

SNACKs for Proof-of-Space Blockchains

Hamza Abusalah 

IMDEA Software Institute, Madrid, Spain.
`hamza.abusalah@imdea.org`

Abstract. SNACKs are succinct non-interactive arguments of chain knowledge. They allow for efficient and generic solutions to blockchain light-client bootstrapping. Abusalah et al. construct SNACKs in the random oracle model for any *single-chain* blockchain from any graph-labeling proof of sequential work (PoSW) scheme. Their SNACK construction is a PoSW-like protocol over the augmented blockchain.

Unlike single-chain blockchains, such as proof-of-work and proof-of-stake blockchains, proof-of-space (PoSpace) blockchains are composed of two chains: a *canonical* proof chain and a data chain. These two chains are related using a signature scheme.

In this work, we construct PoSW-enabled SNACKs for any PoSpace blockchain. Combined with the results of Abusalah et al., this gives the first solution to light-client bootstrapping in PoSpace blockchains. The space cost of our construction is *two* hash values in each augmented PoSpace block. Generating SNACK proofs for a PoSpace blockchain is identical to generating SNACK proofs for single-chain blockchains and amounts to looking up a succinct number of augmented blocks.

1 Introduction

Consider a blockchain protocol Π , say Bitcoin or the Chia Network, and a light client V , which is assumed to hold only minimal information about Π , say its genesis block ψ . A *bootstrapping* protocol [BKLZ20, AFGK22] for a blockchain allows such V to hold a commitment to its stable prefix.

A succinct non-interactive argument of chain knowledge (SNACK) system is a computationally-sound proof system (P, V) that allows a prover P to give a succinct non-interactive proof that convinces a verifier V that P *knows* a *chain* of certain weight. Crucially, the SNACK proof is succinct, i.e., poly-logarithmic in the length of the blockchain.

In [AFGK22], secure bootstrapping is formalized and instantiated for any blockchain protocol Π for which (1) we have a secure SNACK system and (2) a natural and previously used assumption [BKLZ20] on the adversarial mining power holds. This is captured in the (c, ℓ, ϵ) -fork assumption, which informally says that, except with probability ϵ , no adversary can produce a fork containing ℓ consecutive blocks with more than a c -fraction

of them being valid.¹ Furthermore, Abusalah et al. [AFGK22] construct SNACKs for single-chain blockchains, like Bitcoin, generically from any graph-labeling proof of sequential work (PoSW) scheme assuming the (c, ℓ, ϵ) -fork assumption holds for such chains.

SNACKs for PoSpace Blockchains. In this paper, we study SNACKs for PoSpace blockchains [PKF⁺18, CP19]. These blockchains are composed of two chains in tandem: a proof chain and a data chain – see Fig. 2. Both chains are bound by a signature scheme. The proof chain contains only *canonical* data, such as (unique) proofs of space [DFKP15, AAC⁺17] and verifiable-delay function [BBBF18] computations. The data chain contains transactions and any arbitrary data the blockchain protocol allows. The dual nature of these chains and the requirement that the proof chain must remain canonical make designing SNACKs for such chains more involved than their single-chain blockchain counterparts, say Bitcoin.

Contributions. In this work, we extend the framework of [AFGK22] and construct SNACKs for *any* PoSpace blockchain from any graph-labeling PoSW scheme. The cost of our construction is *two* hash values in each augmented block, one in the augmented proof block, and one in the augmented data block. Generating a SNACK proof is as efficient as generating a PoSW proof, which amounts to looking up a succinct² number of blocks.

(We mention that simply defining the PoSpace SNACK to be two SNACK systems, one for the proof chain, and one for the data chain, doesn't result in a secure SNACK, and extra care needs to be exercised in order to prove security of the SNACK and maintain the security of the underlying PoSpace blockchain.)

Therefore, by the results of [AFGK22], by plugging in our PoSpace SNACK construction into their generic bootstrapping protocol, we get, to the best of our knowledge, the first solution to bootstrapping in PoSpace blockchains that avoids setup assumptions. Our protocol, as outlined above, is also practically efficient.³

¹ This assumption was first introduced in [BKLZ20] in the PoW-blockchain setting and adopted in [AFGK22] generically, i.e., without reference to the Sybil-mechanism of the underlying blockchain protocol. Studying the (c, ℓ, ϵ) -fork assumption in various blockchain protocols, and possibly deriving it from their underlying security assumptions, is an interesting open problem that we don't address in this work.

² Depending on the PoSW scheme used, this number maybe $O(t \log n)$ where n is the length of the blockchain and t a security parameter

³ SNACKs are on par with Flyclient in terms of practical efficiency [AFGK22].

1.1 SNACK Systems

Consider a family of weighted DAGs $\Gamma = (\Gamma_n = (G_n, \Omega_n))_{n \in \mathbb{N}}$, where $G_n = ([n]_0, E_n)$ is a DAG on $[n]_0 := \{0, \dots, n\}$ and $\Omega_n : [n]_0 \rightarrow [0, 1]$ is a weight function, i.e., $\sum_{i=0}^n \Omega_n(i) = 1$. We define an NP language $\mathcal{L}_{\Gamma, R, \text{Com}}$ parameterized by Γ , a position-binding commitment scheme Com , and a polynomial-time relation R , which defines validity of labels of vertices in Γ . An element $(\phi, n) \in \mathcal{L}_{\Gamma, R, \text{Com}}$ consists of a Com commitment ϕ to a labeling of G_n that is valid with respect to R . A *labeling* of G_n is any mapping $L : [n]_0 \rightarrow \{0, 1\}^\lambda$.

A SNACK system (P, V) for $\mathcal{L}_{\Gamma, R, \text{Com}}$ in a non-interactive argument system satisfying completeness, (α, ϵ) -knowledge soundness, and succinctness. Completeness guarantees that an honest P holding a witness for a statement $(\phi, n) \in \mathcal{L}_{\Gamma, R, \text{Com}}$ makes V accept. For an $\alpha \in (0, 1]$, (α, ϵ) -knowledge-soundness guarantees that from any convincing prover for a statement (ϕ, n) , one can extract, except with probability ϵ , an R -valid labeling of a path P in Γ_n of weight at least α . Standard knowledge soundness is recovered by setting $\alpha = 1$. However, relaxing α to $\alpha < 1$ allows for more efficient instantiations of SNACKs. Succinctness requires that the proof size as well as its verification time are poly-logarithmic in n and polynomial in the security parameter.

PoSW-Enabled SNACKs. The SNACK construction of [AFGK22] works by carefully augmenting the labels of the graph underlying a (graph-labeling) Proof of Sequential Work (PoSW) with the blockchain data in a way that allows leveraging the security and efficiency guarantees of the PoSW scheme to the SNACK.

A PoSW scheme (P, V) is an (interactive) proof system in which P , on common input an integer parameter n and a statement χ sampled by V , computes a proof π that convinces the verifier that n *sequential* computational steps have been performed since χ was received.

All known PoSW schemes [MMV13, CP18, AKK⁺19, DLM19, AFGK22, AC23] are based on a random-oracle induced labeling of a particular weighted DAG $(G_n = ([n]_0, E_G), \Omega_n)$ for a weight function Ω_n on $[n]_0$.⁴ In particular, upon receiving a statement χ from V , P uses χ to refresh a random oracle $\tau : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ to compute a labeling $L : [n]_0 \rightarrow \{0, 1\}^\lambda$ of G_n , where the label of vertex $i \in [n]_0$ is defined as $L(i) := \tau(\chi, i, L(\text{parents}(i)))$, where $\text{parents}(i)$ denotes the parents of i in G_n in any fixed ordering. Then P sends a vector commitment of L to V ,

⁴ In the literature, where Ω_n is not explicitly defined, one can think of it as the uniform distribution on $[n]_0$.

which, in turn, sends challenges to P sampled according to Ω_n . P then replies by sending, among other things, commitment openings to these challenges. Finally V accepts or rejects. Such an interactive protocol can be made non-interactive in the ROM by applying the Fiat-Shamir transform [FS87].

Completeness guarantees that honest provers make the verifier accept. For $\alpha \in (0, 1]$, (α, ϵ) -knowledge soundness guarantees that from any PoSW successful prover, except with probability ϵ , one can extract a sequentially labeled path P in G_n whose vertices have total weight $\geq \alpha$. For uniform Ω_n , this translates to extracting a sequential path of length $\geq \alpha \cdot n$. Succinctness guarantees that proofs are small and verify fast.

We view blockchains as labeled *chain graphs*, where a chain graph $H_n = ([n]_0, E_H)$ is a DAG with edge set $E_H \supseteq \{(i-1, i) : i \in [n]\}$, i.e., the edge set contains the line graph with potentially more edges. Hence, we use the terms block and label interchangeably. Furthermore, we assume that a blockchain is equipped with a polynomial-time relation R that checks validity, i.e., on input a label/block and the labels of its parents, R outputs a bit indicating validity of the block. For example, the edge set of the underlying chain graph of fixed-difficulty Bitcoin is simply $\{(i-1, i) : i \in [n]\}$, and R takes the i th label and the label of its single parent $i-1$, and outputs 1 if and only if the proof of work checks out. The generality of the notion of a chain graph, which allows for extra edges, is well suited for other blockchains for which the validity of a block may require checking many parent blocks, rather than the immediate parent.

The PoSW-enabled SNACK construction of [AFGK22] works by first defining an *augmented* chain graph $K_n := ([n]_0, E_G \cup E_H)$ where E_G is the edge sets of the underlying chain graph G_n of a PoSW scheme (P, V) , and E_H is the edge set of the chain graph underlying the blockchain in question. An illustration of example graphs is given in Fig. 1. Then, an *augmented mining* algorithm that labels K_n by infusing PoSW-related data into blockchain labels is defined. The validity of labels in K_n is now an augmented validity relation \tilde{R} , which in addition to using R to check blockchain block validity, also checks the consistency of the added PoSW data. The augmentation is carefully done such that running (the Fiat-Shamir-transformed non-interactive counterpart of) (P, V) over the augmented labeled K_n gives rise to a SNACK system for the chain commitment language $\mathcal{L}_{\Gamma, \tilde{R}, \text{Com}}$. Knowledge-soundness of the SNACK follows from that of the PoSW. As we rely on graph-labeling PoSW schemes in the random oracle model, our SNACKs are also secure in the random oracle model.

For blockchain augmentation to work, the blockchain protocol needs to allow it, i.e., the blockchain uses an augmented mining protocol. The cost of such augmentation is minimal: the edge structure of the blockchain changes from E_H to E_K , and each augmented block contains an extra λ -bit hash value, say a 256 bit value. The benefit of this augmentation is that a SNACK system for the augmented blockchain is as efficient as its underlying PoSW scheme. To put things in perspective, a full node miner holding the augmented blockchain generates a SNACK proof by simply looking up a poly-logarithmic number of blocks and sends them to the verifier, which in turn, simply checks their consistency

It is instructive to stress that the SNACK construction works for *any* graph-labeling PoSW scheme whose underlying graph is a chain graph. Hence, if we instantiate the SNACK with the skiplist PoSW from [AFGK22], and the fixed-difficulty Bitcoin whose chain graph has edge set $\{(i-1, i) : i \in [n]\}$, then the SNACK proof size is $O(t \cdot \log^2 n)$ blocks, where t is a security parameter. If we use a variant of the Cohen-Pietrzak [CP18] PoSW as given in [AFGK22], then the proof size drops to $O(t \cdot \log n)$ blocks. In this work, we opt for the skiplist PoSW in our illustrations simply because of its symmetric graph structure.

