . Namely, let  $\pi_{\text{out}} = (\text{pcf}, r, \pi^*)$ , and  $\text{pcf} = (\text{rt}, \text{q}, \text{ridx})$ , and check  $\text{HT.VfyRd}(\text{hk}_{\text{ht}}, \text{rt}, [\text{n}_{\text{out}}], \text{out}, \pi^*) = 1$  and  $\text{PKE.Enc}(\text{pk}, \text{pcf}; r) = \text{epcf}$ ; output 0 otherwise.
3. First verify all the  $\lambda$  proofs w.r.t. input digest  $\text{hz}$ , and then check the consistency of  $\text{v}_i$  values with the starting/ending/intersecting encrypted pseudo-configurations. Track such encrypted pseudo-configurations by  $\text{epcf}^*$  and at the beginning let  $\text{epcf}^* = \text{epcf}$ , i.e. the last encrypted pseudo-configuration in the proof. Then for  $\ell \in [\log T - 1, \lambda]$ , if  $\pi_\ell$  is not empty do the following and output 0 if any of the checks fail:
  - (a) Let  $\pi_\ell = (\text{v}, \rho_1, \rho_2, \text{h}, \hat{\pi})$ . Check  $\text{h} = \text{h}_\ell$ ,  $\text{seBARG.Verify}(\text{crs}_{\text{bp}, \ell}, (\text{v}, \text{h}, i)_{i \in [2]}, \hat{\pi}) = 1$ , and  $\text{SEH.VerifyAcc}((\text{hk}_{\text{cf}, d})_{d \in [\ell]}, \text{v}, \rho_i, i)$  for  $i \in [2]$ .
  - (b) For  $i \in [2]$  parse  $\rho_i = ((\text{v}_{0,i,1}, \text{v}_{0,i,2}), \dots, (\text{v}_{\ell-1,i,1}, \text{v}_{\ell-1,i,2}))$ ,  $\text{v}_{0,i,i} = (\text{epcf}_{i,1}, \text{epcf}_{i,2})$ , and check  $\text{epcf}_{2,2} = \text{epcf}^*$ . Let  $\text{epcf}^* = \text{epcf}_{1,1}$ .
4. Output 1 if  $\text{epcf}^* = \text{PKE.Enc}(\text{pk}, \text{pcf}_{1,0}; 0^{|\text{r}|})$  and 0 otherwise.

<sup>9</sup>Recall that by underlining an algorithm we mean that it can be even more efficiently implemented as a RAM algorithm. We explain this further in Section 3.3.



**UpdateSingleStep**( $\text{crs}, \text{aux}_{\text{in}}, \text{st}_{t-1}, \text{tree}_{t-1}, \Pi_{t-1}$ )  $\rightarrow$  ( $\text{st}_t, \text{tree}_t, \Pi_t$ ). This algorithm does the following:

1. Parse  $\Pi_{t-1} = (\pi_{\text{out}}, \text{epcf}_{t-1}, \pi_0, \dots, \pi_\lambda)$ , and  $\text{aux}_{\text{in}} = (\text{rt}_{\text{in}}, \text{tree}_{\text{in}}, r_{\text{in}}, \text{ct}_{\text{in}})$ .
2. Find  $(\text{cf}_t, \text{pcf}_t)$ , and the reading/writing proofs  $(\text{rop}, \text{wop})$  for RAM transition  $\text{pcf}_{t-1} \rightarrow \text{pcf}_t$  as follows:
  - (a) Get read-proof from input tape:  $(\text{rbit}, \text{rop}) = \text{HT.Read}(\text{hk}_{\text{ht}}, \text{tree}_{\text{in}}, \text{ridx}_{t-1})$  if  $\text{ridx}_{t-1} \leq n$ .
  - (b) Get read-proof from work tape:  $(\text{rbit}, \text{rop}) = \text{HT.Read}(\text{hk}_{\text{ht}}, \text{tree}_{t-1}, \text{ridx}_{t-1} - n)$  if  $\text{ridx}_{t-1} > n$ .
  - (c) Run the RAM step function:  $\mathcal{C}_{\mathcal{M}}(q_{t-1}, \text{rbit}) = (q_t, \text{ridx}_t, \text{widx}, \text{wbit})$ .
  - (d) Get write-proof for work tape:  $(\text{rt}_t, \text{tree}_t, \text{wop}) = \text{HT.Write}(\text{hk}_{\text{ht}}, \text{tree}_{t-1}, \text{widx}, \text{wbit})$ .
  - (e) Update the work tape: write  $\text{wbit}$  at  $\text{widx}^{\text{th}}$  index of  $W_{t-1}$  to get  $W_t$ .
  - (f) Find  $(\text{cf}_t, \text{pcf}_t)$ : Let  $\text{cf}_t = (W_t, q_t, \text{ridx}_t)$  and  $\text{pcf}_t = (\text{rt}_t, q_t, \text{ridx}_t)$ .
3. *Encrypt*  $\text{pcf}_t$ : sample  $r_t$  and compute  $\text{epcf}_t = \text{PKE.Enc}(\text{pk}, \text{pcf}_t; r_t)$ .
4. *Compute a NIZK proof for*  $\text{epcf}_{t-1} \rightarrow \text{epcf}_t$ :  $\pi_{\text{zk}} = \text{NIZK.Prove}(\text{crs}_{\text{zk}}, x = (\text{ct}_{\text{in}}, \text{epcf}_{t-1}, \text{epcf}_t), w)$  for language  $\mathcal{L}_{\text{zk}}$  where  $w = (\text{rt}_{\text{in}}, r_{\text{in}}, \text{pcf}_{t-1}, r_{t-1}, \text{pcf}_t, r_t, \text{rbit}, \text{rop}, \text{wop})$ .
5. *Define the temporary proof and input hash*: let  $\mathbf{t}\text{-}\pi = (\mathbf{t}\text{-}\mathbf{v} = (\text{epcf}_{t-1}, \text{epcf}_t), \mathbf{t}\text{-}\hat{w} = \pi_{\text{zk}})$ , and  $\mathbf{t}\text{-}\mathbf{h} = \text{ct}_{\text{in}}$ . Note that at the moment these values correspond to a level 0 proof in the power-of-2-structure, but they will be treated as proofs for level  $\ell$ , in the next iterative process.
6. *Iteratively merge the proofs*: let  $\ell = 0$  and while  $\pi_\ell \in \Pi$  is not empty run the following iterative process to merge the proofs  $\pi_\ell$  and  $\mathbf{t}\text{-}\pi$  (which are at the same level) and increment  $\ell$  at the end.
  - (a) Parse  $\mathbf{t}\text{-}\pi$  and  $\pi_\ell$  as follows:
    - If  $\ell = 0$  then  $\pi_\ell = (\mathbf{v}_1, \hat{w}_1)$  and  $\mathbf{t}\text{-}\pi = (\mathbf{v}_2, \hat{w}_2)$ .
    - If  $0 < \ell < \log T - 1$  then  $\pi_\ell = (\mathbf{v}_1, \rho_{1,1}, \rho_{1,2}, \hat{\pi}_1)$  and  $\mathbf{t}\text{-}\pi = (\mathbf{v}_2, \rho_{2,1}, \rho_{2,2}, \hat{\pi}_2)$ .
    - If  $\log T - 1 \leq \ell$  then  $\pi_\ell = (\mathbf{v}_1, \rho_{1,1}, \rho_{1,2}, \mathbf{h}_1, \hat{\pi}_1)$  and  $\mathbf{t}\text{-}\pi = (\mathbf{v}_2, \rho_{2,1}, \rho_{2,2}, \mathbf{h}_2, \hat{\pi}_2)$ .
  - (b) *Merge level  $\ell$  hashes to get the next temp  $\mathbf{v}$* :  $\mathbf{t}\text{-}\mathbf{v} = \text{SEH.Hash}(\text{hk}_{\text{cf}, \ell+1}, (\mathbf{v}_1, \mathbf{v}_2))$ .
  - (c) *Merge level  $\ell$  input hashes to get the next temp  $\mathbf{h}$* : if  $\ell \geq \log T - 1$  compute  $\mathbf{t}\text{-}\mathbf{h} = \text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\mathbf{h}_1, \mathbf{h}_2))$ .
  - (d) *Find partial witness  $\hat{w}_i$  for merging BARG proofs*: for  $i \in [2]$ , if  $1 \leq \ell < \log T$  then  $\hat{w}_i = (\rho_{1,2}, \rho_{2,1}, \hat{\pi}_i)$ , and if  $\log T \leq \ell$  then  $\hat{w}_i = (\rho_{1,2}, \rho_{2,1}, \mathbf{h}_i, \hat{\pi}_i)$ .
  - (e) *Merge NIZK/BARG proofs*:  $\mathbf{t}\text{-}\pi^* = \text{seBARG.Prove}\left(\text{crs}_{\text{bp}, \ell+1}, \begin{pmatrix} x_i = \mathbf{t}\text{-}\mathbf{h}, \mathbf{t}\text{-}\mathbf{v}, i \\ w_i = \mathbf{v}_1, \mathbf{v}_2, \hat{w}_i \end{pmatrix}_{i \in [2]}\right)$  for the language  $\mathcal{L}_{\text{bp}, \ell+1}$  in Fig. 4.
  - (f) *Extend the path openings to get the next temp path openings*: for  $i \in [2]$  if  $\ell = 0$  then  $\mathbf{t}\text{-}\rho_i = (\mathbf{v}_1, \mathbf{v}_2)$ , and if  $\ell \geq 1$  then  $\mathbf{t}\text{-}\rho_i = (\rho_{i,i}, (\mathbf{v}_1, \mathbf{v}_2))$ .
  - (g) *Put everything together to get the next temp  $\pi$* : if  $\ell < \log T - 2$  then  $\mathbf{t}\text{-}\pi = (\mathbf{t}\text{-}\mathbf{v}, \mathbf{t}\text{-}\rho_1, \mathbf{t}\text{-}\rho_2, \mathbf{t}\text{-}\pi^*)$ , otherwise,  $\mathbf{t}\text{-}\pi = (\mathbf{t}\text{-}\mathbf{v}, \mathbf{t}\text{-}\rho_1, \mathbf{t}\text{-}\rho_2, \mathbf{t}\text{-}\mathbf{h}, \mathbf{t}\text{-}\pi^*)$ .
7. *Compute output-proof w.r.t.*  $\text{epcf}_t$ :  $(\text{out}_t, \pi_t^*) = \text{HT.Read}(\text{hk}_{\text{ht}}, \text{tree}_t, [n_{\text{out}}])$  and  $\pi_{\text{out}} = (\text{pcf}_t, r_t, \pi_t^*)$ .
8. Output  $(\text{st}_t, \text{tree}_t, \Pi_t = (\pi_{\text{out}}, \text{epcf}_t, \epsilon, \dots, \epsilon, \mathbf{t}\text{-}\pi, \pi_{\ell+1}, \dots, \pi_\lambda))$ .

### 3.3 Analysis

**Theorem 3.5.** Assuming that seBARG is an almost rate-1 seBARG (Definitions 2.6 and 2.8), NIZK is a secure NIZK protocol (Definition 2.15), PKE is a secure PKE system (Definition 2.17), SEH is an almost rate-1 SEH (Definitions 2.2 and 2.4), HT is a hash tree (Definition 2.11), and CRHF is a 2-to-1 collision resistant hash function (Definition 2.10), Construction 3.4 is a zk-IVSc (Definition 3.2).

**Corollary 3.6.** Assuming almost rate-1 seBARGs and PKEs, Construction 3.4 is a zk-IVsC.

**Corollary 3.7.** Assuming either LWE, or DLIN, or sub-exponential DDH, Construction 3.4 is a zk-IVsC.

The zero-knowledge property of our construction follows by NIZK zero-knowledge and PKE security in a straightforward manner, however, the soundness proof is rather a delicate task. Next, we provide an overview of the strong soundness of our construction and refer the reader to Appendix A for a complete analysis of our construction.

**Strong soundness overview.** Define  $\text{exp}_{\alpha,\beta}$  to be an experiment w.r.t. the choice of  $\alpha$  and  $\beta$  (where SEH (resp. seBARG) is binding/extractable on the path  $\alpha$  (resp.  $\beta$ )), that also outputs the trapdoors and secret keys. Note that  $\alpha$  uniquely determines a tuple  $(\text{cf}_{k,t-1}, \text{cf}_{k,t})$  in the computation (and its corresponding path in the proof). Define an extraction algorithm  $\text{Extract}^*$  that by (recursively) using the SEH trapdoors and the PKE secret-key, extracts some  $(\text{pcf}_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)})$  from proof  $\Pi^{(b)}$ . Now note that any  $\text{epcf}_{k,t}$  is in exactly two different path openings, one that branches out to right and one that branches out to left. Now, define two other extraction algorithms  $\text{Extract}_r$  (resp.  $\text{Extract}_l$ ) that first recursively use the seBARG trapdoors, then use path openings from right (resp. left), and finally uses the PKE secret-key to extract some  $\text{pcf}_r^{(b)}$  (resp.  $\text{pcf}_l^{(b)}$ ) from proof  $\Pi^{(b)}$ , where  $(k,t)$  is also defined by the value of  $\alpha$ . Our goal is to inductively show that if  $\text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)}$ , then  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$ , and next increment  $\alpha$  and  $\beta$  while holding the invariant that  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$ . More specifically, we prove the following:

(1) In  $\text{exp}_{1,1}$  we have  $\text{pcf}_{1,0}^{(0)} = \text{pcf}_{1,0}^{(1)}$  (see Claim A.2). We rely on the SEH binding on the path to  $\text{epcf}_{1,0}$ , PKE correctness, and the verifier's equality check with  $\text{pcf}_{1,0} \in \text{crs}$ .

(2) In  $\text{exp}_{\alpha,\alpha}$  if  $\text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)}$  then  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$  (see Claim A.3). To prove this we use an induction on  $\ell$ , the level of the proof  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)})$  that includes a proof for  $\text{pcf}_{k,t-1} \rightarrow \text{pcf}_{k,t}$ . We first extract a witness from  $\hat{\pi}^{(b)}$ , then using SEH binding property argue the extracted witness is consistent with the extracted hash value from  $\mathbf{v}^{(b)}$ . Then starting with hashed inputs  $\mathbf{h}^{(b)}$  that are consistent with  $\text{hz}$ , by relying on collision resistance of CRHF we get that the hashed input values in the witness are also consistent with  $\text{hz}$ . Then we rely on the induction step. For the induction base, we additionally rely on the knowledge soundness of NIZK, together with reading and writing soundness of HT (for  $\text{is-nxt-cnfg}_{\mathcal{M}}$ ), to conclude the claim.

(3) If  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$  in  $\text{exp}_{\alpha,\alpha}$  then  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$  in  $\text{exp}_{\alpha,\alpha+1}$  (see Claim A.4). Note that the statement depends only on SEH trapdoors and PKE secret-key, thus we can rely on the index hiding property of seBARG to change the binding path from  $\alpha$  to  $\alpha+1$ .

(4) If  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$  in  $\text{exp}_{\alpha,\alpha+1}$  then  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$  in  $\text{exp}_{\alpha+1,\alpha+1}$  (see Claim A.5). To prove this we take the following steps – (a) in  $\text{exp}_{\alpha,\alpha+1}$  if  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$  then  $\text{pcf}_l^{(0)} = \text{pcf}_l^{(1)}$ ; this holds as the path on which  $\text{pcf}_l^{(b)}$  is computed is the same as the path on which  $\text{pcf}_{k,t}^{(b)}$  is extracted, (b) in  $\text{exp}_{\alpha,\alpha+1}$  if  $\text{pcf}_l^{(0)} = \text{pcf}_l^{(1)}$  then  $\text{pcf}_r^{(b)} = \text{pcf}_r^{(b)}$ ; this holds either by the verifier's check or the seBARG proof check at some level  $j$ , (c) if  $\text{pcf}_r^{(b)} = \text{pcf}_r^{(b)}$  in  $\text{exp}_{\alpha,\alpha+1}$  then  $\text{pcf}_r^{(b)} = \text{pcf}_r^{(b)}$  in  $\text{exp}_{\alpha+1,\alpha+1}$ ; this holds by the SEH index hiding and the fact that statement depends only on seBARG trapdoors and PKE secret key, (d) in  $\text{exp}_{\alpha+1,\alpha+1}$  if  $\text{pcf}_r^{(b)} = \text{pcf}_r^{(b)}$  then  $\text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)}$ ; this holds as the path on which  $\text{pcf}_r^{(b)}$  is computed is the same as the path on which  $\text{pcf}_{k,t}^{(b)}$  is extracted.

(5) In  $\text{exp}_{KT,KT}$  if  $\text{pcf}_{N,T}^{(0)} = \text{pcf}_{N,T}^{(1)}$  then the adversary cannot generate valid proofs  $\Pi^{(0)}$  and  $\Pi^{(1)}$  for two different final outputs  $\text{out}^{(0)}$  and  $\text{out}^{(1)}$  (see Claim A.6). To prove this claim we rely on the SEH binding and PKE correctness to show  $\text{pcf}_{K,T}^{(b)} = \text{pcf}^{(b)}$  where  $\text{pcf}^{(b)} \in \pi_{\text{out}}^{(b)}$ , which implies that  $\text{pcf}^{(0)} = \text{pcf}^{(1)}$ , and finally rely on reading soundness of HT to prove  $\text{out}^{(0)} = \text{out}^{(1)}$ .

### 3.4 Further Optimizing zk-IVsC

In this section we explain how to further optimize our zk-IVsC.

**What if output is large?** Consider a RAM machine  $\mathcal{M}$  for which the output size  $n_{\text{out}}$  is large, e.g.,  $n_{\text{out}} \simeq S$ . In this scenario, as our proof size and verification time grows with  $n_{\text{out}}$ , our construction becomes inefficient. In this case we alter the proof to not compute  $\pi_{\text{out}}$ , and let the verification to only verify a digest of the final configuration (which is already included in our construction) and conclude the proof by relying on collision resistance of HT. Then the efficiency analysis of our construction becomes independent of  $n_{\text{out}}$ , making the construction efficient even if the output size is large.

**More efficient prover.** Suppose we don't want the running time of the prover (and Update) to grow with the size of  $\text{cf}$ . Note that the only part of the computation that grows with  $|\text{cf}| = S$  is the computation of the hash tree given the configuration  $\text{cf}_{k,0}$ . To avoid this step, we can consider the computation of the RAM machine to keep a tree of the  $\text{cf}$  in its memory. Thus we update the state  $\text{st}$  such that instead of  $\text{st}_i = (\text{cf}_i, \text{pcf}_i, r_i)$ , we have  $\text{st}_i = (\text{cf}_i, \text{pcf}_i, \text{tree}_i, r_i)$ . Note that the size of  $\text{st}$  only suffers from a (multiplicative) logarithmic overhead on the size of  $\text{cf}$ . However, Update algorithm will be relieved from the computation of  $\text{tree}_{k,0}$  each time a new user starts to run the machine. Therefore, the new running time of Update is  $\text{poly}(\lambda, n) + T \cdot \text{poly}(\lambda, \log S, \log n, \log T, \log k, n_{\text{out}})$  and the running time of the new prover is  $K (\text{poly}(\lambda, n) + T \cdot \text{poly}(\lambda, \log S, \log n, \log T, \log K, n_{\text{out}}))$ .

**Succinct communication to the next user.** Note that each time a user finishes running the machine, it has to pass the final state/configuration to the next user (so that they can start running the machine). We highlight that if the amount of information needed to start running the machine by the next user is succinct, then there is a succinct way of sending the last configuration to the next user. Namely, user  $k$  will only write the information that needs to be passed on at the beginning of the work tape and let the rest of the work tape be all zeros (this will only have an additive overhead of size at most  $S$  on the running time of the machine, and w.l.o.g. we can assume  $S \leq T$ ). Then the user will only send the non-zero part of the memory to the next user and the next user will be able to reconstruct the entire state/configuration. Therefore, we can always assume that for all  $k \in [N]$ ,  $\text{st}_{k,T}$  has a short description (e.g., of size  $\text{poly}(\lambda, \log S)$ ).

### 3.5 Predicated zk-IVsC

#### 3.5.1 Definition

In this section, we extend zk-IVsC definition to further impose a well-formedness property on the input to the machine, and then propose a construction that achieves such property.

**Syntax.** A predicated zero-knowledge incrementally verifiable streaming computation (predicated zk-IVsC) scheme for a RAM machine  $\mathcal{M}$  is defined *similarly* to zk-IVsC, except that, first the Gen algorithm additionally takes a predicate size  $n_f$  as input, and second, the primitive also has the following PPT algorithms for proving and verifying the well-formedness of the inputs:

**PredicateUpdate**( $\text{crs}, f_k, z_k, \text{ps}_k, \bar{f}_{k-1}, \text{hz}_{k-1}, \Sigma_{k-1}$ )  $\rightarrow$  ( $\text{hz}_k, \Sigma_k$ ) The predicate update algorithm takes as input a CRS  $\text{crs}$ , a predicate  $f_k$ , an input  $z_k$ , set of predicates  $\bar{f}_{k-1} = (f_j)_{j \in [k-1]}$ , a private state  $\text{ps}_k$ , a digest  $\text{hz}_{k-1}$ , and a proof  $\Sigma_{k-1}$  and outputs an updated digest  $\text{hz}_k$  and an updated proof  $\Sigma_k$ .

**PredicateVerify**( $\text{crs}, \bar{f}, \text{hz}, \Sigma$ )  $\rightarrow$  0/1 The predicate verifier algorithm takes as input a CRS  $\text{crs}$ , a set of predicates  $\bar{f} = (f_k)_{k \in [K]}$ , a digest  $\text{hz}$ , and a proof  $\Sigma$ , and outputs a bit to signal whether the proof is valid or not.

**Remark 3.8** (Predicate Prover). We will define the predicate prover algorithm **PredicateProve**( $\text{crs}, \bar{f}, \bar{z}, \bar{\text{ps}}$ ) for a predicate zk-IVsC inductively as follows (where  $\text{hz}_0, \Sigma_0 = \epsilon$ ):

If  $K = 1$  then output  $(\text{hz}_1, \Sigma_1) \leftarrow \text{PredicateUpdate}(\text{crs}, f_1, z_1, \text{ps}_1, \bar{f}_0, \text{hz}_0, \Sigma_0)$ , otherwise inductively compute  $(\text{hz}_{K-1}, \Sigma_{K-1}) = \text{PredicateProve}(\text{crs}, \bar{f}_{K-1}, \bar{z}_{K-1}, \bar{\text{ps}}_{K-1})$  and output

$$(\text{hz}_K, \Sigma_K) \leftarrow \text{PredicateUpdate}(\text{crs}, f_K, z_K, \text{ps}_K, \bar{f}_{K-1}, \text{hz}_{K-1}, \Sigma_{K-1}).$$

**Definition 3.9** (Predicated zk-IVsC). A predicated zero-knowledge incrementally verifiable streaming computation scheme (Gen, Update, Digest, Verify, PredicateProve, PredicateVerify) for  $\mathcal{M}$  has to satisfy the following properties:

**Completeness.** For every  $\lambda, n, S, T \in \mathbb{N}$ , any  $\bar{z} = (z_k)_{k \in [K]}$  s.t.  $z_k \in \{0, 1\}^n$ ,

$$\Pr \left[ \begin{array}{l} (\bar{z}, \text{out}) \in \mathcal{L}_{\mathcal{M}, T} \wedge f(z_k) = 1 \wedge \\ \text{hz} = \text{hz}' = \text{Digest}(\text{crs}, \bar{z}, \bar{\text{ps}}) \wedge \\ \text{Verify}(\text{crs}, (\text{hz}, \text{out}), \Pi) = 1 \\ \text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}, \Sigma) = 1 \end{array} : \begin{array}{l} \text{crs} \leftarrow \text{Gen}(1^\lambda, n, S, T, n_f) \\ (\text{st}, \text{hz}, \Pi) \leftarrow \text{Prove}(\text{crs}, \bar{z}, \bar{\text{ps}}) \\ \text{find } \text{cf} \in \text{st}, \text{out} = \text{cf}[1, \dots, n_{\text{out}}] \\ (\text{hz}', \Sigma) \leftarrow \text{PredicateProve}(\text{crs}, \bar{f}, \bar{z}, \bar{\text{ps}}) \end{array} \right] = 1.$$

**Efficiency.** For the completeness experiment above,

- The size of the  $\text{crs}$  and setup running time are  $\text{poly}(\lambda, \log S, n, n_f)$ .
- The size of  $\Pi$ , and update/verifier's running time are the same as in IVsC
- The size of  $\Sigma$  is  $\text{poly}(\lambda, n, n_f, \log K)$ , and predicate update/verifier's running time are  $\text{poly}(K) + \text{poly}(\lambda, n, n_f, \log K)$ . (We also consider - and later achieve - a more efficient predicate update/verifier that run in  $\text{poly}(\lambda, n, n_f, \log K)$ , when they are allowed to preprocess the set  $\bar{f}$ ..

**Collision Resistance, Soundness, and Strong Soundness.** Similar to zk-IVsC.

**Predicate Soundness.** There exists PPT extractor  $\mathcal{E}$  and algorithms TGen and HashCheck such that for any stateful PPT adversary  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}, \Sigma) = 1 \wedge \\ (f_{k^*}(z^*) = 0 \vee \\ \text{HashCheck}(\text{crs}, \text{hz}, z^*, k^*, \text{op}) = 0) \end{array} : \begin{array}{l} (n, S, T, n_f, k^*) \leftarrow \mathcal{A}(1^\lambda), \\ (\text{crs}, \text{td}) \leftarrow \text{TGen}(1^\lambda, n, S, T, k^*), \\ (\bar{f}, \text{hz}, \Sigma) \leftarrow \mathcal{A}(\text{crs}) \\ (z^*, \text{op}) \leftarrow \mathcal{E}(\text{td}, \bar{f}, \text{hz}, \Sigma) \end{array} \right] \leq \text{negl}(\lambda),$$

and  $\text{crs}$  and  $\text{crs}' \leftarrow \text{Gen}(1^\lambda, n, S, T)$  are computationally indistinguishable even given  $k^*$ .

**Zero-Knowledge.** There exists a stateful PPT simulator  $\mathcal{S}$  such that for any PPT adversary  $\mathcal{A}$ , there is a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} (n, S, T, n_f) \leftarrow \mathcal{A}(1^\lambda), b \leftarrow \{0, 1\} \\ \text{crs}_0 \leftarrow \text{Gen}(1^\lambda, n, S, T, n_f), \text{crs}_1 \leftarrow \mathcal{S}(1^\lambda, n, S, T, n_f), \\ (\bar{z}, \bar{f}) \leftarrow \mathcal{A}(\text{crs}_b), (f_k(z_k) = 1)_{k \in [K]}, \text{cf} = \mathcal{M}(\bar{z}; 1^T) \\ b = b' : (\text{st}_0, \bar{\text{hz}}_0, \Pi_0) \leftarrow \text{Prove}(\text{crs}_0, \bar{z}, \bar{\text{ps}}) \\ (\text{hz}_{0,K}, \Sigma_0) \leftarrow \text{PredicateProve}(\text{crs}, \bar{f}, \bar{z}, \bar{\text{ps}}) \\ (\text{st}_1, \bar{\text{hz}}_1, \Pi_1, \Sigma_1) \leftarrow \mathcal{S}(\text{crs}_1, \text{cf}) \\ b' \leftarrow \mathcal{A}(\text{st}_b, \bar{\text{hz}}_b, \Pi_b, \Sigma_b) \end{array} \right] \leq \text{negl}(\lambda),$$

where we consider a variant of Prover that additionally outputs all the intermediate  $\text{hz}$  values, i.e.,  $\bar{\text{hz}}_b = (\text{hz}_{b,k})_{k \in [K]}$ .

