Overcoming Data Availability Attacks in Blockchain Systems: LDPC Code Design for Coded Merkle Tree

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

30-08-2021 / 01-02-2022

CITATION

Mitra, Debarnab; Tauz, Lev; Dolecek, Lara (2021): Overcoming Data Availability Attacks in Blockchain Systems: LDPC Code Design for Coded Merkle Tree. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.16532853.v2>

DOI

[10.36227/techrxiv.16532853.v2](https://doi.org/10.36227/techrxiv.16532853.v2)

Overcoming Data Availability Attacks in Blockchain Systems: LDPC Code Design for Coded Merkle Tree

Debnab Mitra, *Student Member, IEEE*, Lev Tauz, *Student Member, IEEE*
and Lara Dolecek, *Senior Member, IEEE*

Abstract

Light nodes in blockchains improve the scalability of the system by storing a small portion of the blockchain ledger. In certain blockchains, light nodes are vulnerable to a *data availability* (DA) attack where a malicious node makes the light nodes accept an invalid block by hiding the invalid portion of the block from the nodes in the system. Recently, a technique based on LDPC codes called Coded Merkle Tree (CMT) was proposed by Yu *et al.* that enables light nodes to detect a DA attack by randomly requesting/sampling portions of the block from the malicious node. However, light nodes fail to detect a DA attack with high probability if a malicious node hides a small stopping set of the LDPC code. To mitigate this problem, Yu *et al.* used random LDPC codes that achieve large minimum stopping set size with high probability. Although effective, these codes are not necessarily optimal for this application, especially at short code lengths, which are relevant for low latency systems, IoT blockchains, etc.. In this paper, we focus on short code lengths and demonstrate that a suitable co-design of specialized LDPC codes and the light node sampling strategy can improve the probability of detection of DA attacks. We consider different adversary models based on their computational capabilities of finding stopping sets in LDPC codes. For a weak adversary model, we devise a new LDPC code construction termed as the *entropy-constrained* PEG (EC-PEG) algorithm which *concentrates* stopping sets to a small group of variable nodes. We demonstrate that the EC-PEG algorithm coupled with a greedy sampling strategy improves the probability of detection of DA attacks. For stronger adversary models, we provide a co-design of a sampling strategy called *linear-programming-sampling* (LP-sampling) and an LDPC code construction called *linear-programming-constrained* PEG (LC-PEG) algorithm. The new co-design demonstrates a higher probability of detection of DA attacks compared to approaches in earlier literature.

Index Terms

Blockchain Systems, Data Availability Attacks, LDPC codes, Coded Merkle Tree

I. INTRODUCTION

Blockchains are tamper-proof ledgers of transaction data maintained by a network of nodes

D. Mitra, L. Tauz, and L. Dolecek are with the ECE Department, UCLA, Los Angeles, CA 90095 USA (e-mail: debarnabucla@ucla.edu, levtau@ucla.edu, dolecek@ee.ucla.edu). A part of this paper was presented at the IEEE Information Theory Workshop 2020 [3]. Research supported in part by the Guru Krupa Foundation and NSF-BSF grant no. 2008728.

in a decentralized manner. They were initially proposed in the field of cryptocurrencies like Bitcoin [4] and Ethereum [5]. However, the decentralized nature of blockchains has led to their application in fields such as supply chains [6], Internet of Things [7], and healthcare [8].

A blockchain is a collection of transaction blocks arranged in the form of a hash-chain. Full nodes in the blockchain network store the entire blockchain ledger and operate on it to validate transactions. However, storing the entire ledger requires a significant storage overhead¹ which prevents resource limited nodes from joining the blockchain system. To alleviate this problem, some blockchain systems also run light nodes [4], [12]. These are nodes that only store the headers corresponding to each block of the blockchain. The header for each block contains a field called a Merkle root which is constructed from the block transactions [4]. Using the Merkle root, light nodes can verify the inclusion of a given transaction in a block via a technique called a Merkle proof. However, they cannot verify the correctness of the transactions in the block.

Assuming that the system has a majority of honest full nodes, light nodes simply accept headers that are a part of the longest header chain because honest full nodes will not mine blocks on chains containing fraudulent transactions (i.e., a longest chain consensus protocol [4] is used). However, when the honest majority assumption is removed, the longest chain protocol becomes insecure for light nodes. As such, researchers were prompted to find methods to provide security even under a dishonest majority of full nodes. One such research endeavor was [1] where authors provided protocols for honest full nodes to broadcast verifiable fraud proofs of invalid transactions. The mechanism allows light nodes, even in the presence of a majority of malicious full nodes, to reject headers of invalid blocks on receiving fraud proofs from an honest full node. However, with a majority of malicious full nodes, the light nodes are still susceptible to data availability (DA) attacks [1], [2]. In this attack, as illustrated in Fig. 1 left panel, a malicious full node generates a block with invalid transactions, publishes the header of the invalid block to the light nodes, and hides the invalid portion of the block from the full nodes. Honest full nodes cannot validate the missing portion of the block and hence are unable to generate fraud proofs to be sent to the light nodes. Since the absence of a fraud proof also corresponds to the situation that the block is valid, light nodes accept the invalid header².

Light nodes can independently detect a DA attack if a request for a portion of the block is rejected by the full node that generates the block. As such, as illustrated in Fig. 1 right panel,

¹At the time of writing, the size of the Bitcoin and Ethereum ledgers are around 380GB [10] and 1150GB [11], respectively.

²In this system, there is no way of identifying honest alarm messages sent by full nodes about block unavailability [2], [13].

Fig. 1: Left: Data Availability (DA) attack; Right: Detection of DA attack via light node sampling

light nodes randomly sample the block, i.e., randomly request for different portions of the block transactions and accept the header if all the requested portions are returned. In this paper, we are interested in reducing the probability of failure for a light node to detect a DA attack for a given sample size, thus improving the security of the system. Since the size of individual transactions is much smaller compared to the entire block, an adversary can hide a very small portion of the block corresponding to the invalid transactions. Such a hiding will result in a high probability of failure for the light nodes using random sampling. To alleviate this problem, authors in [1] proposed coding the block using erasure codes³. When the block is erasure coded, to make the invalid portion of the block unavailable, the malicious block producer must prevent honest full nodes from decoding back the original block. They do so by either 1) hiding a larger portion of the coded block (more than the erasure correcting capability of the code). This hiding can be detected with a high probability by the light nodes using random sampling; 2) incorrectly generating the coded data. In this case, honest full nodes can broadcast verifiable *incorrect-coding (IC) proofs* [1], [2] allowing light nodes to reject the header. To keep the IC proof size small, authors in [1] used 2D Reed-Solomon (RS) codes. 2D-RS codes result in an IC proof size of $O(\sqrt{b} \log b)$, where b is the size of the block. Work in [2] extends the idea into a technique called Coded Merkle Tree (CMT). A CMT uses Low-Density Parity-Check (LDPC) codes for encoding a Merkle tree and it provides the following benefits: 1) small check node (CN) degrees in the LDPC codes reduce the IC proof size to $O(\log b)$ [2]; 2) LDPC codes can be decoded using a linear time peeling decoder [14], thus reducing the decoding complexity compared to Reed-Solomon codes. Despite these benefits, an LDPC code with a peeling decoder leads to certain problematic objects, called stopping sets [14] that allow malicious nodes to successfully

³As with all applications of channel coding, coded redundancy adds extra storage overhead at the full nodes. In this work, we improve the trade-off between the storage overhead and the probability of failure of detecting DA attacks by providing better codes, thus, making channel coding a more viable solution despite the overhead.

hide a smaller portion of the block compared to Reed-Solomon codes. A stopping set of an LDPC code is a set of variable nodes (VNs) that if erased prevents a peeling decoder from fully decoding the original block. If a malicious node hides coded symbols corresponding to a stopping set of the LDPC code, full nodes will not be able to decode the CMT. Since the malicious node can hide the smallest stopping set, the best code design strategy to reduce the probability of failure using random sampling is to construct deterministic LDPC codes with large minimum stopping set size. Constructing such LDPC codes is considered a hard problem [15].

Another important coding parameter for the CMT is the length of the LDPC codes which affect the encoding/decoding complexity and Merkle proof sizes. Short code lengths are beneficial in CMT applications (like low latency blockchains [9] or resource limited IoT blockchains [7]) since they keep the above quantities small. Previous work in [2] have focused on using codes from an LDPC ensemble to construct the CMT. At large code lengths, the LDPC ensemble guarantees, with high probability, a large stopping ratio (the smallest stopping set size divided by the code length [2]) and hence a low probability of failure. However, at short code lengths, the LDPC ensemble is unable to provide good guarantees on the minimum stopping set size. Authors in [2] combat this issue through the use of *bad-code* proofs when codes with a smaller stopping ratio (*bad-codes*) than guaranteed by the ensemble get used. A bad-code proof triggers all nodes in the system to use a newly sampled code from the ensemble. However, at short code lengths, this approach requires many rounds of bad codes until a good code has been found which undermines the security of the system. Thus, the LDPC code design of [2] is inappropriate for short CMT code lengths. Hence, in this paper, we focus on short CMT code lengths and provide deterministic LDPC codes that allow for good detection of DA attacks. For various adversary models, we provide a co-design of specialized LDPC codes and sampling strategies that reduce the probability of failure compared to techniques used in earlier literature.

We can broadly categorize all possible adversaries into three types based on their computational capabilities. The computational complexity is based on how hard it is for a malicious node to find the minimum stopping set in the LDPC code (which is known to be an NP-hard problem [16]). Note that the light node sampling strategy is known by all entities in the system. The first adversary type is termed as a *weak adversary*. A weak adversary does not have the resources to find a large number of stopping sets. It settles for hiding a random one it finds and is unable to take advantage of the light node sampling strategy. The second type is a *medium adversary*. A medium adversary, using more computational resources, can find all stopping sets up to a

certain size and select the stopping set that performs the worst under the posted light node sampling strategy. While the medium adversary has more computational capability than a weak adversary, a medium adversary represents a malicious node with bounded resources and can only find stopping sets up to a certain size within a reasonable time frame. The final type is a *strong adversary* which we assume has unlimited resources and can find all stopping sets (of any size) and hide one among them that performs the worst. These three models represent how much resources we assume an adversary possesses to disrupt our system. As such, our modeling encompasses everything from a single hacker with a standard computer to a small group of hackers with a cluster of computers to a large organization with unlimited resources.

A. Contributions

Our main contributions in this paper are co-design techniques for LDPC codes and coupled light node sampling strategies that result in a low probability of failure under the different adversary models described above. In LDPC codes with no degree-one VNs, all stopping sets are made up of cycles [23]. Since working with stopping sets directly is computationally difficult, in this paper we design LDPC codes by optimizing cycles to indirectly optimize stopping sets. We show that our LDPC codes result in the desired stopping set properties and produce low probability of failures for the different adversary models. The contributions are listed as follows:

- 1) For the weak adversary, we demonstrate that *concentrating* stopping sets in LDPC codes to a small set of VNs and then greedily sampling this small set of VNs results in a low probability of light node failure. We then provide a specialized LDPC code construction technique called the *entropy-constrained* Progressive Edge Growth (EC-PEG) algorithm that is able to concentrate stopping sets in the LDPC code to a small set of VNs. We provide a greedy sampling strategy for the light nodes to sample this small set of VNs. We demonstrate that for a weak adversary, LDPC codes constructed using the EC-PEG algorithm along with the greedy sampling strategy result in a significantly lower probability of failure compared to techniques used in earlier literature.