1.2 Proof of Space Blockchains

We are aware of two PoSpace blockchains: SpaceMint [PKF⁺18] and Chia [CP19] and our treatment covers them both.

Unlike blockchains based on either proofs of work (PoW) or proofs of stake (PoS), proofs of space (PoSpace) based blockchains are composed of two chains: a *canonical* proof chain and a data chain. The proof chain contains unique proofs and hence is canonical. The data chain contains transactions and any arbitrary data that the blockchain permits. The data chain is bound to the proof chain by means of digital signatures.

Without loss of generality, we can view a *PoSpace blockchain* as a tuple of labeled chains whose underlying DAG is $B_n = (C_n, D_n)$ where $C_n = ([n]_0, E_C)$ and $D_n = ([n]_0, E_D)$ are the chain graphs underlying the canonical proof and data chains, respectively. Both C_n and D_n are chain graphs, and we stress that E_C and E_D need not be equal. Furthermore, C_n is bound to D_n by a digital signature scheme $\text{SIG} = (\text{Gen}, \text{Sign}, \text{Vrfy})$ in a simple manner that we explain shortly below.

Example chains graphs for C_n, D_n , and other graphs that we will shortly be discussing are shown in Fig. 2.

Let $(b_i := (c_i, d_i))_{i \in [n]_0}$ denote the labels of these chains, where c_i and d_i denote the i th labels of the canonical and data blocks, respectively.




Fig. 1. Example graphs H_8, G_8, K_8 , where G_8 is the chain graph underlying the PoSW scheme of [AFGK22]. The labeling $(k_i)_{i \in [8]_0}$ of K_8 is computed such that $k_i = (g_i, h_i)$ where g_i is PoSW-related and essentially contains a hash of all parent labels in K_8 , and h_i is a standard blockchain block, where additionally the publicly verifiable proof in it, say PoW in Bitcoin, must depend on all parent labels in K_8 as well as g_i . The augmented blockchain is now the labeled K_8 , and therefore, contains in its i th block, compared to the initial non-augmented blockchain, an additional PoSW-related data g_i , which is typically small, say 256 bits. At this price of augmentation, one can get SNACKs as efficient as PoSW schemes.




Fig. 2. Example DAGs $C_8 = ([8]_0, E_C)$ and $D_8 = ([8]_0, E_D)$.

Although our treatment allows for arbitrary labeling, c_i and d_i , for simplicity of exposition, can be assumed to have the following format (which is faithful to existing PoSpace blockchains):

- $c_i = (i, \pi_i)$ where π_i is a canonical computation that depends on the labels of parent proof blocks $(c_{i_1}, \dots, c_{i_q})$ where $(i_1, \dots, i_q) = \text{parents}_C(i)$. For simplicity, we assume that $\pi = (\delta_i, (\text{VDF}_v^i, \text{VDF}_p^i))$ where δ_i is a proof of space [DFKP15, AAC⁺17] and $(\text{VDF}_v^i, \text{VDF}_p^i)$ is a verifiable delay function [BBBF18] computation/proof pair.
- $d_i = (s_i, \text{data}_i)$ where $s_i \leftarrow \text{Sign}_{\text{sk}}(d_{i_1} \parallel \dots \parallel d_{i_p} \parallel \text{data}_i \parallel c_i)$ is a signature on the parents data blocks $d_{i_1} \parallel \dots \parallel d_{i_p}$ where $(i_1, \dots, i_p) = \text{parents}_D(i)$, the current data data_i , and the current proof block c_i .




Fig. 3. Example of an insecure design: the DAG $U_8 = ([8]_0, E_U := E_C \cup E_D)$ is the union DAG of $C_8 = ([8]_0, E_C)$ and $D_8 = ([8]_0, E_D)$ from Fig. 2, and K_8 augments U_8 by the PoSW G_8 from Fig. 1.

PoSW-Enabled SNACKs for PoSpace Blockchains. Constructing SNACKs for PoSpace blockchains is more subtle than for PoW or PoS blockchains due to the requirement that c_i must be canonical, a condition which is crucial for the security of PoSpace blockchains [PKF⁺18, CP19]. So, simply viewing a PoSpace blockchain as a *single* chain graph, call it U_n , which is the union of both C_n and D_n , i.e., $U_n = ([n]_0, E_U := E_C \cup E_D)$, and then applying the generic SNACK construction outlined above to U_n , would not work. See Fig. 3 for illustrative graphs of this construction.

The reason this natural construction doesn't work is that it violates the canonical nature of the proof chain: in this case, the augmented chain graph is $K_n = ([n]_0, E_K := E_U \cup E_G)$ and its i th augmented label is $k_i = (g_i, u_i)$, where g_i is defined by the underlying PoSW scheme, and depends on all parent labels k_j in K_n , and $u_i := (c_i, d_i)$, where c_i is the proof block and d_i the data block. For SNACK's security [AFGK22], the publicly verifiable proofs in c_i must depend on g_i , which itself depends on (grindable) transaction data contained in d_i , and hence the canonical nature of the proof chain would be lost.

The insecurity of this SNACK construction suggests the following alternative idea: independently of D_n , construct a SNACK generically over C_n which maintains the canonical nature of C_n . However, unless another SNACK over D_n is executed, no sequentiality guarantees can be concluded on the data chain. And by symmetry, having a SNACK on D_n alone doesn't imply sequentiality on C_n . Therefore, to ensure sequentiality on the (labeled) PoSpace blockchain, one has to ensure sequentiality on both (labeled) C_n and D_n simultaneously.

Our SNACK. We construct a SNACK for $B_n = (C_n, D_n)$, by constructing two SNACKs *simultaneously*, one for C_n , call it CS, and one for the

data chain D_n , call it DS. Both CS and DS are generic PoSW-enabled constructions following the blueprint of [AFGK22]. They also satisfy:

1. CS and DS both use the same underlying PoSW scheme PoSW
2. PoSW uses a *deterministic* vector commitment Com, and
3. CS is embedded into DS.

The final PoSpace SNACK is simply DS. The soundness guarantees of DS is that from any convincing prover, we can extract an α -valid path (P, L_P) where $L_P = ((c_{i_1}, d_{i_1}), \dots, (c_{i_k}, d_{i_m}))$ is such that (c_{i_j}, d_{i_j}) is a valid blockchain block and that $(c_{i_1}, \dots, c_{i_m})$ and $(d_{i_1}, \dots, d_{i_m})$ are both sequentially computed. These are the guarantees that a SNACK should provide for a blockchain: sequentiality of its blocks.

Let's justify the design choice made above. Note that assuming the same PoSW chain graph in CS and DS simplifies the final construction, and requiring Com to be deterministic is necessary to preserve the canonical nature of the augmented proof chain.⁵ To see the necessity of embedding CS into DS, let's see what guarantees one would get from these SNACKs individually, and why these guarantees falls short of our goal of ensuring the sequentiality of the combined PoSpace blockchain blocks.

From α -valid paths (P_c, L_{P_c}) and (P_d, L_{P_d}) extracted from CS and DS respectively, we would like to construct an α -valid path (P, L_P) as above. However, as it may be the case that $P_c \neq P_d$, i.e., P_c, P_d may not coincide, constructing (P, L_P) with weight α out of (P_c, L_{P_c}) and (P_d, L_{P_d}) may not be possible.

A natural first idea towards ensuring $P_c = P_d$ would be to fix the same PoSW scheme in both CS and DS. That is, we augment both C_n and D_n with the same PoSW chain graph G_n to arrive at augmented chain graphs K_n^c and K_n^d , respectively. However, this doesn't mean that $K_n^c = K_n^d$ as C_n need not be equal to D_n , and hence, the extracted paths may be such that $P_c \neq P_d$. But as we will show in Lemma 2 below, in any PoSW-enabled SNACK, over an augmented chain graph, say $K_n^c = ([n]_0, E_K = E_C \cup E_G)$, any *extractable* α -valid path (P_c, L_{P_c}) is such that P_c lies in G_n . Still, that P_c and P_d lie in G_n doesn't mean they coincide, but now we are a step closer towards ensuring they do.

To be able to compose an α -valid path (P, L_P) from α -valid paths (P_c, L_{P_c}) and (P_d, L_{P_d}) , not only we want to ensure that $P_c = P_d$, but also that their labels are valid and bound, i.e., let $L_{P_c} = (c_{i_1}, \dots, c_{i_m})$ and $L_{P_d} = (d_{i_1}, \dots, d_{i_m})$, then it must hold that (c_{i_j}, d_{i_j}) is a valid blockchain

⁵ We remark that all PoSW schemes in the ROM [MMV13, CP18, AKK⁺19, DLM19, AFGK22, AC23] use deterministic Com anyway.

block including that d_{i_j} contains a signature on c_{i_j} . Recall that SIG binds C_n to D_n . Now because c_{i_j} is needed to validate d_{i_j} , DS can't simply be independent of CS.

To resolve all issues at once, that is, to make sure $P_c = P_d$ and that $L_P := ((c_{i_1}, d_{i_1}), \dots, (c_{i_m}, d_{i_m}))$ is valid, where $P := P_c = P_d$, we require that CS and DS use the *same* underlying PoSW scheme, and furthermore, *embed* CS into DS. By embedding the augmented labeled proof chain K_n^c into the augmented labeled data chain K_n^d , and relying on Lemma 2 below, we ensure that the same labeled path in K_n^d contains a valid labeling in K_n^c at the same time.

The cost of our construction is the same as that of single-chain blockchains [AFGK22], as recalled in Sect. 1.1, except that we store in the augmented blockchain two, instead of one, PoSW-related data: one in each augmented proof chain block and one in each augmented data chain block. Each PoSW-related data is a λ -bit hash value, say 256 bits.

In Sect. 3.3, we present the PoSpace SNACK construction and prove its security.

2 Preliminaries

In this section, we review the SNACK-related definitions from [AFGK22].

Notation. For a directed acyclic graph (DAG) $G = (V, E)$ on $n+1$ vertices, we always number its vertices $V = [n]_0$ in topological order and often write G_n to make this explicit. For $v \in [n]_0$, we denote the parent vertices of v in G by $\text{parents}_G(v)$, and their number (i.e., the indegree of v) by $\text{deg}_G(v)$; thus, $\text{parents}_G(v) = (v_1, \dots, v_{\text{deg}_G(v)})$. We also let $\text{deg}(G) := \max_{v \in V} \{\text{deg}_G(v)\}$. We drop the subscript G when it's clear from context.

Graph labeling. A *chain graph* is a DAG on $[n]_0$ vertices such that its edge set E contains a path $P := (0, \dots, n)$ which goes through all $[n]_0$.

Definition 1 (Chain graphs). A DAG $G_n = ([n]_0, E_n)$ is a chain graph if $E_n \supseteq \{(i-1, i) : i \in [n]\}$.

A DAG is weighted if its vertices have arbitrary weights that sum to 1. In SNACK constructions, the verifier's challenges are sampled according to the distribution induced by the weights of the underlying DAG.

Definition 2 (Weighted DAGs). We call $\Gamma_n = (G_n, \Omega_n)$ a weighted DAG if $G_n = ([n]_0, E_n)$ is a DAG and $\Omega_n : [n]_0 \rightarrow [0, 1]$ is a function s.t. $\Omega_n([n]_0) = 1$, where for $S \subseteq [n]_0$, $\Omega_n(S) := \sum_{s \in S} \Omega_n(s)$.