### 3.5.2 Construction

**Notations.** We implement our algorithms in the RAM model and to indicate a RAM access, we use underline in the construction. For ease of exposition (and w.l.o.g.) we assume  $T$  is a power of 2. We let  $r_{1,0}$  to be all zeros.

By  $\bar{f}^* = (\bar{f}_1^*, \bar{f}_2^*)$  we denote  $\bar{f}^* = (f_k)_{k \in [1, 2t]}$ ,  $\bar{f}_1^* = (f_k)_{k \in [1, t]}$ , and  $\bar{f}_2^* = (f_k)_{k \in [t+1, 2t]}$ . We say  $\gamma$  corresponds to  $z_k$  (or just  $k$ ) if  $(\gamma_{\log T}, \dots, \gamma_\lambda)$  is the unique path from root to  $k^{\text{th}}$  leaf in a full binary tree of depth  $\lambda - \log T + 1$ .

**Construction overview.** The `PredicateUpdate` works similarly to `Update`, that is, it builds a powers-of-2 structure of proofs, generates proofs at the base level using a NIZK, and recursively merges the proofs using `seBARGs`. Moreover, the construction relies on a decomposition algorithm that decomposes the set of all predicates to their corresponding powers-of-2 structure. More specifically we define  $\text{TreeDecompose}(\bar{f}) = (\bar{f}_{\text{dc},0}, \dots, \bar{f}_{\text{dc},\lambda})$  to be an algorithm that receives as input a set  $\bar{f}$  of size  $k$  and decomposes it to disjoint sets  $\bar{f}_{\text{dc},i}$  as follows: first write  $k$  as a summation of powers of 2, i.e.  $k = \sum_{i=0}^\lambda b_i \cdot 2^i$  where  $b_i \in \{0, 1\}$  and let  $s_i = \sum_{j=0}^i b_j \cdot 2^j$  for  $i \in [-1, \lambda]$ . Finally it outputs  $(\bar{f}_{\text{dc},0}, \dots, \bar{f}_{\text{dc},\lambda})$  where  $\bar{f}_{\text{dc},i} = (f_k)_{k \in [K-s_i+1, K-s_{i-1}]}$  for  $i \in [0, \lambda]$ .

<b>Language <math>\mathcal{L}_{zk,f}</math></b>
<b>Hardwired:</b> $pk, hk_{ht}$
<b>Instance:</b> $x = (f, ct_{in})$
<b>Witness:</b> $w = (z, rt_{in}, r)$
<b>Membership:</b> $w$ is a valid witness for $x \in \mathcal{L}_{zk,f}$ if the following hold:
$ct_{in}$ is the encryption of $rt_{in}$ . Namely, $PKE.Enc(pk, rt_{in}; r) = ct_{in}$ .
$rt_{in}$ is the input root. Namely, $(rt_{in}, tree_{in}) = HT.Hash(hk_{ht}, z)$ .
<b>Input satisfies the predicate.</b> Namely, $Check(f(z)) = 1$ .

Figure 6: Description of language  $\mathcal{L}_{zk,f}$ .

<b>Language <math>\mathcal{L}_{bp,f,\ell}</math></b>
<b>Hardwired:</b> $hk_{ch}, crs_{zk,f}, (crs_{bp,f,j})_{j \in [\ell-1]}$
<b>Instance:</b> $x = (\bar{f}^* = (\bar{f}_1^*, \bar{f}_2^*), h, i)$
<b>Witness:</b> $w = (h_1, h_2, \pi)$
<b>Membership:</b> $w$ is a valid witness for $x \in \mathcal{L}_{bp,f,\ell}$ if the following hold:
$h$ is the hash of $(h_1, h_2)$ . Namely, $CRHF.Hash(hk_{ch}, (h_1, h_2)) = h$ .
$\pi$ certifies well-formedness of $h_i$ w.r.t. $\bar{f}_i^*$ . Namely, if $\ell = \log T$ then $NIZK.Verify(crs_{zk,f}, (\bar{f}_i^*, h_i), \pi) = 1$ , otherwise, $seBARG.Verify(crs_{bp,f,\ell-1}, (\bar{f}_i^*, h_i, j)_{j \in [2]}, \pi) = 1$ .

Figure 7: Description of language  $\mathcal{L}_{bp,f,\ell}$ .

**Building blocks and parameters.** We use the same primitives as in  $zk-IVsC$ . In addition to the sizes in  $zk-IVsC$ , let  $n_f = |f_k|$  for  $k \in [K]$ , and  $n_{bp,f,\ell} = 2^{\ell - \log T + 1} \cdot n_f + \lambda + 1$ . Let  $|r|$  and  $|ct|$  be respectively the size of randomness and ciphertext of PKE, and let  $n_{zk,f,x} = n_f + |ct|$  and  $n_{zk,f,w} = n + \lambda + |r|$ .

**Construction 3.10.** We use the same parsing notation as in  $zk-IVsC$ , except that we modify  $crs$  parsing as  $crs = (pk, hk_{ht}, hk_{ch}, crs_{zk}, crs_{zk,f}, (crs_{bp,\ell}, hk_{cf,\ell})_{\ell \in [\lambda]}, (crs_{bp,f,\ell})_{\ell \in [\log T, \lambda]}, pcf_{1,0})$ , and we additionally parse  $\Sigma = (\pi_{f,\log T-1}, \dots, \pi_{f,\lambda})$  wherever needed. Update, Digest, and Verify algorithms remain the same as in  $zk-IVsC$ .

$Gen(1^\lambda, n, S, T, n_f) \rightarrow crs$ . This algorithm is the same as in  $zk-IVsC$  except that it additionally samples  $crs_{zk,f} = NIZK.Setup(1^\lambda, 1^{n_{zk,f,x}}, 1^{n_{zk,f,w}})$ , lets  $\gamma = 1^\lambda$  and for all  $\ell \in [\log T, \lambda]$ , it samples  $(crs_{bp,f,\ell}, td_{bp,f,\ell}) \leftarrow seBARG.Gen(1^\lambda, 2, n_{bp,f,\ell}, \gamma[\ell - \log T + 1])$ , and outputs  $crs = (pk, hk_{ht}, hk_{ch}, crs_{zk}, crs_{zk,f}, (crs_{bp,\ell}, hk_{cf,\ell})_{\ell \in [\lambda]}, (crs_{bp,f,\ell})_{\ell \in [\log T, \lambda]}, pcf_{1,0})$ .

$PredicateUpdate(crs, f_k, z_k, ps_k, \bar{f}_{k-1}, hz_{k-1}, \Sigma_{k-1}) \rightarrow (hz_k, \Sigma_k)$ . This algorithm does the following:

1. Let  $hz_{k-1} = (h_{\log T-1}, \dots, h_\lambda)$ , and  $\Sigma_{k-1} = (\pi_{f,\log T-1}, \dots, \pi_{f,\lambda})$ .
2. Decompose the predicates  $\bar{f}_{k-1}$  to powers-of-2:  $(\bar{f}_{dc,0}, \dots, \bar{f}_{dc,\lambda}) = TreeDecompose(\bar{f}_{k-1})$ .
3. Encrypt the root of input:  $(rt_{in}, tree_{in}) = HT.Hash(hk_{ht}, z_k)$  and  $ct_{in} = PKE.Enc(pk, rt_{in}; ps_k)$ .
4. Compute a NIZK proof for  $f_k(z_k) = 1$ :  $\pi_{zk,f} = NIZK.Prove(crs_{zk,f}, x = (f_k, ct_{in}), w)$  for the language  $\mathcal{L}_{zk,f}$  where  $w = (z_k, rt_{in}, ps_k)$ .



5. *Define the temporary proof, input hash, and predicates:* Let  $\mathbf{t-h} = \mathbf{ct}_{\text{in}}$ ,  $\mathbf{t-f} = (f_k)$ ,  $\mathbf{t-\pi} = \pi_{\text{zk},f}$ . Note that at the moment these values correspond to a level 0 proof in the power-of-2-structure, but they will be treated as proofs for level  $\ell$ , in the next iterative process.
6. *Iteratively merge the proofs:* let  $\ell = \log T - 1$  and while  $\pi_{f,\ell}$  is not empty run the following iterative process and increment  $\ell$  at the end:
  - (a) Let  $\mathbf{h} = \mathbf{h}_\ell$ ,  $\mathbf{h}' = \mathbf{t-h}$ ,  $\pi_1 = \pi_{f,\ell}$ ,  $\pi_2 = \mathbf{t-\pi}$ ,  $\bar{f}_1^* = \mathbf{t-f}$ ,  $\bar{f}_2^* = \bar{f}_{\text{dc},\ell-\log T+1}$ .
  - (b) Merge level  $\ell$  input hashes to get the next temp  $\mathbf{h}$ :  $\mathbf{t-h} = \text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\mathbf{h}, \mathbf{h}'))$ .
  - (c) Merge level  $\ell$  predicates to get the next temp predicates:  $\mathbf{t-f} = (\bar{f}_1^* \cup \bar{f}_2^*)$ .
  - (d) Merge NIZK/BARG proofs:  $\mathbf{t-\pi} = \text{seBARGProve} \left( \text{crs}_{\text{bp},f,\ell+1}, \begin{matrix} (\mathbf{t-f}, \mathbf{t-h}, i)_{i \in [2]}, \\ (\mathbf{h}, \mathbf{h}', \pi_i)_{i \in [2]} \end{matrix} \right)$  for the language  $\mathcal{L}_{\text{bp},f,\ell+1}$ .
7. Output  $(\mathbf{hz}_k = (\epsilon, \dots, \epsilon, \mathbf{t-h}, \mathbf{h}_{\ell+1}, \dots, \mathbf{h}_\lambda), \Sigma_k = (\epsilon, \dots, \epsilon, \mathbf{t-\pi}, \pi_{f,\ell+1}, \dots, \pi_{f,\lambda}))$ .

$\text{PredicateVerify}(\text{crs}, \bar{f}, \mathbf{hz}, \Sigma) \rightarrow 0/1$ . The verifier first decomposes the predicates  $\bar{f}$  and then verifies NIZK/BARG proofs. Namely, let  $(\bar{f}_{\text{dc},0}, \dots, \bar{f}_{\text{dc},\lambda}) = \text{TreeDecompose}(\bar{f})$ , if  $\pi_{f,\log T-1}$  is not empty then check  $\text{NIZK.Verify}(\text{crs}_{\text{zk},f}, (\bar{f}_{\text{dc},0}, \mathbf{h}_{\log T-1}), \pi_{f,\log T-1}) = 1$ , and for  $\ell \in [\log T, \lambda]$  if  $\pi_{f,\ell}$  is not empty then check  $\text{seBARG.Verify}(\text{crs}_{\text{bp},f,\ell}, (\bar{f}_{\text{dc},\ell-\log T+1}, \mathbf{h}_\ell, i)_{i \in [2]}, \pi_{f,\ell}) = 1$ . Output 1 if all the checks pass and 0 o.w.

### 3.5.3 Analysis

**Theorem 3.11.** Assuming that seBARG is an almost rate-1 seBARG (Definitions 2.6 and 2.8), NIZK is a secure NIZK protocol (Definition 2.15), PKE is a secure PKE system (Definition 2.17), SEH is an almost rate-1 SEH (Definitions 2.2 and 2.4), HT is a hash tree (Definition 2.11), and CRHF is a 2-to-1 collision resistant hash function (Definition 2.10), Construction 3.10 is a predicated zk-IVsC (Definition 3.9).

**Corollary 3.12.** Assuming almost rate-1 seBARGs and PKEs, Construction 3.10 is a predicated zk-IVsC.

**Corollary 3.13.** Assuming either LWE, or DLIN, or sub-exponential DDH, Construction 3.10 is a predicated zk-IVsC.

The analysis of the predicate proof is very similar to the analysis of zk-IVsC. We refer the reader to Appendix B for a complete analysis of our construction.

## 4 Incrementally Computable Zero-Knowledge Arguments (ICZK)

### 4.1 Definition

Consider an NP language  $\mathcal{L} = \{\bar{x} = (x_1, \dots, x_K) \mid \exists \bar{w} = (w_1, \dots, w_K) : \mathcal{R}(\bar{x}, \bar{w}) = 1\}$  defined w.r.t. a relation  $\mathcal{R}$ .



$z_i = (x_i, w_i, y_1, y_2)$  as input but does not use  $y_1$  and  $y_2$  in the computation. Later, we let  $y_1 = \text{ps}_i$  and  $y_2 = r_i$  be randomnesses for different PKE schemes.

**Construction overview.** Let  $\mathcal{M}$  be a RAM machine for relation  $\mathcal{R}$  with space  $S$ , and running time  $T$  (on each input  $z_k = (x_k, w_k)$ ). Our starting point is to consider the computation of  $\mathcal{M}$  as a zk-IVsC where the streaming inputs are  $z_k = (x_k, w_k)$  and the tuple  $(\text{hz}, \Pi)$  is considered as the proof of computation. While this proof is zero-knowledge, one cannot verify that  $\text{hz}$  is an honest digest of  $\bar{z}$ . Our idea is to send a hiding encoding of the witnesses and use the input well-formedness proof to verify the consistency of  $\text{hz}$  with such an encoding. There are two caveats in such an approach: (1) both the encoding of witnesses and encrypted digest of the RAM machine have to be randomized to satisfy any hiding properties, however, our predicated zk-IVsC only supports deterministic computation, and the predicates have to be publicly known, and (2) the predicate proof is sound only at a trapdoor location, thus we have to break the global well-formedness of  $\text{hz}$  to a local well-formedness checks. We get around both of these issues using following blueprint: (1) encrypt the witness and private state  $(w_k, \text{ps}_k)$  under some randomness  $r_k$  to get  $\text{ct}_k$ , (2) use zk-IVsC to prove the computation of a RAM machine  $\mathcal{M}'$ , and (3) use predicate  $f_k$  to prove  $x_k$  is the correct statement,  $(w_k, \text{ps}_k, r_k)$  is consistent with  $\text{ct}_k$ , and  $\text{hz}_{k-1} \xrightarrow{\text{ct}_{\text{in}}} \text{hz}_k$  where  $\text{ct}_{\text{in}}$  is the encryption of  $z_k$ 's root under randomness  $\text{ps}_k$ . Then the final proof consists of  $\overline{\text{ct}}_k$ ,  $\overline{\text{hz}}_k$ , and zk-IVsC computation and predicate proofs.

**Building blocks and parameters.** Let HT be a hash tree with output root of size  $\lambda$  on any input, CRHF be a collision resistant Hash :  $\{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ , seBARG be an almost rate-1 seBARG, SEH be an almost rate-1 SEH, NIZK be a NIZK, PKE be a *randomness-dependent message* public-key encryption (RDM-PKE), and rrPKE be a rate-1 leakage resilient public-key encryption with randomness recovery. In addition to predicated zk-IVsC sizes, we let  $n_{\text{ps}}$  be the size of the randomness  $\text{ps}$  for the PKE above,  $n_r$  be the size of the randomness for the rrPKE above, and  $n = (n_x + n_w + n_{\text{ps}} + n_r)$ . Additionally, we let  $n_f$  be the max size of the function  $f_k$  from Fig. 8.

<b>Predicate <math>f_k</math></b>
<b>Hardwired:</b> $\text{hk}_{\text{ht}}, \text{pk}, \text{pk}', (x_k, \text{ct}_k, \text{hz}_{k-1}, \text{hz}_k)$
<b>Input:</b> $z = (x, w, \text{ps}, r)$
<b>Output:</b> 1 if the following hold:
<b>The RAM input is consistent with NIZK statement.</b> Namely, $x = x_k$ .
<b>The RAM input is consistent with encrypted witness.</b> $\text{PKE.Enc}(\text{pk}, (w, \text{ps}); r) = \text{ct}_k$ .
<b>The RAM input is consistent with <math>\text{hz}_{k-1} \rightarrow \text{hz}_k</math> transition</b> Namely, $(\text{rt}, \text{tree}) = \text{HT.Hash}(\text{hk}_{\text{ht}}, x, w, \text{ps}, r)$ , $\text{ct}_{\text{in}} = \text{PKE.Enc}(\text{pk}', \text{rt}; \text{ps})$ and $\text{nxt-hz}(\text{hk}_{\text{ch}}, \text{hz}_{k-1}, \text{ct}_{\text{in}}) = \text{hz}_k$ using Fig. 5.

Figure 8: Description of predicate  $f_k$  in the ICZK construction.

**Construction 4.3.** In this construction, we always parse  $\Pi = (\overline{\text{hz}}, \Pi_\psi, \Sigma_\psi, \overline{\text{ct}})$  and  $\text{crs} = (\text{crs}_\psi, \text{pk})$ , where we have  $\text{crs}_\psi = (\text{pk}', \text{hk}_{\text{ht}}, \text{hk}_{\text{ch}}, \text{crs}_{\text{zk}}, \text{crs}_{\text{zk}, f}, (\text{crs}_{\text{bp}, \ell}, \text{hk}_{\text{cf}, \ell})_{\ell \in [\lambda]}, (\text{crs}_{\text{bp}, f, \ell})_{\ell \in [\log T, \lambda]}, \text{pcf}_{1,0})$ .

$\text{Setup}(1^\lambda, n_x, n_w) \rightarrow \text{crs}$ . This algorithm samples  $\text{crs}_\psi$  similar to predicated zk-IVsC (except that now PKE is RDM-PKE), and  $(\text{pk}, \text{sk}) \leftarrow \text{rrPKE.Gen}(1^\lambda)$ , and outputs  $\text{crs} = (\text{crs}_\psi, \text{pk})$ .

$\text{Update}(\text{crs}, x_k, w_k, \text{st}_{k-1}, \Pi_{k-1}) \rightarrow (\text{st}_k, \Pi_k)$ . This algorithm does the following:

1. Parse  $\text{crs} = (\text{crs}_\psi, \text{pk})$  and  $\Pi_{k-1} = (\overline{\text{hz}}_{k-1}, \Pi_{\psi, k-1}, \Sigma_{\psi, k-1}, \overline{\text{ct}}_{k-1})$ .
2. *Encrypt the witness and private state:* sample randomness  $\text{ps}_k, r_k$ , and compute  $\text{ct}_k = \text{rrPKE.Enc}(\text{pk}, (w_k, \text{ps}_k); r_k)$ .
3. *Update the computation proof for running  $\mathcal{M}'$  on  $z_k = (x_k, w_k, \text{ps}_k, r_k)$ :*  $(\text{st}_k, \text{hz}_k, \Pi_{\psi, k}) \leftarrow \Psi.\text{Update}(\text{crs}_\psi, z_k = (x_k, w_k, \text{ps}_k, r_k), \text{st}_{k-1}, \text{hz}_{k-1}, \Pi_{\psi, k-1}; \text{ps}_k)$ .
4. *Get the predicates:* find  $\text{hk}_{\text{ht}}, \text{pk}' \in \text{crs}_\psi$ ; then construct  $\overline{f}_k$  using Fig. 8.
5. *Update the predicate well-formedness proof w.r.t.  $f_k(z_k) = 1$ :* compute  $(\text{hz}_k, \Sigma_{\psi, k}) \leftarrow \Psi.\text{PredicateUpdate}(\text{crs}_\psi, f_k, z_k, \text{ps}_k, \overline{f}_{k-1}, \text{hz}_{k-1}, \Sigma_{\psi, k-1})$ .
6. Output  $(\text{st}_k, \Pi_k = (\overline{\text{hz}}_k, \Pi_{\psi, k}, \Sigma_{\psi, k}, \overline{\text{ct}}_k))$ .

$\text{Verify}(\text{crs}, \overline{x}, \Pi) \rightarrow 0/1$ . This algorithm parses  $\text{crs} = (\text{crs}_\psi, \text{pk})$ , and  $\Pi = (\overline{\text{hz}}, \Pi_\psi, \Sigma_\psi, \overline{\text{ct}})$ . It additionally parses  $\text{crs}_\psi$  to find  $\text{hk}_{\text{ht}}$  and  $\text{pk}'$ , and then construct  $\overline{f}$  using Fig. 8. Finally it checks  $\Psi.\text{Verify}(\text{crs}, (\text{hz}_K, 1), \Pi_\psi) = 1$  and  $\Psi.\text{PredicateVerify}(\text{crs}, \overline{f}, \text{hz}_K, \Sigma_\psi) = 1$ . It outputs 1 if both checks pass, and 0 otherwise.

### 4.3 Analysis

**Theorem 4.4.** Assuming that seBARG is an almost rate-1 seBARG (Definitions 2.6 and 2.8), NIZK is a secure NIZK protocol (Definition 2.15), PKE is a secure RDM-PKE system (Definitions 2.17 and 2.24), SEH is an almost rate-1 SEH (Definitions 2.2 and 2.4), HT is a hash tree (Definition 2.11), CRHF is a 2-to-1 collision resistant hash function (Definition 2.10), and rrPKE is a secure rate-1 leakage resilient PKE with randomness recovery (Definitions 2.17 to 2.19 and 2.22), Construction 4.3 is an ICZK (Definition 4.2) with proof size  $K|w_i| + K \cdot \text{poly}(\lambda, \log |x_1|, \log |w_1|, \log K) + \text{poly}(\lambda, \log |x_1|, |w_1|, \log K)$ .

**Corollary 4.5.** Assuming almost rate-1 seBARGs and lossy trapdoor functions, Construction 4.3 is an ICZK.

**Corollary 4.6.** Assuming either LWE, or DLIN, or sub-exponential DDH, Construction 4.3 is an ICZK.

Next, we provide an overview of soundness and zero-knowledge of our construction and refer the reader to Appendix C for a complete analysis of our construction.

**Soundness overview.** To rely on the soundness of zk-IVsC we need to find  $(\overline{z}, \overline{\text{ps}})$  s.t.  $\text{hz}$  is the honest digest of  $(\overline{z}, \overline{\text{ps}})$ . First, note that  $\overline{x}$  is publicly known. Then, we decrypt all the ciphertexts  $\overline{\text{ct}}$  to find  $(\overline{w}, \overline{\text{ps}})$ ; moreover, recall that  $\overline{\text{ct}}$  is generated using PKE with *randomness recovery*, thus, we can recover  $\overline{r}$  as well. Hence, we can reconstruct  $z_k = (x_k, w_k, \text{ps}_k, r_k)$ . Next, we need to prove that  $\text{hz}$  is consistent with  $\overline{z}$  as reconstructed above.

Suppose this doesn't hold; then there is an index  $k$  such that  $\text{nxt-hz}(\text{hz}_{k-1}, \text{ct}_{\text{in}}) \neq \text{hz}_k$  while  $\text{hz}_{k-1} = \text{Digest}(\overline{z}_{k-1}, \overline{\text{ps}}_{k-1})$ . Next we guess index  $k$ , and prove that the above event cannot happen if our guess is correct. Namely, if the predicate proof is extractable on  $k$ , then there is an input  $z^*$  that satisfies the predicate. By the predicate check,  $\text{nxt-hz}(\text{hz}_{k-1}, \text{ct}_{\text{in}}^*) \neq \text{hz}_k$ , where  $\text{ct}_{\text{in}}^*$  is computed from  $z^*$ . Now note that the computation of  $\text{ct}_{\text{in}}^*$  (resp.  $\text{ct}_{\text{in}_k}$ ) from  $z^*$  (resp.  $z_k$ ) is deterministic, and its correctness is checked by the predicate  $f_k$ . Therefore, all that is left to prove is that  $z_k = z^*$ . For this, we crucially rely on the *perfect correctness* of the PKE that is used to compute  $\overline{\text{ct}}$ . Namely,

since we check that each predicate input is consistent with  $\text{ct}_k$ , and the choice of such input is unique by the perfect correctness of PKE, it holds that  $z^*$  is the same as  $z_k$ .

**Incremental zero-knowledge overview.** We first recall the main circularity challenges in proving zero-knowledge of our construction. First, note that the ciphertext  $\text{ct}_k$  is the encryption of  $(w_k, \text{ps}_k)$  and its security relies on the randomness of  $r_k$ , however,  $r_k$  is used as part of the input  $z_k$  everywhere in the zk-IVsC proofs. Moreover, the zero-knowledge of zk-IVsC holds if  $\text{ps}_k$  is random and independent. However, not only there is leakage of  $\text{ps}_k$  in  $\text{ct}_k$ , but also the message encrypted under the randomness  $\text{ps}_k$  (which is the hash of  $z_k$ ), depends on  $\text{ps}_k$ . Our main observation is that we can break down these circularities step by step and bound the leakage at every step. Next, we show how to carefully formulate the hybrids and use the right building blocks to get around such issues.

$\mathcal{H}_0$ : This is the original experiment.

$\mathcal{H}_1$ : Use NIZK simulator to compute  $\text{crs}_{z_k}$ , and  $\pi_{z_k}$  in  $\Pi_\psi$ .  $\mathcal{H}_0 \rightarrow \mathcal{H}_1$  relies on zero-knowledge property of the underlying NIZK. Note that after this step the only information about an intermediate configurations  $\text{cf}_{k,t}$  is in  $\text{epcf}_{k,t}$ , and the only information about the inputs of  $\mathcal{M}$  is in  $(\overline{\text{hz}}, \overline{\text{ct}}, \Sigma_\psi)$  and  $(\text{ct}_{\text{in},k})_{k \in [K]}$  (that is being used in the statements of  $\Pi_\psi$  and  $\Sigma_\psi$ ).

$\mathcal{H}_2$ : Use NIZK simulator to compute  $\text{crs}_{z_k,f}$ , and  $\pi_{z_k,f}$  in  $\Sigma_\psi$ .  $\mathcal{H}_1 \rightarrow \mathcal{H}_2$  also relies on the zero-knowledge property of the underlying NIZK. Note that after this step there is no information about inputs  $(\overline{w}, \overline{\text{ps}}, \overline{r})$  in the proof except the ones in  $(\text{hz}_k, \text{ct}_k, \text{ct}_{\text{in},k})_{k \in [K]}$  where  $\overline{\text{hz}}$  is constructed only using  $(\text{ct}_{\text{in},k})_{k \in [K]}$ . Thus, the only information about  $z_k$  is in  $\text{ct}_k$  and  $\text{ct}_{\text{in},k}$ .

$\mathcal{H}_3$ : For every  $\text{cf}_{k,t}$  (except  $\text{cf}_{1,0}$  and  $\text{cf}_{K,T}$ ), compute  $\text{epcf}_{k,t} \leftarrow \text{PKE.Enc}(\text{pk}, 0^{|\text{pcf}|})$ .  $\mathcal{H}_2 \rightarrow \mathcal{H}_3$  relies on the PKE security. Note that after this step there is no information about the intermediate configurations in the proof.

$\mathcal{H}_4$ : For every  $\text{ct}_k$ , independently sample  $r_k$  and let  $\text{ct}_k = \text{rrPKE.Enc}(\text{pk}, 0^{n_w + |\text{ps}|}; r_k)$ .  $\mathcal{H}_3 \rightarrow \mathcal{H}_4$  relies on the security with leakage of the underlying rrPKE as the only information about randomness  $\overline{r}$  remained in the proof is in  $(\text{ct}_{\text{in},k})_{k \in [K]}$  (that is an encryption of a short digest of the input  $z_k$ ). Note that after this step there's no information about  $\overline{w}$  and  $\overline{\text{ps}}$  in the  $\overline{\text{ct}}$ .