- 2) To secure the light nodes against a medium and a strong adversary, we provide a co-design of a light node sampling strategy called *linear-programming-sampling* (LP-sampling) and an LDPC code construction called *linear-programming-constrained* PEG (LC-PEG) algorithm. LP-sampling is tailor-made for the particular LDPC codes used to construct each layer of the CMT. It is designed by solving a linear program (LP) based on the knowledge of the small stopping sets in the LDPC codes to minimize the probability of failure. We demonstrate that, for a medium and a strong adversary, LDPC codes designed by the LC-PEG algorithm coupled with LP-sampling

result in a lower probability of failure compared to techniques used in earlier literature.

B. Previous Work

In [1], authors proposed to solve DA attacks by encoding the block using 2D-RS codes that results in an IC proof size of $O(\sqrt{b} \log b)$. In [2], authors proposed the CMT and demonstrated that encoding the CMT using LDPC codes results in a small IC proof size of $O(\log b)$. Authors in [2] used codes from a random LDPC ensemble of [18] to construct the CMT to result in a low probability of failure. However, random LDPC ensembles used in [2] were originally designed for other types of channels (i.e., BSC) and we show that they are not the best choice for this specific application at short CMT code lengths. At the same time, as described before, random LDPC ensembles undermine the security of the system, especially at short CMT code lengths. In this work, we demonstrate that the presented co-design techniques result in a lower probability of failure compared to using codes from a random LDPC ensemble and random sampling. Furthermore, to alleviate the security problem, we provide deterministic LDPC code design algorithms in this paper. In [19], authors provide a protocol called *CoVer* based on CMT, which allows a group of light nodes to collectively validate blocks without relying on full nodes. However, [19] still uses random sampling and random LDPC ensembles to mitigate DA attacks.

DA attacks are possible in other blockchain systems as well. *Sharded blockchains* where each node stores a fraction of the entire block are vulnerable to DA attacks that can be solved using the CMT [20]. The LDPC co-design techniques described in this paper can also be used in sharded blockchains. *Side Blockchains* [21] that improve the throughput of block transactions are also vulnerable to DA attacks. The vulnerability is mitigated in [21] by introducing a *DA oracle* that uses the CMT. A similar idea as this paper of co-design to construct specialized LDPC codes to improve the performance of the DA oracle was demonstrated in [22].

While this paper focuses on designing codes to mitigate DA attacks, channel coding has been extensively used to mitigate other scalability issues in blockchain systems: [25] uses network codes to reduce the storage cost associated with full nodes; [26] combines downsampling and erasure coding to reduce the storage cost while allowing nodes to directly use the stored data without decoding; [27] proposes secure fountain codes to reduce the storage and bootstrapping communication cost of full nodes; [28] uses Lagrange coding in sharded blockchains to simultaneously improve storage, computation, and security; [29] proposes using erasure codes to allow light nodes to contribute in storing the blockchain. The proposal in [29] can be combined with techniques proposed in this paper to enable light nodes to ensure data availability.

Fig. 2: Left Panel: Construction process of a CMT. A block of size b is partitioned into k data chunks (data symbols) each of size $\frac{b}{k}$ and a rate R systematic LDPC code is applied to generate n coded symbols. These n coded symbols form the base layer of the CMT. The n coded symbols are then hashed using a hashing function and the hashes of every q coded symbols are concatenated to get one data symbol of the parent layer. The data symbols of this layer are again coded using a rate R systematic LDPC code and the coded symbols are further hashed and concatenated to get the data symbols of its parent layer. This iterative process is continued until there are only t ($t > 1$) hashes in a layer which form the CMT root. Left panel shows a CMT with $n = 16$, $q = 4$, $R = 0.5$ and $t = 4$. The circled symbols in L_1 and L_2 are the Merkle proof of the circled symbol in L_3 . Right panel: DA attack on the CMT.

The rest of this paper is organized as follows. In Section II, we provide the preliminaries and system model. In Section III, we describe the greedy sampling strategy and the EC-PEG algorithm and how they overcome DA attacks against the weak adversary. In Section IV, we present our approach for the medium and strong adversary where we describe the LP-sampling strategy and the LC-PEG algorithm. We discuss system aspects of our co-design in Section V. We provide simulation results in Section VI and concluding remarks in Section VII.

II. PRELIMINARIES AND SYSTEM MODEL

In this section, we first look at the preliminaries of the CMT and LDPC notation. We then present our system, network, and threat model. We use the following notation in the rest of the paper. For $\mathbf{p} = (p_1, \dots, p_t)$ such that $p_i \geq 0$, $\sum_{i=1}^t p_i = 1$, we use the entropy function $\mathcal{H}(\mathbf{p}) = -\sum_{i=1}^t p_i \log(p_i)$. For a vector $\mathbf{a}$, let $\max(\mathbf{a})$ ($\min(\mathbf{a})$) denote the largest (smallest) entry of $\mathbf{a}$ and let a_i denote the i^{th} element of $\mathbf{a}$. For a matrix $\mathbf{M}$ of size $c \times d$, let M_{ki} denote the element of $\mathbf{M}$ on the k^{th} row and i^{th} column, $1 \leq k \leq c$, $1 \leq i \leq d$. Define $x \bmod p := (x)_p$.

A. Coded Merkle Tree (CMT)

1) *CMT construction*: A CMT of a block is built using the block transactions as leaf nodes and the CMT root is included in the block header. It is constructed by encoding each layer of the Merkle tree [4] with an LDPC code and then hashing the layer to generate its parent layer. A simplified description of the CMT construction is shown in Fig. 2 left panel⁴. As shown in Fig. 2 left panel, coded symbols of a layer are interleaved into the data symbols of the parent layer. In this paper, we adopt the interleaving technique introduced in [21]. Let the CMT have

⁴In this paper, we call chunks as symbols where each chunk can be represented as a vector of elements from some finite field.

l layers (except the root), $L_1, L_2, \dots, L_l$, where L_l is the base layer. The root of the CMT is referred to as L_0 . For $1 \leq j \leq l$, let L_j have n_j coded symbols and let the LDPC code used in L_j have a parity check matrix H_j . Let $N_j[i]$, $0 \leq i < n_j$, be the $(i+1)^{th}$ symbol of the j^{th} layer L_j ⁵. Also, let $D_j[i] = N_j[i]$, $0 \leq i < Rn_j$ and $P_j[i] = N_j[i]$, $Rn_j \leq i < n_j$, be the systematic (data) and parity symbols of L_j , respectively. Coded symbols $P_j[i]$, $Rn_j \leq i < n_j$ are obtained from $D_j[i]$, $0 \leq i < Rn_j$ using a rate R systematic LDPC code H_j . In the above CMT, hashes of every q coded symbols of L_j are concatenated to form a data symbol of L_{j-1} . Hence, $n_j = \frac{n_l}{(qR)^{l-j}}$, $j = 1, \dots, l$. The CMT root has $t = n_1$ hashes. Let the number of systematic and parity symbols in L_j be denoted by $s_j = Rn_j$ and $p_j = (1 - R)n_j$, respectively. For $1 \leq j \leq l$, the data symbols of L_{j-1} are formed from the coded symbols of L_j as follows:

$$D_{j-1}[i] = N_{j-1}[i] = \text{concat} \left(\left\{ \text{Hash}(N_j[x]) \mid 0 \leq x < n_j, i = (x)_{s_{j-1}} \right\} \right) \quad \forall 0 \leq i < s_{j-1},$$

where Hash and concat represent the hash and the string concatenation functions, respectively.

2) *Merkle Proof for CMT symbols*: The Merkle proof of a symbol in L_j consists of a data symbol and a parity symbol from each intermediate layer of the tree that is above L_j [21]. An illustration of a Merkle proof is shown in Fig. 2 left panel. In particular, the Merkle proof of the symbol $N_j[i]$, $1 < j \leq l$, is the set of symbols $\{N_{j'}[(i)_{s_{j'}}], N_{j'}[s_{j'} + (i)_{p_{j'}}] \mid 1 \leq j' \leq j-1\}$. Detailed discussion on the properties of Merkle proofs⁶ can be found in [21].

3) *Hash-Aware Peeling decoder*: Using the CMT root and the available symbols of each layer of the CMT, the original block can be decoded using a hash-aware peeling decoder described in [2]. The hash-aware peeling decoder decodes each layer of the CMT (from top to bottom) like a conventional peeling decoder [14]. However, after decoding a symbol in layer j , the decoder matches its hash with the corresponding hash present in layer $j-1$. Matching the hashes allows the decoder to detect IC attacks and generate IC proofs as described in [2]. The IC proof size is proportional to the degree of CNs in the LDPC codes used to build the CMT.

B. Stopping sets and LDPC notation

A stopping set of an LDPC code is a set of VNs such that every CN connected to this set is connected to it at least twice [14]. A stopping set is hidden (made unavailable) by a malicious node if all VNs present in it are hidden. The hash-aware peeling decoder fails to successfully decode layer j of the CMT if a stopping set of H_j is unavailable. Let the parity check matrix

⁵Due to modulo operations, we define $N_j[i]$ starting with index 0 for i . All other variables in the paper start with index 1.

⁶The data part of the Merkle proof of $N_j[i]$ from each layer lie on the path of $N_j[i]$ to the CMT root and can be used to check the integrity of $N_j[i]$ in a manner similar to checking the proofs for regular Merkle trees in [4]. The parity symbols included in the Merkle proof are only for sampling purposes. For each parity symbol, the information provided in the Merkle proof of $N_j[i]$ contain all the symbols in parity symbol's path to root and can be used to check its integrity.

H_j , $1 \leq j \leq l$, of the CMT have n_j columns $\mathcal{V}^{(j)} = \{v_1^{(j)}, v_2^{(j)}, \dots, v_{n_j}^{(j)}\}$. CMT symbol $N_j[i]$, $0 \leq i < n_j$, corresponds to VN $v_{i+1}^{(j)}$. Let the Tanner graph (TG) [14] representation of H_j be denoted by $\mathcal{G}_j$, where $v_i^{(j)}$ is also referred to as the i^{th} VN in $\mathcal{G}_j$ and rows of H_j are referred as CNs in $\mathcal{G}_j$. For a parity check matrix H_j , let $H_j[v_i^{(j)}]$ denote the column of the parity check matrix corresponding to VN $v_i^{(j)}$. A cycle of length g is called a g -cycle. For a set $\mathcal{S}$, let $|\mathcal{S}|$ denote its cardinality. For a cycle (stopping set) in the TG $\mathcal{G}$, we say that a VN v *touches* the cycle (stopping set) iff v is part of the cycle (stopping set). We define the weight of a stopping set as the number of VNs touching it. Let $\omega_{\min}^{(j)}$ denote the minimum stopping set size of H_j , $1 \leq j \leq l$. The girth of a TG is defined as the length of the smallest cycle present in the graph.