SNACK constructions are over labeled chain graphs. An augmented data corresponding to arbitrary blockchain data is infused into the random-oracle-based labeling of chain graphs that underlie SNACKs.

Definition 3 (Oracle-based graph labeling). Let $G_n = ([n]_0, E_n)$ be a DAG and $\tau = (\tau_i)_{i \in [n]_0}$ be a tuple of oracles, with each $\tau_i: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$. For any $X = (x_0, \dots, x_n) \in (\{0, 1\}^*)^{n+1}$ the X -augmented τ -labeling $L^\tau: [n]_0 \rightarrow \{0, 1\}^*$ of G_n is recursively defined as

$$L^\tau(i) := \begin{cases} \tau_i(\varepsilon) \| x_i & \text{if } \text{parents}(i) = \emptyset, \\ \tau_i(L^\tau(\text{parents}(i))) \| x_i & \text{otherwise,} \end{cases} \quad (1)$$

where $L^\tau(\text{parents}(i)) := L^\tau(i_1) \| \dots \| L^\tau(i_k)$ for $(i_1, \dots, i_k) := \text{parents}(i)$. If $X = (\varepsilon, \dots, \varepsilon)$, we call L^τ the τ -labeling of G_n .

SNACKs. A *valid path* is a labeled path whose labels are locally valid according to some poly-time relation R and globally consistent as in (2).

Definition 4 (Valid paths). Let $G_n = ([n]_0, E_n)$ be a DAG, and $R \subseteq \mathbb{N}_0 \times (\{0, 1\}^*)^2$ a relation. Furthermore, let P be a path in G_n , L_P a labeling of P , and $(p_v)_{v \in P} \in (\{0, 1\}^*)^{|P|}$ a $|P|$ -tuple of bitstrings with $p_v = (p_v[1], \dots, p_v[\text{deg}(v)])$. We say that $(P, L_P, (p_v)_{v \in P})$ is an R -valid path in G_n if $\forall v \in P$ with $(v_1, \dots, v_{\text{deg}(v)}) := \text{parents}(v)$, we have

$$R(v, L_P(v), p_v) = 1 \text{ and } \forall i \in [\text{deg}(v)] \text{ if } v_i \in P \text{ then } p_v[i] = L_P(v_i). \quad (2)$$

For a weighted DAG $\Gamma_n = (G_n = ([n]_0, E_n), \Omega_n)$, we say $(P, L_P, (p_v)_{v \in P})$ is (α, R) -valid in Γ_n if in addition $\Omega_n(P) \geq \alpha$.

The language over which SNACKs are defined $\mathcal{L}_{\Gamma, R, \text{Com}}$ is defined via a parameter-dependent ternary polynomial-time (PT) relation $\mathcal{R}_{\Gamma, R, \text{Com}}$ over tuples (prm, η, w) , where prm is generated by a parameter generation G algorithm. A statement $\eta = (\phi, n)$ in $\mathcal{L}_{\Gamma, R, \text{Com}}$ consists of a position-binding Com commitment ϕ to an R -valid labeling of the graph $\Gamma_n \in \Gamma$. The labeling together with an opening of ϕ constitutes a witness w for η .

Definition 5 (Chain commitment language). Let $\Gamma = (\Gamma_n)_{n \geq 0}$ be a family of weighted DAGs and Com a vector commitment scheme, define

$$\mathcal{R}_{\Gamma, R, \text{Com}}^{(\alpha)} := \left\{ \begin{array}{l} (\text{prm}, \eta = (\phi, n), \\ w = (P, L_P, (p_i)_{i \in P}, \rho)) \end{array} : \begin{array}{l} (P, L_P, (p_i)_{i \in P}) \text{ is } (\alpha, R)\text{-valid in } \\ \Gamma_n \wedge \text{Com.ver}(\text{pp}, \phi, L_P, P, \rho) = 1 \end{array} \right\} \quad (3)$$

where $R \subseteq \mathbb{N}_0 \times (\{0, 1\}^*)^2$ is a PT relation that depends on prm . We let $\mathcal{R}_{\Gamma, R, \text{Com}} := \mathcal{R}_{\Gamma, R, \text{Com}}^{(1)}$ and $\mathcal{L}_{\Gamma, R, \text{Com}}$ denote the language defined by $\mathcal{R}_{\Gamma, R, \text{Com}}$.

A SNACK system (P, V) for $\mathcal{L}_{\Gamma, R, \text{Com}}$ in a non-interactive argument system satisfying completeness, (α, ϵ) -knowledge soundness, and succinctness. Completeness guarantees that an honest P holding a witness for a statement $(\phi, n) \in \mathcal{L}_{\Gamma, R, \text{Com}}$ makes V accept. For an $\alpha \in (0, 1]$, (α, ϵ) -knowledge-soundness guarantees that from any convincing prover for a statement (ϕ, n) one can extract, except with probability ϵ , an R -valid labeling of a path P in Γ_n of weight at least α . In our SNACK constructions, due to the use of random-oracle graph-labeling, the labels of R will be guaranteed to be computed *sequentially*. Succinctness requires that the proof size as well as its verification time are poly-logarithmic in n and polynomial in the security parameter.

Definition 6 (SNACK). *A tuple of PPT algorithms (P, V) is a succinct non-interactive argument of chain knowledge (SNACK) for $\mathcal{L}_{\Gamma, R, \text{Com}}$ with parameter generator G from Def. 5 if the following properties hold:*

Completeness: $\forall \lambda \in \mathbb{N}, \text{prm} \leftarrow G(1^\lambda), \eta, w \in \{0, 1\}^*$ with $(\text{prm}, \eta, w) \in \mathcal{R}_{\Gamma, R, \text{Com}}$, we have $\Pr[\pi \leftarrow P(\text{prm}, \eta, w) : V(\text{prm}, \eta, \pi) = 1] = 1$.

(α, ϵ) -Knowledge soundness: For every PPT prover \tilde{P} there exists a PPT extractor E such that

$$\Pr \left[\begin{array}{l} \text{prm} \leftarrow G(1^\lambda); r \xleftarrow{\$} \{0, 1\}^{\text{poly}(\lambda)} \\ (\eta, \pi) := \tilde{P}(\text{prm}; r) \\ w' \leftarrow E(\text{prm}, r) \end{array} : \begin{array}{l} V(\text{prm}, \eta, \pi) = 1 \wedge \\ \mathcal{R}_{\Gamma, R, \text{Com}}^{(\alpha)}(\text{prm}, \eta, w') = 0 \end{array} \right] \leq \epsilon(\lambda), \quad (4)$$

with $\mathcal{R}_{\Gamma, R, \text{Com}}^{(\alpha)}$ from (3).

Succinctness: For all $\text{prm} \leftarrow G(1^\lambda), (\text{prm}, \eta, w) \in \mathcal{R}_{\Gamma, R, \text{Com}}$ and $\pi \leftarrow P(n, t\eta, w)$, we have $|\pi| \leq \text{poly}(\lambda, \log n)$, P runs in time $\text{poly}(\lambda, n)$, and V runs in time $\text{poly}(\lambda, \log n)$.

Our SNACK construction relies on augmented graph-labeling PoSWs. Informally, an (augmented graph-labeling) PoSW scheme $\text{PoSW} := (P := (P.\text{label}, P.\text{open}), V)$ is an (interactive) proof system in which P , on common input a weighted DAG (G_n, Ω_n) , and a statement χ sampled by V , computes a proof that convinces the verifier that a certain number of *sequential* computational steps with weight 1 according to Ω_n have been performed since χ was received. In particular, $P.\text{label}$ computes an augmented τ -labeling L of G_n and sends a vector commitment of L to V , which sends challenges to P , where $P.\text{open}$ replies by giving, among other things, commitment openings to these challenges. Finally V accepts or rejects. (α, ϵ) -knowledge soundness of PoSW guarantees that from any

convincing prover, a sequentially-labeled path of total weight $\geq \alpha$, can be extracted except with probability ϵ .

The formal definition of augmented graph-labeling PoSW schemes utilizes the notion of sequential weights of oracle queries.

Definition 7 (Sequential weight). *Let $Q = (Q_1, \dots, Q_\ell)$ be a sequence of parallel queries to an oracle $\tau = (\tau_i)_{i \in [n]_0}$. We define the sequential weight of Q with respect to a weight function $\Omega_n: [n]_0 \rightarrow [0, 1]$ as*

$$\Omega_{\text{seq}}(Q) := \sum_{i=1}^{\ell} \max \{ \Omega_n(j) : Q_i \text{ contains a query to } \tau_j \} .$$

Note that if Ω_n is uniform, i.e., $\forall i \in [n]_0 : \Omega_n(i) = \frac{1}{n+1}$, then $\Omega_{\text{seq}}(Q) = \frac{\ell}{n+1}$.

Definition 8 (Augmented Graph-Labeling PoSW). *Let $\Gamma = (\Gamma_n = (G_n, \Omega_n))_{n \in \mathbb{N}}$ be a family of weighted DAGs such that for all n , G_n has a unique sink n . A pair of PPT algorithms $(\mathbf{P} := (\mathbf{P}_0, \mathbf{P}_1), \mathbf{V} := (\mathbf{V}_0, \mathbf{V}_1, \mathbf{V}_2))$, with access to an oracle $\tau = (\tau_i)_{i \in \mathbb{N}_0}$ is an augmented (oracle-based) graph-labeling proof of sequential work (GL-PoSW) if it instantiates the template described in Fig. 4 by specifying a vector commitment scheme $\text{Com} = (\text{setup}, \text{commit}, \text{open}, \text{ver})$ and subroutines PoSW.label , PoSW.open and PoSW.ver ; and it satisfies the following properties:*

Completeness: *For all $n, \lambda \in \mathbb{N}$ it holds that*

$$\Pr [(\text{out}_{\mathbf{P}}, \text{out}_{\mathbf{V}}) \leftarrow \langle \mathbf{P}(1^n) \leftrightarrow \mathbf{V}(1^\lambda, n) \rangle : \text{out}_{\mathbf{V}} = 1] = 1 .$$

(α, ϵ) -Soundness: *For all $\lambda \in \mathbb{N}$ and every PPT adversary $(\tilde{\mathbf{P}}', \tilde{\mathbf{P}} = (\tilde{\mathbf{P}}_0, \tilde{\mathbf{P}}_1))$ s.t. $\tilde{\mathbf{P}}$ makes a sequence Q of parallel queries to $\tau = (\tau_j(\cdot))_{j \in [n]_0}$ of sequential weight $\Omega_{\text{seq}}(Q) < \alpha$, it holds that*

$$\Pr \left[\begin{array}{l} (n, \text{st}) \leftarrow \tilde{\mathbf{P}}'(1^\lambda) \\ (\text{out}_{\tilde{\mathbf{P}}}, \text{out}_{\mathbf{V}}) \leftarrow \langle \tilde{\mathbf{P}}(\text{st}) \leftrightarrow \mathbf{V}(1^\lambda, n) \rangle : \text{out}_{\mathbf{V}} = 1 \end{array} \right] \leq \epsilon(\lambda) .$$

Succinctness: *The size of the transcript $|\langle \mathbf{P}(1^n) \leftrightarrow \mathbf{V}(1^\lambda, n) \rangle|$ as a function of λ and n is upper-bounded by $\text{poly}(\lambda, \log n)$. The running time of \mathbf{P} is $\text{poly}(\lambda, n)$ and that of \mathbf{V} is $\text{poly}(\lambda, \log n)$.*