$\mathcal{H}_5$ : For every  $z_k$ , let  $\text{ct}_{\text{in},k} \leftarrow \text{PKE.Enc}(\text{pk}, 0^\lambda)$ .  $\mathcal{H}_4 \rightarrow \mathcal{H}_5$  relies on the RDM security of the underlying PKE as the only information about randomness  $\overline{\text{ps}}$  remained in the proof is in  $(\text{ct}_{\text{in},k})_{k \in [K]}$  (that is the encryption of a short digest of the inputs). Note that after this step there's no information about  $\overline{w}$  in the proof and the proof can be simulated using only  $(\overline{x}, \text{cf})$ .

## 5 From IVsC to Deterministic Distributed Computation

We present this section as PCD for deterministic computation, however, we note that our constructions are not true PCD and we modify the network by adding new communication edges.

**Distributed Computation.** A distributed computation is a tuple  $S = (G, \text{data}, \text{linp}, \mathbb{C})$  where  $G$  denotes the graph of computation,  $\text{data}$  denotes the data that is transferred over the edges,  $\text{linp}$  denotes the local input of any user/vertex, and  $\mathbb{C}$  denotes a compliance function for the inputs and output of any node. More formally we define a distributed computation as follows:

- $G = (V, E)$ .

- $D = \text{Diameter}(G)$ .
- $V = \{v_i^d\}$  where  $d \in D$  and  $i \in [2^{d-1}]$ .
- $E = \{(v_j^{d+1}, v_i^d)\}_{j \in [2^{(i-1)+1, 2 \cdot i}], d \in [D-1]}$ .
- $\text{data}_{\text{inp}}(v_i^d) = (\text{data}(v, v_i^d))_v$  s.t.  $(v, v_i^d) \in E$ .
- $\text{data}_{\text{out}}(v_i^d) = \text{data}(v_i^d, v)$  where  $(v_i^d, v) \in E$ .
- $\mathbb{C}(\text{data}_{\text{inp}}(v), \text{linp}(v), \text{data}_{\text{out}}(v)) = 1$  for all  $v \in V$ .

**Proof Carrying Data (PCD) Systems.** PCD systems are usually defined for non-deterministic computations where the security is w.r.t. a knowledge extractor. However, IVsC is defined for deterministic computation. Therefore, for applications of IVsC to PCD we define PCD for deterministic computations where we assume that the compliance functions  $\mathbb{C}$  for each set of node inputs  $(\text{data}_{\text{inp}}(v), \text{linp}(v))$  outputs 1 only for a unique node output  $\text{data}_{\text{out}}(v)$ . Moreover, this output can be computed using a RAM machine, i.e., there exists a RAM machine  $\mathcal{R}$  s.t.  $\mathcal{R}(\text{data}_{\text{inp}}(v), \text{linp}(v)) = \text{data}_{\text{out}}(v)$ . For the rest of this section by PCD we refer to PCD for deterministic computation.

**Definition 5.1** (PCD for Deterministic Computation). A PCD system  $\Pi = (\text{Gen}, \text{Prove}, \text{Verify})$  for a deterministic compliance function  $\mathbb{C}$  consists of the following algorithms:

- $\text{Gen}(1^\lambda) \rightarrow \text{crs}$
- $\text{Prove}(\text{crs}, \text{linp}(v), \text{data}_{\text{inp}}(v), \pi_{\text{inp}}(v)) \rightarrow (\text{data}_{\text{out}}(v), \pi_{\text{out}}(v))$ .
- $\text{Verify}(\text{crs}, G, (\text{linp}(v))_{v \in V}, \text{data}_{\text{out}}(v^*), \pi_{\text{out}}(v^*)) \rightarrow \{0, 1\}$ .

Where  $v^*$  is the sink of the graph. Note that the Prove algorithm is defined for a single vertex, and we define the  $\text{GlobalProve}(\text{crs}, G, (\text{linp}(v))_{v \in V}) \rightarrow (\text{data}_{\text{out}}(v^*), \pi_{\text{out}}(v^*))$  inductively by writing down a topological order of the graph and running the Prove algorithm for each of the vertices in order.

**Completeness.** For any  $\lambda \in \mathbb{N}$  it holds that:

$$\Pr \left[ \text{Verify}(\text{crs}, X) = 1 \quad : \quad \begin{array}{l} (\text{data}_{\text{out}}(v^*), \pi_{\text{out}}(v^*)) \leftarrow \\ \text{GlobalProve}(\text{crs}, G, (\text{linp}(v))_{v \in V}) \end{array} \right] = 1$$

where  $\text{crs} \leftarrow \text{Gen}(1^\lambda)$ , and  $X = (G, (\text{linp}(v))_{v \in V}, \text{data}_{\text{out}}(v^*), \pi_{\text{out}}(v^*))$ .

**Soundness.** For any PPT stateful admissible adversary  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{crs}, G, (\text{linp}(v))_{v \in V}, \text{out}, \pi) = 1, \\ \wedge \text{out} \neq \text{data}_{\text{out}}(v^*) \end{array} \quad : \quad \begin{array}{l} (G) \leftarrow \mathcal{A}(1^\lambda), \text{crs} \leftarrow \text{Gen}(1^\lambda), \\ ((\text{linp}(v))_{v \in V}, \text{out}, \pi) \leftarrow \mathcal{A}(\text{crs}) \end{array} \right] \leq \text{negl}(\lambda).$$

**Remark 5.2** (From IVsC to Linear PCD). Note that our construction of IVsC (i.e., consider Construction 3.4 without zero-knowledge proofs and properties) implies a PCD for a deterministic computation on a path graph. More specifically, we can use  $(\text{Gen}, \text{Update}, \text{Verify})$  of an IVsC as

(Gen, Prove, Verify) in a PCD for path graphs. Note that we can consider a RAM machine  $\mathcal{R}$  with fixed input size  $n$ , max configuration size  $S$ , and running time  $T$  that computes the compliant function  $\mathbb{C}$ . Then we can instantiate an IVsC using these parameters and let the PCD algorithms be the same as the IVsC algorithms. Note that the way we analyze our construction it is only secure against adversaries that output  $1^{N \cdot T}$ , therefore our PCD constructions from IVsC are also secure only if the adversary outputs the running time of the computation in unary. Additionally, note that our construction achieves a stronger security notion w.r.t. a verifier that only takes a digest of the local inputs. This implies the notion of PCD for path graphs and also enables us to strengthen this notion correspondingly.

In the following, we discuss the set of distributed computations for which we can generate a PCD using our results on IVsC.

### 5.1 Binary Tree to Path Distributed Computation

Here we show how to transform a depth-D binary tree distributed computation into a Path computation with a succinct amount of data (a D overhead on the original data size) carried over the edges. Hence, given a PCD scheme for a path distributed computation, we can generate a PCD scheme for a depth-D (with an acceptable D, say  $\text{poly}(\lambda)$  at most) binary tree distributed computation.

**Construction 5.3.** Given  $S = (G, \text{data}, \text{linp}, \mathbb{C})$  construct  $S' = (G', \text{data}', \text{linp}, \mathbb{C}')$  as follows:

1.  $G' = (V, E')$ .
2.  $E' = E_1 \cup E_2 \cup E_3$  where
  - $E_1 = \{(v_{2i-1}^D, v_{2i}^D)\}_{i \in [2^{d-2}]}$ .
  - $E_2 = \{(v_{2i}^{d+1}, v_i^d)\}_{d \in [D-1], i \in [2^{d-1}]}$ .
  - $E_3 = \{(v_{2i-1}^d, v_{(2i-1)(2^{D-d}+1)}^D)\}_{d \in [D-1], i \in [2^{d-2}]}$ .
3.  $\text{linp}'(v_i^d) = \text{linp}(v_i^d)$ .
4. Let  $\text{data}'_{\text{out}}(v) = (\text{data}_{\text{out}}^1(v), \dots, \text{data}_{\text{out}}^D(v))$  and  $\text{data}'_{\text{inp}}(v) = (\text{data}_{\text{inp}}^1(v), \dots, \text{data}_{\text{inp}}^D(v))$  be as follows:
  - Let  $\text{data}'_{\text{inp}}(v_1^D) = \epsilon$ , and for  $(d, i) \neq (D, 1)$ ,  $(v_i^{d'}, v_i^d) \in E'$ , let  $\text{data}'_{\text{inp}}(v_i^d) = \text{data}'_{\text{out}}(v_i^{d'})$ .
  - For any edge  $(v_{2i-1}^D, v_{2i}^D) \in E_1$ , let  $\text{data}'_{\text{out}}(v_{2i-1}^D) = \text{data}'_{\text{inp}}(v_{2i}^D)$  and append  $\text{data}_{\text{out}}(v_{2i-1}^D)$  to  $\text{data}_{\text{out}}^D(v_{2i-1}^D)$  where  $\mathbb{C}(\epsilon, \text{linp}(v_{2i-1}^D), \text{data}_{\text{out}}(v_{2i-1}^D)) = 1$ .
  - For any edge  $(v_{2i}^{d+1}, v_i^d) \in E_2$ , let  $\text{data}'_{\text{out}}(v_{2i}^{d+1}) = \text{data}'_{\text{inp}}(v_i^d)$ ,  $\text{data}_{\text{out}}^{d+2}(v_{2i}^{d+1}) = \epsilon$ , and append  $\text{data}_{\text{out}}(v_{2i}^{d+1})$  to  $\text{data}_{\text{out}}^{d+1}(v_{2i}^{d+1})$  where  $\mathbb{C}(\text{data}_{\text{inp}}^{d+2}(v_{2i}^{d+1}), \text{linp}(v_{2i}^{d+1}), \text{data}_{\text{out}}(v_{2i}^{d+1})) = 1$ .
  - For any edge  $(v_{2i-1}^d, v_{(2i-1)(2^{D-d}+1)}^D) \in E_3$ , let  $\text{data}'_{\text{out}}(v_{2i-1}^d) = \epsilon$  and append  $\text{data}_{\text{out}}(v_{2i-1}^d)$  to  $\text{data}_{\text{out}}^d(v_{2i-1}^d)$  where  $\mathbb{C}(\text{data}_{\text{inp}}^{d+1}(v_{2i-1}^d), \text{linp}(v_{2i-1}^d), \text{data}_{\text{out}}(v_{2i-1}^d)) = 1$ .
5.  $\mathbb{C}'(\text{data}'_{\text{inp}}(v), \text{linp}(v), \text{data}'_{\text{out}}(v)) = 1$  for all  $v \in V$ .



**Correctness.** From the construction it is clear that for any  $v \in V$ ,  $\text{data}_{\text{inp}}(v) \subseteq \text{data}'_{\text{inp}}(v)$ , therefore, user  $v$  has all the data it needs to perform its computation. Additionally, it is clear from the construction of  $S'$  that every  $v$  has exactly one input and output data, so it is a path computation.

**Succinctness.** Let  $\max_{(v,v') \in E} |\text{data}(v,v')| = S_{\text{PCD}}$ , then for any user/vertex  $v \in V$ , we have:  $|\text{data}_{\text{inp}}| = D \cdot S_{\text{PCD}}$  since for any  $d \in D$  (except one at most),  $\text{data}_{\text{inp}}^d(v)$  includes one data at most.

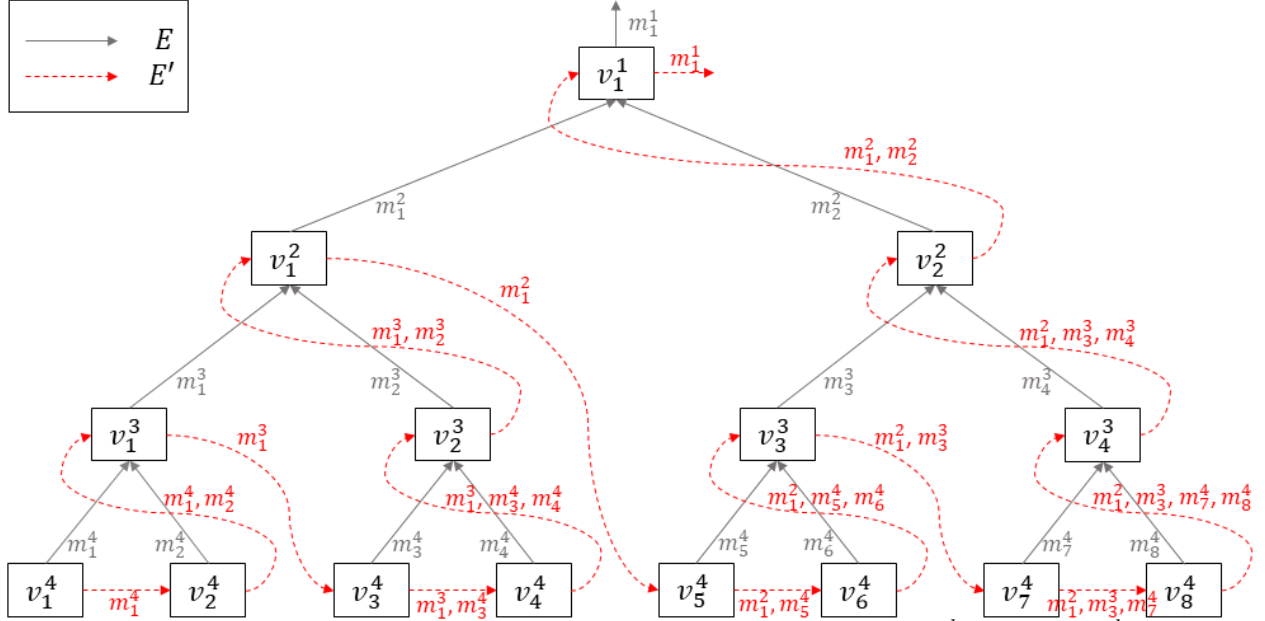


Figure 9: Example of Construction 5.3 for a depth 4 tree where  $m_i^d = \text{data}_{\text{out}}(v_i^d)$  in the original distributed computation. Here  $E'$  and  $\text{data}_{\text{out}}(v_i^d)$  are determined by red.

## 5.2 Class of Succinctly Transformable Distributed Computations to a Path

Here we generalize the approach in the previous subsection and discuss a bigger class of distributed computations that can be transformed into a path distributed computation. First note that the main challenge in such a transformation is to keep a succinct set of all the data that is needed to do the rest of the computation. More specifically, in this approach, we have to find a topological order of a directed graph such that at any point in the order, the amount of data that needs to be kept to do the rest of the computation is succinct. We can transform any graph  $G$  for which such a topological order exists, to a path of computation with a succinct amount of data carried over the edges.

More formally, let  $(v_1, \dots, v_N)$  be a topological order of a graph  $G$ . Denote  $\mathbb{L}$  to be a labeling function such that if the computation for  $v_i$  is done then  $\mathbb{L}(v_i) = 1$ , and otherwise  $\mathbb{L}(v_i) = 0$ . Define a node  $v_i$  to be open,  $\mathbb{O}(v_i) = 1$ , if  $\mathbb{L}(v_i) = 1$  and  $\mathbb{L}(v_j) = 0$  where  $(v_i, v_j) \in E$ . Suppose we start the computation from  $v_1$  and stop right after some arbitrary point  $v_i$ . Define the set of open nodes as follows  $\mathcal{O}_i = \{v_j \mid \mathbb{O}(v_j) = 1\}$ . Now note that at any point  $v_i$ , we only need to keep  $\text{data}_{\text{out}}$  for the open nodes. This is because if a node  $v_k$  is not open either (1)  $\mathbb{L}(v_k) = 0$  hence  $\text{data}_{\text{out}}(v_k)$  hasn't

been computed yet, or (2)  $\mathbb{L}(v_k) = 1$  and  $\mathbb{L}(v_{k'}) = 1$  where  $(v_k, v_{k'}) \in E$ , therefore,  $\text{data}_{\text{out}}(v_k)$  is already used as part of  $\text{data}_{\text{inp}}(v_{k'})$  and it will not be used elsewhere, so we don't need to keep it. Putting this all together means that if for a graph  $G$ , there exists a topological order  $(v_1, \dots, v_N)$  such that  $|\mathcal{O}_i| = \text{poly}(\lambda)$  for any  $i \in [N]$ , then there exists a transformation from  $\mathbf{S} = (G, \dots)$  to a linear distributed computation with a succinct amount of data carried over the edges.

**Zero-Knowledge for Distributed Computation.** We call a graph *good* if it is succinctly transformable to a path. Defining zero-knowledge for these graphs is a challenging task let alone constructing a zero-knowledge PCD for these graphs. Here we informally discuss a high-level idea to achieve zero-knowledge and then analyze the kind of zero-knowledge that it potentially achieves. After transforming a graph into a path, we generate a pair of public and secret keys of an FHE encryption system for the sink node and publish the public key. Then every user encrypts its local input under that public key and the computation is defined as a homomorphic computation over the inputs. This approach achieves zero-knowledge for the inputs for every user against all the other users except the sink user. Note that the same approach can be taken to achieve privacy in a path graph as well but our construction has two advantages over this. First, it is based on simpler assumptions than FHE. Second, the sink user learns nothing about the other users' inputs other than the last configuration of the machine.

Now if the sink user of a good graph is not predetermined, e.g., the graph is still evolving and more users are joining, then we can use the same idea with a leveled set of public and secret keys, where every user has a pair of secret and public keys, and the public keys are published. The first user starts by encrypting its message under its parent's public key and performs the homomorphic computation and sends the result to the next user. Every user upon receiving an input has two choices, either (1) the computation is done under some other user's public key, in which case the user will continue the computation under the same public key, or (2) the computation is done under that user's public key, in which case it uses its secret key to decrypt the computation, then encrypts the computation using its parent's public key and performs the rest of the computation homomorphically under that public key. This approach achieves zero-knowledge for inputs for user  $v$  against any user  $v'$  if there is no path from  $v$  to  $v'$  in the graph.

We leave the formalization and constructions of zero-knowledge PCD for general good graphs (and without FHE) as an interesting open problem.

## 6 Benefits and Applications

**The streaming scenario.** In many modern applications, the data for computation might not be available *all at once*, or it could be in an ongoing process of generation. Additionally, datasets are often too large and it is difficult to store the entire data simultaneously, or with a single data processor. E.g., consider a cloud service provider running an ML process on a massive dataset collected by a trusted agency such as the National Institute of Health (NIH). The data could be a large corpus of *sensitive* patient data from many different sources (e.g., healthcare providers, research facilities, etc), where the data is signed by each contributing source. *The question is how can we verify that the cloud provider trained the model correctly?*

NIH might ask the cloud provider to use IVC, but this would require the training data to be *fixed at the beginning*. Therefore, if new data arrives (e.g., new routine blood sample results), then the cloud provider has to restart the IVC computation. Also, if the ML algorithm is itself a

streaming algorithm (e.g., reinforcement learning), then the cloud cannot take advantage of this. Furthermore, the proof might leak information about the sensitive data, thus cannot be shared with third-party clients (e.g., an insurance provider, or research lab) who are not authorized to look at the entire data but only the parameters of the trained ML model.

One could consider many more such applications (e.g., online video-processing of traffic-videos/user-videos by National Highway Institute/Youtube, etc). We believe that IVsC protocols will be a very valuable tool for all such applications. A cloud provider can create incremental proofs over streaming (sensitive) data. Moreover, as long as there is a publicly checkable predicate associated with the data (e.g., every medical record is signed by a trusted entity), then the proof can be verified without the entire dataset, but just the additional predicated proof. Thus, it will solve: (a) the streaming problem, and (b) data privacy problem.

**Step-by-step zero-knowledge.** Similarly, the notion of step-by-step zero-knowledge could be a great addition to practical deployments of proof systems. For instance, consider any sensitive computation performed by one user/device, or even multiple. A popular threat in many deployed cryptographic systems is of total system corruption. Unfortunately, standard zero-knowledge completely fails in such settings. However, step-by-step zero-knowledge offers a new approach for capturing such practical threats by defining a *best-possible* zero-knowledge property. E.g., consider any application where a user wants to prove that it correctly computed a PRF on a public input  $x$ . That is, it wants to prove  $y = F_K(x)$ , where  $K$  is the user’s secret. Such proofs of PRF evaluations routinely occur in many blockchain applications (e.g., for leader election in proof-of-stake) or anonymous credentials/ blind signatures (for generating blinded credentials/signatures). Using a step-by-step ZK protocol, we can prove that the user’s PRF key  $K$  will never be fully compromised even an attacker gets a hold of their system during the proof generation process. Basically, if one use GGM-style PRFs [GGM86], then after each computation step, the user can delete the PRG seed after it correctly expands it (depending upon the next input bit). This merely sketches how step-by-step ZK could go beyond conventional zero-knowledge formulations to provide better protections, and we believe further exploration would be quite meaningful for applications.

**Acknowledgments.** We would like to thank Elahe Sadeghi and Saikumar Yadugiri for their valuable suggestions and assistance in refining the content of this paper. We also thank anonymous reviewers for valuable feedback. Their inputs greatly improved the clarity and presentation of the ideas discussed.

## References

- [ADKL19] Prabhanjan Ananth, Apoorvaa Deshpande, Yael Tauman Kalai, and Anna Lysyanskaya. Fully homomorphic nizk and niwi proofs. In *Theory of Cryptography Conference*, pages 356–385. Springer, 2019.
- [BBB<sup>+</sup>18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE symposium on security and privacy (SP)*, pages 315–334. IEEE, 2018.
- [BC23] Benedikt Bünz and Binyi Chen. Protostar: Generic efficient accumulation/folding for special sound protocols. *Cryptology ePrint Archive*, 2023.

- [BC24] Dan Boneh and Binyi Chen. Latticefold: A lattice-based folding scheme and its applications to succinct proof systems. *Cryptology ePrint Archive*, 2024.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for snarks and proof-carrying data. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 111–120, 2013.
- [BCL<sup>+</sup>21] Benedikt Bünz, Alessandro Chiesa, William Lin, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data without succinct arguments. In *Advances in Cryptology—CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*, pages 681–710. Springer, 2021.
- [BCMS20] Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data from accumulation schemes. *Cryptology ePrint Archive*, 2020.
- [BCPT13] Eleanor Birrell, Kai-Min Chung, Rafael Pass, and Sidharth Telang. Randomness-dependent message security. In *Theory of Cryptography: 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013. Proceedings*, pages 700–720. Springer, 2013.
- [BCTV14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 276–294. Springer, Heidelberg, August 2014.
- [BDFG21] Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Halo infinite: Proof-carrying data from additive polynomial commitments. In *Advances in Cryptology—CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*, pages 649–680. Springer, 2021.
- [BGH19] Sean Bowe, Jack Grigg, and Daira Hopwood. Recursive proof composition without a trusted setup. *Cryptology ePrint Archive*, 2019.
- [BKP<sup>+</sup>23] Nir Bitansky, Chethan Kamath, Omer Paneth, Ron Rothblum, and Prashant Nalini Vasudevan. Batch proofs are statistically hiding. *Cryptology ePrint Archive*, 2023.
- [BMNW24] Benedikt Bünz, Pratyush Mishra, Wilson Nguyen, and William Wang. Accumulation without homomorphism. *Cryptology ePrint Archive*, 2024.
- [BMRS20] Joseph Bonneau, Izaak Meckler, Vanishree Rao, and Evan Shapiro. Coda: Decentralized cryptocurrency at scale. *Cryptology ePrint Archive*, 2020.
- [BR94] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *CRYPTO’93*, volume 773 of *LNCS*, pages 232–249. Springer, Heidelberg, August 1994.
- [BSCTV17] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. *Algorithmica*, 79:1102–1160, 2017.

- [But22] Vitalik Buterin. The different types of zk-evms. <https://vitalik.eth.limo/general/2022/08/04/zkevm.html>, 2022.
- [BWW23] Eli Bradley, Brent Waters, and David J Wu. Batch arguments to nizks from one-way functions. *Cryptology ePrint Archive*, 2023.
- [CCDW20] Weikeng Chen, Alessandro Chiesa, Emma Dauterman, and Nicholas P Ward. Reducing participation costs via incremental verification for ledger systems. *Cryptology ePrint Archive*, 2020.
- [CCG<sup>+</sup>23] Megan Chen, Alessandro Chiesa, Tom Gur, Jack O’Connor, and Nicholas Spooner. Proof-carrying data from arithmetized random oracles. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 379–404. Springer, 2023.
- [CCS22] Megan Chen, Alessandro Chiesa, and Nicholas Spooner. On succinct non-interactive arguments in relativized worlds. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 336–366. Springer, 2022.
- [CG] Jiaqi Cheng and Rishab Goyal. Personal communication.
- [CGSY24] Alessandro Chiesa, Ziyi Guan, Shahar Samocha, and Eylon Yogev. Security bounds for proof-carrying data from straightline extractors. In *Theory of Cryptography Conference*, pages 464–496. Springer, 2024.
- [CJJ21] Arka Rai Choudhuri, Abhihek Jain, and Zhengzhong Jin. Snargs for  $\mathcal{P}$  from LWE. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 68–79. IEEE, 2021.
- [COS20] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. In *Advances in Cryptology—EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part I 39*, pages 769–793. Springer, 2020.
- [CT10] Alessandro Chiesa and Eran Tromer. Proof-carrying data and hearsay arguments from signature cards. In *ICS*, volume 10, pages 310–331, 2010.
- [CTV13] Stephen Chong, Eran Tromer, and Jeffrey A Vaughan. Enforcing language semantics using proof-carrying data. *Cryptology ePrint Archive*, 2013.
- [CTV15] Alessandro Chiesa, Eran Tromer, and Madars Virza. Cluster computing in zero knowledge. In *Advances in Cryptology—EUROCRYPT 2015: 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26–30, 2015, Proceedings, Part II 34*, pages 371–403. Springer, 2015.
- [CW23] Jeffrey Champion and David J Wu. Non-interactive zero-knowledge from non-interactive batch arguments. *Cryptology ePrint Archive*, 2023.

- [CY21] Alessandro Chiesa and Eylon Yogev. Subquadratic snargs in the random oracle model. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*, pages 711–741. Springer, 2021.
- [DGKV22] Lalita Devadas, Rishab Goyal, Yael Kalai, and Vinod Vaikuntanathan. Rate-1 non-interactive arguments for batch-np and applications. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1057–1068. IEEE, 2022.
- [FGK<sup>+</sup>10] David Mandell Freeman, Oded Goldreich, Eike Kiltz, Alon Rosen, and Gil Segev. More constructions of lossy and correlation-secure trapdoor functions. In *Public Key Cryptography–PKC 2010: 13th International Conference on Practice and Theory in Public Key Cryptography, Paris, France, May 26–28, 2010. Proceedings 13*, pages 279–295. Springer, 2010.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.
- [GHMR18] Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, and Ahmadreza Rahimi. Registration-based encryption: removing private-key generator from ibe. In *Theory of Cryptography: 16th International Conference, TCC 2018, Panaji, India, November 11–14, 2018, Proceedings, Part I 16*, pages 689–718. Springer, 2018.
- [GKP<sup>+</sup>13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. How to run turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2013.
- [GV20] Rishab Goyal and Satyanarayana Vusirikala. Verifiable registration-based encryption. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 621–651. Springer, Heidelberg, August 2020.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 99–108, 2011.
- [HN23] Mathias Hall-Andersen and Jesper Buus Nielsen. On valiant’s conjecture: impossibility of incrementally verifiable computation from random oracles. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 438–469. Springer, 2023.
- [HW15] Pavel Hubacek and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 163–172, 2015.
- [KB23] Assimakis Kattis and Joseph Bonneau. Proof of necessary work: succinct state verification with fairness guarantees. In *International Conference on Financial Cryptography and Data Security*, pages 18–35. Springer, 2023.