C. System and Network Model

We consider a blockchain system similar to [1] and [2] that has full nodes and light nodes. One of the full nodes acts as a *block producer* of a new block. We consider the same blockchain network model as [2]. In particular, we assume a synchronous network where every honest node (full or light) is connected to at least one honest full node and the messages sent on the network are anonymous. The network can have a dishonest majority of full nodes, but each light node is connected to at least one honest full node (thus preventing eclipse attacks [1]). Nodes broadcast a message (fraud proofs, IC proofs, and CMT symbols) by sending the message to all its connected nodes. The connected nodes check the message correctness (Merkle proofs) and forward valid messages to their neighbors⁷. In the following, we describe actions performed by the block producer, other full nodes, and light nodes. We also mention the items included in the publicly available protocol that is designed by a blockchain system designer to be used by nodes in the system. In Section V, we provide a discussion on the blockchain system designer.

1) *Items included in the protocol:* Parity check matrices H_j , $1 \leq j \leq l$, systematic generator matrix of each H_j , and the light node sampling strategy (a rule to sample CMT symbols).

2) *Block Producer:* A full node that produces (mines) a new block (see Fig. 1). On producing a new block, the block producer encodes the block to construct its CMT using the systematic generator matrices specified in the protocol. It then broadcasts all the coded symbols in the CMT (including the root) to other full nodes and the root of the CMT to the light nodes. On receiving a sampling request from the light nodes, it returns the requested symbols along with their Merkle proofs. The block producer can be malicious and can act arbitrarily.

⁷Since messages are communicated only to connected nodes, the cost of broadcasting is not high. Moreover, honest nodes prevent fake communication from malicious nodes by forwarding only valid messages.

3) *Full nodes that are not the block producer*: These nodes perform Merkle proof checks on the coded symbols of the CMT that they receive from a block producer, other full nodes, or light nodes (see Fig. 1). They forward symbols that satisfy the Merkle proofs to other connected full nodes. Using the symbols that they received, they decode each layer of the CMT with a hash-aware peeling decoder using the parity check matrices H_j , $1 \leq j \leq l$, specified in the protocol. After decoding the base layer of the CMT, which contains transaction data, they verify all the transactions. They store a local copy of all blocks (i.e., its CMT) that they verify to be valid (i.e., fully available, having no fraudulent transactions and no incorrect-coding at any layer). They declare the availability of this valid block to all other nodes and respond to sample requests from the light nodes. If they find a certain block to be invalid, either due to fraudulent transactions or incorrect coding, they broadcast a fraud proof or an IC proof for other nodes to reject the block. If they find a certain layer of the CMT to be unavailable (i.e., having coded symbols missing that prevent decoding), they reject the block. A malicious full node need not follow the above protocol and can act arbitrarily.

4) *Light nodes*: These nodes are storage constrained and only store the CMT root corresponding to each block (see Fig. 1). They download only a small portion of the block and perform tasks like fraud and IC proof checks. Additionally, light nodes check the availability of each layer of the CMT. They do so by making sampling requests for coded symbols of the CMT base layer from the block producer (or any other full node that declares the block to be available). They make sample requests using the sampling strategy specified in the protocol. They perform Merkle proof checks on the returned symbols and broadcast symbols that satisfy the Merkle proofs to other connected full nodes. Upon receiving all the requested symbols and verifying their Merkle proofs, light nodes accept the block as available and store the block header. On receiving fraud proofs or IC proofs sent out by a full node, light nodes verify the proof and reject the header if the proof is correct. We assume that each light node is honest.

Remark 1. *In this paper, we provide co-design of LDPC codes and sampling strategies (that are included in the protocol) to reduce the probability of failure. As such, we do not compromise on other performance metrics considered in [2]: the CMT root has a fixed size t which does not grow with the blocklength; the hash-aware peeling decoder has a decoding complexity linear in the blocklength; we empirically show that the IC proof size for our codes is similar to [2].*

D. Threat Model

A blockchain system involves two aspects: block generation and block verification. The block

generation depends on the consensus algorithm used in the blockchain e.g., Proof of Work (PoW) [4], Proof of Stake (PoS) [30], etc.. However, a DA attack caused by an adversary with dishonest majority (in terms of work, stake, etc.) affects the block verification process. Hence, the exact consensus algorithm used by the blockchain is not relevant to our work. Similar to [1] and [2], we focus on block verification and propose LDPC codes to mitigate DA attacks⁸.

Similar to [1] and [2], we model the security of our system in terms of two properties: *i) Soundness*: If a light node thinks that a block is available and accepts the block, then at least one honest full node in the system will be able to fully decode all layers of the CMT corresponding to the block; *ii) Agreement*: If one light node determines that a block is available, all light nodes in the system determine that the block is available. Let $P_f^{S,A}$ be the probability that soundness or agreement fails in the system due to a DA attack. In Section V, we show that in our proposed co-design, $P_f^{S,A}$ is reduced by reducing the probability of failure of a single light node to detect DA attacks when there is a sufficiently large number of light nodes in the system. Thus, in the rest of the paper, we focus on reducing the probability of failure of a single light node.

We consider an adversary that conducts a DA attack by hiding coded symbols of the CMT. An illustration of a DA attack is shown in Fig. 2 right panel. On receiving sampling requests from the light nodes, the adversary only returns coded symbols that it has not hidden and ignores other requests⁹. The adversary conducts a DA attack at layer j of the CMT by 1) generating coded symbols of layer j , that satisfy their Merkle proof, for the light nodes to accept these coded symbols as valid, and 2) hiding a small portion of the coded symbols of layer j , corresponding to a stopping set of H_j , such that honest full nodes are not able to decode the layer. A DA attack at layer j prevents an honest full node from generating a fraud proof of fraudulent transactions (if $j = l$) or an IC proof for incorrect coding at layer j . Since an incorrect coding can occur at any layer, for the full nodes to be able to send IC proofs, light nodes must detect a DA attack at any layer j that the adversary may perform. They do so by sampling few base layer coded symbols. For each intermediate layer j , $1 \leq j < l$, the symbols of layer j collected as part of the Merkle proofs of the base layer samples are used to check the availability of layer j .

Light nodes fail to detect a DA attack if none of the base samples requested or the symbols in their Merkle proofs are hidden. Let $P_f^{(j)}(s)$, $1 \leq j \leq l$, be the probability of failure of detecting

⁸Note that forking-based double spending attacks (related to block generation) where an adversary generates an invalid longest chain are still possible with a dishonest majority of full nodes [1] but are not necessary to launch a DA attack.

⁹Using Merkle proofs, light nodes can verify the integrity of the returned samples with respect to the Merkle root, preventing the adversary from returning incorrect data. Since the light nodes also broadcast the returned samples, the adversary cannot return the hidden portion which prevented honest full nodes from decoding back the unavailable layer.

a DA attack at layer j by a single light node when it samples s base layer coded symbols. Also, let $J^{\max} = \underset{1 \leq j \leq l}{\operatorname{argmax}} P_f^{(j)}(s)$. To maximize the probability of failure, we assume that the adversary is able to perform a DA attack at layer $J^{\max}$. We now provide precise mathematical definitions of the three adversary models discussed in Section I based on their computational capabilities:

1) *Weak Adversary*: For each layer j , $1 \leq j \leq l$, they hide stopping sets of size $< \mu_j$ for the parity check matrix H_j (for some integer μ_j). Moreover, they do not exhaustively find all stopping sets of a particular size of a given parity check matrix or perform a tailored search for stopping sets. Instead, we assume that to conduct a DA attack at layer j , for all stopping sets of H_j of a particular size, they randomly choose one of them to hide.

2) *Medium Adversary*: For each layer j , $1 \leq j \leq l$, they hide stopping sets of size $< \mu_j$ for the parity check matrix H_j . However, they use the knowledge of the sampling strategy employed by the light nodes to hide the worst case stopping set that has the lowest probability of being sampled by the light nodes. Let Ψ_j be set of all stopping sets of H_j of size $< \mu_j$. Also, let $P_f^{(j)}(s) = \max_{\psi \in \Psi_j} P_f^{(j)}(s; \psi)$, where $P_f^{(j)}(s; \psi)$ is the probability of failure for the light nodes to detect a DA attack at layer j under the light node sampling strategy when the adversary hides the stopping set ψ of H_j . For $J^{\max} = \underset{1 \leq j \leq l}{\operatorname{argmax}} P_f^{(j)}(s)$, the medium adversary conducts a DA attack at layer $J^{\max}$ by hiding a stopping set ψ from $\Psi_{J^{\max}}$ with the highest $P_f^{J^{\max}}(s; \psi)$.

3) *Strong Adversary*: They can find the worst case stopping sets of any size of H_j , $1 \leq j \leq l$. Let Ψ_j^∞ be the set of all stopping sets of H_j . Similar to the medium adversary, define $P_f^{(j)}(s) = \max_{\psi \in \Psi_j^\infty} P_f^{(j)}(s; \psi)$ and $J^{\max} = \underset{1 \leq j \leq l}{\operatorname{argmax}} P_f^{(j)}(s)$. The strong adversary conducts a DA attack at layer $J^{\max}$ by hiding a stopping set ψ from $\Psi_{J^{\max}}^\infty$ with the highest $P_f^{J^{\max}}(s; \psi)$.

The co-design that we provide to mitigate DA attacks against weak adversaries, i.e., the EC-PEG algorithm and the greedy sampling strategy, has the advantage of being computationally cheap and does not involve finding stopping sets. In order to mitigate DA attacks against a medium and a strong adversary we provide LP-sampling and the LC-PEG algorithm. LP-sampling uses stopping sets of size $< \mu_j$ from layer j of the CMT and is more computationally expensive. It is an overkill for the weak adversary which can be mitigated using cheaper techniques. Factors such as the choice of the consensus algorithm, area of deployment, etc. can give an idea about the expected computational capabilities of full nodes in the system and allow the system designer to choose the adversary model. For example, in PoS [30] and PoSpace [31] consensus blockchains, full nodes need not have a high computational power and a weak adversary would be a reasonable model to follow. For PoW blockchains [4], full nodes are expected to have high computational

power and a strong and medium adversary model would be a suitable design choice. Another example is small scale IoT-blockchains where the blockchain nodes are IoT devices [7]. Here, full nodes have low computational power and a weak adversary model would be appropriate.

In our co-design to mitigate a DA attack against a medium and a strong adversary, we assume that a blockchain system designer decides the value of μ_j , $1 \leq j \leq l$, and is able to find all stopping sets of H_j of size $< \mu_j$, that is used to design LP-sampling. Although finding all stopping sets of H_j of size $< \mu_j$ is NP-hard, since we focus on short code lengths in this paper, the set of stopping sets can be found in a reasonable amount of time using Integer Linear Programming (ILP) methods demonstrated in [32]. Note that μ_j and the set of all stopping sets of H_j of size $< \mu_j$ that the designer uses to design LP-sampling is not publicly released. Only the final design output, i.e., the LP-sampling strategy is included in the protocol. Here, we have made a trusted set up assumption of a blockchain system designer to design the items included in the protocol. In Section V, we will discuss potential ways to prevent security attacks by a malicious designer and how some attacks are naturally handled by our co-design method.

Given the above adversary models, we provide LDPC code construction and sampling strategies to minimize the probability of failure for a single light node to detect DA attacks. In the next section, we discuss the techniques to mitigate DA attacks conducted by a weak adversary.