We say that (\mathbf{P}, \mathbf{V}) is (α, ϵ) -knowledge-sound we additionally have:

Verifier $V = (V_0, V_1, V_2)$:	Prover $P = (P_0, P_1)$:
<p>Stage V_0: On input $(1^\lambda, n)$:</p> <ol style="list-style-type: none"> 1. $\chi \xleftarrow{\\$} \{0, 1\}^\lambda$ 2. $\text{pp} \leftarrow \text{Com.setup}(1^\lambda)$ 3. send $\text{prm} := (\chi, \text{pp})$ to P_0 <p>Stage V_1: On input ϕ_L:</p> <ol style="list-style-type: none"> 1. $\forall i \in [t]$ do $\iota_i \xleftarrow{\\$} \Omega_n$ 2. send $\iota = (\iota_i)_{i=1}^t$ to P_1 <p>Stage V_2: On input $(\gamma_i = (o_i, \rho_i))_{i=1}^t$:</p> <ol style="list-style-type: none"> 1. $\forall i \in [t]$ do <ol style="list-style-type: none"> (a) $b_i^{(1)} := \text{PoSW.ver}(\chi, \iota_i, o_i)$ (b) $b_i^{(2)} := \text{Com.ver}(\text{pp}, \phi_L, L(\iota_i), \iota_i, \rho_i)$ 2. output $\bigwedge_{i=1}^t (b_i^{(1)} \wedge b_i^{(2)})$ 	<p>Stage P_0: On input 1^n and $\text{prm} := (\chi, \text{pp})$:</p> <ol style="list-style-type: none"> 1. $L := \text{PoSW.label}(\chi, 1^n)$ <i>Use χ to sample oracles $\tau := (\tau_i(\cdot))_{i \in [n]_0}$ and compute a (possibly augmented) τ-labeling L of G_n satisfying Def. 3.</i> 2. $(\phi_L, \text{aux}) \leftarrow \text{Com.commit}(\text{pp}, L)$ 3. send ϕ_L to V_1 <p>Stage P_1: On input $\iota = (\iota_i)_{i=1}^t$:</p> <ol style="list-style-type: none"> 1. $\forall i \in [t]$ do <ol style="list-style-type: none"> (a) $o_i \leftarrow \text{PoSW.open}(\chi, \text{pp}, \phi_L, \text{aux}, L, \iota_i)$ <i>We assume that o_i contains $L(\iota_i)$</i> (b) $\rho_i \leftarrow \text{Com.open}(\text{pp}, \phi_L, \text{aux}, L(\iota_i), \iota_i)$ (c) $\gamma_i := (o_i, \rho_i)$ 2. send $(\gamma_1, \dots, \gamma_t)$ to V_2

Fig. 4. The template of a GL-PoSW, parametrized by a family of weighted DAGs $(\Gamma_n)_{n \in \mathbb{N}}$, a vector commitment scheme Com and the number of challenges t .

(α, ϵ) -Knowledge soundness: *There exists a PPT extractor E such that for every PPT adversary $(\tilde{P}', \tilde{P} = (\tilde{P}_0, \tilde{P}_1))$ we have*

$$\Pr \left[\begin{array}{l} r \xleftarrow{\$} \{0, 1\}^{\text{poly}(\lambda)}; \\ (n, \text{st}) := \tilde{P}'(1^\lambda; r); \\ (\text{out}_{\tilde{P}}, \text{out}_V) \leftarrow \langle \tilde{P}(\text{st}; r) \leftrightarrow V(n) \rangle(1^\lambda); \\ w' \leftarrow E^{\tilde{P}}(1^\lambda, r) \end{array} : \begin{array}{l} \text{out}_V = 1 \wedge \\ \mathcal{R}_{\Gamma, R, \text{Com}}^{(\alpha)}(\text{prm}, \\ (\text{out}_{\tilde{P}_0}, n), w') = 0 \end{array} \right] \leq \epsilon(\lambda),$$

where prm is as sampled by V_0 , $\text{out}_{\tilde{P}_0}$ is the output ϕ_L of \tilde{P}_0 and relation $\mathcal{R}_{\Gamma, R, \text{Com}}^{(\alpha)}$ is as in (3) with $R := R_\chi$ defined as

$$R(i, L(i), p_i) = 1 \quad \text{iff} \quad L(i) = \tau_i(p_i) \parallel x_i \text{ for some } x_i \in \{0, 1\}^*. \quad (5)$$

The connection between SNACKs and PoSW schemes is captured by the following lemma.

Lemma 1 ([AFGK22]). *Let Π be an (interactive) (α, ϵ) -knowledge-sound GL-PoSW based on a family of weighted DAGs Γ and a commitment scheme Com . Then applying the Fiat-Shamir [FS87] transformation Π results in an (α, ϵ) -knowledge-sound SNACK system for $\mathcal{R}_{\Gamma, R, \text{Com}}^{(\alpha)}$ when R is defined as in (5).*

3 SNACKs for Proof-of-Space Blockchains

We extend the PoSW-enabled SNACK construction of [AFGK22] to the context of PoSpace blockchains. Our construction in a nutshell follows the simple outline of (1) defining an appropriate DAG (2) labeling it and (3) running a PoSW-like protocol over it.

3.1 Proof-of-Space Blockchains

We are aware of two PoSpace blockchains: SpaceMint [PKF⁺18] and Chia [CP19] and our treatment covers them both.

Unlike blockchains based on either proofs of work (PoW) or proofs of stake (PoS), proofs of space (PoSpace) based blockchains are composed of two chains: a *canonical* proof chain and a data chain. The proof chain contains unique proofs and hence is canonical. The data chain contains transactions and any arbitrary data that the blockchain permits. The data chain is bound to the proof chain by means of digital signatures.

Without loss of generality we can view a *PoSpace blockchain* as a tuple of labeled chains whose underlying DAG is $B_n = (C_n, D_n)$ where $C_n = ([n]_0, E_C)$ and $D_n = ([n]_0, E_D)$ are the chain graphs underlying the canonical proof and data chains, respectively. Both C_n and D_n are chain graphs in the sense of Def. 1, and we stress that E_C and E_D need not be equal. Furthermore, C_n is bound to D_n by a digital signature scheme $\text{SIG} = (\text{Gen}, \text{Sign}, \text{Vrfy})$ in a simple manner that we explain shortly below. Example chain graphs for C_n, D_n are shown in Fig. 2.

We view *blockchain mining* as the process of labeling the vertices of these chains. We let $(b_i := (c_i, d_i))_{i \in [n]_0}$ denote the labels of these chains, where c_i and d_i denote the i th labels of the canonical and data blocks, respectively. Although our treatment allows for arbitrary labeling, c_i and d_i , for simplicity of exposition, can be assumed to have the following format (which is faithful to existing PoSpace blockchains):

- $c_i = (i, \pi_i)$ where π_i is a canonical computation that depends on the labels of parent proof blocks $(c_{i_1}, \dots, c_{i_q})$ where $(i_1, \dots, i_q) = \text{parents}_C(i)$. For simplicity, we assume that $\pi = (\delta_i, (\text{VDF}_v^i, \text{VDF}_p^i))$ where δ_i is a proof of space [DFKP15, AAC⁺17] and $(\text{VDF}_v^i, \text{VDF}_p^i)$ is a verifiable delay function [BBBF18] computation/proof pair.
- $d_i = (s_i, \text{data}_i)$ where $s_i \leftarrow \text{Sign}_{\text{sk}}(d_{i_1} \parallel \dots \parallel d_{i_p} \parallel \text{data}_i \parallel c_i)$ is a signature on the parents data blocks $d_{i_1} \parallel \dots \parallel d_{i_p}$ where $(i_1, \dots, i_p) = \text{parents}_D(i)$, the current data data_i , and the current proof block c_i .

These simplifying assumptions are without loss of generality. For example, in both Chia and SpaceMint, π_i contains a *unique* PoSpace δ_i . The PoSpace challenge chal_i for δ_i is uniquely determined by the labels of its parents $\text{parents}_C(i)$. The value of δ_i is defined by chal_i and the public key pk associated with the signing key sk of SIG. In Chia, π_i additionally contains $(\text{VDF}_v^i, \text{VDF}_p^i)$ where VDF_v^i is a verifiable delay function evaluation on input x_i for a time parameter t_i and VDF_p^i is a unique proof of correctness of VDF_v^i ; both x_i and t_i are uniquely defined by $\text{parents}_C(i)$.⁶

3.2 SNACKs for PoSpace Blockchains: An Overview

Constructing SNACKs for PoSpace blockchains is more subtle than for PoW blockchains, mainly due to the requirement that proof chain blocks must remain canonical. That the proof chain must be canonical (non-grindable) is crucial for the security of PoSpace blockchains [PKF⁺18, CP19]. In Sect. A, we give an intuition on why proofs in PoSpace blockchains must remain canonical.

We give here a high-level overview of the SNACK construction of [AFGK22]. Let $H_n = ([n]_0, E_H)$ be the underlying chain graph of the blockchain in question and $G_n = ([n]_0, E_G)$ the chain graph of any (graph-labeling) PoSW scheme. Then, the SNACK construction works by first defining an augmented chain graph $K_n = ([n]_0, E_K = E_H \cup E_G)$ whose i th augmented label is $k_i = (g_i, h_i)$ where g_i is defined by the underlying PoSW scheme and h_i contains the actual content of the block including the publicly verifiable (say PoW) proof π_i . The SNACK then would essentially be a non-interactive augmented PoSW on this labeled K_n . The (α, ϵ) -knowledge soundness guarantees imply that from any successful prover, we can extract, except with probability ϵ , an (α, R) -valid path $(P, L_P, (p_v)_{v \in P})$ as defined in Def. 4, such that the labels of L_P are sequentially computed and have total weight α .

For notational simplicity, we refer to (α, R) -valid paths $(P, L_P, (p_v)_{v \in P})$ by (P, L_P) , and when R is either clear from the context or irrelevant for the discussion, we call an (α, R) -valid path, α -valid.

Our SNACK. We construct a SNACK for $B_n = (C_n, D_n)$, by constructing two SNACKs *simultaneously*, one for C_n , call it CS, and one for the

⁶ In fact, in Chia, the pair $(\text{VDF}_v^i, \text{VDF}_p^i)$ is a pair of tuples, i.e., $\text{VDF}_v^i = (y_{i_1}, \dots, y_{i_k})$ and $\text{VDF}_p^i = (\rho_{i_1}, \dots, \rho_{i_k})$ where (y_{i_j}, ρ_{i_j}) is a VDF evaluation/proof pair on a challenge and time parameter pair (x_{i_j}, t_{i_j}) , which is uniquely defined by the proof chain so far (c_0, \dots, c_{i-1}) .

data chain D_n , call it DS. Both CS and DS are generic PoSW-enabled constructions following the blueprint of [AFGK22]. They also satisfy:

1. CS and DS both use PoSW, i.e., the same underlying PoSW scheme
2. PoSW uses a *deterministic* Com, and
3. CS is embedded into DS.

The final PoSpace SNACK is simply DS. The soundness guarantees of DS is that from any convincing prover, we can extract an α -valid path (P, L_P) where $L_P = ((c_{i_1}, d_{i_1}), \dots, (c_{i_k}, d_{i_m}))$ is such that (c_{i_j}, d_{i_j}) is a valid blockchain block and that $(c_{i_1}, \dots, c_{i_m})$ and $(d_{i_1}, \dots, d_{i_m})$ are both sequentially computed. These are the guarantees that a SNACK should provide for a blockchain: sequentiality of its blocks.