- [KLVW23] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Boosting batch arguments and ram delegation. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1545–1552, 2023.
- [KRR14] Yael Tauman Kalai, Ran Raz, and Ron D Rothblum. How to delegate computations: the power of no-signaling proofs. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 485–494, 2014.
- [KS22] Abhiram Kothapalli and Srinath Setty. Supernova: Proving universal machine executions without universal circuits. *Cryptology ePrint Archive*, 2022.
- [KST22] Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla. Nova: Recursive zero-knowledge arguments from folding schemes. In *Annual International Cryptology Conference*, pages 359–388. Springer, 2022.
- [Mer87] Ralph C Merkle. A digital signature based on a conventional encryption function. In *Conference on the theory and application of cryptographic techniques*, pages 369–378. Springer, 1987.
- [Min17] O(1) labs. mina cryptocurrency. <https://minaprotocol.com/>, 2017.
- [Nao03] Moni Naor. On cryptographic assumptions and challenges. In *Annual International Cryptology Conference*, pages 96–109. Springer, 2003.
- [NT16] Assa Naveh and Eran Tromer. Photoproof: Cryptographic image authentication for any set of permissible transformations. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 255–271. IEEE, 2016.
- [OPWW15] Tatsuaki Okamoto, Krzysztof Pietrzak, Brent Waters, and Daniel Wichs. New realizations of somewhere statistically binding hashing and positional accumulators. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 121–145. Springer, Heidelberg, November / December 2015.
- [PP22] Omer Paneth and Rafael Pass. Incrementally verifiable computation via rate-1 batch arguments. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1045–1056. IEEE, 2022.
- [PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 187–196, 2008.
- [Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *Theory of Cryptography: Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008. Proceedings 5*, pages 1–18. Springer, 2008.
- [XZC<sup>+</sup>22] Tiancheng Xie, Jiaheng Zhang, Zerui Cheng, Fan Zhang, Yupeng Zhang, Yongzheng Jia, Dan Boneh, and Dawn Song. zkbridge: Trustless cross-chain bridges made practical. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 3003–3017, 2022.



[ZGSX23] Tianyu Zheng, Shang Gao, Yubo Song, and Bin Xiao. Leaking arbitrarily many secrets: Any-out-of-many proofs and applications to ringct protocols. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 2533–2550. IEEE, 2023.

## A zk-IVsC Complete Analysis.

*Proof of Theorem 3.5.* We prove completeness, efficiency, strong soundness, and zero-knowledge properties of our construction.

**Completeness.** The completeness follows by the construction and the completeness of the underlying seBARG, NIZK, PKE, H, HT and SEH.

**Efficiency.** Recall that we are applying almost rate-1 SEH, and almost rate-1 seBARG. Moreover, note that the size of a ciphertext only grows polynomially with the message size and the size of a NIZK proof only grows polynomially with the size witness size and the statement size. Hence the size of each proof  $\pi_i$  is computed as follows (where  $c$  and  $c'$  are constants):

- $|\pi_{\text{out}}| = \text{poly}(\lambda, \log S, n_{\text{out}})$ , and  $|\text{epcf}|, |\text{pcf}|, |\pi_0| = \text{poly}(\lambda, \log S, \log n)$ .
- $|\pi_1| = 2n_{\text{cf},0} + (1 + c/\lambda)|\pi_0| + \text{poly}(\lambda) \leq (1 + c/\lambda) \cdot \text{poly}(\lambda, \log S, \log n)$
- For  $\ell \leq \log T - 1$ ,  $|\pi_\ell| = (1 + c'/\lambda)n_{\text{cf},\ell-1} + 2|\rho| + (1 + c/\lambda)|\pi_{\ell-1}| + \text{poly}(\lambda) \leq (1 + (c + c')/\lambda)^\ell \cdot \text{poly}(\lambda, \log S, \log n, \ell)$ .
- For  $\log(T \cdot K) \geq \ell \geq \log T$ ,  $|\pi_\ell| = (1 + c'/\lambda)n_{\text{cf},\ell-1} + 2|\rho| + 2\lambda + (1 + c/\lambda)|\pi_{\ell-1}| + \text{poly}(\lambda) \leq (1 + (c + c')/\lambda)^\ell \cdot \text{poly}(\lambda, \log S, \log n, \ell) \leq e^{c+c'} \cdot \text{poly}(\lambda, \log S, \log n, \ell)$ .

Therefore it holds that  $|\Pi| = \text{poly}(\lambda, \log S, \log n, \log T, \log K)$ . Note that here we consider RAM machines with polynomially bounded outputs. If the output size  $n_{\text{out}}$  grows larger (say, grows with  $S$ ), we slightly modify the construction so that the proof size does not grow with  $n_{\text{out}}$  (see Section 3.4).

By the same argument since verifier only checks seBARG proofs and SEH path openings validity, and since the verification circuit size for BARGs at the base level (which grows with the NIZK verification size which in turn grows with the size of  $\text{is-nxt-cnfg}_{\mathcal{M}}$  and with encryption) is  $\text{poly}(\lambda, \log S, \log n, n_{\text{out}})$ , the verifier's running time grows with  $\text{poly}(\lambda, \log S, \log n, \log T, \log K)$ .

The setup running time grows with the batch argument, the NIZK and SEH setup running time, and the pseudo-configuration computation time. The first three grow at most with  $\text{poly}(\lambda, \log S, \log n, \ell)$  for  $\ell \leq \lambda$ , thus the total running time of the  $\text{crs}$  and  $\text{hk}$  generations is  $\text{poly}(\lambda, \log S, \log n)$ . For pseudo-configuration computation, the size of the work tape can be large, but since it is initialized to all zeros, the pseudo-configuration computation can be done efficiently by taking advantage of the symmetric structure of the computation when run on all zeros input. Thus, we only need an order of tree-depth steps to compute the root, and hence the pseudo-configuration computation time is  $\text{poly}(\lambda, \log S)$  making the total running time of the setup  $\text{poly}(\lambda, \log S, \log n)$ . With the same argument, the  $\text{crs}$  size is  $\text{poly}(\lambda, \log S, \log n)$ .

Finally, assuming RAM access to the memory (in the underlined parts in the construction) that includes the entire tree of the input and configurations, one can generate reading and writing openings by only reading the path from root to leaf of the tree. Now since each element on the path is

$\text{poly}(\lambda)$ -size and the path length is logarithmic in  $n$  and  $S$ , the total running time of a single reading or writing generation is  $\text{poly}(\lambda, \log S, \log n)$ . Thus, the running time of a single step of update only grows with  $\text{poly}(\lambda, \log S, \log n, \log T, \log k, n_{\text{out}})$ . Hence, the running time of the update grows with  $T$  single steps plus the hash tree computation of the beginning and input digests. Therefore, the total running time of **Update** grows with  $\text{poly}(\lambda, n) + \text{poly}(\lambda, S) + T \cdot \text{poly}(\lambda, \log S, \log n, \log T, \log k, n_{\text{out}})$ .

**Strong Soundness.** Define the original experiment  $\text{exp}_{\text{orig}}$  and the event **Bad** as follows:

$$\text{exp}_{\text{orig}} := \left[ \begin{array}{l} (n, S, T) \leftarrow \mathcal{A}(1^\lambda), \\ \text{crs} \leftarrow \text{Gen}(1^\lambda, n, S, T), \\ (\text{hz}, (\text{out}^{(b)}, \Pi^{(b)})_{b \in \{0,1\}}, 1^{K \cdot T}) \leftarrow \mathcal{A}(\text{crs}) \end{array} \right],$$

$$\text{Bad} := \left[ \begin{array}{l} \text{Verify}(\text{crs}, (\text{hz}, \text{out}^{(0)}), \Pi^{(0)}) = 1 \wedge \\ \text{Verify}(\text{crs}, (\text{hz}, \text{out}^{(1)}), \Pi^{(1)}) = 1 \wedge \\ \text{out}^{(0)} \neq \text{out}^{(1)} \end{array} \right].$$

Our goal is to prove that for any stateful PPT adversary  $\mathcal{A}$  and any  $\lambda \in \mathbb{N}$ ,

$$\Pr [\text{Bad} : \text{exp}_{\text{orig}}] \leq \text{negl}(\lambda).$$

**Notation.** For  $t = T$ , we define  $(k, t, t+1) := (k+1, 0, 1)$  (e.g.,  $(\text{cf}_{k,t}, \text{cf}_{k,t+1}) := (\text{cf}_{k+1,0}, \text{cf}_{k+1,1})$ ). Let  $\text{pcf}_{k,t}^{(b)}$  (resp.  $\text{epcf}_{k,t}^{(b)}$ ) be the pseudo-configuration (resp. encrypted pseudo-configuration) of  $t^{\text{th}}$  configuration on the  $k^{\text{th}}$  input in the proof  $\Pi^{(b)}$  for  $b \in \{0,1\}$ . Furthermore, let  $\text{Gen}_{\alpha,\beta}$  be a generation algorithm with respect to  $\alpha, \beta$ , where the NIZK extractor (instead of NIZK setup) is used to generate  $\text{crs}_{\text{zk}}$ , i.e.  $(\text{crs}_{\text{zk}}, \text{td}_{\text{zk}}) \leftarrow \text{NIZK}.\mathcal{E}(1^\lambda, 1^{n_{\text{zk},x}}, 1^{n_{\text{zk},w}})$ , and let  $\text{Gen}_{\alpha,\beta}$  in addition to the  $\text{crs}$  output  $\text{td} = (\text{td}_{\text{bp}}, \text{td}_{\text{cf}}, \text{td}_{\text{zk}}, \text{sk})$  (where  $\text{td}_{\text{bp}} = (\text{td}_{\text{bp},\ell})_{\ell \in [\lambda]}$  and  $\text{td}_{\text{cf}} = (\text{td}_{\text{cf},\ell})_{\ell \in [\lambda]}$ ). Define  $\text{depth}(\cdot)$  for a hash value  $v$  (resp. a proof  $\pi^*$ ) be  $\ell$  if  $\text{hk}_{\text{cf},\ell}$  (resp.  $\text{crs}_{\text{bp},\ell}$ ) is used to generate that value. For any hash value  $v$  of depth 0 let  $v = (\text{epcf}_1, \text{epcf}_2)$ , then we denote by  $v[b] = \text{epcf}_b$ . We say  $\alpha$  or  $\beta$  corresponds to  $(k, t-1, t)$  if it indicates the unique path to  $v^* = (\text{epcf}_{k,t-1}, \text{epcf}_{k,t})$  from a root of depth  $\lambda$  (a complete tree). Additionally, define the following extraction algorithms:

- $\text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)}) \rightarrow (\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)})$ . It finds  $(k, t-1, t)$  corresponding to  $\alpha$  (w.l.o.g.  $\text{td}_{\text{cf},\ell}$  includes  $\alpha[\ell]$  for all  $\ell \in [\lambda]$ ). Then locates  $v_d^{(b)} \in \Pi^{(b)}$  (where  $d = \text{depth}(v_d)$ ) such that it includes  $(\text{pcf}_{k,t-1}, \text{pcf}_{k,t})$ . Then recursively uses  $v_{\ell-1}^{(b)} = \text{SEH}.\text{Extract}(\text{td}, v_\ell^{(b)})$  for  $\ell \in [d]$  to find  $(\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)})$ . Since this algorithm only uses  $v^{(b)} \in \pi_\ell^{(b)} \in \Pi^{(b)}$ , we can w.l.o.g. let  $\text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)}) \equiv \text{Extract}_{\text{cf}}(\text{td}, \pi_\ell^{(b)}) \equiv \text{Extract}_{\text{cf}}(\text{td}, v^{(b)})$ .
- $\text{Extract}_{\text{bp}}(\text{td}, \Pi^{(b)}) \rightarrow w^{(b)}$ . It finds  $(k, t-1, t)$  corresponding to  $\beta$ , then locates  $\pi_\ell^{(b)} = (v^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, h^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$  such that it includes a proof for  $(\text{pcf}_{k,t-1}, \text{pcf}_{k,t})$ . Then it computes  $w^{(b)} \leftarrow \text{seBARG}.\mathcal{E}(\text{td}_{\text{bp},\ell}, (h^{(b)}, v^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)})$ . Since this algorithm only uses  $\hat{\pi}^{(b)} \in \pi_\ell^{(b)} \in \Pi^{(b)}$ , we can w.l.o.g. let  $\text{Extract}_{\text{bp}}(\text{td}, \Pi^{(b)}) \equiv \text{Extract}_{\text{bp}}(\text{td}, \pi_\ell^{(b)}) \equiv \text{Extract}_{\text{bp}}(\text{td}, \hat{\pi}^{(b)})$ .
- $\text{Extract}_r(\text{td}, \Pi^{(b)}) \rightarrow (\text{epcf}_r^{(b)})$ . It finds  $(k, t, t+1)$  corresponding to  $\beta$ . Let the path to  $(k, t, t+1)$  be  $(1, \dots, 1, 2, \beta[j+1], \dots, \beta[\lambda])$  for some  $j \in [\lambda]$  (number of 1's at the beginning can be 0). In words this algorithm extracts using batch proofs until finding a proof at level  $j$  and

its witness, then it uses path opening from *right* to find an encrypted pseudo-configuration  $\text{epcf}_r^{(b)}$  corresponding to  $\text{cf}_{k,t}$ .

Formally,  $\text{Extract}_r$  finds a proof  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$  such that it contains a poof for  $(\text{pcf}_{k,t}, \text{pcf}_{k,t+1})$ . Now find  $\text{epcf}_r^{(b)}$  as follows:

- If  $(k, t)$  is the leftmost config in the tree with root  $\mathbf{v}^{(b)}$  (e.g.,  $\ell < j$ ), then let  $\rho_1^{(b)} = ((\mathbf{v}_{0,1}^{(b)}, \mathbf{v}_{0,2}^{(b)}), \dots, (\mathbf{v}_{\ell-1,1}^{(b)}, \mathbf{v}_{\ell-1,2}^{(b)}))$ ,  $\mathbf{v}_{0,1}^{(b)} = (\text{epcf}_r^{(b)}, \text{epcf}')$ .
- Otherwise, it holds that  $\ell \geq j$ , then recursively use  $\text{seBARG.Extract}(\text{td}_{\text{bp},d}, (x_{d,1}^{(b)}, x_{d,2}^{(b)})) = w_d^{(b)}$  for  $d \in [j, \ell]$  where  $x_{d,i}^{(b)} = (\tilde{\mathbf{h}}_d^{(b)}, \tilde{\mathbf{v}}_d^{(b)}, i)$  for  $i \in [2]$  constructed as follows, (1)  $(\tilde{\mathbf{h}}_\ell^{(b)}, \tilde{\mathbf{v}}_\ell^{(b)}) = (\mathbf{v}^{(b)}, \mathbf{h}^{(b)})$ , (2) find  $(\tilde{\mathbf{v}}_{d-1,1}^{(b)}, \tilde{\mathbf{v}}_{d-1,2}^{(b)}) \in w_d^{(b)}$  and let  $\tilde{\mathbf{v}}_{d-1}^{(b)} = \tilde{\mathbf{v}}_{d-1, \beta[d]}^{(b)}$ , and (3) if  $d < \log T$  let  $\tilde{\mathbf{h}}_{d-1}^{(b)} = \tilde{\mathbf{h}}_d^{(b)}$ , otherwise Find  $(\tilde{\mathbf{h}}_{d-1,1}^{(b)}, \tilde{\mathbf{h}}_{d-1,2}^{(b)}) \in w_d^{(b)}$  and let  $\tilde{\mathbf{h}}_{d-1}^{(b)} = \tilde{\mathbf{h}}_{d-1, \beta[d]}^{(b)}$ . Finally find  $(\tilde{\mathbf{v}}_{j-1,2}^{(b)}, \tilde{\rho}_{j-1,2}^{(b)}) \in w_j^{(b)}$  and do the following, (1) if  $j = 1$  then let  $\tilde{\mathbf{v}}_{j-1,2}^{(b)} = (\text{epcf}_r^{(b)}, \text{epcf}')$ , (2) otherwise let  $\rho_{j-1,2}^{(b)} = ((\mathbf{v}_{0,1}^{(b)}, \mathbf{v}_{0,2}^{(b)}), \dots, (\mathbf{v}_{j-2,1}^{(b)}, \mathbf{v}_{j-2,2}^{(b)}))$ , and  $\mathbf{v}_{0,1}^{(b)} = (\text{epcf}_r^{(b)}, \text{epcf}')$ .

- $\text{Extract}_l(\text{td}, \Pi^{(b)}) \rightarrow (\text{epcf}_l^{(b)})$ . This is similar to  $\text{Extract}_r$  except that in the end it uses path opening from *left* to find an encrypted pseudo-configuration  $\text{epcf}_l^{(b)}$ . Namely, it finds  $(k, t, t+1)$  corresponding to  $\beta$ . Let the path to  $(k, t, t+1)$  be  $(1, \dots, 1, 2, \beta[j+1], \dots, \beta[\lambda])$  for some  $j \in [\lambda]$  (number of 1's at the beginning can be 0).

Formally,  $\text{Extract}_l$  finds a proof  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$  such that it contains a poof for  $(\text{pcf}_{k,t-1}, \text{pcf}_{k,t})$ . Now find  $\text{epcf}_l^{(b)}$  as follows:

- If  $(k, t)$  is the rightmost config in the tree with root  $\mathbf{v}^{(b)}$  (e.g.,  $\ell < j$ ), then let  $\rho_2^{(b)} = ((\mathbf{v}_{0,1}^{(b)}, \mathbf{v}_{0,2}^{(b)}), \dots, (\mathbf{v}_{\ell-1,1}^{(b)}, \mathbf{v}_{\ell-1,2}^{(b)}))$ ,  $\mathbf{v}_{0,2}^{(b)} = (\text{epcf}', \text{epcf}_l^{(b)})$ .
- Otherwise similar to  $\text{Extract}_r$  find  $w_j^{(b)}$ . Finally find  $(\tilde{\mathbf{v}}_{j-1,1}^{(b)}, \tilde{\rho}_{j-1,1}^{(b)}) \in w_j^{(b)}$  and do the following, (1) if  $j = 1$  then let  $\tilde{\mathbf{v}}_{j-1,1}^{(b)} = (\text{epcf}', \text{epcf}_l^{(b)})$ , (2) otherwise let  $\rho_{j-1,1}^{(b)} = ((\mathbf{v}_{0,1}^{(b)}, \mathbf{v}_{0,2}^{(b)}), \dots, (\mathbf{v}_{j-2,1}^{(b)}, \mathbf{v}_{j-2,2}^{(b)}))$ ,  $\mathbf{v}_{0,2}^{(b)} = (\text{epcf}', \text{epcf}_l^{(b)})$ .

For any stateful PPT adversary  $\mathcal{A}$  and any  $\lambda \in \mathbb{N}$  define the experiment  $\text{exp}_{\alpha, \beta}$  (where  $\alpha$

corresponds to  $(k, t - 1, t)$  as follows:

$$\text{exp}_{\alpha, \beta} := \left[ \begin{array}{l} n, S, T \leftarrow \mathcal{A}(1^\lambda), \\ (\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha, \beta}(1^\lambda, n, S, T), \\ (\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T}) \leftarrow \mathcal{A}(\text{crs}), \\ (\text{epcf}_{k, t-1}^{(b)}, \text{epcf}_{k, t}^{(b)}) \leftarrow \text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)}) \text{ for } b \in \{0, 1\}, \\ \text{pcf}_{k, t-b'}^{(b)} = \text{PKE.Dec}(\text{sk}, \text{epcf}_{k, t-b'}^{(b)}) \text{ for } b, b' \in \{0, 1\}, \\ w^{(b)} \leftarrow \text{Extract}_{\text{bp}}(\text{td}, \Pi^{(b)}) \text{ for } b \in \{0, 1\}, \\ \text{epcf}_r^{(b)} \leftarrow \text{Extract}_r(\text{td}, \Pi^{(b)}) \text{ for } b \in \{0, 1\}, \\ \text{pcf}_r^{(b)} = \text{PKE.Dec}(\text{sk}, \text{epcf}_r^{(b)}) \text{ for } b \in \{0, 1\}, \\ \text{epcf}_l^{(b)} \leftarrow \text{Extract}_l(\text{td}, \Pi^{(b)}) \text{ for } b \in \{0, 1\}, \\ \text{pcf}_l^{(b)} = \text{PKE.Dec}(\text{sk}, \text{epcf}_l^{(b)}) \text{ for } b \in \{0, 1\} \end{array} \right]$$

Note that if  $\alpha = \beta = 1^\lambda$ , then  $\text{exp}_{\alpha, \beta}$  is the same as the strong soundness experiment except that the setup algorithm uses NIZK extractor to generate  $\text{crs}_{\text{zk}}$  and the experiment performs some additional computations on  $\mathcal{A}$ 's output.

We define the following hybrids.

$\mathcal{H}_0$ : This is the original experiment.

$\mathcal{H}_{k, t, 1}$ : This is  $\text{exp}_{\alpha, \beta}$  where  $\alpha$  and  $\beta$  correspond to  $(k, t - 1, t)$ .

$\mathcal{H}_{k, t, 2}$ : This is  $\text{exp}_{\alpha, \beta}$  where  $\alpha$  corresponds to  $(k, t - 1, t)$ , and  $\beta$  corresponds to  $(k, t, t + 1)$ .

$\mathcal{H}_{k, t, 3}$ : This is  $\text{exp}_{\alpha, \beta}$  where  $\alpha$  and  $\beta$  correspond to  $(k, t, t + 1)$ . Note that this hybrid is equal to  $\mathcal{H}_{k, t+1, 1}$  (or  $\mathcal{H}_{k+1, 1, 1}$  if  $t = T$ ).

If proofs  $\Pi^{(0)}$  and  $\Pi^{(1)}$  verify, we prove the followings hold except with a negligible probability:

- (Claim A.1) For any adversary  $\mathcal{A}$  the probability that  $\text{out}^{(0)} \neq \text{out}^{(1)}$  in  $\mathcal{H}_0$  and  $\mathcal{H}_{1, 1, 1}$  is negligibly close (henceforth, our goal is to prove  $\Pr[\text{Bad} : \text{exp}_{\alpha, \beta}] \leq \text{negl}(\lambda)$ ).
- (Claim A.2) In hybrid  $\mathcal{H}_{1, 1, 1}$   $\text{pcf}_{1, 0}^{(0)} = \text{pcf}_{1, 0}^{(1)}$ .
- (Claim A.3) In hybrid  $\mathcal{H}_{k, t, 1}$  if  $\text{pcf}_{k, t-1}^{(0)} = \text{pcf}_{k, t-1}^{(1)}$  then  $\text{pcf}_{k, t}^{(0)} = \text{pcf}_{k, t}^{(1)}$ .
- (Claim A.4)  $\mathcal{H}_{k, t, 1}$  and  $\mathcal{H}_{k, t, 2}$  are indistinguishable.
- (Claim A.5)  $\mathcal{H}_{k, t, 2}$  and  $\mathcal{H}_{k, t, 3}$  are indistinguishable.
- (Claim A.6) In hybrid  $\mathcal{H}_{N, T, 1}$  if  $\text{pcf}_{N, T}^{(0)} = \text{pcf}_{N, T}^{(1)}$  then the adversary cannot generate valid proofs  $\Pi^{(0)}$  and  $\Pi^{(1)}$  for two different final outputs  $\text{out}^{(0)}$  and  $\text{out}^{(1)}$ .

**Claim A.1.** If NIZK extractor satisfies crs indistinguishability and  $\alpha = \beta = 1^\lambda$ , then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\text{Bad} : \text{exp}_{\alpha, \beta}] - \Pr[\text{Bad} : \text{exp}_{\text{orig}}]| \leq \text{negl}(\lambda).$$

*Proof.* The proof directly follows from the crs indistinguishability of the NIZK extractor.  $\square$

**Claim A.2.** If SEH is somewhere binding w.r.t. path opening, PKE is correct, and  $\alpha = \beta = 1^\lambda$ , then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{1,0}^{(0)} \neq \text{pcf}_{1,0}^{(1)} : \exp_{\alpha,\beta}] \leq \text{negl}(\lambda).$$

*Proof.* We will prove that for  $b \in \{0, 1\}$ ,  $\text{pcf}_{1,0}^{(b)} = \text{pcf}_{1,0}$  (as given in the crs). First let  $\ell \in [\lambda]$  be the biggest index such that  $\pi_\ell^{(b)} \neq \epsilon$ , then parse  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)})$ ,  $\rho_1^{(b)} = ((\mathbf{v}_{0,1}^{(b)}, \mathbf{v}_{0,2}^{(b)}), \dots, (\mathbf{v}_{\ell-1,1}^{(b)}, \mathbf{v}_{\ell-1,2}^{(b)}))$ , and  $\mathbf{v}_{0,1}^{(b)} = (\text{epcf}_{1,0}^{(b)}, \text{epcf}_{1,1}^{(b)})$ . Note that since  $\alpha = \beta = 1^\lambda$ , then  $\text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)}) = (\text{epcf}_{1,0}^{(b)}, \text{epcf}_{1,1}^{(b)})$  by somewhere binding w.r.t. path opening of SEH. Also, note that by the verifier's check  $\text{epcf}_{1,0}^{(b)} = \text{PKE.Enc}(\text{pk}, \text{pcf}_{1,0}; 0^{|r|})$  for  $b \in \{0, 1\}$ . Then by PKE correctness it holds that  $\text{pcf}_{1,0}^{(b)} = \text{pcf}_{1,0}$ , and hence we have  $\text{pcf}_{1,0}^{(0)} = \text{pcf}_{1,0}^{(1)}$ .  $\square$

**Claim A.3.** If seBARG is somewhere extractable, NIZK has knowledge extractor, PKE is correct, SEH is somewhere binding, HT has reading and writing soundness, and CRHF is collision resistant, and  $\alpha$  and  $\beta$  correspond to  $(k, t-1, t)$ , and for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{k,t-1}^{(0)} \neq \text{pcf}_{k,t-1}^{(1)} : \exp_{\alpha,\beta}] \leq \text{negl}(\lambda),$$

then for any such adversary, it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta}] \leq \text{negl}(\lambda).$$

*Proof.* To prove the claim, by total law of the probability we only need to show that for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$\delta_0 = \Pr[\text{Bad} \wedge \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta}] \leq \text{negl}(\lambda).$$

We prove this by induction on the depth  $\ell$  of the proof  $\pi_\ell^{(b)}$  that contains  $(\text{pcf}_{k,t-1}, \text{pcf}_{k,t})$  (the depth is the same for both proofs as it only depends on the time  $T$  and number of inputs  $N$ ). Note that this is a two-step induction, first on the depth  $\ell \leq \log T - 1$ , and then on the depth  $\ell \geq \log T$ . Also note that if  $\ell < \log T - 1$  then there is no  $\mathbf{h}^{(b)}$  in the proof  $\pi_\ell^{(b)}$  that contains  $(\text{pcf}_{k,t-1}, \text{pcf}_{k,t})$ , thus in this case we let  $\mathbf{h}^{(b)} = \mathbf{h}_{\log T-1}$  where  $\mathbf{h}_{\log T-1} \in \text{hz}$ .