III. LDPC CODE AND SAMPLING CO-DESIGN FOR WEAK ADVERSARY

In this section, we demonstrate our novel design idea of *concentrating* stopping sets in LDPC codes to reduce the probability of failure against a weak adversary. Since working with stopping sets directly is computationally difficult, we focus on *concentrating* cycles to indirectly *concentrate* stopping sets. It is well known that codes with irregular VN degree distributions are prone to small stopping sets. Thus, we consider VN degree regular LDPC codes of VN degree $d_v \geq 3$ in this paper. In the following, we first look at the effect of the light node sampling strategy on the probability of failure when a DA attack occurs on the base layer of the CMT. This will motivate the LDPC code construction for the base layer. Later, we demonstrate how the LDPC code construction for the base layer can be used in all layers by aligning the columns of the parity check matrices before constructing the CMT. For simplicity of notation, we denote H_l by H having n VNs $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ and TG $\mathcal{G}$. Consider the following definition.

Definition 1. For a parity check matrix H , let $ss^\kappa = (ss_1^\kappa, ss_2^\kappa, \dots, ss_n^\kappa)$ denote the VN-to-stopping-set of weight κ distribution where ss_i^κ is the fraction of stopping sets of H of weight κ touched by v_i . Similarly, for a parity check matrix H , let $\zeta^g = (\zeta_1^g, \zeta_2^g, \dots, \zeta_n^g)$ be the VN-to-

g -cycle distribution where ζ_i^g is the fraction of g -cycles of H touched by v_i .

We informally say that distribution ss^κ (ζ^g) is concentrated if a small set of VNs have high corresponding stopping set (g -cycle) fractions ss_i^κ (ζ_i^g). The following lemma demonstrates that LDPC codes with concentrated ss^κ results in a smaller probability of light node failure when a weak adversary conducts a DA attack (on the base layer). The proof is straightforward and we omit it due to space limitations. It can be found in [3] and references therein.

Lemma 1. *Let $\mathcal{SS}_\kappa$ denote the set of all weight κ stopping sets of H . For a weak adversary that randomly hides a stopping set from $\mathcal{SS}_\kappa$, the probability of failure at the base layer, $P_f^{(l)}(s)$, when the light nodes use s samples and any sampling strategy satisfies $P_f^{(l)}(s) \geq 1 - \max_{S \subseteq \mathcal{V}, |S|=s} \tau(S, \kappa)$. Here, $\tau(S, \kappa)$ is the fraction of stopping sets of weight κ touched by the subset of VNs S of H . The lower bound in the above equation is achieved when light nodes sample, with probability one, the set $\mathcal{S}_\kappa^{opt} = \operatorname{argmax}_{S \subseteq \mathcal{V}, |S|=s} \tau(S, \kappa)$.*

Lemma 1 suggests that for a sample size s , the lowest probability of failure is $1 - \tau(\mathcal{S}_\kappa^{opt}, \kappa)$ and is achieved when the light nodes sample the set $\mathcal{S}_\kappa^{opt}$. Now, $\tau(\mathcal{S}_\kappa^{opt}, \kappa)$ is large if a majority of stopping sets of weight κ are touched by a small subset of VNs. This goal is achieved if the distributions ss^κ are concentrated towards a small set of VNs. Thus, designing LDPC codes with concentrated ss^κ increases $\tau(\mathcal{S}_\kappa^{opt}, \kappa)$ and reduces the probability of failure. In Section III-B, we design the EC-PEG algorithm that achieves concentrated stopping set distributions.

Algorithm 1 Light node sampling strategy for weak adversary: `greedy-set`($\mathcal{G}, g_{\min}, g_{\max}, s$)

- 1: **Inputs:** TG $\mathcal{G}$, $g_{\min}$, $g_{\max}$, s , **Output:** $S_{greedy}^{(s)}$, **Initialize:** $S_{greedy}^{(s)} = \emptyset$, $g = g_{\min}$, $\hat{\mathcal{G}} = \mathcal{G}$
  - 2: **while** $|S_{greedy}^{(s)}| < s$ **do**
  - 3: $v = \text{VN that touches the maximum number of } g\text{-cycles in } \hat{\mathcal{G}}$ (ties broken randomly)
  - 4: $S_{greedy}^{(s)} = S_{greedy}^{(s)} \cup \{v\}$, Purge v and all its incident edges from $\hat{\mathcal{G}}$
  - 5:   **if** $\hat{\mathcal{G}}$ has no g -cycles **then** $g = g + 2$
  - 6:   **if** $g \geq g_{\max}$ **then**
  - 7: $\mathcal{V}_r = \text{randomly select } s - |S_{greedy}^{(s)}| \text{ VNs from } \hat{\mathcal{G}} \text{ (ordered arbitrarily); } S_{greedy}^{(s)} = S_{greedy}^{(s)} \cup \mathcal{V}_r$
-

We are unaware of an efficient method to find $\mathcal{S}_\kappa^{opt}$. Instead, we use a greedy algorithm using cycles to find the light node samples, provided in Algorithm 1. Algorithm 1 takes as input the TG $\mathcal{G}$, its girth $g_{\min}$, an upper bound cycle length $g_{\max}$, and the sample size s . It outputs a set of VNs $S_{greedy}^{(s)}$ that the light nodes will sample. The probability of failure using this strategy when a weak adversary randomly hides a stopping set of size κ from the base layer is $P_f^{(l)}(s) = 1 - \tau(S_{greedy}^{(s)}, \kappa)$ (see proof of Lemma 1 in [3]). We refer to the VNs provided by `greedy-set`($\mathcal{G}, g_{\min}, g_{\max}, s$) as *greedy samples*. At the end of this section, we empirically show that concentrating the cycle

distributions ζ^g also concentrates the stopping set distributions. Thus, the EC-PEG algorithm aims to concentrate the cycle distributions to improve the probability of failure. It is easy to see that the complexity of Algorithm 1 is dominated by the complexity of finding cycles (of worst case length $g_{\max}$) and is $O(n^{g_{\max}/2})$ using brute force.

Remark 2. (*Overall Greedy Sampling Strategy*) In the above sampling strategy, some coded symbols may never get sampled which can affect the soundness of the system. We alleviate this problem without affecting the probability of failure by modifying the sampling strategy as follows: Let ρ be a fixed parameter where $0 < \rho < 1$. For a total of s samples, the light nodes select ρs greedy samples $S_{\text{greedy}}^{(\rho s)} = \text{greedy-set}(\mathcal{G}, g_{\min}, g_{\max}, \rho s)$ and randomly select $s - \rho s$ base layer coded symbols for the remaining samples. We discuss the soundness and agreement of this modified strategy in Section V. For this strategy, $P_f^{(l)}(s) = [1 - \tau(S_{\text{greedy}}^{(\rho s)}, \kappa)] \left(1 - \frac{\omega_{\min}^{(l)}}{n_l}\right)^{(s-\rho s)}$.

A. Aligning the parity check matrices of the CMT

In the above discussion, we demonstrated how to mitigate a DA attack conducted by a weak adversary on the base layer of the CMT using greedy sampling. Now, we extrapolate the idea of greedy sampling to the intermediate layers. Since the intermediate layers are sampled via the Merkle proofs of the base layers samples, we align the base and intermediate layer symbols such that the intermediate layers are also sampled greedily. We do so by aligning (permuting) the columns of the parity check matrices used in different CMT layers. We align the columns such that the samples of an intermediate layer j collected from the Merkle proofs of the base layer samples coincide with the greedy samples for layer j provided by $\text{greedy-set}(\mathcal{G}_j, g_{\min}^{(j)}, g_{\max}^{(j)}, \tilde{s})$. Here, $g_{\min}^{(j)}$ is the girth of $\mathcal{G}_j$ and $g_{\max}^{(j)}$ is the upper cycle length for layer j .

We assume that the output $S_{\text{greedy}}^{(s)}$ of Algorithm 1 is ordered according to the order VNs were added to $S_{\text{greedy}}^{(s)}$. Let $S_{\text{ordered}}^{(j)} = \text{greedy-set}(\mathcal{G}_j, g_{\min}^{(j)}, g_{\max}^{(j)}, n_j)$, $1 \leq j \leq l$. VNs in $S_{\text{ordered}}^{(j)}$ are all the VNs of H_j ordered (permuted) according to the order they were added to $S_{\text{ordered}}^{(j)}$. Hence, we denote $S_{\text{ordered}}^{(j)}[i]$ as the i^{th} VN in this ordered list of VNs. The procedure to align the columns of the parity check matrices of different layers of the CMT is provided in Algorithm 2. In the algorithm, we first permute the columns of the base layer parity check matrix H_l (to obtain $\widetilde{H}_l$) such the the VNs in $S_{\text{ordered}}^{(l)}$ appear as columns $1, 2, \dots, n_l$ in $\widetilde{H}_l$ (line 3). Recall that when the base layer symbol corresponding to $v_i^{(l)}$ is sampled, then for every intermediate layer j , the VNs with with subscript indices $\{1 + (i - 1)_{s_j}, 1 + s_j + (i - 1)_{p_j}\}$ get sampled. We assign columns of $\widetilde{H}_j$ at these indices (starting from $i = 1$) the columns of H_j correspond to the greedy samples in $S_{\text{ordered}}^{(j)}$ from start to end (lines 4-7). We continue this process until

all columns of $\widetilde{H}_j$ have been assigned. The complexity of Algorithm 2 is dominated by the complexity of finding $S_{ordered}^{(j)}$ using Algorithm 1 and has a complexity of $O(\sum_{j=1}^l n_j^{g_{max}^{(j)}/2})$.

Algorithm 2 Aligning parity check matrices of CMT for greedy sampling

---

- 1: **Inputs:** $H_j, S_{ordered}^{(j)}, 1 \leq j \leq l$, **Outputs:** $\widetilde{H}_j, 1 \leq j \leq l$
  - 2: **Initialize:** $\widetilde{H}_j$: matrix with unassigned columns, $1 \leq j \leq l$, counter = 1
  - 3: $\widetilde{H}_l[i] = H_l[S_{ordered}^{(l)}[i]], 1 \leq i \leq n_l$
  - 4: **for** $j = 1, 2, \dots, l-1$ **do for** $i = 1, 2, \dots, n_l$ **do**
  - 5: $d = 1 + (i-1)_{s_j}, p = 1 + s_j + (i-1)_{p_j}$
  - 6:   **if** $\widetilde{H}_j[d]$ is not assigned before **then** $\widetilde{H}_j[d] = H_j[S_{ordered}^{(j)}[\text{counter}]]$, counter += 1
  - 7:   **if** $\widetilde{H}_j[p]$ is not assigned before **then** $\widetilde{H}_j[p] = H_j[S_{ordered}^{(j)}[\text{counter}]]$, counter += 1
  - 8:   **if** all columns of $\widetilde{H}_j$ have been assigned **then** Break i for loop
-

Remark 3. The parity check matrices $\widetilde{H}_j, 1 \leq j \leq l$, after the alignment are included in the protocol. Recall that a CMT is built using systematic LDPC codes. Under the assumption of full rank, for the parity check matrices $\widetilde{H}_j, 1 \leq j \leq l$, the corresponding generator matrices are constructed in a systematic form which are then included in the protocol for CMT construction. Also, after the alignment, the overall greedy sampling strategy as described in Remark 2 becomes: sample the first ρs coded symbols of the base layer of the CMT and then randomly sample $s - \rho s$ base layer coded symbols. This sampling rule is included in the protocol.