Let's justify the design choice made above. Note that assuming the same PoSW chain graph in CS and DS simplifies the final construction, and requiring Com to be deterministic is necessary to preserve the canonical nature of the augmented proof chain.⁷ To see the necessity of embedding CS into DS, let's see what guarantees one would get from these SNACKs individually, and why these guarantees falls short of our goal of ensuring the sequentiality of the combined PoSpace blockchain blocks.

From α -valid paths (P_c, L_{P_c}) and (P_d, L_{P_d}) extracted from CS and DS respectively, we would like to construct an α -valid path (P, L_P) as above. However, as it may be the case that $P_c \neq P_d$, i.e., P_c, P_d may not coincide, constructing (P, L_P) with weight α out of (P_c, L_{P_c}) and (P_d, L_{P_d}) may not be possible.

A natural first idea towards ensuring $P_c = P_d$ would be to fix the same PoSW scheme in both CS and DS. That is, we augment both C_n and D_n with the same PoSW chain graph G_n to arrive at augmented chain graphs K_n^c and K_n^d , respectively. However, this doesn't mean that $K_n^c = K_n^d$ as C_n need not be equal to D_n , and hence, the extracted paths may be such that $P_c \neq P_d$. But as we will show in Lemma 2 below, in any PoSW-enabled SNACK, over an augmented chain graph, say $K_n^c = ([n]_0, E_K = E_C \cup E_G)$, any *extractable* α -valid path (P_c, L_{P_c}) is such that P_c lies in G_n . Still, that P_c and P_d lie in G_n doesn't mean they coincide, but now we are a step closer towards ensuring they do.

To be able to compose an α -valid path (P, L_P) from α -valid paths (P_c, L_{P_c}) and (P_d, L_{P_d}) , not only we want to ensure that $P_c = P_d$, but also that their labels are valid and bound, i.e., let $L_{P_c} = (c_{i_1}, \dots, c_{i_m})$ and $L_{P_d} = (d_{i_1}, \dots, d_{i_m})$, then it must hold that (c_{i_j}, d_{i_j}) is a valid blockchain

⁷ We remark that all PoSW schemes in the ROM [MMV13, CP18, AKK⁺19, DLM19, AFGK22, AC23] use deterministic Com anyway.

block including that d_{i_j} contains a signature on c_{i_j} . Recall that SIG binds C_n to D_n . Now because c_{i_j} is needed to validate d_{i_j} , DS can't simply be independent of CS.

To resolve all issues at once, that is, to make sure $P_c = P_d$ and that $L_P := ((c_{i_1}, d_{i_1}), \dots, (c_{i_m}, d_{i_m}))$ is valid, where $P := P_c = P_d$, we require that CS and DS use the *same* underlying PoSW scheme, and furthermore, *embed* CS into DS. By embedding the augmented labeled proof chain K_n^c into the augmented labeled data chain K_n^d , and relying on Lemma 2 below, we ensure that the same labeled path in K_n^d contains a valid labeling in K_n^c at the same time.

3.3 SNACK for PoSpace Blockchains: The Main Construction

For simplicity, fix an integer n and let PoSW be any (graph-labeling) PoSW scheme and $\Gamma_n = (G_n = ([n]_0, E_G), \Omega_n)$ its underlying weighted chain graph, where G_n is a chain graph and $\Omega_n : [n]_0 \rightarrow [0, 1]$ s.t. $\Omega_n([n]_0) = 1$ is a weight function. We will use the PoSW from [AKK⁺19, AFGK22] in our *illustrative examples*. Its underlying DAG is depicted in Fig. 1. (We emphasize the our SNACK construction works for *any* graph-labeling PoSW scheme whose underlying graph is a chain graph.)

Furthermore, let $B_n = (C_n, D_n)$ be a PoSpace blockchain with ψ being its genesis block, and R_ψ^c and R_ψ^d be the polynomial-time validity relations for the (labeled) proof and data chains, respectively. That is, let $(i_1, \dots, i_p) := \text{parents}_C(i)$, then

$$R_\psi^c(i, c_i, (c_{i_1}, \dots, c_{i_p})) = 1 \quad (6)$$

iff c_i is a valid proof chain block, and for $(i_1, \dots, i_q) := \text{parents}_D(i)$:

$$R_\psi^d(i, (c_i, d_i := (s_i, \text{data}_i)), (d_{i_1}, \dots, d_{i_q})) = 1 \quad (7)$$

iff d_i is a valid data chain block. In particular, (7) implies that

$$\text{Vrfy}_{\text{pk}}(d_{i_1} \parallel \dots \parallel d_{i_q} \parallel \text{data}_i \parallel c_i, s_i) = 1 \quad (8)$$

The validity relation for B_n is simply the relation that checks that both (6) and (7) hold simultaneously. These validity relations are blockchain specific and they can be augmented or redefined to suite the specific instantiation of the PoSpace blockchain in question.




Fig. 5. Example K_8^c and K_8^d defined in (9) and (10) where C_n and D_n are from Fig. 2 and G_8 is from Fig. 1. CS and DS are w.r.t. K_8^c and K_8^d , respectively. The dashed arrows indicate the labels of the source are embedded into the labels of the target. The labels of these graphs are computed by SInit and SMine of Fig. 6.

Augmented Blockchains. We define *augmented chain graphs* K_n^c and K_n^d that respectively underlie CS and DS as follows:

$$K_n^c = ([n]_0, E_{K^c}) \quad \text{with} \quad E_{K^c} = E_G \cup E_C \quad (9)$$

$$K_n^d = ([n]_0, E_{K^d}) \quad \text{with} \quad E_{K^d} = E_G \cup E_D \cup E_C \quad (10)$$

Note that while K_n^c augments C_n with G_n , K_n^d augments the union of C_n and D_n with G_n . The reason for this is that we would like the DS extractor to succeed in extracting a labeled path in (the labeled) K_n^d such that it contains a labeled path in K_n^c , and for this to be possible, we make sure that (1) the (augmented) labels of K_n^c are embedded (as data) in the (augmented) labels of K_n^d and that (2) $E_{K^d} \supseteq E_{K^c}$. Examples of these graphs are depicted in Fig. 5.

Let τ be an oracle and χ a bitstring to be defined later, we define oracles $(\tau_i)_{i \in [n]_0}$ as $\tau_i(\cdot) := \tau(i, \chi, \cdot)$ and use these oracles to label our augmented blockchain.

The *augmented blockchain* is obtained by *labeling* the augmented chain graphs K_n^c and K_n^d at once, and furthermore, embedding K_n^c in K_n^d as data. This labeling is done using oracles $(\tau_i)_{i \in [n]_0}$, where $\tau_i(\cdot) := \tau(i, \chi, \cdot)$ for a random oracle τ and random bitstring χ .

This labeling is formalized by the augmented mining algorithms `Init` and `SMine`, in Fig. 6. In particular, from an initial genesis block ψ , we define in `Init`, an augmented genesis block $\sigma := L_K(0)$, which contains, in addition to ψ and pk_0 , PoSW-related data such as χ and pp .

<p>Algorithm SInit : On input 1^λ and ψ:</p> <p>Computing k_0^c:</p> <ol style="list-style-type: none"> 1. $\chi \leftarrow \{0, 1\}^\lambda$ 2. $\ell_0^c := \tau_0(\varepsilon)$ 3. $\text{pp} \leftarrow \text{Com.setup}(1^\lambda)$ 4. $(\phi_0^c, \text{aux}_0^c) := \text{Com.commit}(\text{pp}, \ell_0^c)$ 5. $g_0^c := (\ell_0^c, \phi_0^c)$ 6. $(\text{sk}_0, \text{pk}_0) \leftarrow \text{Gen}(1^\lambda)$ 7. $\pi_0 := (\psi, \chi, \text{pp}, \text{pk}_0)$ By definition $R_\psi^c(0, (g_0^c, \pi_0), \varepsilon) = 1$ 8. $c_0 := (0, \pi_0)$ 9. $L_{K^c}(0) := k_0^c := (g_0^c, c_0)$ <p>Computing k_0^d:</p> <ol style="list-style-type: none"> 1. $\ell_0^d := \tau_0(\varepsilon)$ 2. $(\phi_0^d, \text{aux}_0^d) := \text{Com.commit}(\text{pp}, \ell_0^d)$ 3. $g_0^d := (\ell_0^d, \phi_0^d)$ 4. $\text{data}_0 := (\psi, k_0^c)$ 5. $s_0 \leftarrow \text{Sign}_{\text{sk}_0}(g_0^d \parallel \text{data}_0)$ 6. $d_0 := (s_0, \text{data}_0)$ 7. $k_0^d := (g_0^d, d_0)$ 8. $L_{K^d}(0) := k_0^d$ <p>return $(\sigma := L_{K^d}(0)), \text{aux}_0^c, \text{aux}_0^d$</p>	<p>Algorithm SMine: On input $(\text{sk}, \text{pk}, \text{data}_i, (k_j^d := (g_j^d, d_j))_{j \in [i-1]_0})$:</p> <p>Parse $(k_j^c)_{j \in [i-1]_0}$ out of $(k_j^d)_{j \in [i-1]_0}$ See Line 4 in k_i^d below.</p> <p>Computing k_i^c:</p> <ol style="list-style-type: none"> 1. $\ell_i^c := \tau_i(L_{K^c}(\text{parents}_{K^c}(i)))$ 2. $(\phi_i^c, \text{aux}_i^c) := \text{Com.commit}(\text{pp}, (k_j^c)_{j \in [i-1]_0} \parallel \ell_i^c)$ 3. $g_i^c := (\ell_i^c, \phi_i^c)$ 4. Let π_i be s.t. $R_\sigma^c(i, (g_i^c, \pi_i), L_{K^c}(\text{parents}_{K^c}(i))) = 1$ π_i is associated with pk for $(\text{sk}, \text{pk}) \in [\text{Gen}(1^\lambda)]$ 5. $c_i := (i, \pi_i)$ 6. $k_i^c := (g_i^c, c_i)$ <p>Computing k_i^d:</p> <ol style="list-style-type: none"> 1. $\ell_i^d := \tau_i(L_{K^d}(\text{parents}_{K^d}(i)))$ 2. $(\phi_i^d, \text{aux}_i^d) := \text{Com.commit}(\text{pp}, (k_j^d)_{j \in [i-1]_0} \parallel \ell_i^d)$ 3. $g_i^d := (\ell_i^d, \phi_i^d)$ 4. $\text{data}_i := (\text{data}_i, k_i^c)$ 5. $s_i \leftarrow \text{Sign}_{\text{sk}}(L_{K^d}(\text{parents}_{K^d}(i)) \parallel g_i^d \parallel \text{data}_i)$ 6. $d_i := (s_i, \text{data}_i)$ 7. $L_{K^d}(i) := k_i^d := (g_i^d, d_i)$ $k_i^d = (\ell_i^d, \phi_i^d, s_i, \text{data}_i, k_i^c)$ Note: $R_\sigma^d(i, L_{K^d}(i), L_{K^d}(\text{parents}_{K^d}(i))) = 1$ <p>return $(L_{K^d}(i), \text{aux}_i^c, \text{aux}_i^d)$</p>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 6. The mining algorithms Init and SMine for PoSpace-augmented blockchains.