Let  $(\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, \hat{\pi}^{(b)}) \in \pi_\ell^{(b)}$  and define  $\text{Ver}_\ell$  be the following event:

$$\text{Ver}_\ell := \left[ \begin{array}{l} \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell}, (\mathbf{h}^{(0)}, \mathbf{v}^{(0)}, j)_{j \in [2]}, \hat{\pi}^{(0)}) = 1 \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell}, (\mathbf{h}^{(1)}, \mathbf{v}^{(1)}, j)_{j \in [2]}, \hat{\pi}^{(1)}) = 1 \end{array} \right]$$

and

$$\delta_{1,\ell} = \Pr[\text{Ver}_\ell \wedge \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta}].$$

Since  $\text{Ver}_\ell$  is implied by  $\text{Bad}$ , it holds that  $\delta_{1,\ell} \geq \delta_0$ , thus it suffices to prove that  $\delta_{1,\ell} \leq \text{negl}(\lambda)$ .

In the following induction let  $\text{hz} = (\mathbf{h}_{\log T-1}, \dots, \mathbf{h}_\lambda)$  and note that  $\mathbf{h}^{(b)} = \mathbf{h}_{\log T-1}$  by the verifier's check.

**Induction Base.** Now suppose  $\pi_1^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$  is the proof that contains  $(\text{pcf}_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)})$  (i.e.  $\ell = 1$ ). This means  $\text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)}) = \text{SEH.Extract}(\text{td}_{\text{cf},1}, \mathbf{v}^{(b)})$  and  $w^{(b)} = \text{Extract}_{\text{bp}}(\text{td}, \Pi^{(b)}) = \text{seBARG.Extract}(\text{td}_{\text{bp},1}, (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)}) = (\tilde{\mathbf{v}}_1^{(b)}, \tilde{\mathbf{v}}_2^{(b)}, \pi_{\text{zk}}^{(b)})$ , is a witness for  $x^{(b)} = (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, \beta[1])$ . Let  $\tilde{\mathbf{v}}_{\beta[1]}^{(b)} = (\tilde{\text{epcf}}_{k,t-1}^{(b)}, \tilde{\text{epcf}}_{k,t}^{(b)})$ .

Now let  $\text{exp}_{\alpha,\beta}$  additionally compute  $\text{NIZK.E}(\text{td}_{\text{zk}}, x_{\text{zk}}^{(b)} = (\mathbf{h}^{(b)}, \tilde{\text{epcf}}_{k,t-1}^{(b)}, \tilde{\text{epcf}}_{k,t}^{(b)}, \pi_{\text{zk}}^{(b)}) = w_{\text{zk}}^{(b)}$  and parse  $w_{\text{zk}}^{(b)} = (\text{rt}_{\text{in}}^{(b)}, r_{\text{in}}^{(b)}, \tilde{\text{pcf}}_{k,t-1}^{(b)}, r_{k,t-1}^{(b)}, \tilde{\text{pcf}}_{k,t}^{(b)}, r_{k,t}^{(b)}, \text{rbit}^{(b)}, \text{rop}^{(b)}, \text{wop}^{(b)})$ .

Recall that  $\mathcal{L}_{\text{bp},1}$  from Fig. 4 is the corresponding language for  $(x^{(b)}, w^{(b)})$  and  $\mathcal{L}_{\text{zk}}$  from Fig. 3 is the corresponding language for  $(x_{\text{zk}}^{(b)}, w_{\text{zk}}^{(b)})$ .

Now consider the following probabilities:

$$\delta_2 = \Pr \left[ \begin{array}{l} (x^{(0)}, w^{(0)}) \in \mathcal{L}_{\text{bp},1} \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},1}, (\mathbf{h}^{(1)}, \mathbf{v}^{(1)}, j)_{j \in [2]}, \hat{\pi}^{(1)}) = 1 \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \text{exp}_{\alpha,\beta} \right].$$

$$\delta_3 = \Pr \left[ \begin{array}{l} (x^{(b)}, w^{(b)}) \in \mathcal{L}_{\text{bp},1} \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \text{exp}_{\alpha,\beta} \right].$$

$$\delta_4 = \Pr \left[ \begin{array}{l} (\tilde{\text{epcf}}_{k,t-1}^{(0)}, \tilde{\text{epcf}}_{k,t}^{(0)}) = (\text{epcf}_{k,t-1}^{(0)}, \text{epcf}_{k,t}^{(0)}) \wedge \\ \text{NIZK.Verify}(\text{crs}_{\text{zk}}, x_{\text{zk}}^{(0)}, \pi_{\text{zk}}^{(0)}) = 1 \\ (x^{(1)}, w^{(1)}) \in \mathcal{L}_{\text{bp},1} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \text{exp}_{\alpha,\beta} \right].$$

$$\delta_5 = \Pr \left[ \begin{array}{l} (\tilde{\text{epcf}}_{k,t-1}^{(b)}, \tilde{\text{epcf}}_{k,t}^{(b)}) = (\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}) \text{ for } b \in \{0, 1\} \wedge \\ \text{NIZK.Verify}(\text{crs}_{\text{zk}}, x_{\text{zk}}^{(b)}, \pi_{\text{zk}}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \text{exp}_{\alpha,\beta} \right].$$

$$\delta_6 = \Pr \left[ \begin{array}{l} (\tilde{\text{epcf}}_{k,t-1}^{(b)}, \tilde{\text{epcf}}_{k,t}^{(b)}) = (\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}) \text{ for } b \in \{0, 1\} \wedge \\ (x_{\text{zk}}^{(0)}, w_{\text{zk}}^{(0)}) \in \mathcal{L}_{\text{zk}} \wedge \\ \text{NIZK.Verify}(\text{crs}_{\text{zk}}, x_{\text{zk}}^{(1)}, \pi_{\text{zk}}^{(1)}) = 1 \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \text{exp}_{\alpha,\beta} \right].$$

$$\delta_7 = \Pr \left[ \begin{array}{l} (\tilde{\text{epcf}}_{k,t-1}^{(b)}, \tilde{\text{epcf}}_{k,t}^{(b)}) = (\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}) \text{ for } b \in \{0, 1\} \wedge \\ (x_{\text{zk}}^{(b)}, w_{\text{zk}}^{(b)}) \in \mathcal{L}_{\text{zk}} \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \text{exp}_{\alpha,\beta} \right].$$

$$\delta_8 = \Pr \left[ \begin{array}{l} (\tilde{\text{pcf}}_{k,t-1}^{(b)}, \tilde{\text{pcf}}_{k,t}^{(b)}) = (\text{pcf}_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)}) \text{ for } b \in \{0, 1\} \wedge \text{rt}_{\text{in}}^{(0)} = \text{rt}_{\text{in}}^{(1)} \wedge \\ \text{is-nxt-cnfg}_{\mathcal{M}}(\text{hk}_{\text{ht}}, \text{rt}_{\text{in}}^{(b)}, \tilde{\text{pcf}}_{k,t-1}^{(b)}, \tilde{\text{pcf}}_{k,t}^{(b)}, \text{rbit}^{(b)}, \text{rop}^{(b)}, \text{wop}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \text{exp}_{\alpha,\beta} \right].$$

$$\delta_9 = \Pr \left[ \begin{array}{l} \text{rbit}^{(0)} = \text{rbit}^{(1)} \wedge \\ \mathbb{C}\mathcal{M}(\tilde{\mathbf{q}}_{k,t-1}^{(b)}, \text{rbit}^{(b)}) = (\hat{\mathbf{q}}_{k,t}^{(b)}, \hat{\text{rid}}_{k,t}^{(b)}, \text{wid}\mathbf{x}^{(b)}, \text{wbit}^{(b)}) \wedge \\ (\tilde{\mathbf{q}}_{k,t}^{(b)}, \tilde{\text{rid}}_{k,t}^{(b)}) = (\hat{\mathbf{q}}_{k,t}^{(b)}, \hat{\text{rid}}_{k,t}^{(b)}) \wedge \\ \text{HT.VfyWt}(\text{hk}_{\text{ht}}, \tilde{\text{rt}}_{k,t-1}^{(b)}, \text{wid}\mathbf{x}^{(b)}, \text{wbit}^{(b)}, \tilde{\text{rt}}_{k,t}^{(b)}, \text{wop}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ (\tilde{\text{pcf}}_{k,t-1}^{(b)}, \tilde{\text{pcf}}_{k,t}^{(b)}) = (\text{pcf}_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)}) \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} \right] : \exp_{\alpha,\beta}.$$

We let  $\delta_1 = \delta_{1,1}$ , and  $\epsilon_i = \delta_i - \delta_{i+1}$  for  $i \in [8]$  and  $\epsilon_9 = \delta_9$ . Since  $\delta_1 = \sum_{i=1}^6 \epsilon_i$  we only need to show that for  $i \in [9]$ ,  $\epsilon_i$  is negligible. Now the following Eqs. (3) to (8) holds.

**Case 1.** For  $b \in \{0, 1\}$  it holds that:

$$\Pr[\text{seBARG.Verify}(\text{crs}_{\text{bp},1}, (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)}) = 1 \wedge (x^{(b)}, w^{(b)}) \notin \mathcal{L}_{\text{bp},1} : \exp_{\alpha,\beta}] \geq \epsilon_{b+1} \quad (3)$$

Suppose towards the contradiction that for some  $b^* \in \{0, 1\}$ ,  $\epsilon_{b^*+1}(\cdot)$  is non-negligible. Then we construct  $\mathcal{B}_{\text{bp}}$  s.t.  $\mathcal{B}_{\text{bp}} = (\alpha, \beta, 1, \mathcal{A}, b^*)$  breaks the seBARG security. Note that the advantage of  $\mathcal{B}_{\text{bp}}(\alpha, \beta, 1, \mathcal{A}, b^*)$  is greater than that of  $\mathcal{A}$  in Eq. (3) contradicting the seBARG security.

$\mathcal{B}_{\text{bp}}(\alpha, \beta, \ell, \mathcal{A}, b)$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{crs}^*$  from the challenger where  $(\text{crs}^*, \text{td}^*) \leftarrow \text{seBARG.Gen}(1^\lambda, 2, n_{\text{bp},\ell}, \beta[\ell])$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha,\beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{crs}_{\text{bp},\ell}$  with  $\text{crs}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .
4. Finds  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ , and output  $((\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)})$ .

**Case 2.** For  $b \in \{0, 1\}$  it holds that:

$$\Pr \left[ \begin{array}{l} \text{SEH.Hash}(\text{hk}_{\text{cf},1}, (\tilde{\mathbf{v}}_1^{(b)}, \tilde{\mathbf{v}}_2^{(b)})) = \mathbf{v}^{(b)} \wedge \\ (\tilde{\text{epcf}}_{k,t-1}^{(b)}, \tilde{\text{epcf}}_{k,t}^{(b)}) \neq (\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}) \end{array} : \exp_{\alpha,\beta} \right] \geq \epsilon_{b+3} \quad (4)$$

Suppose towards the contradiction that for some  $b^* \in \{0, 1\}$ ,  $\epsilon_{b^*+3}(\cdot)$  is non-negligible. Then we construct  $\mathcal{B}_{\text{seh}}$  s.t.  $\mathcal{B}_{\text{seh}} = (\alpha, \beta, 1, \mathcal{A}, b^*)$  breaks the SEH security. Note that the advantage of  $\mathcal{B}_{\text{seh}}(\alpha, \beta, 1, \mathcal{A}, b^*)$  is greater than that of  $\mathcal{A}$  in Eq. (4) contradicting the SEH security.

$\mathcal{B}_{\text{seh}}(\alpha, \beta, \ell, \mathcal{A}, b)$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{hk}^*$  from the challenger where  $(\text{hk}^*, \text{td}^*) \leftarrow \text{SEH.Gen}(1^\lambda, \Sigma_{\text{cf}}^{\ell-1}, 2, \alpha[\ell])$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha,\beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{hk}_{\text{cf},\ell}$  with  $\text{hk}^*$ .



3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .
4. Finds  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ .
5. Computes  $w^{(b)} = \text{seBARG.Extract}(\text{td}_{\text{bp}, \ell}, (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)})$  and finds  $(\tilde{\mathbf{v}}_1^{(b)}, \tilde{\mathbf{v}}_2^{(b)}) \in w^{(b)}$ .
6. Outputs  $((\tilde{\mathbf{v}}_1^{(b)}, \tilde{\mathbf{v}}_2^{(b)}), \mathbf{v}^{(b)})$

**Case 3.** For  $b \in \{0, 1\}$  it holds that:

$$\Pr[\text{NIZK.Verify}(\text{crs}_{\text{zk}}, (\mathbf{h}^{(b)}, \text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}, \pi_{\text{zk}}^{(b)}) = 1 \wedge (x_{\text{zk}}^{(b)}, w_{\text{zk}}^{(b)}) \notin \mathcal{L}_{\text{zk}} : \exp_{\alpha, \beta}] \geq \epsilon_{b+5} \quad (5)$$

Suppose towards the contradiction that for some  $b^* \in \{0, 1\}$ ,  $\epsilon_{b^*+1}(\cdot)$  is non-negligible. Then we construct  $\mathcal{B}_{\text{zk}}$  s.t.  $\mathcal{B}_{\text{zk}} = (\alpha, \beta, 1, \mathcal{A}, b^*)$  breaks the NIZK knowledge extractor property. Note that the advantage of  $\mathcal{B}_{\text{zk}}(\alpha, \beta, 1, \mathcal{A}, b^*)$  is greater than that of  $\mathcal{A}$  in Eq. (3) contradicting NIZK security.

$\mathcal{B}_{\text{zk}}(\alpha, \beta, \ell, \mathcal{A}, b)$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{crs}^*$  from the challenger where  $(\text{crs}^*, \text{td}^*) \leftarrow \text{NIZK.E}(1^\lambda, 1^{n_x}, 1^{n_w})$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha, \beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{crs}_{\text{zk}}$  with  $\text{crs}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .
4. Finds  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ .
5. Compute  $w^{(b)} = \text{seBARG.Extract}(\text{td}_{\text{bp}, 1}, (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)})$ .
6. Parse  $w^{(b)} = (\tilde{\mathbf{v}}_1^{(b)}, \tilde{\mathbf{v}}_2^{(b)}, \pi_{\text{zk}}^{(b)})$ , and let  $\tilde{\mathbf{v}}_{\beta[1]}^{(b)} = (\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)})$ .
7. Output  $((\mathbf{h}^{(b)}, \text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}), \pi_{\text{zk}}^{(b)})$ .

**Case 4.** We claim  $\epsilon_7 = 0$  (6).

Note that if for  $b \in \{0, 1\}$ ,  $(x_{\text{zk}}^{(b)}, w_{\text{zk}}^{(b)}) \in \mathcal{L}_{\text{zk}}$  (where  $x_{\text{zk}}^{(b)} = (\mathbf{h}^{(b)}, \text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)})$  and  $w_{\text{zk}}^{(b)} = (\text{rt}_{\text{in}}^{(b)}, r_{\text{in}}^{(b)}, \text{pcf}_{k,t-1}^{(b)}, r_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)}, r_{k,t}^{(b)}, \text{rbit}^{(b)}, \text{rop}^{(b)}, \text{wop}^{(b)})$ ), then by the definition of  $\mathcal{L}_{\text{zk}}$  (see Fig. 3) for  $b \in \{0, 1\}$  it holds that  $\text{PKE.Enc}(\text{pk}, \text{rt}_{\text{in}}^{(b)}; r_{\text{in}}^{(b)}) = \mathbf{h}^{(b)}$ ,  $\text{PKE.Enc}(\text{pk}, \text{pcf}_{k,t-1}^{(b)}; r_{k,t-1}^{(b)}) = \text{epcf}_{k,t-1}^{(b)}$ , and  $\text{PKE.Enc}(\text{pk}, \text{pcf}_{k,t}^{(b)}; r_{k,t}^{(b)}) = \text{epcf}_{k,t}^{(b)}$ .

Now first note that, as discussed in the beginning of the claim's proof, it holds that  $\mathbf{h}^{(0)} = \mathbf{h}^{(1)} = \mathbf{h}_{\log T-1}$ , thus by the perfect correctness of PKE for  $b \in \{0, 1\}$  we have  $\text{rt}_{\text{in}}^{(0)} = \text{rt}_{\text{in}}^{(1)}$ . Moreover, by the definition of  $\exp_{\alpha, \beta}$  for  $b, b' \in \{0, 1\}$  it holds that  $\text{pcf}_{k,t-b'}^{(b)} = \text{PKE.Dec}(\text{sk}, \text{epcf}_{k,t-b'}^{(b)})$ . Thus if for  $b \in \{0, 1\}$  it holds that  $(\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}) = (\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)})$ , then by the perfect correctness of PKE for  $b \in \{0, 1\}$  we have  $(\text{pcf}_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)}) = (\text{pcf}_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)})$ . By the same argument, the other direction can be proved, therefore, we conclude that  $\delta_7 = \delta_8$  and  $\epsilon_7 = 0$ .

**Case 5.** If  $\text{ridx}_{k,t-1}^{(0)} \leq n$  then let  $(m^*, i^*) = (h_{\log T-1}, \text{ridx}_{k,t-1}^{(0)})$ , otherwise let  $(m^*, i^*) = (\text{rt}_{k,t-1}^{(0)}, \text{ridx}_{k,t-1}^{(0)})$ . Then it holds that:

$$\Pr \left[ \begin{array}{l} \text{rbit}^{(0)} \neq \text{rbit}^{(1)} \wedge \\ \text{HT.VfyRd}(\text{hk}_{\text{MT}}, m^*, i^*, \text{rbit}^{(b)}, \text{rop}^{(b)}) = 1 \wedge \text{ for } b \in \{0, 1\} \end{array} : \exp_{\alpha, \beta} \right] \geq \epsilon_8 \quad (7)$$

Suppose towards the contradiction that  $\epsilon_8(\cdot)$  is non-negligible. Then we construct  $\mathcal{B}_{\text{ht}}^{\text{read}}$  s.t.  $\mathcal{B}_{\text{ht}}^{\text{read}} = (\alpha, \beta, 1, \mathcal{A})$  breaks the hash tree security. Note that the advantage of  $\mathcal{B}_{\text{ht}}^{\text{read}}(\alpha, \beta, 1, \mathcal{A})$  is greater than that of  $\mathcal{A}$  in Eq. (7) contradicting the hash tree reading soundness.

$\mathcal{B}_{\text{ht}}^{\text{read}}(\alpha, \beta, \ell, \mathcal{A})$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{hk}^*$  from the challenger where  $\text{hk}^* \leftarrow \text{HT.Gen}(1^\lambda)$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha, \beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{hk}_{\text{ht}}$  with  $\text{hk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .
4. For  $b \in \{0, 1\}$ , first finds  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ , then computes  $w^{(b)} = \text{seBARG.Extract}(\text{td}_{\text{bp}, \ell}, (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)})$ ,  $w_{\text{zk}}^{(b)} = \text{NIZK.E}(\text{td}_{\text{zk}}, (\mathbf{h}^{(b)}, \text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}), \pi_{\text{zk}}^{(b)})$ , and finally finds  $(\text{rbit}^{(b)}, \text{rop}^{(b)}) \in w_{\text{zk}}^{(b)}$ .
5. Outputs  $(m^*, i^*, \text{rbit}^{(0)}, \text{rop}^{(0)}, \text{rbit}^{(1)}, \text{rop}^{(1)})$  where  $(m^*, i^*) = (h_{\log T-1}, \text{ridx}_{k,t-1}^{(0)})$  if  $\text{ridx}_{k,t-1}^{(0)} \leq n$ , and  $(m^*, i^*) = (\text{rt}_{k,t-1}^{(0)}, \text{ridx}_{k,t-1}^{(0)})$  otherwise.

**Case 6.** It holds that:

$$\Pr \left[ \begin{array}{l} \mathbb{C}_{\mathcal{M}}(\mathbf{q}_{k,t-1}^{(0)}, \text{rbit}^{(0)}) = (\mathbf{q}_{k,t}^{(b)}, \text{ridx}_{k,t}^{(b)}, \text{wbit}^{(0)}, \text{widx}^{(0)}) \wedge \text{rt}_{k,t}^{(0)} \neq \text{rt}_{k,t}^{(1)} \wedge \\ \text{HT.VfyWt}(\text{hk}_{\text{MT}}, \text{rt}_{k,t-1}^{(0)}, \text{widx}^{(0)}, \text{wbit}^{(0)}, \text{rt}_{k,t}^{(b)}, \text{wop}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \end{array} : \exp_{\alpha, \beta} \right] \geq \epsilon_9 \quad (8)$$

Suppose towards the contradiction that  $\epsilon_9(\cdot)$  is non-negligible. Then we construct  $\mathcal{B}_{\text{ht}}^{\text{write}}$  s.t.  $\mathcal{B}_{\text{ht}}^{\text{write}} = (\alpha, \beta, 1, \mathcal{A})$  breaks the hash tree security. Note that the advantage of  $\mathcal{B}_{\text{ht}}^{\text{read}}(\alpha, \beta, 1, \mathcal{A})$  is greater than that of  $\mathcal{A}$  in Eq. (8) contradicting the hash tree writing soundness.

$\mathcal{B}_{\text{ht}}^{\text{write}}(\alpha, \beta, \ell, \mathcal{A})$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{hk}^*$  from the challenger where  $\text{hk}^* \leftarrow \text{HT.Gen}(1^\lambda)$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha, \beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{hk}_{\text{ht}}$  with  $\text{hk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .
4. For  $b \in \{0, 1\}$ , first finds  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ , then computes  $w^{(b)} = \text{seBARG.Extract}(\text{td}_{\text{bp}, \ell}, (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)})$ ,  $w_{\text{zk}}^{(b)} = \text{NIZK.E}(\text{td}_{\text{zk}}, (\mathbf{h}^{(b)}, \text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}), \pi_{\text{zk}}^{(b)})$ , and finally finds  $(\text{rbit}^{(b)}, \text{wop}^{(b)}) \in w_{\text{zk}}^{(b)}$ .
5. Computes  $\mathbb{C}_{\mathcal{M}}(\mathbf{q}_{k,t-1}^{(0)}, \text{rbit}^{(0)}) = (\mathbf{q}_{k,t}^{(b)}, \text{ridx}_{k,t}^{(b)}, \text{wbit}^{(0)}, \text{widx}^{(0)})$ .
6. Outputs  $(\text{rt}_{k,t-1}^{(0)}, \text{widx}^{(0)}, \text{wbit}^{(0)}, \text{rt}_{k,t}^{(0)}, \text{wop}^{(0)}, \text{rt}_{k,t}^{(1)}, \text{wop}^{(1)})$ .

**Induction Step (1).** Suppose the claim holds for any  $\ell' < \ell \leq \log T - 1$ , namely,  $\delta_{1,\ell'} \leq \text{negl}(\lambda)$ . Let  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$  be the proof that contains  $(\text{pcf}_{k,t-1}^{(b)}, \text{pcf}_{k,t}^{(b)})$ . Now the followings hold:

- $\text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)}) = \text{Extract}_{\text{cf}}(\text{td}, \mathbf{v}^{(b)}) = \text{Extract}_{\text{cf}}(\text{td}, \mathbf{v}_{\alpha[\ell]}^{(b)})$  where  $\mathbf{v}_{\alpha[\ell]}^{(b)} = \text{SEH.Extract}(\text{td}_{\text{cf},\ell}, \mathbf{v}^{(b)})$ . Let  $\text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)})$  additionally output  $\mathbf{v}_{\alpha[\ell]}^{(b)}$ .
- $w^{(b)} = \text{Extract}_{\text{bp}}(\text{td}, \Pi^{(b)}) = \text{seBARG.Extract}(\text{td}_{\text{bp},\ell}, (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)}) = (\tilde{\mathbf{v}}_1^{(b)}, \tilde{\mathbf{v}}_2^{(b)}, \tilde{\rho}_1^{(b)}, \tilde{\rho}_2^{(b)}, \tilde{\pi}^{(b)})$ , is a witness for  $x^{(b)} = (\mathbf{h}^{(b)}, \mathbf{v}^{(b)}, \beta[\ell])$

Let  $\mathcal{L}_{\text{bp},\ell}$  be the corresponding language for  $(x^{(b)}, w^{(b)})$  defined in the *Fig. 4*.

Now consider the following probabilities:

$$\begin{aligned} \delta_2 &= \Pr \left[ \begin{array}{l} (x^{(0)}, w^{(0)}) \in \mathcal{L}_{\text{bp},\ell} \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell}, (\mathbf{h}^{(1)}, \mathbf{v}^{(1)}, j)_{j \in [2]}, \hat{\pi}^{(1)}) = 1 \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \exp_{\alpha,\beta} \right]. \\ \delta_3 &= \Pr \left[ \begin{array}{l} (x^{(b)}, w^{(b)}) \in \mathcal{L}_{\text{bp},\ell} \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \exp_{\alpha,\beta} \right]. \\ \delta_4 &= \Pr \left[ \begin{array}{l} \tilde{\mathbf{v}}_{\alpha[\ell]}^{(0)} = \mathbf{v}_{\alpha[\ell]}^{(0)} \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell-1}, (\mathbf{h}^{(0)}, \tilde{\mathbf{v}}_{\alpha[\ell]}^{(0)}, j)_{j \in [2]}, \tilde{\pi}^{(0)}) = 1 \wedge \\ (x^{(1)}, w^{(1)}) \in \mathcal{L}_{\text{bp},1} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \exp_{\alpha,\beta} \right]. \\ \delta_5 &= \Pr \left[ \begin{array}{l} \tilde{\mathbf{v}}_{\alpha[\ell]}^{(b)} = \mathbf{v}_{\alpha[\ell]}^{(b)} \text{ for } b \in \{0, 1\} \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell-1}, (\mathbf{h}^{(b)}, \tilde{\mathbf{v}}_{\alpha[\ell]}^{(b)}, j)_{j \in [2]}, \tilde{\pi}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} : \exp_{\alpha,\beta} \right]. \end{aligned}$$

We let  $\delta_1 = \delta_{1,\ell}$ , and  $\epsilon_i = \delta_i - \delta_{i+1}$  for  $i \in [4]$  and  $\epsilon_5 = \delta_5$ . Since  $\delta_1 = \sum_{i=1}^5 \epsilon_i$  we only need to show that for  $i \in [5]$ ,  $\epsilon_i$  is negligible.

**Case 1.** Similar to the case 1 in the induction base if  $\epsilon_{b+2}$  is non-negligible then  $\mathcal{B}_{\text{bp}}(\alpha, \beta, \ell, \mathcal{A}, b)$  breaks the security of seBARG.

**Case 2.** Similar to the case 2 in the induction base if  $\epsilon_{b+4}$  is non-negligible then  $\mathcal{B}_{\text{seh}}(\alpha, \beta, \ell, \mathcal{A}, b)$  breaks the security of SEH.

**Case 3.** Notice that  $\epsilon_5 \leq \delta_{1,\ell-1}$ , therefore, by the induction hypothesis,  $\epsilon_5 \leq \text{negl}(\lambda)$ .