For a CMT built using $\widetilde{H}_j, 1 \leq j \leq l$, provided by Algorithm 2, greedy sampling of the base layer of the CMT according to Algorithm 1 ensures that all intermediate layers of the CMT are greedily sampled according to Algorithm 1 through the Merkle proofs of the base layer samples. Next, we provide a design strategy to construct LDPC codes with concentrated stopping set distributions that result in a low probability of failure under greedy sampling. Note that codes produced in the next subsection are aligned by Algorithm 2 and then included in the protocol.

B. Entropy-Constrained PEG (EC-PEG) Algorithm

The EC-PEG algorithm is based on minimizing the entropy of cycle distribution ζ^g . The intuition behind our algorithm is using the fact that uniform distributions have high entropy and distributions that are concentrated have low entropy. Thus, we construct LDPC codes using the PEG algorithm [17] by making CN selections that minimize the entropy of the cycle distributions. Algorithm 3 presents the EC-PEG algorithm for constructing a TG $\widetilde{\mathcal{G}}$ with n VNs, m CNs, and VN degree d_v that concentrates distributions $\zeta^{g'}, \forall g' < g_c$. Choice of g_c is a complexity constraint of how many cycles we keep track in the algorithm. All ties in the algorithm are broken randomly.

The PEG algorithm builds a TG by iterating over the set of VNs and for each VN v_j in the TG, establishing d_v edges to it. For establishing the k^{th} edge to VN v_j , the PEG algorithm

encounters two situations: i) addition of the edge is possible without creating cycles; ii) addition of the edge creates cycles. In both situations, the PEG algorithm finds a set of *candidate* CNs that it proposes to connect to v_j , to maximize the girth. We abstract out the steps followed in [17] to find the set of *candidate* CNs by a procedure $\text{PEG}(\tilde{\mathcal{G}}, v_j)$. The procedure returns the set of *candidate* CNs $\mathcal{K}$ for establishing a new edge to VN v_j under the TG setting $\tilde{\mathcal{G}}$ according to the PEG algorithm in [17]. For ii), the procedure returns the cycle length g of the smallest cycles formed when an edge is added between any CN in $\mathcal{K}$ and v_j . For i), it returns $g = \infty$. $\mathcal{K}$ is the set of all CNs in $\tilde{\mathcal{G}}$ that create new g -cycles when an edge is added between the CN and v_j . When $g = \infty$, $\mathcal{K}$ is the set of all CNs in $\tilde{\mathcal{G}}$ that if connected to v_j create no cycles.

Algorithm 3 EC-PEG Algorithm

```

1: Inputs:  $n, m, d_v, g_c$ , Outputs:  $\tilde{\mathcal{G}}, g_{\min}$ , Initialize  $\tilde{\mathcal{G}}$  to  $n$  VNs,  $m$  CNs and no edges
2: Initialize  $\Lambda_i^{(g')} = 0$ , for all  $g' < g_c$  and  $1 \leq i \leq n$ ,  $T = |\{4, 6, \dots, g_c - 2\}|$ 
3: for  $j = 1$  to  $n$  do
4:   for  $k = 1$  to  $d_v$  do
5:      $[\mathcal{K}, g] = \text{PEG}(\tilde{\mathcal{G}}, v_j)$ 
6:     if  $g \geq g_c$  then
7:        $c^{sel} = \text{Select a CN from } \mathcal{K} \text{ with the minimum degree under the current TG setting } \tilde{\mathcal{G}}$ 
8:     else  $\triangleright (g\text{-cycles, } g < g_c, \text{ are created})$ 
9:       for each  $c$  in  $\mathcal{K}$  do
10:         $\lambda_i^{(g',c)} = \Lambda_i^{(g')}, g' < g_c, 1 \leq i \leq n$ 
11:         $\mathcal{L}_{cycles} = \text{new } g\text{-cycles formed in } \tilde{\mathcal{G}} \text{ due to the addition of edge between } c \text{ and } v_j$ 
12:        for all  $v$  in  $\tilde{\mathcal{G}}$  do  $\lambda_v^{(g,c)} = \lambda_v^{(g,c)} + |\{\mathcal{O} \in \mathcal{L}_{cycles} \mid v \text{ is part of } \mathcal{O}\}|$ 
13:         $\alpha^{(g')} = (\alpha_1^{(g')}, \alpha_2^{(g')}, \dots, \alpha_n^{(g')})$ , where  $\alpha_i^{(g')} = \frac{\lambda_i^{(g',c)}}{\sum_{i=1}^n \lambda_i^{(g',c)}}$ ,  $g' < g_c$  (define  $\frac{0}{0} = 0$ )
14:         $\alpha_{g_c} = (\sum_{g' < g_c} \frac{\alpha_1^{(g')}}{T}, \sum_{g' < g_c} \frac{\alpha_2^{(g')}}{T}, \dots, \sum_{g' < g_c} \frac{\alpha_n^{(g')}}{T})$ ; Entropy $[c] = \mathcal{H}(\alpha_{g_c})$ 
15:         $c^{sel} = \text{CN in } \mathcal{K} \text{ with minimum Entropy}[c]$ 
16:         $\Lambda_i^g = \lambda_i^{(g,c^{sel})}, 1 \leq i \leq n$ 
17:       $\tilde{\mathcal{G}} = \tilde{\mathcal{G}} \cup \text{edge}\{c^{sel}, v_j\}$ 

```

Thus, when the $\text{PEG}(\tilde{\mathcal{G}}, v_j)$ procedure returns $g \geq g_c$, either no new cycles are created or the cycles created have length $\geq g_c$. In both these situations, similar to the original PEG algorithm in [17], we select a CN from $\mathcal{K}$ with the minimum degree under the current TG setting $\tilde{\mathcal{G}}$ (line 7). When $\text{PEG}(\tilde{\mathcal{G}}, v_j)$ returns $g < g_c$, we modify the CN selection procedure so that the resultant cycle distributions get concentrated. We explain the modified CN selection procedure next.

While progressing through the EC-PEG algorithm, for all g' -cycles, $g' < g_c$, we maintain VN-to- g' -cycle counts $\Lambda^{(g')} = (\Lambda_1^{(g')}, \Lambda_2^{(g')}, \dots, \Lambda_n^{(g')})$, where $\Lambda_i^{(g')}$ is the number of g' -cycles that are touched by VN v_i . When the $\text{PEG}(\tilde{\mathcal{G}}, v_j)$ procedure returns $g < g_c$, for each candidate CN $c \in \mathcal{K}$, new g -cycles are formed in the TG when an edge is established between c and v_j .

Fig. 3: Results for LDPC codes with $R = 0.5$, $d_v = 4$, $n = 128$ using different PEG algorithms. The x-axis in all the plots are the VN indices v_i in the decreasing order of the 6-cycle fractions ζ_i^6 (for the respective codes); Left panel: cycle distributions ζ^6 and ζ^8 ; Middle Panel: stopping set distribution ss^{13} ; Right Panel: stopping set distribution ss^{14} . The lines in the middle and right panels are the best fit lines for ss^k indicating the graph slope.

These cycles are listed in $\mathcal{L}_{cycles}$ (line 11). For these new g -cycles, we calculate the resultant VN -to- g -cycle counts $\lambda_i^{(g,c)}$, $1 \leq i \leq n$, if an edge is established between c and v_j (line 12). Using $\lambda_i^{(g,c)}$, we calculate the VN -to- g' -cycle normalized counts $\alpha^{g'} = (\alpha_1^{g'}, \alpha_2^{g'}, \dots, \alpha_n^{g'})$ (line 13) and then the joint normalized cycle counts α_{g_c} for g' -cycles, $g' < g_c$ (line 14). The joint normalized cycle counts α_{g_c} is simply the average of the normalized cycle counts across all the cycle lengths. Using α_{g_c} , we calculate the entropy $\mathcal{H}(\alpha_{g_c})$ for each CN c in $\mathcal{K}$ (line 14). Our modified CN selection procedure is to select a CN from $\mathcal{K}$ with minimum Entropy[] (line 15). We then update the VN -to- g -cycle counts for the new g -cycles that get created (line 16) to be used in future iterations. Minimizing the entropy of the joint normalized cycle counts ensures that the different cycle distributions are concentrated towards the same set of VNs.

We now mention the complexity of the EC-PEG algorithm. Note that the complexity of the original PEG algorithm is $O(mn)$ [17]. The EC-PEG algorithm differs from the original PEG algorithm in steps 8-16. Step 14 has the largest complexity which results in the complexity of the EC-PEG algorithm to be at most $O(mn^2) = O(n^3)$. Note that $\mathcal{L}_{cycles}$ in step 11 is obtained during step 5 as a by-product and does not incur additional complexity.

Fig. 3 demonstrates the effectiveness of the EC-PEG algorithm in concentrating the stopping set distribution. In Fig. 3 left panel, we plot the cycle distributions generated by the PEG and EC-PEG algorithms. From the figure, we see that the EC-PEG algorithm generates significantly concentrated distributions ζ^6 and ζ^8 compared to the original PEG algorithm. Fig. 3 middle and right panels show the corresponding stopping set distributions ss^k . We see that for the EC-PEG algorithm, the VNs towards the left (right) on the x-axis have high (low) stopping set fraction. Thus, concentrating the cycle distributions concentrates the stopping set distributions towards the same set of VNs as the cycles. In Section VI, we demonstrate that such concentrated distributions result in a low probability of failure using the greedy sampling strategy in Algorithm 1.

IV. LDPC CODE AND SAMPLING CO-DESIGN FOR MEDIUM AND STRONG ADVERSARY

For the medium and strong adversary, the EC-PEG algorithm and greedy sampling is insufficient to secure the system and requires stronger code and sampling design. In this section, we focus on overcoming these stronger adversaries that hide the worst case stopping set. Similar to Section III, we first look at a medium and a strong adversary who conduct a DA attack on the base layer of the CMT and propose a sampling strategy for the light nodes to sample the base layer to minimize the probability of failure. This will motivate the construction of LDPC codes for the base layer. Finally, we will generalize the sampling strategy and LDPC construction for the situation when the adversary conducts a DA attack at any layer of the CMT.

Recall that for each layer j , $1 \leq j \leq l$, the medium adversary hides stopping sets of H_j of size $< \mu_j$. Let $\Psi_j = \{\psi_1^{(j)}, \psi_2^{(j)}, \dots, \psi_{|\Psi_j|}^{(j)}\}$ be the set of all stopping sets of H_j of size $< \mu_j$, $1 \leq j \leq l$. For Ψ_j , let $\Pi^{(j)}$ denote the *VN-to-stopping-set* adjacency matrix of size $|\Psi_j| \times n_j$, where $\Pi_{ki}^{(j)} = 1$ iff $v_i^{(j)}$ touches stopping set $\psi_k^{(j)}$, else $\Pi_{ki}^{(j)} = 0$, $1 \leq i \leq n$, $1 \leq k \leq |\Psi_j|$.