SMine, on input $(k_j^d := (g_j^d, d_j))_{j \in [i-1]_0}$, i.e., the augmented labels of the first i vertices of K_n^d , as well as the current data data_i including transactions and the signing/verification key pair (sk, pk) of an arbitrary space farmer, and some auxiliary information $\text{aux}_{i-1}^c, \text{aux}_{i-1}^d$ related to the commitment opening (which we explicitly state in Fig. 6, but ignore in this informal discussion for simplicity), computes the augmented labels of both the proof and data chain as follows.

As for the i th augmented proof chain label, SMine computes the PoSW label ℓ_i^c using the random oracle τ_i and the graph structure of K_n^c , computes a *deterministic* commitment ϕ_i^c of the labels $(k_j^c := (g_j^c, c_j))_{j \in [i-1]_0}$ and ℓ_i^c , and defines the label $g_i^c := (\ell_i^c, \phi_i^c)$. We stress that the commitment ϕ_i^c must be deterministic, for otherwise the proof chain becomes grindable. The label of the proof chain is defined as $c_i := (i, \pi_i)$ and the augmented label is defined as $L_{K^c}(i) := k_i^c := (g_i^c, c_i)$. For the sequentiality guarantees of the PoSW to carry on to the SNACKs, we must ensure that in the

augmented blockchain, π_i is computed *after* ϕ_i^c – see [AFGK22] for a detailed discussion on this point. However, note that Line 4 in Fig. 6 doesn't make this explicit, as ensuring this condition is blockchain-specific.

As for the i th augmented data chain label, SMine computes the PoSW label ℓ_i^d using τ_i and the graph structure of K_n^d , computes a commitment ϕ_i^d of the labels $(k_j^d := (g_j^d, d_j))_{j \in [i-1]_0}$ and ℓ_i^d , and defines the label $g_i^d := (\ell_i^d, \phi_i^d)$. A space farmer/miner who generates a PoSpace proof δ_i using pk , computes a digital signature s_i , using the corresponding signing key sk , on $(L_{K^d}(\text{parents}_D(i)) \| g_i^d \| \text{data}_i)$, where $L_{K^d}(\text{parents}_D(i))$ are the augmented labels of the parents of i in D , data_i is the current data including (the now embedded) k_i^c , transactions, and any arbitrary data that the original mining protocol allows. Finally set $d_i := (s_i, \text{data}_i)$ and $L_{K^d} := k_i^d := (g_i^d, d_i)$.

Blockchain validity must be adapted to accommodate the augmentation. Therefore, we define the *augmented validity relations* R_σ^c and R_σ^d of the proof and data chains, by overriding R_ψ^c and R_ψ^d , respectively. Both R_σ^c and R_σ^d still consider the same graph structure as R_ψ^c and R_ψ^d but expect augmented labels. Concretely, we override R_ψ^c from (6) as

$$R_\sigma^c(i, L_{K^c}(i), L_{K^c}(\text{parents}_C(i))) = 1 \quad , \quad (11)$$

iff the augmented block is valid. Note that σ is defined by Slnit and is the augmented genesis block. As before, this validity is blockchain specific. Similarly, we override R_ψ^d from (7) as

$$R_\sigma^d(i, L_{K^d}(i), L_{K^d}(\text{parents}_D(i))) = 1 \quad (12)$$

iff the augmented data block is valid, which in particular, implies that

$$\text{Vrfy}_{\text{pk}}(L_{K^d}(\text{parents}_D(i)) \| g_i^d \| \text{data}_i, s_i) = 1 \quad . \quad (13)$$

Note that R_σ^d doesn't verify transaction consistency in data_i . The consistency of data_i is assumed, i.e., we assume honest miners would not finalize block i that contains data_i that is inconsistent with $\text{data}_0, \dots, \text{data}_{i-1}$. The consistency of data is orthogonal to the SNACK construction.

As the i th augmented (proof/data) block contains, not only blockchain-specific, but also PoSW-specific data, we define augmented validity relations \tilde{R}_σ^c and \tilde{R}_σ^d that check the validity of (a) the blockchain-specific data using R_σ^c and R_σ^d , respectively, and check (b) the PoSW data. Concretely, define $\tilde{R}_\sigma^c, \tilde{R}_\sigma^d$ and R_σ :

$$\begin{aligned} \tilde{R}_\sigma^c(i, L_{K^c}(i), L_{K^c}(\text{parents}_{K^c}(i))) &= 1 \Leftrightarrow \exists x_i \text{ s.t.} & (14) \\ R_\sigma^c(i, L_{K^c}(i), L_{K^c}(\text{parents}_C(i))) &= 1 \wedge L_{K^c}(i) = \tau_i(L_{K^c}(\text{parents}_{K^c}(i))) \| x_i. \end{aligned}$$

$$\begin{aligned} \tilde{R}_\sigma^d(i, L_{K^d}(i), L_{K^d}(\text{parents}_{K^d}(i))) &= 1 \Leftrightarrow \exists x_i \text{ s.t.} & (15) \\ L_{K^d}(i) &= \tau_i(L_{K^d}(\text{parents}_{K^d}(i))) \| x_i \wedge R_\sigma^d(i, L_{K^d}(i), L_{K^d}(\text{parents}_D(i))) = 1. \end{aligned}$$

$$\begin{aligned} R_\sigma(i, L_{K^d}(i), L_{K^d}(\text{parents}_{K^d}(i))) &= 1 \Leftrightarrow & (16) \\ \tilde{R}_\sigma^c(i, L_{K^c}(i), L_{K^c}(\text{parents}_{K^c}(i))) &= \tilde{R}_\sigma^d(i, L_{K^d}(i), L_{K^d}(\text{parents}_{K^d}(i))) = 1. \end{aligned}$$

where by construction (see Fig. 6), $L_{K^c}(i)$ is contained in $L_{K^d}(i)$ as part of data_i and $\text{parents}_{K^c}(i)$ is contained in $\text{parents}_{K^d}(i)$ by definition of K_n^d .

Arguments for Augmented PoSpace Blockchains. We first define the witness relation \mathcal{R} with respect to which we construct our SNACK.

Definition 9 (Augmented PoSpace chain language). *We define the augmented PoSpace chain relation $\mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(1)}$ for any $\alpha \in (0, 1]$ as*

$$\mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(\alpha)} := \left\{ \begin{array}{l} ((\text{prm} := (\sigma, \text{pp}), \eta := (\phi^c, \phi^d, n), \\ w := (P, L_P^d, (p_i^d)_{i \in P}, \rho^c, \rho^d)) \text{ s.t.} \\ (P, L_P^d, (p_i^d)_{i \in P}) \text{ is an } (\alpha, R_\sigma)\text{-valid path in } \Gamma_n^d \\ \wedge \text{Com.ver}(\text{pp}, \phi^c, L_P^c, P, \rho^c) = 1 \\ \wedge \text{Com.ver}(\text{pp}, \phi^d, L_P^d, P, \rho^d) = 1 \end{array} \right\}. \quad (17)$$

where R_σ is defined in (16) and $\Gamma_n^d := (K_n^d, \Omega_n)$ for K_n^d as in (10) and Ω_n from the underlying PoSW scheme PoSW. We let $\mathcal{L}_{\Gamma_n^d, R_\sigma, \text{Com}}$ be the language associated with $\mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(1)}$.

Having formalized $\mathcal{L}_{\Gamma_n^d, R_\sigma, \text{Com}}$, we now build for it a succinct Argument of Chain Knowledge⁸ (ACK) system $\text{ACK} := (\text{ACK.P}, \text{ACK.V})$, which we later Fiat-Shamir to get our final DS. The ACK is given in Fig. 8. Note that the prover takes as input the labeling of K_n^d as its witness. Additionally, the SNACK parameter generator G outputs prm . Syntactically, ACK can be seen as two copies of the respective ACK construction of [AFGK22], one for the proof chain and one for the data chain. The difference is that we embed K_n^c into K_n^d , use the same underlying PoSW scheme as well as the same verifier challenges in these two copies.

⁸ A SNACK is a succinct and non-interactive ACK.

<p>Algorithm PoSW.ver_{K^c} : On input (χ, ι_i, o_i):</p> <ol style="list-style-type: none"> 1. Run $b_i := \text{PoSW.ver}(\chi, \iota_i, o_i)$ modified as follows: whenever it queries $\tau_j(L_{K^c}(\text{parents}_G(j)))$ for some j, issue query $\tau_j(L_{K^c}(\text{parents}_{K^c}(j)))$ instead. (Missing labels are provided in $o_{i,2}$.) 2. return b_i 	<p>Algorithm PoSW.open_{K^c} : On input $(\chi, \text{pp}, \phi_n, \text{aux}_n, L, \iota_i)$:</p> <ol style="list-style-type: none"> 1. $o_{i,1} \leftarrow \text{PoSW.open}(\chi, \text{pp}, \phi_n, \text{aux}_n, L, \iota_i)$ (PoSW.open acts based on edges E_G.) 2. $\mathcal{J} := \{j \in [n]_0 : L_{K^c}(\text{parents}_G(j))$ which appear in $o_{i,1}\}$ 3. $o_{i,2} := \{(j, L_{K^c}(\text{parents}_G(j)))\}_{j \in \mathcal{J}}$ 4. return $o_i := (o_{i,1}, o_{i,2})$
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 7. Algorithms PoSW.open_{K^c} and PoSW.ver_{K^c} defined based on PoSW.open and PoSW.ver, respectively. Algorithms PoSW.open_{K^d} and PoSW.ver_{K^d} are defined analogously by changing every occurrence of c to d and C to D .

Theorem 1. *Let $\text{SNACK} := (\text{SNACK.P}, \text{SNACK.V})$ be the non-interactive counterpart of ACK from Fig. 8, then in the ROM, SNACK is an (α, ϵ) -knowledge-sound SNACK for $\mathcal{L}_{\Gamma^d, R_\sigma, \text{Com}}$ as in Def. 9 if PoSW is an (α, ϵ) -knowledge-sound τ -based GL-PoSW as in [AFGK22] with Com being its underlying deterministic commitment scheme, $(G_n = ([n]_0, E_G), \Omega_n)_{n \in \mathbb{N}}$ its weighted graph family, and τ modeled as a random oracle.*

SNACK’s Cost and Guarantees. We started with an underlying chain $B_n = (C_n, D_n)$ of a PoSpace blockchain, and augmented it to (K_n^c, K_n^d) , on which we run ACK, whose non-interactive counterpart SNACK is the SNACK construction for the PoSpace blockchain. The space cost of SNACK is storing in each block in K_n^c a PoSW label and a commitment pair (ℓ_i^c, ϕ_i^c) . The same holds for K_n^d . (Note that the embedding of $L_{K^c}(i)$ into $L_{K^d}(i)$ doesn’t mean that we actually store $L_{K^c}(i)$ twice, in K_n^c and K_n^d ; all what it means is that computing and verifying $L_{K^d}(i)$ requires having $L_{K^c}(i)$ explicitly given as input.) If we instantiate PoSW from either PoSW schemes in [AFGK22], then $\phi_i^c = \ell_i^c$ and $\phi_i^d = \ell_i^d$ and hence the space cost is ℓ_i^c and ℓ_i^d per (PoSpace) block. Setting $|\ell_i^c| = |\ell_i^d| = 256$ bits is a reasonable instantiation. Furthermore, modeling a hash function as a RO, ℓ_i^c and ℓ_i^d are efficient hash computations.