**Induction Step (2).** Suppose the claim holds for any  $\ell' < \ell$  (where  $\ell \geq \log T$ ), namely,  $\delta_{1,\ell'} \leq \text{negl}(\lambda)$ . We follow the same approach as induction step (1) except that  $w^{(b)}$  additionally includes  $(\tilde{h}_1^{(b)}, \tilde{h}_2^{(b)})$ . Thus we only need to show that  $\tilde{h}_{\alpha[\ell]}^{(0)} = \tilde{h}_{\alpha[\ell]}^{(1)}$  to use the induction step. We let  $\delta_2$  and  $\delta_3$  to be the same as in induction step (1), and define the rest as follows:

$$\delta_4 = \Pr \left[ \begin{array}{l} \tilde{v}_{\alpha[\ell]}^{(0)} = v_{\alpha[\ell]}^{(0)} \wedge \\ \text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\tilde{h}_1^{(0)}, \tilde{h}_2^{(0)})) = h^{(0)} \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell-1}, (\tilde{h}_{\alpha[\ell]}^{(0)}, \tilde{v}_{\alpha[\ell]}^{(0)}, j)_{j \in [2]}, \tilde{\pi}^{(0)}) = 1 \wedge \\ (x^{(1)}, w^{(1)}) \in \mathcal{L}_{\text{bp},1} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} \quad : \quad \text{exp}_{\alpha,\beta} \right].$$

$$\delta_5 = \Pr \left[ \begin{array}{l} \tilde{v}_{\alpha[\ell]}^{(b)} = v_{\alpha[\ell]}^{(b)} \text{ for } b \in \{0,1\} \wedge \\ \text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\tilde{h}_1^{(b)}, \tilde{h}_2^{(b)})) = h^{(b)} \text{ for } b \in \{0,1\} \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell-1}, (\tilde{h}_{\alpha[\ell]}^{(b)}, \tilde{v}_{\alpha[\ell]}^{(b)}, j)_{j \in [2]}, \tilde{\pi}^{(b)}) = 1 \text{ for } b \in \{0,1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} \quad : \quad \text{exp}_{\alpha,\beta} \right].$$

$$\delta_6 = \Pr \left[ \begin{array}{l} \tilde{v}_{\alpha[\ell]}^{(b)} = v_{\alpha[\ell]}^{(b)} \text{ for } b \in \{0,1\} \wedge \tilde{h}_{\alpha[\ell]}^{(0)} = \tilde{h}_{\alpha[\ell]}^{(1)} \wedge \\ \text{seBARG.Verify}(\text{crs}_{\text{bp},\ell-1}, (\tilde{h}_{\alpha[\ell]}^{(b)}, \tilde{v}_{\alpha[\ell]}^{(b)}, j)_{j \in [2]}, \tilde{\pi}^{(b)}) = 1 \text{ for } b \in \{0,1\} \wedge \\ \text{pcf}_{k,t-1}^{(0)} = \text{pcf}_{k,t-1}^{(1)} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} \end{array} \quad : \quad \text{exp}_{\alpha,\beta} \right].$$

Similar to induction step (1) define  $\delta_1$ , and  $\epsilon_i$  for  $i \in [6]$ .

**Case 1 and 2.** Similar to the case 1 and 2 in the induction step (1).

**Case 3.** It holds that:

$$\Pr \left[ \begin{array}{l} \tilde{h}_{\alpha[\ell]}^{(0)} \neq \tilde{h}_{\alpha[\ell]}^{(1)} \wedge \\ \text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\tilde{h}_1^{(b)}, \tilde{h}_2^{(b)})) = h_\ell \text{ for } b \in \{0,1\} \end{array} \quad : \quad \text{exp}_{\alpha,\beta} \right] \geq \epsilon_5 \quad (9)$$

Suppose towards the contradiction that  $\epsilon_5(\cdot)$  is non-negligible. Then we construct  $\mathcal{B}_{\text{ch}}$  s.t.  $\mathcal{B}_{\text{ch}}(\alpha, \beta, \ell, \mathcal{A})$  breaks the CRHF security. Note that the advantage of  $\mathcal{B}_{\text{ch}}(\alpha, \beta, \ell, \mathcal{A})$  is greater than that of  $\mathcal{A}$  in Eq. (9) contradicting the CRHF security.

$\mathcal{B}_{\text{ch}}(\alpha, \beta, \ell, \mathcal{A})$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{hk}^*$  from the challenger where  $\text{hk}^* \leftarrow \text{CRHF.Gen}(1^\lambda)$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha,\beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{hk}_{\text{ch}}$  with  $\text{hk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .
4. Finds  $\pi_\ell^{(b)} = (v^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, h^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ .
5. Computes  $w^{(b)} = \text{seBARG.Extract}(\text{td}_{\text{bp},\ell}, (h^{(b)}, v^{(b)}, j)_{j \in [2]}, \hat{\pi}^{(b)})$  and finds  $(\tilde{h}_1^{(b)}, \tilde{h}_2^{(b)}) \in w^{(b)}$ .
6. Outputs  $((\tilde{h}_1^{(0)}, \tilde{h}_2^{(0)}), (\tilde{h}_1^{(1)}, \tilde{h}_2^{(1)}))$ .

**Case 4.** Notice that  $\epsilon_6 \leq \delta_{1,\ell-1}$ , therefore, by the induction hypothesis,  $\epsilon_6 \leq \text{negl}(\lambda)$ . □

**Claim A.4.** If seBARG is index hiding,  $\alpha$  and  $\beta$  correspond to  $(k, t - 1, t)$ ,  $\beta'$  corresponds to  $(k, t, t + 1)$ , and for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta}] \leq \text{negl}(\lambda),$$

then for any such adversary, it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta'}] \leq \text{negl}(\lambda).$$

*Proof.* There is some index  $j$  s.t.  $\beta = (2, \dots, 2, 1, \beta[j + 1], \dots, \beta[\lambda])$  and  $\beta' = (1, \dots, 1, 2, \beta[j + 1], \dots, \beta[\lambda])$ . For  $d \in [0, j]$  define  $\beta_d$  as  $\beta_d[\ell] = \beta'[\ell]$  for  $\ell \in [d]$  and  $\beta_d[\ell] = \beta[\ell]$  at every other index. Therefore  $\beta_0 = \beta$  (correspond to  $(k, t - 1, t)$ ) and  $\beta_j = \beta'$  (corresponds to  $(k, t, t + 1)$ ). Note that every  $\beta_{d-1}$  and  $\beta_d$  differ only on index  $d$ .

If the claim doesn't hold, then there is an adversary  $\mathcal{A}$  s.t. for some  $d \in [j]$  there is a non-negligible function  $\epsilon(\cdot)$  s.t.:

$$\Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta_d}] - \Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta_{d-1}}] \geq \epsilon(\lambda) \quad (10)$$

Now we construct  $\mathcal{B}_{\text{bp}}^{\text{ih}}$  s.t.  $\mathcal{B}_{\text{seh}}^{\text{ih}}(\alpha, \beta_d, d, \mathcal{A})$  breaks the index-hiding property of seBARG on two indices  $e_1 = \beta_d[d]$  and  $e_2 = \beta_{d-1}[d]$ .

$\mathcal{B}_{\text{bp}}^{\text{ih}}(\alpha, \beta, \ell, \mathcal{A})$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{crs}^*$  from the challenger where  $(\text{crs}^*, \text{td}^*) \leftarrow \text{seBARG.Gen}(1^\lambda, 2, n_{\text{bp},\ell-1}, e_b)$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha,\beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{crs}_{\text{bp},\ell}$  with  $\text{crs}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .
4. If for any  $b \in \{0, 1\}$ ,  $\text{Verify}(\text{crs}, (\text{hz}, \text{out}^{(b)}), \Pi^{(b)}) = 0$  or  $\text{out}^{(0)} = \text{out}^{(1)}$  then abort.
5. Finds  $\pi_\ell^{(b)} = (v^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ .
6. Compute  $(\text{epcf}_{k,t-1}^{(b)}, \text{epcf}_{k,t}^{(b)}) \leftarrow \text{Extract}_{\text{cf}}(\text{td}, \Pi^{(b)})$  for  $b \in \{0, 1\}$ .
7. Compute  $\text{pcf}_{k,t}^{(b)} = \text{PKE.Dec}(\text{sk}, \text{epcf}_{k,t}^{(b)})$  for  $b \in \{0, 1\}$ .
8. Output 1 if  $\text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)}$  and 2 otherwise.

In the reduction above  $b = 1$  corresponds to  $\exp_{\alpha,\beta_d}$  and  $b = 2$  corresponds to  $\exp_{\alpha,\beta_{d-1}}$ . Therefore,  $\mathcal{B}_{\text{bp}}^{\text{ih}}(\alpha, \beta_d, d, \mathcal{A})$ 's distinguishing advantage is greater than that of  $\mathcal{A}$  in Eq. (10) contradicting the seBARG security. □

**Claim A.5.** If seBARG, SEH, HT, and CRHF are secure,  $\alpha$  corresponds to  $(k, t - 1, t)$ ,  $\alpha'$  and  $\beta'$  correspond to  $(k, t, t + 1)$ , and for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta'}] \leq \text{negl}(\lambda),$$

then for any such adversary, it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha',\beta'}] \leq \text{negl}(\lambda).$$

*Proof.* Consider the following probabilities:

$$\delta_0 = \Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta'}]$$

$$\delta_1 = \Pr[\text{Bad} \wedge \text{pcf}_l^{(0)} \neq \text{pcf}_l^{(1)} : \exp_{\alpha,\beta'}]$$

$$\delta_2 = \Pr[\text{Bad} \wedge \text{pcf}_r^{(0)} \neq \text{pcf}_r^{(1)} : \exp_{\alpha,\beta'}]$$

$$\delta_3 = \Pr[\text{Bad} \wedge \text{pcf}_r^{(0)} \neq \text{pcf}_r^{(1)} : \exp_{\alpha',\beta'}]$$

$$\delta_4 = \Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(0)} \neq \text{pcf}_{k,t}^{(1)} : \exp_{\alpha',\beta'}]$$

We let  $\epsilon_i = \delta_i - \delta_{i-1}$  for  $i \in [4]$ . Since  $\delta_4 = \delta_0 + \sum_{i=1}^4 \epsilon_i$ , and by the claim assumption  $\delta_0 \leq \text{negl}(\lambda)$ , we only need to show that for  $i \in [4]$ ,  $\epsilon_i$  is negligible.

**Case 1.** It holds that:

$$\epsilon_1 \leq \Pr[\text{Bad} \wedge \text{pcf}_l^{(0)} \neq \text{pcf}_l^{(1)} \wedge \text{pcf}_{k,t}^{(0)} = \text{pcf}_{k,t}^{(1)} : \exp_{\alpha,\beta'}]$$

Thus for some  $b \in \{0, 1\}$  it holds that:

$$\epsilon_1/2 \leq \Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(b)} \neq \text{pcf}_l^{(b)} : \exp_{\alpha,\beta'}]$$

Now two cases could happen:

- $\text{pcf}_l^{(b)}$  is computed using some  $\rho_2^{(b)} \in \pi_\ell^{(b)} \in \Pi^{(b)}$ . In this case by the verifier's check it holds that  $\text{SEH.VerifyAcc}((\text{hk}_{\text{cf},d})_{d \in [\ell]}, \mathbf{v}^{(b)}, \rho_2^{(b)}, 2) = 1$  where  $\mathbf{v}^{(b)} \in \pi_\ell^{(b)}$ . Now since  $\alpha$  corresponds to  $(k, t - 1, t)$ , it is binding on the path  $\rho_2^{(b)}$ , thus by the SEH security and PKE correctness it holds that  $\Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(b)} \neq \text{pcf}_l^{(b)} : \exp_{\alpha,\beta'}] \leq \text{negl}(\lambda)$ .
- Find  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$  that contains a proof for  $(k, t - 1, t)$ . Note that  $j < \ell$  thus  $\alpha$  and  $\beta$  both are binding on the same path from a level  $\ell$  proof until a level  $j$  proof. Additionally recall that to find  $\text{pcf}_l^{(b)}$ , seBARG extraction is recursively used to find witness  $w_j^{(b)}$ . Then  $w_j^{(b)}$  is parsed to get  $(\tilde{\mathbf{v}}_{j-1,1}^{(b)}, \tilde{\rho}_{j-1,1}^{(b)}) \in w_j^{(b)}$  and  $\text{pcf}_l^{(b)}$  is found using them. Let  $\mathbf{v}_d^{(b)}$  be the hash of pseudo-configurations at level  $d$  that is extracted by recursively running SEH extraction starting from  $\mathbf{v}^{(b)}$ . Now two cases can happen:

- $\tilde{v}_{j-1}^{(b)} \neq v_{j-1}^{(b)}$ . This contradicts the seBARG and SEH security since both are binding on the same path from  $\ell$  to  $j$ .
- $\tilde{v}_{j-1}^{(b)} = v_{j-1}^{(b)} \wedge \text{pcf}_{k,t}^{(b)} \neq \text{pcf}_l^{(b)}$ . This contradicts either the somewhere statistically binding of SEH on the path from  $v_{j-1}^{(b)}$  to  $\text{epcf}_l^{(b)}$  or the PKE correctness.

**Case 2.** Similar to case 1, for some  $b \in \{0, 1\}$  it holds that:

$$\epsilon_2/2 \leq \Pr[\text{Bad} \wedge \text{pcf}_r^{(b)} \neq \text{pcf}_l^{(b)} : \exp_{\alpha, \beta'}]$$

Two cases could happen:

- $\text{pcf}_l^{(b)}$  (resp.  $\text{pcf}_r^{(b)}$ ) is computed using some  $\rho_2^{(b)} \in \pi_\ell^{(b)} \in \Pi^{(b)}$  (resp.  $\rho_1^{(b)} \in \pi_{\ell'}^{(b)} \in \Pi^{(b)}$ ) where  $\ell > \ell'$ . In this case by the verifier's check, the leftmost encrypted pseudo-configuration in  $\pi_{\ell'}^{(b)}$  is the same as the rightmost encrypted pseudo-configuration in  $\pi_\ell^{(b)}$ , namely,  $\text{epcf}_l^{(b)} = \text{epcf}_r^{(b)}$ , thus by PKE correctness we have  $\text{pcf}_l^{(b)} = \text{pcf}_r^{(b)}$ .
- Otherwise, recall that for both  $\text{pcf}_l^{(b)}$  and  $\text{pcf}_r^{(b)}$ , first seBARG extractor is recursively being using to find  $w_j^{(b)}$  for a batch proof at level  $j$ . Let  $(\tilde{v}_{j-1,1}^{(b)}, \tilde{v}_{j-1,2}^{(b)}) \in w_j^{(b)}$ , then by the definition of  $\mathcal{L}_{\text{bp},j}$ , validity of  $w_j^{(b)}$  (which is implied by the extraction correctness of batch argument) implies that the leftmost encrypted pseudo-configuration in the tree with root  $\tilde{v}_{j-1,2}^{(b)}$  is equal to the rightmost encrypted pseudo-configuration in  $\tilde{v}_{j-1,1}^{(b)}$ , namely,  $\text{epcf}_l^{(b)} = \text{epcf}_r^{(b)}$ , thus by PKE correctness we have  $\text{pcf}_l^{(b)} = \text{pcf}_r^{(b)}$ .

Therefore in both cases it holds that  $\Pr[\text{Bad} \wedge \text{pcf}_r^{(b)} \neq \text{pcf}_l^{(b)} : \exp_{\alpha, \beta'}] \leq \text{negl}(\lambda)$ .

**Case 3.** Similar to the proof of Claim A.4 we prove  $\epsilon_3 \leq \text{negl}(\lambda)$ . There is some index  $j$  s.t.  $\alpha = (2, \dots, 2, 1, \alpha[j+1], \dots, \alpha[\lambda])$  and  $\alpha' = (1, \dots, 1, 2, \alpha[j+1], \dots, \alpha[\lambda])$ . For  $d \in [0, j]$  define  $\alpha_d$  as  $\alpha_d[\ell] = \alpha'[\ell]$  for  $\ell \in [d]$  and  $\alpha_d[\ell] = \alpha[\ell]$  at every other index. Therefore  $\alpha_0 = \alpha$  (correspond to  $(k, t-1, t)$ ) and  $\alpha_j = \alpha'$  (corresponds to  $(k, t, t+1)$ ). Note that every  $\alpha_{d-1}$  and  $\alpha_d$  differ only on index  $d$ . There is an adversary  $\mathcal{A}$  s.t. for some  $d \in [j]$  it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_r^{(0)} \neq \text{pcf}_r^{(1)} : \exp_{\alpha_d, \beta'}] - \Pr[\text{Bad} \wedge \text{pcf}_r^{(0)} \neq \text{pcf}_r^{(1)} : \exp_{\alpha_{d-1}, \beta'}] \geq \epsilon_3(\lambda)/\lambda \quad (11)$$

Now we construct  $\mathcal{B}_{\text{seh}}^{\text{ih}}$  s.t.  $\mathcal{B}_{\text{seh}}^{\text{ih}}(\alpha_d, \beta', d, \mathcal{A})$  breaks the index-hiding property of SEH on two indices  $e_1 = \alpha_d[d]$  and  $e_2 = \alpha_{d-1}[d]$ .

$\mathcal{B}_{\text{seh}}^{\text{ih}}(\alpha, \beta, \ell, \mathcal{A})$  This PPT attacker does the following:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$  and  $\text{hk}^*$  from the challenger where  $(\text{hk}^*, \text{td}^*) \leftarrow \text{seBARG.Gen}(1^\lambda, 2, \Sigma_{\text{cf}}^{\ell-1}, e_b)$ .
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}_{\alpha, \beta}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{hk}_{\text{cf}, \ell}$  with  $\text{hk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$  and receives  $(\text{hz}, (\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}), 1^{N \cdot T})$ .



4. If for any  $b \in \{0, 1\}$ ,  $\text{Verify}(\text{crs}, (\text{hz}, \text{out}^{(b)}), \Pi^{(b)}) = 0$  or  $\text{out}^{(0)} = \text{out}^{(1)}$  then abort.
5. Finds  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$ .
6. Compute  $\text{epcf}_r^{(b)} \leftarrow \text{Extract}_r(\text{td}, \Pi^{(b)})$  and  $\text{pcf}_r^{(b)} = \text{PKE.Dec}(\text{sk}, \text{epcf}_r^{(b)})$  for  $b \in \{0, 1\}$ .
7. Output 1 if  $\text{pcf}_r^{(0)} \neq \text{pcf}_r^{(1)}$  and 2 otherwise.

In the reduction above  $b = 1$  corresponds to  $\text{exp}_{\alpha, \beta_d}$  and  $b = 2$  corresponds to  $\text{exp}_{\alpha, \beta_{d-1}}$ . Therefore,  $\mathcal{B}_{\text{seh}}^{\text{ih}}(\alpha_d, \beta', d, \mathcal{A})$ 's distinguishing advantage is greater than that of  $\mathcal{A}$  in Eq. (11) contradicting the SEH security.

**Case 4.** Similar to case 1, for some  $b \in \{0, 1\}$  it holds that:

$$\epsilon_4/2 \leq \Pr[\text{Bad} \wedge \text{pcf}_{k,t}^{(b)} \neq \text{pcf}_r^{(b)} : \text{exp}_{\alpha', \beta'}]$$

Note that  $\alpha'$  and  $\beta'$  are the same path, thus the proof is similar to case one.  $\square$

**Claim A.6.** If SEH is somewhere binding w.r.t. path opening, HT has reading soundness, PKE is correct, and  $\alpha$  and  $\beta$  correspond to  $(K, T-1, T)$ , and for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$\Pr[\text{Bad} \wedge \text{pcf}_{K,T}^{(0)} \neq \text{pcf}_{K,T}^{(1)} : \text{exp}_{\alpha, \beta}] \leq \text{negl}(\lambda),$$

then for any such adversary, it holds that:

$$\Pr[\text{Bad} : \text{exp}_{\alpha, \beta}] \leq \text{negl}(\lambda).$$

*Proof.* It suffices to show that for any PPT adversary  $\mathcal{A}$  it holds that:

$$\delta_0 = \Pr[\text{Bad} \wedge \text{pcf}_{K,T}^{(0)} = \text{pcf}_{K,T}^{(1)} : \text{exp}_{\alpha, \beta}] \leq \text{negl}(\lambda). \quad (12)$$

Let  $\pi_\ell^{(b)} = (\mathbf{v}^{(b)}, \rho_1^{(b)}, \rho_2^{(b)}, \mathbf{h}^{(b)}, \hat{\pi}^{(b)}) \in \Pi^{(b)}$  be the proof that contains  $(\text{pcf}_{K, T-1}^{(b)}, \text{pcf}_{K, T}^{(b)})$  and For  $b \in \{0, 1\}$  parse  $\rho_2^{(b)} = ((\mathbf{v}_{0,1}^{(b)}, \mathbf{v}_{0,2}^{(b)}), \dots, (\mathbf{v}_{\ell-1,1}^{(b)}, \mathbf{v}_{\ell-1,2}^{(b)}))$ , and  $\mathbf{v}_{0,2}^{(b)} = (\text{epcf}_1^{(b)}, \text{epcf}_2^{(b)})$ , and  $\pi_{\text{out}}^{(b)} = (\text{pcf}_2^{(b)}, r^{(b)}, \hat{\pi}_{\text{out}}^{(b)})$ , and  $\text{pcf}_2^{(b)} = (\text{rt}_2^{(b)}, \text{q}_2^{(b)}, \text{ridx}_2^{(b)})$ . Now, define the following probabilities:

$$\delta_1 = \Pr \left[ \begin{array}{l} \text{HT.VfyRd}(\text{hk}_{\text{ht}}, \text{rt}_2^{(b)}, [n_{\text{out}}], \text{out}^{(b)}, \hat{\pi}_{\text{out}}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ \text{PKE.Enc}(\text{pk}, \text{pcf}_2^{(b)}; r^{(b)}) = \text{epcf}_2^{(b)} \text{ for } b \in \{0, 1\} \wedge \\ \text{SEH.VerifyAcc}((\text{hk}_{\text{cf}, d})_{d \in [\ell]}, \mathbf{v}^{(b)}, \rho_2^{(b)}, 2) \text{ for } b \in \{0, 1\} \wedge \\ \text{out}^{(0)} \neq \text{out}^{(1)} \wedge \text{pcf}_{K,T}^{(0)} = \text{pcf}_{K,T}^{(1)} \end{array} : \text{exp}_{\alpha, \beta} \right].$$

$$\delta_2 = \Pr \left[ \begin{array}{l} \text{HT.VfyRd}(\text{hk}_{\text{ht}}, \text{rt}_2^{(b)}, [n_{\text{out}}], \text{out}^{(b)}, \hat{\pi}_{\text{out}}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ \text{PKE.Enc}(\text{pk}, \text{pcf}_2^{(b)}; r^{(b)}) = \text{epcf}_2^{(b)} \text{ for } b \in \{0, 1\} \wedge \\ \text{epcf}_{K,T}^{(0)} = \text{epcf}_2^{(0)} \wedge \\ \text{SEH.VerifyAcc}((\text{hk}_{\text{cf}, d})_{d \in [\ell]}, \mathbf{v}^{(1)}, \rho_2^{(1)}, 2) \wedge \\ \text{out}^{(0)} \neq \text{out}^{(1)} \wedge \text{pcf}_{K,T}^{(0)} = \text{pcf}_{K,T}^{(1)} \end{array} : \text{exp}_{\alpha, \beta} \right].$$

$$\delta_3 = \Pr \left[ \begin{array}{l} \text{HT.VfyRd}(\text{hk}_{\text{ht}}, \text{rt}_2^{(b)}, [n_{\text{out}}], \text{out}^{(b)}, \hat{\pi}_{\text{out}}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ \text{PKE.Enc}(\text{pk}, \text{pcf}_2^{(b)}; r^{(b)}) = \text{epcf}_2^{(b)} \text{ for } b \in \{0, 1\} \wedge \\ \text{epcf}_{K,T}^{(b)} = \text{epcf}_2^{(b)} \text{ for } b \in \{0, 1\} \wedge \\ \text{out}^{(0)} \neq \text{out}^{(1)} \wedge \text{pcf}_{K,T}^{(0)} = \text{pcf}_{K,T}^{(1)} \end{array} \quad : \quad \exp_{\alpha, \beta} \right].$$

$$\delta_4 = \Pr \left[ \begin{array}{l} \text{HT.VfyRd}(\text{hk}_{\text{ht}}, \text{rt}_2^{(b)}, [n_{\text{out}}], \text{out}^{(b)}, \hat{\pi}_{\text{out}}^{(b)}) = 1 \text{ for } b \in \{0, 1\} \wedge \\ \text{pcf}_{K,T}^{(b)} = \text{pcf}_2^{(b)} \text{ for } b \in \{0, 1\} \wedge \\ \text{out}^{(0)} \neq \text{out}^{(1)} \wedge \text{pcf}_{K,T}^{(0)} = \text{pcf}_{K,T}^{(1)} \end{array} \quad : \quad \exp_{\alpha, \beta} \right].$$

First note that  $\delta_0 \leq \delta_1$  since all the events are implied by the verification of the proofs. Also,  $\delta_3 = \delta_4$  by the perfect correctness of PKE. Moreover, the advantage of an adversary  $\mathcal{B}$  attacking the reading soundness of HT by outputting  $(\text{rt}_2^{(0)}, [n_{\text{out}}], \text{out}^{(0)}, \pi_{\text{out}}^{(0)}, \text{out}^{(1)}, \pi_{\text{out}}^{(1)})$  when  $\mathcal{A}$  generates  $((\text{out}^{(0)}, \Pi^{(0)}), (\text{out}^{(1)}, \Pi^{(1)}))$  is at least  $\delta_4$ , thus  $\delta_4 \leq \text{negl}(\lambda)$ . Therefore, we only need to show that  $\delta_{b+1} - \delta_{b+2} \leq \text{negl}(\lambda)$  for  $b \in \{0, 1\}$ . It holds that:

$$\delta_{b+1} - \delta_{b+2} \leq \epsilon_b = \Pr \left[ \text{SEH.VerifyAcc}((\text{hk}_{\text{cf},d})_{d \in [\ell]}, \mathbf{v}^{(b)}, \rho_2^{(b)}, 2) \wedge \text{pcf}_{K,T}^{(b)} = \text{pcf}_2^{(b)} \quad : \quad \exp_{\alpha, \beta} \right].$$

Now by somewhere binding w.r.t. path opening of SEH it holds that  $\epsilon_b \leq \text{negl}(\lambda)$  concluding the proof. □

□

□

**Zero-Knowledge.** We construct the simulator  $\mathcal{S}(\text{crs}, \text{cf})$  as follows:

- It generates  $\text{crs}$  similar to  $\text{Gen}$  except that it runs the NIZK simulator to generate  $\text{crs}_{\text{zk}}$ .
- To generate the proof it does the following:
  1. Compute  $\text{epcf}_{k,t}$  for  $k \in [K], t \in [T]$  as follows: (1) let  $\text{epcf}_{1,0} = \text{PKE.Enc}(\text{pk}, \text{pcf}_{1,0}; 0^{|r|})$ , (2) let  $\text{cf} = (W, q, \text{ridx})$  and compute  $(\text{rt}, \text{tree}) = \text{HT.Hash}(\text{hk}_{\text{ht}}, W)$ , and  $\text{pcf}_{K,T} = (\text{rt}, q, \text{ridx})$ , then sample randomness  $r$  and compute  $\text{epcf}_{K,T} = \text{PKE.Enc}(\text{pk}, \text{pcf}_{K,T})$ , (3) for every other  $(k, t)$  independently (under different randomness) compute  $\text{epcf}_{k,t} = \text{PKE.Enc}(\text{pk}, 0^{|\text{pcf}|})$ .
  2. Let  $\text{st} = (\text{cf}, \text{pcf}_{K,T}, r)$ .
  3. Generate  $\text{ct}_{\text{in},k}$  for  $k \in [K]$  by independently encrypting  $0^\lambda$ .
  4. To generate  $\text{hz}$  use the same iterative process as in [Digest](#) except that in [Item 2a](#) instead of computing  $\text{ct}_{\text{in}}$  use  $\text{ct}_{\text{in},k}$  computed above.
  5. To generate the proof use the same iterative process as in [Update](#) except that do not compute  $(\text{pcf}_{t-1}, \text{pcf}_t, \text{rbit}, \text{rop}, \text{wop})$ , and to generate  $\pi_{\text{zk}}$  use NIZK simulator. Also, never compute  $\pi_{\text{out}}$  except when  $(k, t) = (K, T)$  in which case compute  $\pi_{\text{out}}$  similar to the prover.