Definition 2. A sampling (with replacement) strategy $(\mathbf{x}, \beta^{(l)})$ is a $n_l \times 1$ vector $\mathbf{x} = [x_1 \cdots x_{n_l}]^T$, where x_i is the probability that a light node requests for the i^{th} base layer symbol (i.e., $v_i^{(l)}$) for every sample request and $\beta^{(l)}$ controls the minimum probability of requesting a given CMT base layer symbol. $(\mathbf{x}, \beta^{(l)})$ satisfy $0 \leq \beta^{(l)} \leq x_i \leq 1$, $\sum_{i=1}^{n_l} x_i = 1$.

Let $P_{(f,med)}^{(l)}(s)$ ($P_{(f,str)}^{(l)}(s)$) be the probability of failure against a medium (strong) adversary for a DA attack on the base layer of the CMT. (Define similarly $P_{(f,med)}^{(j)}(s)$ and $P_{(f,str)}^{(j)}(s)$ for DA attack on layer j). We have the following lemma (all proofs are deferred to the Appendix).

Lemma 2. For a sampling strategy $(\mathbf{x}, \beta^{(l)})$, $P_{(f,med)}^{(l)}(s) = [\max(1 - \Pi^{(l)}\mathbf{x})]^s$. Define $P_{(f,str-bnd)}^{(l)}(s) := (1 - \beta^{(l)}\mu_l)^s$. Then, $P_{(f,str)}^{(l)}(s) \leq \max(P_{(f,med)}^{(l)}(s), P_{(f,str-bnd)}^{(l)}(s))$.

In the rest of the paper, we assume that $P_{(f,str)}^{(l)}(s)$ is equal to the upper bound provided in Lemma 2. We find the light node sampling strategy by formulating a linear program (LP) in $(\mathbf{x}, \beta^{(l)})$ to minimize the probabilities in Lemma 2. The optimization problem (which can be easily converted into an LP by introducing additional variables) is provided below:

$$\underset{\mathbf{x}, \beta^{(l)}}{\text{minimize}} \quad \max\left(\max(1 - \Pi^{(l)}\mathbf{x}), \theta \times [1 - \beta^{(l)}\mu_l]\right) \quad (1)$$

$$\text{subject to} \quad \beta^{(l)} \leq x_i \leq 1, \quad i = 1, \dots, n_l; \quad \beta^{(l)} \geq 0; \quad \sum_{i=1}^{n_l} x_i = 1,$$

where θ , $0 \leq \theta \leq 1$, is a parameter that controls the trade-off between $P_{(f,med)}^{(l)}(s)$ and $P_{(f,str)}^{(l)}(s)$.

A. Linear-programming-sampling (LP-sampling) for DA attacks on any layer of the CMT

In this subsection, we modify LP (1) to take into effect a DA attack conducted on any layer

of the CMT and derive the sampling strategy based on the modified LP. We first align the columns of the parity check matrices of all the CMT layers as described in Section III-A. Assume that the stopping sets and VNs in the following are based on the aligned parity check matrices.

Recall that a base layer symbol samples, via its Merkle proof, two symbols from every intermediate layer of the CMT. Hence, the events of sampling intermediate layer symbols are not disjoint. To calculate the probability that each intermediate layer symbol is sampled, we define for each j , $1 \leq j \leq l-1$, a matrix $A^{(j)}$ of size $n_j \times n_l$ whose entries are as follows: 1) if $(k \leq s_j \text{ and } 1 + (i-1)_{s_j} = k)$ then $A_{ki}^{(j)} = 1$; 2) if $(k > s_j \text{ and } 1 + s_j + (i-1)_{p_j} = k)$ then $A_{ki}^{(j)} = 1$; 3) $A_{ki}^{(j)} = 0$ for all other cases. For simplicity, assume that $A^{(l)}$ is an $n_l \times n_l$ identity matrix. Also, define for $1 \leq j \leq l$, the matrices $\Delta^{(j)} = \min(\Pi^{(j)} A^{(j)}, 1)$ where the minimum is element wise. Using the above matrices, we calculate $P_{(f,med)}^{(j)}(s)$ and $P_{(f,str)}^{(j)}(s)$ in Lemma 3. First, consider the following definition.

Definition 3. A sampling (with replacement) strategy $(x, \beta^{(1)}, \beta^{(2)}, \dots, \beta^{(l)})$ is a sampling strategy $(x, \beta^{(l)})$, such that for $x^{(j)} = A^{(j)}x$, $1 \leq j \leq l-1$, $x_i^{(j)}$'s satisfy $x_i^{(j)} \geq \beta^{(j)}$, $1 \leq i \leq n_j$, $1 \leq j \leq l-1$. Parameter $\beta^{(j)}$, $1 \leq j \leq l$, is a non-negative real number and controls the minimum probability of requesting a given symbol from layer j of the CMT.

Lemma 3. For a sampling strategy $(x, \beta^{(1)}, \dots, \beta^{(l)})$, let $x^{(j)} = A^{(j)}x$, $1 \leq j \leq l$. $x_k^{(j)}$ is the probability that $v_k^{(j)}$ is sampled and $P_{(f,med)}^{(j)}(s) = [\max(1 - \Delta^{(j)}x)]^s$. Also, for $1 \leq j < l$, let $P_{(f,str-bnd)}^{(j)}(s) := (1 - \frac{1}{2}\beta^{(j)}\mu_j)^s$. Then, $P_{(f,str)}^{(j)}(s) \leq \max(P_{(f,str-bnd)}^{(j)}(s), P_{(f,med)}^{(j)}(s))$.

Using Lemmas 2 and 3, we formulate the following LP to find the light node sampling strategy:

$$\begin{aligned} \underset{x, \beta^{(1)}, \dots, \beta^{(l)}}{\text{minimize}} \quad & \max \left(\max_{1 \leq j \leq l} \max(1 - \Delta^{(j)}x), \max_{1 \leq j \leq l} \theta^{(j)} \times [1 - \xi^{(j)}\beta^{(j)}\mu_j] \right) \end{aligned} \quad (2a)$$

$$\text{subject to} \quad \beta^{(l)} \leq x_i \leq 1, \quad i = 1, \dots, n_l, \quad \sum_{i=1}^{n_l} x_i = 1, \quad (2b)$$

$$\beta^{(j)} \leq \min(A^{(j)}x), \quad j = 1, \dots, l-1; \quad \beta^{(j)} \geq 0, \quad j = 1, \dots, l, \quad (2c)$$

where $\xi^{(j)} = \frac{1}{2}$ for $1 \leq j < l$ and $\xi^{(l)} = 1$. The first and second term in the outer maximum above corresponds to the probability of failure against the medium and strong adversary for a DA attack on different layers of the CMT. $\theta^{(j)}$'s are trade-off parameters and control the importance given to a strong adversary on layer j of the CMT compared to a medium adversary.

The sampling strategy $(x, \beta^{(1)}, \dots, \beta^{(l)})$ obtained as the optimal solution of LP (2) is called LP-sampling and is included in the protocol. To reduce the probability of failure against a medium and a strong adversary under LP-sampling, we design LDPC codes aimed towards minimizing

the probability for each layer. The complexity of LP-sampling is determined by the complexity of finding all stopping sets of H_j of size $< \mu_j$. Although stopping set enumeration is NP-hard, they can be found in a reasonable time for small code lengths using an ILP [32]. However, it is difficult to obtain an analytical complexity expression for stopping set enumeration using ILP.

B. Linear-programming-Constrained PEG (LC-PEG) Algorithm

In this section, we design LDPC codes that perform well under LP-sampling. We design such codes by modifying the CN selection procedure in the PEG algorithm. We call our construction *linear-programming-constrained* PEG or LC-PEG algorithm since it is trying to minimize the optimal objective value of an LP. Codes designed in this section are aligned by Algorithm 2 and then included in the protocol. Similar to the EC-PEG algorithm, we optimize cycles instead of stopping sets. The motivation for focusing on cycles is the following: for lists $\mathcal{C}$ and Ψ of cycles and stopping sets, respectively, such that for every $\psi \in \Psi$, there exists a $\mathcal{O} \in \mathcal{C}$ which is part of ψ , we have $\max_{\psi \in \Psi} (1 - \sum_{v_i: v_i \in \psi} x_i) \leq \max_{\mathcal{O} \in \mathcal{C}} (1 - \sum_{v_i: v_i \in \mathcal{O}} x_i)$. Thus, the optimal objective value of LP (1) can be upper bounded by the optimal objective value of a modified version of LP (1) which is based on cycles. We select CNs in the PEG algorithm depending on the optimal objective value they produce on the modified LP. Algorithm 4 presents our LC-PEG algorithm for constructing a TG $\tilde{\mathcal{G}}$ with n VNs, m CNs, and VN degree d_v . All ties are broken randomly.

Algorithm 4 LC-PEG Algorithm

```

1: Inputs:  $n, m, d_v, g_c, T_{th}, \hat{\theta}, \hat{\mu}$ ; Outputs:  $\tilde{\mathcal{G}}, g_{\min}$ 
2: Initialize  $\tilde{\mathcal{G}}$  to  $n$  VNs,  $m$  CNs and no edges,  $\mathcal{L} = \emptyset$ 
3: for  $j = 1$  to  $n$  do
4:   for  $k = 1$  to  $d_v$  do
5:      $[\mathcal{K}, g] = \text{PEG}(\tilde{\mathcal{G}}, v_j)$ 
6:      $\mathcal{K}_{\mindeg} = \text{CNs in } \mathcal{K} \text{ with the minimum degree under the current TG setting } \tilde{\mathcal{G}}$ 
7:     if  $g \geq g_c$  then  $c^{sel} = \text{Select a CN randomly from } \mathcal{K}_{\mindeg}$ 
8:     else  $\triangleright (g\text{-cycles, } g < g_c, \text{ are created})$ 
9:        $\mathcal{K}_{mincycles} = \text{CNs in } \mathcal{K}_{\mindeg} \text{ that result in the minimum number of new } g\text{-cycles due}$ 
 $\text{to the addition of edge between the CN and } v_j$ 
10:      for each  $c$  in  $\mathcal{K}_{mincycles}$  do
11:         $\mathcal{L}_{cycles}^c = \text{new } g\text{-cycles formed in } \tilde{\mathcal{G}} \text{ due to the addition of edge between } c \text{ and } v_j$ 
12:         $\text{cost}[c] = \text{LP-objective}(\mathcal{L} \cup \mathcal{L}_{cycles}^c, \tilde{\mathcal{G}})$ 
13:         $c^{sel} = \text{CN in } \mathcal{K}_{mincycles} \text{ with minimum cost}[c]$ 
14:         $\mathcal{L}^{sel} = \text{cycles in } \mathcal{L}_{cycles}^{c^{sel}} \text{ that have EMD} \leq T_{th}; \mathcal{L} = \mathcal{L} \cup \mathcal{L}^{sel}$ 
15:       $\tilde{\mathcal{G}} = \tilde{\mathcal{G}} \cup \text{edge}\{c^{sel}, v_j\}$ 