Generating SNACK proofs is identical to generating PoSW proofs in the underlying PoSW scheme when $E_C = E_D = \{(i-1, i) : i \in [n]\}$. The more edges E_C and E_D contain, the bigger the proof size. Furthermore optimizations similar to those given in [AFGK22] are possible.

At this storage and computation costs, we get sequentiality guarantees on the blockchain: Fix prm and a statement $(\phi^c, \phi^d, n) \in \mathcal{L}_{\Gamma^d, R_\sigma, \text{Com}}$, then from any convincing SNACK prover, a witness w can be extracted, and such a witness contains an (α, R_σ) -valid path in Γ_n^d , which is sequentially

Verifier ACK.V = (V ₁ , V ₂)	Prover ACK.P:
<p>Stage V₁: On input $(1^\lambda, \eta)$:</p> <ol style="list-style-type: none"> 1. $\forall i \in [t]$ do $\iota_i \xleftarrow{\\$} \Omega_n$ 2. $\iota_0 := 0$ 3. send $\iota := (\iota_i)_{i \in [t]_0}$ to P <p>Stage V₂: On input $\gamma := (\gamma_i)_{i \in [t]_0}$:</p> <ol style="list-style-type: none"> 1. $\forall i \in [t]$ do: <ol style="list-style-type: none"> (a) $b_{i,1}^c := \tilde{R}_\sigma^c(\iota_i, k_{\iota_i}^c, p_{\iota_i}^c)$ (b) $b_{i,2}^c := \text{PoSW.ver}_{K^c}(\chi, \iota_i, o_i^c)$ (c) $b_{i,3}^c := \text{Com.ver}(\text{pp}, \phi_n^c, k_{\iota_i}^c, \iota_i, \rho_i^c)$ (d) $b_{i,1}^d := \tilde{R}_\sigma^d(\iota_i, k_{\iota_i}^d, p_{\iota_i}^d)$ (e) $b_{i,2}^d := \text{PoSW.ver}_{K^d}(\chi, \iota_i, o_i^d)$ (f) $b_{i,3}^d := \text{Com.ver}(\text{pp}, \phi_n^d, k_{\iota_i}^d, \iota_i, \rho_i^d)$ 2. output $\bigwedge_{i=0}^t b_{i,1}^c \wedge b_{i,2}^c \wedge b_{i,3}^c \wedge b_{i,1}^d \wedge b_{i,2}^d \wedge b_{i,3}^d$ 	<p>On input $(1^\lambda, (\eta, (k_j^d)_{j \in [n]_0}, (\text{aux}_n^c, \text{aux}_n^d), \iota))$:</p> <ol style="list-style-type: none"> 1. Parse η as $((\sigma, \text{pp}), (\phi_n^c, \phi_n^d, n))$ 2. Parse $(k_j^c)_{j \in [n]_0}$ out of $(k_j^d)_{j \in [n]_0}$ 3. $\forall i \in [t]_0$ do: <ol style="list-style-type: none"> (a) $o_i^c \leftarrow \text{PoSW.open}_{K^c}(\chi, \text{pp}, \phi_n^c, \text{aux}_n^c, (k_j^c)_{j \in [n]_0}, \iota_i)$ (b) $o_i^d \leftarrow \text{PoSW.open}_{K^d}(\chi, \text{pp}, \phi_n^d, \text{aux}_n^d, (k_j^d)_{j \in [n]_0}, \iota_i)$ <i>We assume o_i^d contains both $L_{K^d}(\iota_i)$ and $p_{\iota_i}^d := L_{K^d}(\text{parents}_{K^d}(\iota_i))$. Similarly for o_i^c.</i> (c) $\rho_i^c \leftarrow \text{Com.open}(\text{pp}, \phi_n^c, \text{aux}_n^c, k_{\iota_i}^c, \iota_i)$ (d) $\rho_i^d \leftarrow \text{Com.open}(\text{pp}, \phi_n^d, \text{aux}_n^d, k_{\iota_i}^d, \iota_i)$ (e) $\gamma_i := ((o_i^c, o_i^d), (\rho_i^c, \rho_i^d))$ 4. send $\gamma := (\gamma_i)_{i \in [t]_0}$ to V₂

Fig. 8. The interactive proof system ACK which underlies our SNACK construction.

computed. That the extracted path lies in K_n^d is clear, but that it also lies in K_n^c is not. In the sequel, we show that extracted paths must be in K_n^c as well, and hence this shows that the extracted path contains valid blocks in the combined augmented blockchain. This, in turn, allows us to talk of the PoSpace blockchain as a single sequentially mined chain.

Lemma 2 says that if P is a path extracted by the knowledge-soundness of SNACK from Theorem 1, then the edges of P lie in G_n .

Lemma 2. *Let SNACK be as in Theorem 1 and let*

$$(\sigma, \eta := (\phi^c, \phi^d, n), w := (P, L_P^d, (p_i^d)_{i \in P}, \rho^c, \rho^d)) \in \mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(\alpha)}$$

be such that w is output by the extractor guaranteed by the (α, ϵ) -knowledge soundness of SNACK, then P is a path in G_n .

The following lemma shows that SNACK contains an embedded SNACK system for $\mathcal{L}_{\Gamma_n^c, \tilde{R}_\sigma^c, \text{Com}}$, where $\mathcal{L}_{\Gamma_n^c, \tilde{R}_\sigma^c, \text{Com}}$ is as in Def. 5 and \tilde{R}_σ^c as in (14).

Lemma 3. *Let $(\text{prm}, (\phi^c, \phi^d, n), (P, L_P^d, (p_i^d)_{i \in P}, \rho^c, \rho^d)) \in \mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(\alpha)}$ and P be a path in G_n , then $(\text{prm}, (\phi^c, n), (P, L_P^c, (p_i^c)_{i \in P}, \rho^c)) \in \mathcal{R}_{\Gamma_n^c, \tilde{R}_\sigma^c, \text{Com}}^{(\alpha)}$, where $\mathcal{R}_{\Gamma_n^c, \tilde{R}_\sigma^c, \text{Com}}^{(\alpha)}$ is defined as in Def. 5 and $L^c(i), p_i^c$ are embedded in and extracted from $L^d(i), p_i^d$.*

Theorem 1 and Lemma 3 imply Corollary 1, which shows that SNACK proves knowledge of sequentially computed and valid PoSpace blockchain.

Corollary 1. *Let SNACK be as in Theorem 1 and let*

$$\left(\sigma, \eta := (\phi^c, \phi^d, n), w := (P, L_P^d, (p_i^d)_{i \in P}, \rho^c, \rho^d)\right) \in \mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(\alpha)}$$

be s.t. w is output by the extractor guaranteed by the (α, ϵ) -knowledge soundness of SNACK, then (P, L_P^d) is a sequentially-computed path containing valid PoSpace blocks, i.e., let $P = (i_0, \dots, i_k)$, then $\forall j \in [k]_0$, $L_P^d(i_j)$ contains a valid augmented data chain block which contains a signature on a valid augmented proof chain block $L_P^c(i_j)$, and for every $j \in [k]$, $L_P^c(i_j)$ is computed after $L_P^c(i_{j-1})$ and $L_P^d(i_j)$ after $L_P^d(i_{j-1})$.

Proofs of Theorem 1 and Lemmas 2 and 3 are in Sect. B.

Acknowledgements. I want to thank Karen Klein for a fruitful discussion at the beginning of the project and for her feedback on an early version of this paper. Parts of this work were done while the first author was at TUWien funded by the Vienna Science and Technology Fund (WWTF)[10.47379/VRG18002]. The project has also received funding in part from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under project PICOCRYPT (grant agreement No. 101001283), the Spanish Government under projects SCUM (ref. RTI2018-102043-B-I00), the Madrid Regional Government under project BLOQUES (ref. S2018/TCS-4339), and a research grant from Nomadic Labs and the Tezos Foundation.

References

- AAC⁺17. Hamza Abusalah, Joël Alwen, Bram Cohen, Danylo Khilko, Krzysztof Pietrzak, and Leonid Reyzin. Beyond hellman’s time-memory trade-offs with applications to proofs of space. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 357–379. Springer, Heidelberg, December 2017.
- AC23. Hamza Abusalah and Valerio Cini. An incremental posw for general weight distributions. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part II*, volume 14005 of *Lecture Notes in Computer Science*, pages 282–311. Springer, 2023.
- AFGK22. Hamza Abusalah, Georg Fuchsbauer, Peter Gaži, and Karen Klein. Snacks: Leveraging proofs of sequential work for blockchain light clients. *Cryptology ePrint Archive*, Paper 2022/240, 2022. <https://eprint.iacr.org/2022/240>.
- AKK⁺19. Hamza Abusalah, Chethan Kamath, Karen Klein, Krzysztof Pietrzak, and Michael Walter. Reversible proofs of sequential work. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 277–291. Springer, Heidelberg, May 2019.
- BBBF18. Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 757–788. Springer, Heidelberg, August 2018.
- BKLZ20. Benedikt Bünz, Lucianna Kiffer, Loi Luu, and Mahdi Zamani. Flyclient: Super-light clients for cryptocurrencies. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*, pages 928–946. IEEE, 2020.
- CP18. Bram Cohen and Krzysztof Pietrzak. Simple proofs of sequential work. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 451–467. Springer, Heidelberg, April / May 2018.
- CP19. Bram Cohen and Krzysztof Pietrzak. The chia network blockchain, July 9, 2019. <https://www.chia.net/assets/ChiaGreenPaper.pdf>.
- DFKP15. Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 585–605. Springer, Heidelberg, August 2015.
- DLM19. Nico Döttling, Russell W. F. Lai, and Giulio Malavolta. Incremental proofs of sequential work. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 292–323. Springer, Heidelberg, May 2019.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.
- MMV13. Mohammad Mahmoody, Tal Moran, and Salil P. Vadhan. Publicly verifiable proofs of sequential work. In Robert D. Kleinberg, editor, *ITCS 2013*, pages 373–388. ACM, January 2013.

PKF⁺18. Sunoo Park, Albert Kwon, Georg Fuchsbauer, Peter Gazi, Joël Alwen, and Krzysztof Pietrzak. SpaceMint: A cryptocurrency based on proofs of space. In Sarah Meiklejohn and Kazuo Sako, editors, *FC 2018*, volume 10957 of *LNCS*, pages 480–499. Springer, Heidelberg, February / March 2018.

A On the Canonical Nature of Proof Chains in Proof of Space Blockchains

We give here an intuition of why proofs must be canonical in PoSpace blockchains [PKF⁺18, CP19], and refer the reader to the respective papers for details. In a PoSpace blockchain, given a proof chain of length i , each PoSpace farmer can propose at most one valid proof label c_{i+1} for the next block, which either gets adopted or abandoned by the network of miners. This is in contrast to PoW blockchains, where a miner, given a blockchain of length i , can in principle generate infinitely many valid next blocks by tweaking transaction data and using different PoW nonces. As the mining resource in a PoW blockchain is exactly this PoW computation, and a miner’s success in appending the next valid block to the chain needs to be proportional to its mining resource, this process of generating many PoWs is the *honest* mining strategy. However, in PoSpace blockchains, the mining resource is *space*, and therefore, being able to generate more than one valid PoSpace value δ_{i+1} for the next block for a single unit of space is a *malicious* mining strategy that must be prevented. To make sure that a space miner/farmer when given a proof chain of length i can only propose at most a single next block, it must be ensured that each PoSpace farmer, for each space unit, gets a single PoSpace challenge chal_i and that all used primitives such as the PoSpace scheme itself, signatures (if used in the proof chain), VDFs, etc. are *unique*. This ensures that a miner can generate and propose at most a single block for each space unit per index i .