We define the following hybrids.

- $\mathcal{H}_0$ : This is the original experiment. Namely, the adversary sends  $(n, S, T)$ , challenger sends  $\text{crs}_0 \leftarrow \text{Gen}(1^\lambda, n, S, T)$ , adversary sends  $\bar{z}$ , challenger computes  $\text{cf} = \mathcal{M}(\bar{z}; 1^T)$  and sends  $(\text{st}_0, \text{hz}_0, \Pi_0) \leftarrow \text{Prove}(\text{crs}_0, \bar{z})$  to the adversary and finally adversary outputs  $b'$ .
- $\mathcal{H}_1$ : Similar to the original experiment except that NIZK simulator is used to generate  $\text{crs}_{\text{zk}}$ , and  $\pi_{\text{zk}}$  at every iteration.
- $\mathcal{H}_{2,k,t}$ : Similar to  $\mathcal{H}_1$  except that for every  $\text{cf}_{i,j}$  until  $\text{cf}_{k,t}$  except  $\text{cf}_{1,0}$  and  $\text{cf}_{K,T}$ , we let  $\text{epcf}_{i,j} = \text{PKE.Enc}(\text{pk}, 0^{|\text{pcf}|})$  (and for the rest  $\text{epcf}_{i,j}$  is honestly generated). Note that  $\mathcal{H}_{2,1,0} = \mathcal{H}_1$  and  $\mathcal{H}_{2,K,T-1} = \mathcal{H}_{2,K,T}$ .
- $\mathcal{H}_{3,k}$ : Similar to  $\mathcal{H}_{2,K,T}$  except that for every  $z_i$  until  $z_k$ , we let  $\text{ct}_{\text{in},i} = \text{PKE.Enc}(\text{pk}, 0^\lambda)$  (and for the rest  $\text{ct}_{\text{in},i}$  is honestly generated). Note that  $\mathcal{H}_{3,0} = \mathcal{H}_{2,K,T}$  and  $\mathcal{H}_{3,K}$  is the output of the simulator. Namely in  $\mathcal{H}_{3,K}$  the adversary sends  $(n, S, T)$ , the simulator sends  $\text{crs}_1 \leftarrow \mathcal{S}(1^\lambda, n, S, T)$ , adversary sends  $\bar{z}$ , simulator computes  $\text{cf} = \mathcal{M}(\bar{z}; 1^T)$  and sends  $(\text{st}_1, \text{hz}_1, \Pi_1) \leftarrow \mathcal{S}(\text{crs}_0, \text{cf})$  to the adversary and finally adversary outputs  $b'$ .

For any adversary let  $\mathcal{H}^*(\mathcal{A})$  be the output of the adversary in hybrid  $\mathcal{H}^*$ .

**Claim A.7.** If NIZK is zero-knowledge then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_1(\mathcal{A}) = 1] - \Pr[\mathcal{H}_0(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Towards the contradictions suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  s.t.  $|\Pr[\mathcal{H}_1(\mathcal{A}) = 1] - \Pr[\mathcal{H}_0(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ . We construct adversary  $\mathcal{B}_{\text{zk}}$  that breaks the zero-knowledge property of the underlying NIZK protocol as follows:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$ , and  $\text{crs}^*$  from NIZK zero-knowledge challenger.
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{crs}_{\text{zk}}$  with  $\text{crs}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$ , receives  $\bar{z}$ .
4. Computes  $(\text{st}, \text{hz}, \Pi)$  similar to the prover except that at every single step, instead of computing  $\pi_{\text{zk}}$ , queries the challenger on  $(x, w)$  and receive  $\pi_{\text{zk}}$ .
5. Outputs whatever  $\mathcal{A}$  outputs.

If the NIZK challenger generates  $\text{crs}$  and proof using  $\text{Gen}$  and  $\text{Prove}$ , then  $\mathcal{B}_{\text{zk}}$  simulates  $\mathcal{H}_0$  for  $\mathcal{A}$  and if the challenger uses NIZK simulator, then  $\mathcal{B}_{\text{zk}}$  simulates  $\mathcal{H}_1$  for  $\mathcal{A}$ . Thus  $\mathcal{B}_{\text{zk}}$  has a non-negligible advantage  $\epsilon(\lambda)$  in breaking the zero-knowledge property of NIZK.  $\square$

In the following claim we let  $k \in [K]$ ,  $t \in [T]$ , and  $(k, t) \neq (K, T)$ .

**Claim A.8.** If PKE is secure then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_{2,k,t}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{2,k,t-1}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Towards the contradictions suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  s.t.  $|\Pr[\mathcal{H}_{2,k,t}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{2,k,t-1}(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ . We construct adversary  $\mathcal{B}_{\text{pke}}$  that breaks the security of the underlying PKE protocol as follows:

1. Receives  $(n, S, T)$  from  $\mathcal{A}$ , and  $\text{pk}^*$  from PKE challenger.
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}(1^\lambda, n, S, T)$  and updates  $\text{crs}$  by replacing  $\text{pk}$  with  $\text{pk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$ , receives  $\bar{z}$ .
4. Computes  $(\text{st}, \text{hz}, \Pi)$  similar to  $\mathcal{H}_{2,k,t-1}$  except that instead of computing  $\text{epcf}_{k,t}$ , queries the challenger on  $\text{pcf}_{k,t}$  and receive  $\text{epcf}_{k,t}$ .
5. Outputs whatever  $\mathcal{A}$  outputs.

If the PKE challenger generates the ciphertext honestly, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{2,k,t-1}$  for  $\mathcal{A}$  and if it responds with a ciphertext of all zeros, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{2,k,t}$  for  $\mathcal{A}$ . Thus  $\mathcal{B}_{\text{pke}}$  has a non-negligible advantage  $\epsilon(\lambda)$  in breaking the security of PKE.  $\square$

In the following claim we let  $k \in [K]$ .

**Claim A.9.** If PKE is secure then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_{3,k}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{3,k-1}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Similar to the proof of Claim A.8.  $\square$

Therefore we conclude the proof of zero-knowledge property by a hybrid argument.

## B Predicated zk-IVsC Complete Analysis.

*Proof of Theorem 3.11.* The completeness, and strong soundness proofs are similar to zk-IVsC. We only need to prove efficiency, predicate proof soundness, and zero-knowledge properties of our construction.

**Efficiency.** The  $\text{crs}$  size and Setup running time in addition to the other elements (which is similar to IVsC) grows with  $\text{crs}_{\text{zk},f}$  size and generation time which grow with  $\text{poly}(\lambda, n_{\text{zk},f,x}, n_{\text{zk},f,w}) = \text{poly}(\lambda, n, n_f)$ , thus in total grows with  $\text{poly}(\lambda, \log S, n, n_f)$ .

Now note the size of  $\pi_{f, \log T-1}$  and its generation/verification time grows with NIZK statement and witness and the the circuit that decides the language, thus all three grow with  $\text{poly}(\lambda, n, n_f)$ . Additionally, similar to the IVsC analysis the proof size and its generation/verification (given the tree decomposition of  $\bar{f}_k$ ) grows with  $\text{poly}(\lambda, n, n_f, \log K)$ . Finally tree decomposition of  $\bar{f}_k$  grows at most with  $\text{poly}(K)$ .

**Predicate Proof Soundness.** First note that the crs indistinguishability of Gen and TGen follows directly from the crs indistinguishability of the seBARG, and crs indistinguishability of NIZK extractor. Next we construct algorithms TGen, HashCheck and the extractor  $\mathcal{E}$ .

The  $\text{TGen}(1^\lambda, n, S, T, n_f, k^*) \rightarrow \text{crs}$  algorithm is the same as Gen except that  $\gamma$  corresponds to  $k^*$ , and NIZK extractor  $\mathcal{E}$  is used to generate  $\text{crs}_{\text{zk},f}$ . It computes  $\text{td} = ((\text{crs}_{\text{bp},f,\ell})_{\ell \in [\log T, \lambda]}, \text{td}_{\text{zk},f}, k^*)$ .

The  $\text{HashCheck}(\text{crs}, \text{hz}, z^*, k^*, \text{op}) \rightarrow 0/1$  algorithm does the following:

1. Let  $\text{h}_j \in \text{hz}$  be a hash that includes  $k^{*\text{th}}$  leaf and  $(b_{j-1}, \dots, b_{\log T-1})$  (where  $b_i \in \{0, 1\}$ ) be the path from  $\text{h}_j$  to the  $k^{*\text{th}}$  leaf. Then parse opening as  $\text{op} = (r, y_{\log T-1,1}, y_{\log T-1,2}, \dots, y_{j-1,1}, y_{j-1,2})$ .
2. For  $i \in [\log T, j-1]$  check  $\text{CRHF.Hash}(\text{hk}_{\text{ch}}, (y_{i-1,1}, y_{i-1,2})) = y_{i,b_i}$ .
3. If  $j \geq \log T$  check  $\text{CRHF.Hash}(\text{hk}_{\text{ch}}, (y_{j-1,1}, y_{j-1,2})) = \text{h}_j$  otherwise let  $y_{\log T-1, b_{\log T-1}} = \text{h}_j$ .
4. Compute  $(\text{rt}, \text{tree}) = \text{HT.Hash}(\text{hk}_{\text{ht}}, z^*)$  and check  $y_{\log T-1, b_{\log T-1}} = \text{PKE.Enc}(\text{pk}, \text{rt}; r)$ .
5. Output 1 if all the checks pass and 0 otherwise.

The extractor  $\mathcal{E}$  does the following:

1. Find a proof  $\pi_{f,\ell}$  s.t. it contains a proof for  $z_{k^*}$  and let  $\pi_{f,\ell}^* = \pi_{f,\ell}$ .
2. Let  $\text{hz}_{k-1} = \{\text{h}_{\log T-1}, \dots, \text{h}_\lambda\}$  and  $\text{h}_{\ell, \gamma[\ell+1]}^* = \text{h}_\ell$ .
3. Compute  $\bar{f}_{\text{dc}, \ell - \log T + 1}$  using  $\text{TreeDecompose}(\bar{f})$  and let  $\bar{f}_{\ell, \gamma[\ell+1]}^* = \bar{f}_{\text{dc}, \ell - \log T + 1}$ .
4. Recursively extract from the proof  $\pi_{f,\ell}$ . Namely for  $i \in [\log T, \ell]$  Compute  $(\text{h}_{i-1,1}^*, \text{h}_{i-1,2}^*, \pi_{f,i-1}^*) = \text{seBARG.}\mathcal{E}(\text{td}_{\text{bp},f,i}, (\bar{f}_{i,\gamma[i+1]}^* = (\bar{f}_{i-1,1}^*, \bar{f}_{i-1,2}^*), \text{h}_{i,\gamma[i+1]}^*, j)_{j \in [2]}, \pi_{f,i}^*)$ .
5. Compute  $(z^*, \text{h}^*, r^*) = \text{NIZK.}\mathcal{E}(\text{td}_{\text{zk},f}, (\bar{f}_{\log T-1, \gamma[\log T]}^*, \text{h}_{\log T-1, \gamma[\log T]}^*), \pi_{f, \log T-1}^*)$ .
6. Compute  $\text{op} = (r^*, \text{h}_{\log T-1,1}^*, \text{h}_{\log T-1,2}^*, \dots, \text{h}_{\ell-1,1}^*, \text{h}_{\ell-1,2}^*)$  and output  $(z^*, \text{op})$ .

Additionally let Extract be an algorithm that performs only one extraction (either BARG or NIZK depending on  $\ell$ ) from proof  $\pi_{f,\ell}$ .

For any stateful PPT adversary  $\mathcal{A}$  and any  $\lambda \in \mathbb{N}$  define the experiment  $\text{exp}_{\alpha,\beta}$  (where  $\alpha$  corresponds to  $(k, t-1, t)$ ) as follows:

$$\text{exp} := \left[ \begin{array}{l} n, S, T, n_f, k^* \leftarrow \mathcal{A}_1(1^\lambda), \\ (\text{crs}, \text{td}) \leftarrow \text{TGen}(1^\lambda, n, S, T, n_f, k^*), \\ (\bar{f}, \text{hz}, \Sigma) \leftarrow \mathcal{A}_2(\text{crs}), \\ (\bar{f}_{\text{dc},0}, \dots, \bar{f}_{\text{dc},\lambda}) = \text{TreeDecompose}(\bar{f}) \\ (z^*, \text{op}) \leftarrow \mathcal{E}(\text{td}, \bar{f}, \text{hz}, \Sigma) \\ w \leftarrow \text{Extract}(\text{td}, \bar{f}, \text{hz}, \Sigma) \end{array} \right]$$

Moreover, let Bad be the following event:

$$\text{Bad} := \left[ \begin{array}{l} \text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}, \Sigma) = 1 \wedge \\ (f_{k^*}(z^*) = 0 \vee \text{HashCheck}(\text{crs}, \text{hz}, z^*, k^*, \text{op}) = 0) \end{array} \right].$$

To prove the claim, we need to show that for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$\delta_0 = \Pr[\text{Bad} : \text{exp}] \leq \text{negl}(\lambda).$$

We prove this by induction on the depth  $\ell$  of the proof  $\pi_{f,\ell}$  that contains a proof for  $z_{k^*}$ . First note that by the construction of the  $\mathcal{E}$  it holds that for any  $\mathbf{h}_\ell \in \text{hz}$  that includes  $k^*$ th leaf,  $(\gamma_\ell, \dots, \gamma_{\log T})$  is a path from  $\mathbf{h}_\ell$  to  $k^*$ th leaf. Thus we just need to check the CRHF hash correctness on the path to  $k^*$ th leaf and the encryption + hash correctness on the base level. Given  $\text{op} = (r^*, \mathbf{h}_{\log T-1,1}^*, \mathbf{h}_{\log T-1,2}^*, \dots, \mathbf{h}_{\ell-1,1}^*, \mathbf{h}_{\ell-1,2}^*)$  Let  $\text{op}_j = (r^*, \mathbf{h}_{\log T-1,1}^*, \mathbf{h}_{\log T-1,2}^*, \dots, \mathbf{h}_{j-1,1}^*, \mathbf{h}_{j-1,2}^*)$  and  $\text{HashCheck}'(\text{crs}, \mathbf{h}_j, z^*, \text{op}_j)$  partially run  $\text{HashCheck}$  as follows: 1) check the CRHF hash correctness in the  $\text{op}_j$  on path  $(\gamma_j, \dots, \gamma_{\log T})$  from root  $\mathbf{h}_j$ , 2) check the base level encryption + hash correctness.

Define  $\text{Ver}_\ell$  be the following event:

- If  $\ell \geq \log T$  then

$$\text{Ver}_\ell := \left[ \begin{array}{l} \text{seBARG.Verify}(\text{crs}_{\text{bp},f,\ell}, (\bar{f}_{\ell,\gamma[\ell+1]}^*, \mathbf{h}_\ell, j)_{j \in [2]}, \pi_{f,\ell}) = 1 \wedge \\ (f_{k^*}(z^*) = 0 \vee \text{HashCheck}'(\text{crs}, \mathbf{h}_\ell, z^*, \text{op}_\ell) = 0) \end{array} \right].$$

- Otherwise

$$\text{Ver}_\ell := \left[ \begin{array}{l} \text{NIZK.Verify}(\text{crs}_{\text{zk},f}, (\bar{f}_{\log T-1,\gamma[\log T]}^*, \mathbf{h}_\ell), \pi_{f,\ell}) = 1 \wedge \\ (f_{k^*}(z^*) = 0 \vee \text{HashCheck}'(\text{crs}, \mathbf{h}_\ell, z^*, \text{op}_\ell) = 0) \end{array} \right].$$

and define  $\delta_{1,\ell} = \Pr[\text{Ver}_\ell : \text{exp}]$ . Note that  $\delta_{1,\ell} \geq \delta_0$  for all  $\ell$ , thus it suffices to prove that  $\delta_{1,\ell} \leq \text{negl}(\lambda)$ .

**Induction Base.** Suppose  $\ell = \log T - 1$ , thus it holds that  $\mathcal{E}(\text{td}, \bar{f}, \text{hz}, \Sigma) = (z^*, \text{op} = r^*)$  and  $w = \text{Extract}(\text{td}, \bar{f}, \text{hz}, \Sigma) = \text{NIZK.E}(\text{td}_{\text{zk},f}, x = (\bar{f}_{\log T-1,\gamma[\log T]}^*, \mathbf{h}_\ell), \pi_{f,\ell}) = (z^*, \mathbf{h}^*, r^*)$ . Now we define the following probability:

$$\delta_2 = \Pr \left[ (x, w) \in \mathcal{L}_{\text{zk},f} \wedge (f_{k^*}(z^*) = 0 \vee \text{HashCheck}'(\text{crs}, \mathbf{h}_\ell, z^*, \text{op}_\ell) = 0) : \text{exp} \right].$$

Note that:

$$\delta_{1,\ell} - \delta_2 \leq \Pr[\text{NIZK.Verify}(\text{crs}_{\text{zk},f}, x, \pi_{f,\ell}) = 1 \wedge (x, w) \notin \mathcal{L}_{\text{zk},f} : \text{exp}] \leq \text{negl}(\lambda)$$

where the last inequality holds by the argument of knowledge property of the NIZK. Additionally note that  $\delta_2 = 0$  as  $(x, w) \in \mathcal{L}_{\text{zk},f}$  implies (a)  $f_{k^*}(z^*) = 0$ , and (b)  $(\mathbf{h}^*, \text{tree}^*) = \text{HT.Hash}(\text{hk}_{\text{ht}}, z^*)$  and  $\text{PKE.Enc}(\text{pk}, \mathbf{h}^*; r^*) = \mathbf{h}_\ell$ , thus  $\text{HashCheck}'(\text{crs}, \mathbf{h}_\ell, z^*, \text{op}_\ell) = 1$ . So it holds that  $\delta_{1,\log T-1} \leq \text{negl}(\lambda)$ .

**Induction Step.** Suppose  $\delta_{1,i} \leq \text{negl}(\lambda)$  for all  $i < \ell$ , and we want to prove that  $\delta_{1,\ell} \leq \text{negl}(\lambda)$ . Then it holds that  $w = \text{seBARG.E}(\text{td}_{\text{bp},f,\ell}, (\bar{f}_{\ell,\gamma[\ell+1]} = (\bar{f}_1^*, \bar{f}_2^*), \mathbf{h}_\ell, j)_{j \in [2]}, \pi_{f,\ell}) = (\mathbf{h}_1^*, \mathbf{h}_2^*, \pi_f^*)$ .

$$\vee \text{HashCheck}'(\text{crs}, \mathbf{h}_{\ell-1,\gamma[\ell]}^*, z^*, \text{op}_{\ell-1}) = 0$$

Now for  $b = \gamma[\ell]$  we define the following probability:

$$\delta_2 = \Pr[(x = (\bar{f}_b^*, \mathbf{h}_b^*, b), w) \in \mathcal{L}_{\text{bp},f,\ell} \wedge (f_{k^*}(z^*) = 0 \vee \text{HashCheck}'(\text{crs}, \mathbf{h}_\ell, z^*, \text{op}_\ell) = 0) : \text{exp}].$$

Note that:

$$\delta_{1,\ell} - \delta_2 \leq \Pr[\text{seBARG.Verify}(\text{crs}_{\text{bp},f,\ell}, ((\bar{f}_1^*, \bar{f}_2^*), \mathbf{h}_\ell, j)_{j \in [2]}, \pi_{f,\ell}) = 1 \wedge (x, w) \notin \mathcal{L}_{\text{bp},f,\ell} : \text{exp}] \leq \text{negl}(\lambda)$$

where the last inequality holds by the knowledge extractor property of the seBARG. Additionally  $(x, w) \in \mathcal{L}_{\text{bp},f,\ell}$  implies that 1)  $\text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\mathbf{h}_1^*, \mathbf{h}_2^*)) = \mathbf{h}_\ell$ , 2) if  $\ell = \log T$  then  $\text{NIZK.Verify}(\text{crs}_{\text{zk},f}, (\bar{f}_b^*, \mathbf{h}_b^*), \pi_f^*) = 1$ , and if  $\ell > \log T$  then  $\text{seBARG.Verify}(\text{crs}_{\text{bp},f,\ell-1}, (\bar{f}_b^*, \mathbf{h}_b^*, j)_{j \in [2]}, \pi_f^*) = 1$ .

Now since  $\text{HashCheck}'(\text{crs}, \mathbf{h}_\ell, z^*, \text{op}_\ell) = 0$  is equivalent with

$$\text{CRHF.Hash}(\text{hk}_{\text{ch}}, (\mathbf{h}_1^*, \mathbf{h}_2^*)) \neq \mathbf{h}_\ell \vee \text{HashCheck}'(\text{crs}, \mathbf{h}_b^*, z^*, \text{op}_{\ell-1}) = 0,$$

it holds that  $\delta_2 \leq \delta_{1,\ell-1} \leq \text{negl}(\lambda)$  where the last inequality holds by the induction assumption. Thus it holds that  $\delta_{1,\ell} \leq \text{negl}(\lambda)$  concluding the proof.

**Zero-Knowledge.** We construct the simulator  $\mathcal{S}(\text{crs}, \text{cf})$  as follows:

- It generates  $\text{crs}$  similar to  $\text{Gen}$  except that it runs the NIZK simulators to generate  $\text{crs}_{\text{zk}}$  and  $\text{crs}_{\text{zk},f}$ .
- To generate the proof it does the following:
  1. Compute  $\text{cf}_{k,t}$ ,  $\text{st}$ ,  $\text{ct}_{\text{in},k}$ ,  $\bar{\text{hz}}$ , and  $\Pi$  similar to the zero-knowledge simulator for  $\text{zk-IVsC}$  except that for  $\bar{\text{hz}}$  output all the intermediate  $\text{hz}$  values (i.e.,  $\text{hz}_k$  for  $k \in [K]$ ) instead of the final  $\text{hz}$ .
  2. Compute  $\Sigma$  similar to the predicate prover except that use NIZK simulator to compute  $\pi_{\text{zk},f}$  where we use the previously computed  $\text{ct}_{\text{in},k}$  in the statement.

We define the following hybrids (note that the hybrids are similar to the proof of zero-knowledge property of  $\text{zk-IVsC}$  except that we additionally have a  $\mathcal{H}_2$  to simulate the predicate proofs).

$\mathcal{H}_0$ : This is the original experiment. Namely, the adversary sends  $(n, S, T, n_f)$ , challenger sends  $\text{crs}_0 \leftarrow \text{Gen}(1^\lambda, n, S, T, n_f)$ , adversary sends  $(\bar{z}, \bar{f})$  s.t.  $f_k(z_k) = 1$ . Then Challenger computes  $\text{cf} = \mathcal{M}(\bar{z}; 1^T)$ ,  $(\text{st}_0, \bar{\text{hz}}_0, \bar{\text{ps}}, \Pi_0) \leftarrow \text{Prove}(\text{crs}_0, \bar{z})$ ,  $(\text{hz}_{0,K}, \Sigma_0) \leftarrow \text{PredicateProve}(\text{crs}, \bar{f}, \bar{z}, \bar{\text{ps}})$  and sends  $(\text{st}_0, \bar{\text{hz}}_0, \Pi_0, \Sigma_0)$  to the adversary and finally adversary outputs  $b'$ .

$\mathcal{H}_1$ : Similar to the original experiment except that NIZK simulator is used to generate  $\text{crs}_{\text{zk}}$ , and  $\pi_{\text{zk}}$  at every iteration of computing  $\Pi$ .

$\mathcal{H}_2$ : Similar to  $\mathcal{H}_1$  except that NIZK simulator is used to generate  $\text{crs}_{\text{zk},f}$ , and  $\pi_{\text{zk},f}$  at every iteration of computing  $\Sigma$ .

$\mathcal{H}_{3,k,t}$ : Similar to  $\mathcal{H}_2$  except that for every  $\text{cf}_{i,j}$  until  $\text{cf}_{k,t}$  except  $\text{cf}_{1,0}$  and  $\text{cf}_{K,T}$ , we let  $\text{epcf}_{i,j} = \text{PKE.Enc}(\text{pk}, 0^{\text{Pcf}})$  (and for the rest  $\text{epcf}_{i,j}$  is honestly generated). Note that  $\mathcal{H}_{3,1,0} = \mathcal{H}_2$  and  $\mathcal{H}_{3,K,T-1} = \mathcal{H}_{3,K,T}$ .

$\mathcal{H}_{4,k}$ : Similar to  $\mathcal{H}_{3,K,T}$  except that for every  $z_i$  until  $z_k$ , we let  $\text{ct}_{\text{in},i} = \text{PKE.Enc}(\text{pk}, 0^\lambda)$  (and for the rest  $\text{ct}_{\text{in},i}$  is honestly generated). Note that  $\mathcal{H}_{4,0} = \mathcal{H}_{3,K,T}$  and  $\mathcal{H}_{4,K}$  is the output of the simulator. Namely in  $\mathcal{H}_{4,K}$  the adversary sends  $(n, S, T, n_f)$ , the simulator sends  $\text{crs}_1 \leftarrow \mathcal{S}(1^\lambda, n, S, T, n_f)$ , adversary sends  $\bar{z}(\bar{z}, \bar{f})$  s.t.  $f_k(z_k) = 1$ . The simulator computes  $\text{cf} = \mathcal{M}(\bar{z}; 1^T)$  and sends  $(\text{st}_1, \bar{\text{hz}}_1, \Pi_1, \Sigma_1) \leftarrow \mathcal{S}(\text{crs}_0, \text{cf})$  to the adversary and finally adversary outputs  $b'$ .

For any adversary let  $\mathcal{H}^*(\mathcal{A})$  be the output of the adversary in hybrid  $\mathcal{H}^*$ . The proof closely follows the proof of zero-knowledge property in zk-IVsC except that we additionally have to prove the indistinguishability of hybrids  $\mathcal{H}_1$  and  $\mathcal{H}_2$ . Thus we have the following claim:

**Claim B.1.** If NIZK is zero-knowledge then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_2(\mathcal{A}) = 1] - \Pr[\mathcal{H}_1(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Similar to the proof of Claim A.7. □

□

## C ICZK Complete Analysis.

*Proof of Theorem 4.4.* Here we prove completeness, adaptive soundness and incremental zero-knowledge of our construction.