```

In the LC-PEG algorithm, we use the concept of the extrinsic message degree (EMD) of a set of VNs that allows us to rank the harm a cycle may have in creating stopping sets. EMD of a set

of VNs is the number of CN neighbors singly connected to the set [23] and is calculated using the method in [24]. EMD of a cycle is the EMD of the VNs involved in the cycle. Low EMD cycles are more likely to form stopping sets and we term cycles with EMD below a threshold T_{th} as *bad cycles*. We use bad cycles to form the modified linear program below:

$$\min_{\hat{x}, \hat{\beta}} \max \left(\max(1 - C\hat{x}), \hat{\theta} \times [1 - \hat{\beta}\hat{\mu}] \right) \text{ s.t. } \hat{\beta} \leq \hat{x}_i \leq 1, i = 1, \dots, \hat{n}; \hat{\beta} \geq 0; \sum_{i=1}^{\hat{n}} \hat{x}_i = 1. \quad (3)$$

The LC-PEG algorithm uses LP (3) via the procedure $\text{LP-objective}(\hat{\mathcal{L}}, \hat{\mathcal{G}})$ which outputs its optimal objective value. The procedure has inputs of a list $\hat{\mathcal{L}} = \{\mathcal{O}_1, \dots, \mathcal{O}_{|\hat{\mathcal{L}}|}\}$ of cycles and a TG $\hat{\mathcal{G}}$. Let $\hat{\mathcal{G}}$ have $\hat{n}$ VNs $\{\hat{v}_1, \dots, \hat{v}_{\hat{n}}\}$. Here, C is a matrix of size $|\hat{\mathcal{L}}| \times \hat{n}$, such that $C_{ki} = 1$ if $\hat{v}_i$ touches $\mathcal{O}_k$, else $C_{ki} = 0$, $1 \leq i \leq \hat{n}$, $1 \leq k \leq |\hat{\mathcal{L}}|$. Also, $\hat{\theta}$, $0 \leq \hat{\theta} \leq 1$, is a parameter.

In the LC-PEG algorithm, we use the procedure $\text{PEG}()$ defined in Section III-B for the EC-PEG algorithm. The LC-PEG algorithm proceeds exactly as the EC-PEG algorithm when the $\text{PEG}()$ procedure returns cycle length $g \geq g_c$. When the $\text{PEG}()$ procedure returns cycle length $g < g_c$, we select a CN from the set of *candidate* CNs $\mathcal{K}$ such that the resultant LDPC codes have a low optimal objective value of LP (1). We explain the CN selection procedure next.

While progressing through the LC-PEG algorithm, we maintain a list $\mathcal{L}$ of cycles. $\mathcal{L}$ contains cycles of length $g < g_c$ that had EMD less than or equal to threshold T_{th} when they were formed. Cycles in $\mathcal{L}$ are considered *bad* cycles and we base our CN selection procedure on these cycles. When the $\text{PEG}()$ procedure returns candidate CNs $\mathcal{K}$, we first select the set of CNs $\mathcal{K}_{mindeg}$ that have the minimum degree under the current TG setting $\tilde{\mathcal{G}}$ (line 6). Of the CNs in $\mathcal{K}_{mindeg}$, we select the set of CNs $\mathcal{K}_{mincycles}$ that form the minimum number of new g -cycles if an edge is established between the CN and v_j (line 9). Now for every CN c in $\mathcal{K}_{mincycles}$, we find the list $\mathcal{L}_{cycles}^c$ of new g -cycles formed due to the addition of an edge between c and v_j (line 11) and compute $\text{LP-objective}(\mathcal{L} \cup \mathcal{L}_{cycles}^c, \tilde{\mathcal{G}})$ to get $\text{cost}[c]$ (line 12). Our modified CN selection procedure is to select a CN in $\mathcal{K}_{mincycles}$ that has the minimum $\text{cost}[c]$ (line 13). After selecting c^{sel} using the above criteria, we update $\mathcal{L}$ as follows: let $\mathcal{L}^{sel}$ be the list of g -cycles in $\mathcal{L}_{cycles}^{c^{sel}}$ that have $\text{EMD} \leq T_{th}$. We add $\mathcal{L}^{sel}$ to $\mathcal{L}$ (line 14). Finally, we update the TG $\tilde{\mathcal{G}}$ (line 15).

Remark 4. We empirically observed that reducing the number of cycles in the TG (and hence the number of stopping sets) reduces the probability of failure against the medium and strong adversary when LP-sampling is employed. The above holds even if the size of the smallest stopping set remains unchanged. This is in contrast to random sampling where the probability of failure only depends on the size of the smallest stopping set and is agnostic to the number of stopping sets of small size present in the code. Thus, based on this observation, we have added

line 9 in our LC-PEG algorithm which selects CNs $\mathcal{K}_{\text{mincycles}}$ that form the minimum number of cycles when a new edge is established. However, we further make an informed choice among the CNs in $\mathcal{K}_{\text{mincycles}}$ to select a CN that has the minimum optimal objective value of LP (3).

We now discuss the complexity of the LC-PEG algorithm. Note that it differs from the original PEG algorithm (that has complexity $O(mn)$ [17]) in steps 8-14. Of these steps, step 12 has the largest complexity due to solving LP (3). An LP $\min_{Az \leq b} c^T z$ with d variables and t constraints can be solved with complexity $\tilde{O}((nnz(A) + d^2)\sqrt{d})$ [33] where $nnz(A)$ is the number of non-zero entries in A and $\tilde{O}$ hides factors poly-logarithmic in d and t . In our case, LP (3) has n variables and at most mnd_v constraints (since step 11 in the algorithm can result in at most m cycles) and hence $nnz(A) \leq g_c mnd_v$. Thus, the overall complexity of the LC-PEG algorithm is at most $\tilde{O}(mn\sqrt{n}(g_c mnd_v + n^2)) = \tilde{O}(n^{4.5})$.

V. SYSTEM ASPECTS

1) *Security Performance*: Here, we discuss how *soundness* and *agreement* defined in Section II-D are affected by our co-design. Let M be the total number of light nodes in the system and $\eta_{\text{rec}} = \left(\max_{1 \leq j \leq l} \frac{n_j - \omega_{\min}^{(j)} + 1}{n_j} \right)$. We have the following lemmas (we defer the proofs to the Appendix).

Lemma 4. *For a weak adversary, when light nodes sample according to the overall greedy sampling strategy, the probability of soundness or agreement failure in the system $P_f^{S,A}$ satisfies*

$$P_f^{S,A} \leq \max \left(\max_{1 \leq j \leq l, \omega^{(j)} < \mu_j} \left[1 - \tau(S_{\text{greedy}}^{(\rho s, j)}, \omega^{(j)}) \right] \left(1 - \frac{\omega^{(j)}}{n_j} \right)^{s - \rho s} \right], 2^{\lceil \mathcal{H}(\eta_{\text{rec}}, 1 - \eta_{\text{rec}}) n_l - Ms(1 - \rho) \log(\frac{1}{\eta_{\text{rec}}}) \rceil} \right).$$

Here, $S_{\text{greedy}}^{(\rho s, j)}$ is the samples of layer j , $1 \leq j \leq l$, collected when the light nodes request for the first ρs coded symbols from the base layer of the CMT.

Lemma 5. *For a medium and a strong adversary, when light nodes sample according to LP-sampling $\mathbf{x}$, the probability of soundness or agreement failure in the system $P_f^{S,A}$ satisfies*

$$P_f^{S,A} \leq \max \left(\max_{1 \leq j \leq l} P_f^{(j)}(s), 2^{\lceil \mathcal{H}(\eta_{\text{rec}}, 1 - \eta_{\text{rec}}) n_l - Ms \log \left(\frac{1}{\sum_{i=1}^l \eta_{\text{rec}} n_i x_{[i]} \right)} \rceil} \right).$$

Here, $P_f^{(j)}(s) = P_{(f, \text{med})}^{(j)}(s)$ and $P_f^{(j)}(s) = P_{(f, \text{str})}^{(j)}(s)$ for the medium and strong adversary, respectively, as defined in Section IV-A and $x_{[i]}$ is the i^{th} largest entry in vector $\mathbf{x}$.

The first term in the maximum in Lemma 4 and 5 is the probability of failure of a single light node against different adversaries. Thus, when the number of light nodes M is large, $P_f^{S,A}$ is affected by the probability of failure of a single light node, which we minimize in this paper.

2) *Blockchain System Designer*: In the system model in Section II-C, we have made a trusted set up assumption of a blockchain system designer who designs the parity checks matrices and the LP-sampling strategy. Note that for greedy sampling, after the overall sampling rule described in

Remark 3, nothing more needs to be designed by the system designer. Additionally, as previously mentioned in Section II-D, only the final LP-sampling strategy obtained by solving LP (2) is included in the protocol and inputs to LP (2) (μ_j and set of all stopping sets of H_j of size $< \mu_j$) are not part of the protocol. Existing examples of blockchain systems that rely on trusted set up assumptions include [30], [34], [35]. In our system, there are two attacks possible by a compromised designer: i) incorrect protocol design (i.e., the designed sampling strategy and LDPC codes do not result in the claimed probability of failure. Here, the probability of failure can be thought of as an output of the protocol design computation task and nodes join the system based on the published probability of failure performance); ii) the designer acts as the adversary and launches a DA attack using the known stopping sets of H_j of size $< \mu_j$.

A possible direction to remove the first attack is as follows. A cryptographic tool called zk-STARK [36] can be used by the blockchain protocol designer to create verifiable proofs of correct computation of the LDPC codes, the LP-sampling strategy, and the probability of failure. This is a one time proof that can be verified by nodes (full and light) before joining the blockchain system to ensure that the protocol is correctly designed. The proof created using zk-STARK has the following important properties: it has a small size, it can be verified using significantly less computational complexity compared to the actual computation, it is secure against quantum computers, it reveals no information about the secrets involved in the computation (in this case μ_j and all stopping sets of H_j of size $< \mu_j$).

In the second attack, the system designer acts as the adversary (medium) to launch a DA attack using the knowledge of the stopping sets (which it enumerated while correctly designing LP-sampling x). However, this DA attack will be detected by the light nodes with a probability of failure $P_{(f,med)}^{(j)}(s)$ which is guaranteed by the protocol. Also, to launch this DA attack, the system designer spends the same amount of computational power as a medium adversary who doesn't have the knowledge of the stopping sets and wishes to attack the system. Thus, the system designer is not at an advantage to launch DA attacks due to the knowledge of the secret.