B Omitted Proofs

Proof (of Theorem 1). The proof strategy is simple: we use (ACK.P, ACK.V) (Fig. 8) and Alg. SMine (Fig. 6) to build an augmented PoSW $\bar{II} := (\bar{\text{PoSW.P}}, \bar{\text{PoSW.V}})$ whose knowledge-soundness implies that of (ACK.P, ACK.V). The non-interactive (SNACK.P, SNACK.V) is simply the Fiat-Shamir transformation [FS87] applied to (ACK.P, ACK.V).

Concretely, in Fig. 9 we define an augmented PoSW \bar{II} whose

1. labeling is X -augmented τ -labeling, for an X that we specify later,
2. underlying weighted graph family is $(K_n^d, \Omega_n)_{n \in \mathbb{N}}$ where K_n^d is defined in (10),
3. witness relation $\mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(1)}$ as in Def. 9 is defined with respect to a relation R_σ as defined in (16).

We will then show that $\bar{\Pi}$ is (α, ϵ) -knowledge sound if $\Pi := (\text{PoSW.P}, \text{PoSW.V})$ is (α, ϵ) -knowledge sound, i.e., we need to argue that the security guarantees of Π are preserved under the modifications in the above Items 1, 2, and 3.

As for Items 1 and 2, note that the X -augmented labeling L_{K^d} differs from the labeling L_G of $(\text{PoSW.P}, \text{PoSW.V})$ in that the $\text{parents}(\cdot)$ function used in the τ -labeling is now w.r.t. graph K_n^d rather than $G_n = ([n]_0, E_G)$. Recall that $K_n^d = ([n]_0, E_{K^d})$ has the same vertex set $[n]_0$ of G_n but differs in that its edge set is expanded as $E_{K^d} = E_G \cup E_D \cup E_C$ where E_D and E_C are the edge sets of D_n and C_n , respectively. However, when τ is modeled as a random oracle, it is easy to see that if Π is a τ -based (α, ϵ) -knowledge-sound PoSW for the weighted graph family $(G_n, \Omega_n)_{n \in \mathbb{N}}$, then $\bar{\Pi}$ is an X -augmented τ -based (α, ϵ) -knowledge-sound PoSW for the weighted graph family $(K_n^d, \Omega_n)_{n \in \mathbb{N}}$ and $X = (x_0, \dots, x_n)$, where x_i can be any bitstring, in our case $x_i = (\phi_i^d, s_i, \text{data}_i)$. Note that $k_i^d = (\ell_i^d, \phi_i^d, s_i, \text{data}_i)$ as computed by SMine (and Slnit). To argue this we need to argue that appending x_i to the random oracle τ does not affect soundness of the PoSW. This however follows because (1) τ is a random oracle and the extra input x_i simply defines an x_i -salted random oracle, and crucially (2) the new edge structure E_{K^d} does not give a malicious prover $\overline{\text{PoSW.P}}$ any more power than its counterpart PoSW.P : this is ensured by ACK.V by making sure that for each challenge $\iota_i \in [n]_0$, the PoSW-related responses of any prover are verified w.r.t. the edge structure imposed by E_G , which suffices for preserving the underlying soundness of Π . (The edge structure E_{K^d} is used to verify the validity of the augmented blockchain using R_σ .)⁹

As for Item 3, recall that the witness relation for Π is $\mathcal{R}_{\Gamma, R, \text{Com}}^{(\alpha)}$ as in Def. 5 and R as in (5), while the witness relation for $\bar{\Pi}$ is $\mathcal{R}_{\Gamma_n^d, R_\sigma, \text{Com}}^{(\alpha)}$ as in Def. 9 and R_σ is defined in (16). This imposes the extra checks in ACK.V_2 of Fig. 8: Lines 1a, 1b, 1c, and 1d. However, these extra checks don't affect soundness: if a protocol Π is sound when the verifier executes a check C_1 , then it is clearly sound if the verifier executes an additional check C_2 and accepts if and only if both checks pass. However, $\bar{\Pi}$ is not guaranteed to remain complete due to the extra checks: clearly the verifier would reject a proof π that verifies with respect to C_1 but fails with respect C_2 . To make sure that completeness is also preserved, we need to make sure that an

⁹ Note that if the PoSW-related responses are verified w.r.t. E_{K^d} , one can no longer prove soundness of the PoSW scheme as its underlying graph would then be K_n^d and now G_n . Soundness of any PoSW scheme relies on its graph structure and manipulating its structure may allow for shortcuts.

Verifier $\overline{\text{PoSW.V}} = (V_0, V_1, V_2)$: Stage V_0: On input $(1^\lambda, n, \psi)$: 1. $(\sigma, \text{aux}_0^c, \text{aux}_0^d) \leftarrow \text{Init}(1^\lambda, \psi)$ 2. send σ to P_0 Stage V_1: On input (ϕ_n^c, ϕ_n^d) : 1. $\eta := (\sigma, (\phi_n^c, \phi_n^d), n)$ 2. $\iota \leftarrow \text{ACK.V}_1(1^\lambda, \eta)$ 3. send ι to P_1 Stage V_2: On input γ : output $\text{ACK.V}_2(\gamma)$	Prover $\overline{\text{PoSW.P}} = (P_0, P_1)$: Stage P_0: On input $(1^\lambda, 1^n, (\text{sk}_i, \text{pk}_i, \text{data}_i)_{i \in [n]})$ and σ : 1. $L_{K^d}(0) := \sigma$ 2. $\forall i \in [n]$ do $(L_{K^d}(i), \text{aux}_i^c, \text{aux}_i^d) \leftarrow \text{SMine}(\text{sk}_i, \text{pk}_i, \text{data}_i, (L_{K^d}(j))_{j \in [i-1]_0})$ 3. parse (ϕ_n^c, ϕ_n^d) out of $L_{K^d}(n)$ 4. send (ϕ_n^c, ϕ_n^d) to V_1 Stage P_1: On input ι : 1. $\gamma \leftarrow \text{ACK.P}(1^\lambda, \eta, (L_{K^d}(j))_{j \in [n]_0}, \text{aux}_n^c, \text{aux}_n^d, \iota)$ 2. send γ to V_2
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 9. The augmented PoSW scheme constructed from mining algorithms and an ACK system $\text{ACK.V} = (\text{ACK.V}_1, \text{ACK.V}_2), \text{ACK.P}$.

honest augmented prover $\overline{\text{PoSW.P}}$ gets extra information that allows it to pass these extra checks: in fact, $\eta, (L_{K^d}(j))_{j \in [n]_0}, \text{aux}_n^c, \text{aux}_n^d$ provided as input to $\overline{\text{PoSW.P}}$ by the output of (the honest) SMine ensure this: to show this, we show that Line 2 of Fig. 8 holds, i.e., that we can parse $(k_j^c)_{j \in [n]_0}$ out of $(k_j^d)_{j \in [n]_0}$: it suffices to observe that, by construction (see SMine and Sinit), $\forall j \in [n]_0$, $L_{K^d}(j)$ contains $L_{K^c}(j)$, and that, by construction of (10), $E_{K^c} \subseteq E_{K^d}$. This allows $\overline{\text{PoSW.P}}$ via invoking ACK.P to run all subroutines corresponding to the extra checks.

This establishes that $\overline{\Pi}$ is an (α, ϵ) -knowledge-sound X -augmented τ -based (with labeling L_{K^d}) PoSW for the weighted graph family (K_n^d, Ω_n) . By absorbing the computation of $\overline{\text{PoSW.V}}_0$ into ACK.V_1 and $\overline{\text{PoSW.P}}_0$ into ACK.P , the pair $(\text{ACK.P}, \text{ACK.V})$ is syntactically an augmented PoSW.¹⁰

Now by Lemma 1 it holds that $(\text{SNACK.P}, \text{SNACK.V})$, the Fiat-Shamir transform of $(\text{ACK.P}, \text{ACK.V})$, is an (α, ϵ) -knowledge-sound SNACK for the language $\mathcal{L}_{\Gamma^d, R_\sigma, \text{Com}}$ as in Def. 9. \square

Proof (of Lemma 2). The SNACK construction starts with a PoSW scheme Π with an underlying weighted DAG is (G_n, Ω_n) and validity relation is R ,¹¹ and augments it into X -augmented PoSW $\overline{\Pi}$ whose underlying

¹⁰ To justify this syntactic manipulation, note that $\overline{\text{PoSW.V}}_0$ outputs σ which contains χ and pp which are generated identically to any graph-labeling PoSW scheme. Similarly we assume that the (honest) input to ACK.P was computed by an (honest) $\overline{\text{PoSW.P}}_0$, which is the computation of the honest mining Mine , which is in turn an honest PoSW computation.

¹¹ This R is the PoSW relation in (5).

weighted DAG is (K_n^d, Ω_n) where K_n^d is as in (10) and validity relation R_σ . Now the (α, ϵ) -knowledge soundness of Π implies that from any convincing prover an (α, R) -valid path in G_n can be extracted. Similarly for $\bar{\Pi}$, i.e., its (α, ϵ) -knowledge soundness implies that from any convincing prover an (α, R_σ) -valid path in K_n^d can be extracted. However, by construction, $\bar{\Pi}$ runs (see Fig. 8) the augmented PoSW procedures PoSW.open_{K^d} , PoSW.ver_{K^d} , PoSW.open_{K^c} , PoSW.ver_{K^c} .¹² Although these procedures consider augmented labelings L_{K^d} and L_{K^c} , they still, and crucially so,¹³ consider the edge structure E_G of G_n , and not E_{K^c} or E_{K^d} . Therefore, the extracted path must still lie in G_n even though its labels are augmented. This dependency on E_G is highlighted in Fig. 7, which is invoked by these augmented PoSW algorithms. This implies that the extracted (α, R_σ) -valid path in K_n^d must actually lie G_n . \square

Proof (of Lemma 3). We need to show that an (α, R_σ) -valid path in Γ_n^d contains an $(\alpha, \tilde{R}_\sigma^c)$ -valid path in Γ_n^c , where R_σ and \tilde{R}_σ^c are defined in (16) and (14), respectively. By design, $L^d(i)$ contains $L^c(i)$ as a substring – see Fig. 6. By (2), it holds that $p^d(i) = L_P^d(i)$ and therefore $p^c(i)$ is contained in $p^d(i)$ as a substring. This, together with the fact that R_σ internally checks \tilde{R}_σ^c , implies that $(P, L_P^c, (p_i^c)_{i \in P})$ is an $(\alpha, \tilde{R}_\sigma^c)$ -valid in Γ_n^c . Furthermore, by (17), it holds that $\text{Com.ver}(\text{pp}, \phi^c, L_P^c, P, \rho^c) = 1$, where pp is contained in σ . Consequently, $(\sigma, (\phi^c, n), w) \in \mathcal{R}_{\Gamma_n^c, \tilde{R}_\sigma^c, \text{Com}}^{(\alpha)}$. \square

¹² Note that these augmented PoSW algorithms work identically to their underlying PoSW.open , PoSW.ver from Π , except that the labels are augmented and some extra parent labels need to be provided by PoSW.open_{K^d} and PoSW.open_{K^c} for R_σ^c and R_σ^d to be checked – and these extra parents are defined by K^d and K^c .

¹³ See the proof of Theorem 1.