**Completeness.** The completeness follows by the construction and the completeness of the underlying  $\Psi$ , and PKE.

**Efficiency.** Note that the size of  $\text{hz}_k$  values and  $\Pi_{\psi,k}$  are  $\text{poly}(\lambda, \log |x_1|, \log |w_1|, \log K)$ , and the size of  $\Sigma_{\psi,k}$  is  $\text{poly}(\lambda, \log |x_1|, |w_1|, \log K)$ . Additionally since rrPKE is rate-1, the size of each  $\text{ct}_k$  is  $|w_i| + \text{poly}(\lambda)$ . Therefore, the total proof size is  $K|w_i| + K \cdot \text{poly}(\lambda, \log |x_1|, \log |w_1|, \log K) + \text{poly}(\lambda, \log |x_1|, |w_1|, \log K)$ .

**Adaptive Soundness.** Let  $\text{Exp}_{\text{ICZK}}$  be the following experiment:

$$\text{Exp}_{\text{ICZK}} := \left[ \begin{array}{l} (\text{crs}, \text{sk}) \leftarrow \text{Setup}(1^\lambda, 1^{n_x}, 1^{n_w}), \\ (\bar{x}, \Pi = (\overline{\text{hz}}, \Pi_\psi, \Sigma_\psi, \overline{\text{ct}})) \leftarrow \mathcal{A}(\text{crs}) \\ |x_k| = n_x, ((w_k, \text{ps}_k; r_k) \leftarrow \text{PKE.Dec}(\text{sk}, \text{ct}_k))_{k \in [K]} \\ (z_k = x_k, w_k, \text{ps}_k, r_k)_{k \in [K]} \end{array} \right]$$

and **Bad** be the following event:

$$\text{Bad} := \left[ \begin{array}{l} \bar{x} \notin \mathcal{L} \wedge \Psi.\text{Verify}(\text{crs}, (\text{hz}_K, 1), \Pi_\psi) = 1 \wedge \\ \Psi.\text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}_K, \Sigma_\psi) = 1 \end{array} \right].$$

Note that to prove the adaptive soundness of Construction 4.3 we only need to prove that  $\delta_1 = \Pr[\text{Bad} : \text{Exp}_{\text{ICZK}}] \leq \text{negl}(\lambda)$ . Now define the following probabilities:

$$\delta_2 = \Pr \left[ \begin{array}{l} \bar{x} \notin \mathcal{L} \wedge \Psi.\text{Verify}(\text{crs}, (\text{hz}_K, 1), \Pi_\psi) = 1 \wedge \\ \text{hz}_K = \Psi.\text{Digest}(\text{crs}, \bar{z}, \bar{\text{ps}}) \end{array} : \text{Exp}_{\text{ICZK}} \right].$$

$$\delta_3 = \Pr \left[ \bar{x} \notin \mathcal{L} \wedge \Psi.\text{Verify}(\text{crs}, (\text{hz}_K, 1), \Pi_\psi) = 1 : \text{Exp}_{\text{ICZK}}, \text{hz}_K = \Psi.\text{Digest}(\text{crs}, \bar{z}, \bar{\text{ps}}) \right].$$

Now note that  $\bar{x} \notin \mathcal{L}$  implies  $\mathcal{M}(\bar{x}, \bar{w}) = 0$  that implies  $\mathcal{M}'(\bar{z}) = 0$  thus by the soundness property that we prove in predicated zk-IVsC it holds that  $\delta_3 \leq \text{negl}(\lambda)$ . Additionally note that



$\delta_2 \leq \delta_3$ . Hence, it suffices to show that  $\delta_1 - \delta_2 \leq \text{negl}(\lambda)$ , namely we need to show that  $\delta^* \leq \text{negl}(\lambda)$  where:

$$\delta^* = \Pr \left[ \begin{array}{l} \Psi.\text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}_K, \Sigma_\psi) = 1 \wedge \\ \text{hz}_K \neq \Psi.\text{Digest}(\text{crs}, \bar{z}, \bar{f}) \end{array} : \text{Exp}_{\text{ICZK}} \right].$$

Suppose for some non-negligible function  $\epsilon(\cdot)$  it holds that  $\delta^* \geq \epsilon(\lambda)$ , then for some  $k^* \in [K]$  and some non-negligible function  $\epsilon'(\cdot)$  it holds that:

$$\Pr \left[ \begin{array}{l} \Psi.\text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}_K, \Sigma_\psi) = 1 \wedge \\ \text{hz}_{k^*} \neq \Psi.\text{Digest}(\text{crs}, \bar{z}_{k^*}, \bar{\text{ps}}_{k^*}) \wedge \\ \text{hz}_{k^*-1} = \Psi.\text{Digest}(\text{crs}, \bar{z}_{k^*-1}, \bar{\text{ps}}_{k^*-1}) \end{array} : \text{Exp}_{\text{ICZK}} \right] \geq \epsilon'(\lambda). \quad (13)$$

Now we construct an adversary  $\mathcal{B}$  that given  $\mathcal{A}$  breaks the predicate soundness property that we proved in predicated zk-IVsC.  $\mathcal{B}$  receives  $n_x$  and  $n_w$  from  $\mathcal{A}$ , samples a random index  $k' \in [K]$ , sends  $(n, S, T, n_f, k')$  to the challenger, receives  $\text{crs}_\psi$  from the challenger, samples  $(\text{pk}, \text{sk})$ , send  $\text{crs} = (\text{crs}_\psi, \text{pk})$  to  $\mathcal{A}$ , receives  $(\bar{x}, \Pi)$  from  $\mathcal{A}$ , computes  $(\bar{w}, \bar{\text{ps}}, \bar{r})$  by decrypting  $\bar{\text{ct}}$  using  $\text{sk}$ , lets  $z_k = (x_k, w_k, \text{ps}_k, r_k)$  for  $k \in [K]$ , finds  $\text{pk}'$  and  $\text{hk}_{\text{ht}}$  in  $\text{crs}_\psi$ , computes  $\bar{f}$ , and sends  $(\bar{f}, \text{hz}_K, \Sigma_\psi)$  to the challenger. Note that this is the same as  $\text{Exp}_{\text{ICZK}}$  except that  $\text{crs}_\psi$  is generated using  $\Psi.\text{TGen}$  on index  $k'$ . Let this experiment be  $\text{Exp}_{\text{ICZK}, k'}$ . Note that  $\mathcal{B}$ 's advantage can be computed as follows:

$$\delta_{\mathcal{B}} = \sum_{k'=1}^K 1/K \cdot \Pr \left[ \begin{array}{l} \Psi.\text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}_K, \Sigma_\psi) = 1 \wedge \\ f_{k'}(z^*) = 0 \end{array} : \begin{array}{l} \text{Exp}_{\text{ICZK}, k'}, \\ z^* \leftarrow \Psi.\mathcal{E}(\text{td}, \bar{f}, \text{hz}_K, \Sigma_\psi) \end{array} \right],$$

and it holds that,

$$\delta_{\mathcal{B}} \geq 1/K \cdot \Pr \left[ \begin{array}{l} \Psi.\text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}_K, \Sigma_\psi) = 1 \wedge \\ f_{k^*}(z^*) = 0 \end{array} : \begin{array}{l} \text{Exp}_{\text{ICZK}, k^*}, \\ z^* \leftarrow \Psi.\mathcal{E}(\text{td}, \bar{f}, \text{hz}_K, \Sigma_\psi) \end{array} \right]. \quad (14)$$

Now let  $z^* = (x^*, w^*, \text{ps}^*, r^*)$ , and consider the following cases:

- If  $x^* \neq x_{k^*}$  then by the the description of  $f_{k^*}$  it holds that  $f_{k^*}(z^*) = 0$ .
- If  $\text{PKE}.\text{Enc}(\text{pk}, (w^*, \text{ps}^*); r^*) \neq \text{ct}_{k^*}$  then by the the description of  $f_{k^*}$  it holds that  $f_{k^*}(z^*) = 0$  (otherwise by perfect correctness of PKE it holds that  $(w^*, \text{ps}^*, r^*) = (w_{k^*}, \text{ps}_{k^*}, r_{k^*})$ ).
- If  $z^* = z_{k^*}$ , then  $f_{k^*}(z^*) = 0$  is equivalent with  $f_{k^*}(z_{k^*}) = 0$ .

Let  $(\text{rt}_{k^*}, \text{tree}_{k^*}) = \text{HT}.\text{Hash}(\text{hk}_{\text{ht}}, (x_{k^*}, w_{k^*}, \text{ps}_{k^*}, r_{k^*}))$  and  $\hat{\text{ct}}_{k^*} = \text{PKE}.\text{Enc}(\text{pk}', \text{rt}_{k^*}; \text{ps}_{k^*})$ . Thus by Eq. (14) and the description of  $f_{k^*}$  it holds that:

$$\delta_{\mathcal{B}} \geq \Pr \left[ \begin{array}{l} \Psi.\text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}_K, \Pi_\psi) = 1 \wedge \\ \text{nxt-hz}(\text{hz}_{k-1}, \hat{\text{ct}}_{k^*}) \neq \text{hz}_k \end{array} : \text{Exp}_{\text{ICZK}, k^*} \right] \geq \epsilon(\lambda)/K - \text{negl}(\lambda).$$

which means

$$\delta_{\mathcal{B}} \geq \Pr \left[ \begin{array}{l} \Psi.\text{PredicateVerify}(\text{crs}, \bar{f}, \text{hz}_K, \Pi_\psi) = 1 \wedge \\ \text{hz}_{k^*} \neq \Psi.\text{Digest}(\text{crs}, \bar{z}_{k^*}, \bar{\text{ps}}_{k^*}) \wedge \\ \text{hz}_{k^*-1} = \Psi.\text{Digest}(\text{crs}, \bar{z}_{k^*-1}, \bar{\text{ps}}_{k^*-1}) \end{array} : \text{Exp}_{\text{ICZK}, k^*} \right] \geq \epsilon(\lambda)/K - \text{negl}(\lambda).$$

there by Eq. (13) it holds that  $\delta_{\mathcal{B}} \geq \epsilon'(\lambda)$  which contradicts the predicate soundness property that we proved for predicated zk-IVsC concluding the soundness of Construction 4.3.

**Incremental zero-Knowledge.** We construct the simulator  $\mathcal{S}(\text{crs}, \text{cf})$  as follows:

- It generates  $\text{crs}$  similar to the zero-knowledge simulator in the predicated zk-IVsC.
- To generate the proof it does the following:
  1. Compute  $\Pi_\psi$ ,  $\Sigma_\psi$ , and  $\bar{\text{hz}}$  similar to the zero-knowledge simulator in the predicated zk-IVsC.
  2. Compute  $\bar{\text{ct}}$  by independently encrypting  $0^{n_w + |\text{ps}|}$  for every  $k \in [K]$ .

We define the following hybrids:

$\mathcal{H}_0$ : This is the original experiment. Namely, the challenger on input  $(n_x, n_w)$ , sends  $\text{crs}_0 \leftarrow \text{Gen}(1^\lambda, n_x, n_w)$  to  $\mathcal{A}$ ,  $\mathcal{A}$  sends  $(\bar{x}, \bar{w})$  s.t.  $\mathcal{R}(\bar{x}, \bar{w}) = 1$  and  $\mathcal{M}(\bar{x}, \bar{w}) = \text{cf}$ . Then the challenger sends  $(\text{st}_0, \Pi_0) = \text{Prove}(\text{crs}, \bar{x}, \bar{w})$  to  $\mathcal{A}$ , and finally  $\mathcal{A}$  outputs  $b'$ .

$\mathcal{H}_1$ : Similar to  $\mathcal{H}_0$  except that NIZK simulator is used to generate  $\text{crs}_{\text{zk}}$ , and  $\pi_{\text{zk}}$  at every iteration of computing  $\Pi_\psi$ .

This step only uses the zero-knowledge property of the underlying NIZK. Note that after this step the only information about the intermediate configurations is in  $\text{epcf}_{k,t}$  (for  $k \in [K]$  and  $t \in [T]$ ) that is the encryption of the root of the pseudo-configuration. Additionally, the only information about the inputs of the machine is in  $(\bar{\text{hz}}, \bar{\text{ct}}, \Sigma_\psi)$  and  $(\text{ct}_{\text{in},k})_{k \in [K]}$  (that is being used in the statements of  $\Pi_\psi$ ).

$\mathcal{H}_2$ : Similar to  $\mathcal{H}_1$  except that NIZK simulator is used to generate  $\text{crs}_{\text{zk},f}$ , and  $\pi_{\text{zk},f}$  at every iteration of computing  $\Sigma_\psi$ .

This step only uses the zero-knowledge property of the underlying NIZK. Note that after this step there is no information about inputs  $(\bar{w}, \bar{\text{ps}}, \bar{r})$  in the proof except the ones in  $(\bar{\text{hz}}, \bar{\text{ct}})$  and  $(\text{ct}_{\text{in},k})_{k \in [K]}$  (that is being used in the statements of  $\Pi_\psi$  and  $\Sigma_\psi$ ). Additionally note that the only information about the input  $z_k$  other than the ones in  $\text{ct}_k$ , is a short (of size  $\lambda$  bits) digest of the input  $z_k$  that is encrypted in  $\text{ct}_{\text{in},k}$  (since  $\bar{\text{hz}}$  is constructed only using  $(\text{ct}_{\text{in},k})_{k \in [K]}$ ).

$\mathcal{H}_{3,k,t}$ : Similar to  $\mathcal{H}_2$  except that for every  $\text{cf}_{i,j}$  until  $\text{cf}_{k,t}$  except  $\text{cf}_{1,0}$  and  $\text{cf}_{K,T}$ , and in the generation of  $\Pi_\psi$  we let  $\text{epcf}_{i,j} \leftarrow \text{PKE}.\text{Enc}(\text{pk}, 0^{|\text{pcf}|})$  (and for the rest  $\text{epcf}_{i,j}$  is honestly generated). Note that  $\mathcal{H}_{3,1,0} = \mathcal{H}_2$  and  $\mathcal{H}_{3,K,T-1} = \mathcal{H}_{3,K,T}$ .

This step only uses the security of the underlying PKE. Note that after this step there is no information about the intermediate configurations in the proof.

$\mathcal{H}_{4,k}$ : Similar to  $\mathcal{H}_{3,K,T}$  except that for every  $\text{ct}_i$  until  $\text{ct}_k$ , we independently sample  $r_i$  and let  $\text{ct}_i = \text{rrPKE.Enc}(\text{pk}, 0^{n_w + |\text{ps}|}; r_i)$  (and for the rest  $\text{ct}_i$  is honestly generated). Note that  $\mathcal{H}_{4,0} = \mathcal{H}_{3,K,T}$ .

This step only uses the security with leakage of the underlying  $\text{rrPKE}$  as the only information about randomness  $\bar{r}$  remained in the proof is in  $(\text{ct}_{\text{in},k})_{k \in [K]}$  (that is the encryption of a short (of size  $\lambda$ ) bits of the inputs). Note that after this step there's no information about  $\bar{w}$  and  $\bar{\text{ps}}$  in the  $\bar{\text{ct}}$ . The only information left in the proof is in  $(\text{ct}_{\text{in},k})_{k \in [K]}$ .

$\mathcal{H}_{5,k}$ : Similar to  $\mathcal{H}_{4,K}$  except that for every  $z_i$  until  $z_k$ , we let  $\text{ct}_{\text{in},i} \leftarrow \text{PKE.Enc}(\text{pk}, 0^\lambda)$  (and for the rest  $\text{ct}_{\text{in},i}$  is honestly generated). Note that  $\mathcal{H}_{5,0} = \mathcal{H}_{4,K}$  and  $\mathcal{H}_{5,K}$  is the output of the simulator. Namely in  $\mathcal{H}_{5,K}$  the simulator on input  $(n_x, n_w)$ , sends  $\text{crs}_1 \leftarrow \mathcal{S}(1^\lambda, n_x, n_w)$  to  $\mathcal{A}$ ,  $\mathcal{A}$  sends  $(\bar{x}, \bar{w})$  s.t.  $\mathcal{R}(\bar{x}, \bar{w}) = 1$  and  $\mathcal{M}(\bar{x}, \bar{w}) = \text{cf}$ . Then the simulator sends  $(\text{st}_1, \Pi_1) = \mathcal{S}(\text{crs}, \bar{x}, \text{cf})$  to  $\mathcal{A}$ , and finally  $\mathcal{A}$  outputs  $b'$ .

This step only uses the randomness-dependent message security of the underlying  $\text{PKE}$  as the only information about randomness  $\bar{\text{ps}}$  remained in the proof is in the message of  $(\text{ct}_{\text{in},k})_{k \in [K]}$  (that is the encryption of a short (of size  $\lambda$ ) bits of the inputs). Note that after this step there's no information about  $\bar{w}$  in the proof and the proof can be simulated using only  $(\bar{x}, \text{cf})$ .

For any adversary let  $\mathcal{H}^*(\mathcal{A})$  be the output of the adversary in hybrid  $\mathcal{H}^*$ . We have the following claims:

**Claim C.1.** If NIZK is zero-knowledge then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_1(\mathcal{A}) = 1] - \Pr[\mathcal{H}_0(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* hi Towards the contradictions suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  s.t.  $|\Pr[\mathcal{H}_1(\mathcal{A}) = 1] - \Pr[\mathcal{H}_0(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ . We construct adversary  $\mathcal{B}_{\text{zk}}$  that breaks the zero-knowledge property of the underlying NIZK protocol as follows:

1. Receives  $(1^{n_x}, 1^{n_w})$  as input, and  $\text{crs}^*$  from NIZK zero-knowledge challenger.
2. Computes  $(\text{crs}, \text{td}) \leftarrow \text{Gen}(1^\lambda, 1^{n_x}, 1^{n_w})$  and updates  $\text{crs}$  by replacing  $\text{crs}_{\text{zk}} \in \text{crs}_\psi$  with  $\text{crs}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$ , receives  $\bar{x}, \bar{w}$ .
4. Computes  $(\text{st}, \Pi)$  similar to the prover except that at every step of updating  $\text{hz}_{k-1}$  to  $\text{hz}_k - 1$   $\Pi_\psi$  (every call of `UpdateSingleStep` in subroutine  $\Psi.\text{Update}$ ), instead of computing  $\pi_{\text{zk}}$ , queries the challenger on the correct computation of corresponding  $(x, w)$  and receive  $\pi_{\text{zk}}$ .
5. Outputs whatever  $\mathcal{A}$  outputs.

If the NIZK challenger generates  $\text{crs}$  and proof using  $\text{Gen}$  and  $\text{Prove}$ , then  $\mathcal{B}_{\text{zk}}$  simulates  $\mathcal{H}_0$  for  $\mathcal{A}$  and if the challenger uses NIZK simulator, then  $\mathcal{B}_{\text{zk}}$  simulates  $\mathcal{H}_1$  for  $\mathcal{A}$ . Thus  $\mathcal{B}_{\text{zk}}$  has a non-negligible advantage  $\epsilon(\lambda)$  in breaking the zero-knowledge property of NIZK.  $\square$

**Claim C.2.** If NIZK is zero-knowledge then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_2(\mathcal{A}) = 1] - \Pr[\mathcal{H}_1(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Towards the contradictions suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  s.t.  $|\Pr[\mathcal{H}_2(\mathcal{A}) = 1] - \Pr[\mathcal{H}_1(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ . We construct adversary  $\mathcal{B}_{\text{zk}}$  that breaks the zero-knowledge property of the underlying NIZK protocol as follows:

1. Receives  $(1^{n_x}, 1^{n_w})$  as input, and  $\text{crs}^*$  from NIZK zero-knowledge challenger.
2. Computes  $(\text{crs}, \text{td})$  similar to  $\mathcal{H}_1$  and updates  $\text{crs}$  by replacing  $\text{crs}_{\text{zk},f} \in \text{crs}_\psi$  with  $\text{crs}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$ , receives  $\bar{x}, \bar{w}$ .
4. Computes  $(\text{st}, \Pi)$  similar to the  $\mathcal{H}_1$  except that at every update of computing  $\Sigma_\psi$  (every call of subroutine  $\Psi.\text{PredicateUpdate}$ ), instead of computing  $\pi_{\text{zk},f}$ , queries the challenger on the correct computation of the corresponding  $(x, w)$  and receive  $\pi_{\text{zk},f}$ .
5. Outputs whatever  $\mathcal{A}$  outputs.

If the NIZK challenger generates  $\text{crs}$  and proof using  $\text{Gen}$  and  $\text{Prove}$ , then  $\mathcal{B}_{\text{zk}}$  simulates  $\mathcal{H}_1$  for  $\mathcal{A}$  and if the challenger uses NIZK simulator, then  $\mathcal{B}_{\text{zk}}$  simulates  $\mathcal{H}_2$  for  $\mathcal{A}$ . Thus  $\mathcal{B}_{\text{zk}}$  has a non-negligible advantage  $\epsilon(\lambda)$  in breaking the zero-knowledge property of NIZK.  $\square$

In the following claim we let  $k \in [K]$ ,  $t \in [T]$ , and  $(k, t) \neq (K, T)$ .

**Claim C.3.** If PKE is secure then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_{3,k,t}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{3,k,t-1}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Towards the contradictions suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  s.t.  $|\Pr[\mathcal{H}_{3,k,t}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{3,k,t-1}(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ . We construct adversary  $\mathcal{B}_{\text{pke}}$  that breaks the security of the underlying PKE protocol as follows:

1. Receives  $(1^{n_x}, 1^{n_w})$  as input, and  $\text{pk}^*$  from PKE challenger.
2. Computes  $(\text{crs}, \text{td})$  similar to  $\mathcal{H}_2$  and updates  $\text{crs}$  by replacing  $\text{pk} \in \text{crs}_\psi$  with  $\text{pk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$ , receives  $(\bar{x}, \bar{w})$ .
4. Computes  $(\text{st}, \Pi)$  similar to  $\mathcal{H}_2$  except that when running single step of computing  $\Pi_\psi$  (calling  $\text{UpdateSingleStep}$  in subroutine  $\Psi.\text{Update}$ ) on  $(k, t)$ , instead of computing  $\text{epcf}_{k,t}$ , queries the challenger on  $\text{pcf}_{k,t}$  and receive  $\text{epcf}_{k,t}$ .
5. Outputs whatever  $\mathcal{A}$  outputs.

If the PKE challenger generates the ciphertext honestly, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{3,k,t-1}$  for  $\mathcal{A}$  and if it responds with a ciphertext of all zeros, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{3,k,t}$  for  $\mathcal{A}$ . Thus  $\mathcal{B}_{\text{pke}}$  has a non-negligible advantage  $\epsilon(\lambda)$  in breaking the security of PKE.  $\square$

In the following claim we let  $k \in [K]$ .

**Claim C.4.** If rrPKE is secure then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_{4,k}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{4,k-1}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Towards the contradictions suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  s.t.  $|\Pr[\mathcal{H}_{4,k}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{4,k-1}(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ . We construct adversary  $\mathcal{B}_{\text{pke}}$  that breaks the security with leakage of the underlying PKE protocol as follows:

1. Receives  $(1^{n_x}, 1^{n_w})$  as input, and  $\text{pk}^*$  from PKE challenger.
2. Computes  $(\text{crs}, \text{td})$  similar to  $\mathcal{H}_{3,K,T}$  and updates  $\text{crs}$  by replacing  $\text{pk}$  with  $\text{pk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$ , receives  $(\bar{x}, \bar{w})$ .
4. Computes  $(\text{st}, \Pi)$  similar to  $\mathcal{H}_{3,K,T}$  except that instead of computing  $\text{ct}_k \in \bar{\text{ct}}$ , queries the challenger on  $((w_k, \text{ps}_k), g_k(\cdot))$ , and receive  $(\text{ct}_k, g(r_k))$ , where we define  $g_k(r)$  to compute  $\text{HT.Hash}(\text{hk}_{\text{ht}}, z_k = (x_k, w_k, \text{ps}_k, r_k))$  and output only  $\text{rt}_k$  (and not  $\text{tree}_k$ ). Note that  $g_k(\cdot)$  has  $(\text{hk}_{\text{ht}}, x_k, w_k, \text{ps}_k)$  hard-coded.
5. Outputs whatever  $\mathcal{A}$  outputs.

If the PKE challenger generates the ciphertext honestly, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{4,k-1}$  for  $\mathcal{A}$  and if it responds with a ciphertext of all zeros, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{4,k}$  for  $\mathcal{A}$ . Additionally note that  $|\text{rt}_k| = \lambda$ , thus there is only  $\lambda$  bits of leakage about the randomness  $r_k$  that is used to encrypt  $(w_k, \text{ps}_k)$ . Therefore  $\mathcal{B}_{\text{pke}}$  has a non-negligible advantage  $\epsilon(\lambda)$  in breaking the security with leakage of PKE.  $\square$

In the following claim we let  $k \in [K]$ .

**Claim C.5.** If PKE is secure then for any stateful PPT adversary  $\mathcal{A}$  it holds that:

$$|\Pr[\mathcal{H}_{5,k}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{5,k-1}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

*Proof.* Towards the contradictions suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  s.t.  $|\Pr[\mathcal{H}_{5,k}(\mathcal{A}) = 1] - \Pr[\mathcal{H}_{5,k-1}(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ . We construct adversary  $\mathcal{B}_{\text{pke}}$  that breaks the randomness-dependent message security of the underlying PKE protocol as follows:

1. Receives  $(1^{n_x}, 1^{n_w})$  as input, and  $\text{pk}^*$  from PKE challenger.
2. Computes  $(\text{crs}, \text{td})$  similar to  $\mathcal{H}_{4,K}$  and updates  $\text{crs}$  by replacing  $\text{pk}' \in \text{crs}_\psi$  with  $\text{pk}^*$ .
3. Sends  $\text{crs}$  to  $\mathcal{A}$ , receives  $(\bar{x}, \bar{w})$ .
4. Computes  $(\text{st}, \Pi)$  similar to  $\mathcal{H}_{4,K,T}$  except that instead of computing  $\text{ct}_{\text{in},k}$ , interacts with the challenger as follows: sends  $g_k(\cdot)$ , receives  $g_k(\text{ps}_k)$ , sends  $g_k(\text{ps}_k)$ , and receive  $\text{ct}_{\text{in},k}$ . We define  $g_k(\text{ps}_k)$  to compute  $\text{HT.Hash}(\text{hk}_{\text{ht}}, z_k = (x_k, w_k, \text{ps}_k, r_k))$  and output only  $\text{rt}_k$  (and not  $\text{tree}_k$ ). Note that  $g_k(\cdot)$  has  $(\text{hk}_{\text{ht}}, x_k, w_k, r_k)$  hard-coded.
5. Outputs whatever  $\mathcal{A}$  outputs.

If the PKE challenger generates the ciphertext honestly, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{5,k-1}$  for  $\mathcal{A}$  and if it responds with a ciphertext of all zeros, then  $\mathcal{B}_{\text{pke}}$  simulates  $\mathcal{H}_{5,k}$  for  $\mathcal{A}$ . Additionally note that  $|\text{rt}_k| = \lambda$ , thus there is only  $\lambda$  bits of leakage about the randomness  $\text{ps}_k$  that is used to encrypt  $\text{rt}_k$ . Therefore  $\mathcal{B}_{\text{pke}}$  has a non-negligible advantage  $\epsilon(\lambda)$  in breaking the randomness-dependent message security PKE.  $\square$

Therefore we conclude the proof of zero-knowledge property by a hybrid argument.

□