VI. SIMULATION RESULTS

In this section, we compare the performance of our co-design techniques with that of codes designed by the original PEG algorithm and the performance of [2] using random LDPC codes and random sampling (RS). The different CMTs used for simulation are parametrized by $\mathcal{T} = (n_l, R, q, l)$ (individual parameters are defined in Section II-A). For a CMT $\mathcal{T}$, in order to compare the performance of different PEG based codes, we choose $\mu_j = \omega_{\min}^{(j),PEG} + \gamma$, $1 \leq j \leq l$, for the

Fig. 4: The probability of light node failure for various coding schemes and sampling strategies for CMT $\mathcal{T}_1 = (128, 0.5, 4, 4)$; Left panel: probability of failure for a DA attack on the base layer for different stopping set sizes. The black curve is achieved using stopping ratio $\nu^* = 0.064353$. The value ν^* is the best stopping ratio obtained for a rate 0.5 code following the method in [2, section 5.3] using parameters $(c, d) = (8, 16)$. GS refers to the overall greedy sampling strategy described in Section V 1) where we have used $\rho = 0.9$. $P_{f,\omega}^{(j)}(s)$ for GS is calculated as $[1 - \tau(S_{greedy}^{(\rho s, j)}, \omega)](1 - \frac{\omega}{n_j})^{s - \rho s}$, where $S_{greedy}^{(\rho s, j)}$ is described in Lemma 4; Right panel: probability of failure across different layers of the CMT. $P_f^{(j)}(s)$ is calculated as $P_f^{(j)}(s) = \max_{\omega < \mu_j} P_{f,\omega}^{(j)}(s)$, where $\mu_j = \omega_{\min}^{(j), PEG} + 6$.

various adversary models described in Section II-D. Here, $\omega_{\min}^{(j), PEG}$ is the minimum stopping set size for an LDPC code constructed using the original PEG algorithm for layer j of the CMT $\mathcal{T}$ and γ is a parameter. We calculate the probability of failure when the light nodes request for s base layer samples using random sampling for various scenarios as follows: for the base layer when the adversary hides a stopping set of size ω , $P_{f,\omega}^{(l)}(s) = (1 - \frac{\omega}{n_l})^s$; for intermediate layers, we calculate the probability of failure for the medium and strong adversary by substituting $\mathbf{x} = \frac{1_{n_l}}{n_l}$ in the probability of failure expressions provided in Section IV-A, where 1_{n_l} is a vector of ones of length n_l ; for an LDPC code with a stopping ratio ν^* we calculate the probability of failure at the base layer using random sampling as $P_f^{(l)}(s) = (1 - \nu^*)^s$. The LDPC codes at different layers of the CMTs are aligned using Algorithm 2 where we use $g_{\max} = g_c$ (observed cycles in the code constructions) and $g_{\min}$ is set to the girth of the respective codes.

Fig. 4 demonstrates the performance of the EC-PEG algorithm and the greedy sampling strategy for CMT $\mathcal{T}_1 = (128, 0.5, 4, 4)$ and a weak adversary. For the EC-PEG algorithm, we have used the parameters: $d_v = 4$ for all layers, $R = 0.5$, $g_c^{(4)} = 10$ and $g_c^{(j)} = 8$ for $j = 1, 2, 3$. For the adversary model we have chosen $\gamma = 6$. Note that $\omega_{\min}^{(4), PEG} = 9$ and thus $\mu_4 = 9 + 6 = 15$. In Fig. 4 left panel, we plot $P_{f,\omega}^{(4)}(s)$ for various coding algorithms and sampling strategies when a weak adversary conducts a DA attack on the base layer of the CMT by hiding stopping sets of size $\omega < \mu_4$. The codes designed by the original PEG and EC-PEG algorithms have a minimum stopping set size of 9 and 10, respectively. For these algorithms, $P_{f,\omega}^{(4)}(s)$ quickly becomes zero for

Fig. 5: Weak adversary performance plots for $N_l = 200$, $R = 0.5$, $d_v = 4$. Left panel: probability of failure for a DA attack on the base layer for different stopping set sizes (see Fig. 4 left panel for plot properties). We have parameters $\omega_{\min}^{PEG} = 13$ and $\gamma = 5$, $\rho = 0.9$; Right panel: stopping set distribution ss^{17} .

$\omega = 9, 10$ using greedy sampling as s increases. Hence, we have not included these stopping set sizes in Fig. 4 left panel. The figure demonstrates three benefits of our co-design. The first benefit is due to the use of deterministic LDPC codes that provide larger stopping set sizes than random ensembles, as can be seen when comparing the black and green curves. The second benefit comes from using greedy sampling as opposed to random sampling, which can be observed by comparing the green and red curves. The final benefit is provided by the EC-PEG algorithm, as can be seen by comparing the red and blue curves. These benefits combine to significantly reduce $P_{f, \omega}^{(4)}(s)$ compared to the black curve which was proposed in earlier literature¹⁰.

In Fig. 4 right panel, we plot the probability of failure $P_f^{(j)}(s)$ when the weak adversary conducts a DA attack on layer j of CMT $\mathcal{T}_1$. From Fig. 4 right panel, we see that the base layer of the CMT (L_4) has a larger probability of failure compared to other layers and the probability of failures for the intermediate layers quickly become very small. This is due to the alignment of the columns of the parity check matrices, which ensures that each intermediate layer is greedily sampled¹¹. We next observe that the EC-PEG algorithm with greedy sampling results in a lower $P_f^{(j)}(s)$ compared to the original PEG algorithm for all layers of the CMT. Moreover, for the base layer, $P_f^{(4)}(s)$ (for both EC-PEG and original PEG coupled with greedy sampling) is lower than the probability of failure using random sampling for $\omega = 14$ (green curve) and the probability of failure achieved by random LDPC codes and random sampling (black curve). Thus, in combination, the co-design of concentrated LDPC codes and greedy sampling results

¹⁰The singularities in some plots in Fig. 4 (e.g., Original PEG + GS $\omega = 11$) is because $P_{f, \omega}^{(4)}(s)$ becomes zero after certain number of greedy samples. This situation happens when all the stopping sets of weight ω get touched by the greedy samples.

¹¹The plots for $P_{f, \omega}^{(4)}(s)$ and $P_f^{(j)}(s)$ in Fig. 4 sometimes exhibit floors (i.e., they remain constant for different values of s). This is due to i) the new greedy samples that are selected (on increasing s) do not increase the number of stopping sets that are touched; ii) the number of random samples in the overall greedy sampling strategy remain same on increasing s .

Fig. 6: For CMT $\mathcal{T}_1 = (128, 0.5, 4, 4)$, the probability of failure for a DA attack at different layers when the base layer is randomly sampled and the probability of failure when a strong and a medium adversary conduct a DA attack on different layers and the base layer is sampled using LP-sampling. For all layers, we plot $P_{(f, str-bnd)}^{(j)}(s)$ instead of $P_{(f, str)}^{(j)}(s)$. Left panel: Original PEG; Middle panel: LC-PEG; Right panel: MC-PEG.

in a significantly lower $P_{f, \omega}^{(4)}(s)$ compared to methods proposed in [2]. To illustrate the benefits of the EC-PEG algorithm and the greedy sampling strategy, we provide plots similar to Fig. 3 and Fig. 4 for a different choice of code parameters in Fig. 5. From the figure, we see similar stopping set concentration and probability of failure improvement as in Fig. 3 and Fig. 4.

In Figs. 6, 7 and Table I, we demonstrate the performance of the LC-PEG algorithm and LP-sampling (LS) against a medium and a strong adversary. Fig. 6 and 7 correspond to CMT $\mathcal{T}_1 = (128, 0.5, 4, 4)$ where we have used $\gamma = 4$, thus $\mu_j = \omega_{\min}^{(j), PEG} + 4$. Table II lists $\omega_{\min}^{(j), PEG}$ for different j . Additionally, for LP-sampling, we have used $\theta^{(4)} = 0.993$, $\theta^{(j)} = 1$, $j = 1, 2, 3$. For the LC-PEG algorithm, we have used $d_v = 4$ for all layers, $R = 0.5$, $g_c^{(4)} = 10$ and $g_c^{(j)} = 8$ for $j = 1, 2, 3$, $T_{th}^{(j)} = 3$ for $j = 1, 2$, $T_{th}^{(j)} = 4$ for $j = 3, 4$, $\hat{\theta}^{(j)} = 0.997$, $j = 1, 2, 3, 4$ (we tested with $T_{th} = 3, 4$ and $\hat{\theta} = 0.995, 0.996, 0.997, 0.998$ and picked the codes that provide the lowest $P_f^{J^{\max}}(s)$) and $\hat{\mu}_j = \mu_j$, $j = 1, 2, 3, 4$. To demonstrate the effectiveness of the LC-PEG algorithm, we also plot the performance of an algorithm termed as the *Minimum-Cycles* PEG (MC-PEG) algorithm. It is the same as the LC-PEG algorithm but instead of the CN selection steps in lines 10-13 of Algorithm 4, the MC-PEG algorithm selects a CN randomly from $\mathcal{K}_{mincycles}$ as c^{sel} .

We first look at the improvements provided by LP-sampling. Fig. 6 shows the performance of LP-sampling for a DA attack at different layers of the CMT constructed using the PEG, LC-PEG and MC-PEG algorithms. We see that while the probability of failure for some layers worsens in comparison to random sampling, for the worst layer, which is the base layer, the probability of failure improves for both the strong and medium adversary. We generally find that the base layer is the worst layer so we focus on the base layer in the subsequent simulations.

We plot $P_f^{(4)}(s)$ vs. s for the PEG, MC-PEG, and LC-PEG algorithms using LP-sampling in Fig. 7 left panel. From the figure we see the following improvements. The first improvement is

Fig. 7: The probability of light node failure for a DA attack on the base layer of CMT $\mathcal{T}_1 = (128, 0.5, 4, 4)$; Left panel: comparison of different coding schemes and sampling strategies. The black curve uses $\nu^* = 0.064353$; Right panel: variation in $P_f^{(4)}(s = 30)$ for the strong and medium adversary as a function of $\theta^{(4)}$ for $\theta^{(j)} = 1, j = 1, 2, 3$. TABLE I: $P_f^{(l)}(s = 0.25n_l)$ for a DA attack on the base layer of CMT for various CMT parameters, coding schemes, and sampling strategies. The parameters used for the different CMTs is listed in Table II. For $R = 0.4$, we get stopping ratio $\nu^* = 0.085059$ following the method in [2, Section 5.3] using parameters $(c, d) = (6, 10)$.

CMT $\mathcal{T} = (n_l, R, q, l)$	Random Sampling				LP-Sampling					
	Ensemble	PEG	MC-PEG	LC-PEG	Strong Adversary			Medium Adversary		
					PEG	MC-PEG	LC-PEG	PEG	MC-PEG	LC-PEG
(128, 0.5, 4, 4)	0.1190	0.0970	0.0740	0.0428	0.04914	0.04602	0.04106	0.03925	0.03675	0.03279
(208, 0.5, 4, 4)	0.0314	0.0204	0.0155	0.0267	0.0304	0.02427	0.02294	0.00815	0.00651	0.00615
(200, 0.5, 4, 3)	0.0356	0.0347	0.0202	0.0202	0.02778	0.02028	0.01781	0.00606	0.00447	0.00388
(200, 0.4, 5, 4)	0.0117	0.0089	0.0067	0.0052	0.00558	0.00513	0.00484	0.00225	0.00207	0.00195

between the black and magenta curves due to using deterministic LDPC codes that produce larger stopping set sizes. The second improvement is due to using LP-sampling compared to random sampling. Compared to random sampling (magenta curve), LP-sampling with the original-PEG algorithm results in a lower probability of failure for the medium (red-solid curve) and strong adversary (red-dotted curve). The third improvement (between the red and light blue curves) comes from utilizing the key characteristic of the MC-PEG algorithm to reduce the number of small cycles as discussed in Remark 4. The final improvement comes from the informed check node selection in the LC-PEG algorithm to create tailored codes for LP-sampling as seen by comparing the dark and light blue curves. Overall, the co-design of the LC-PEG algorithm and LP-sampling results in a lower $P_f^{(4)}(s)$ compared to other techniques.

In Fig. 7 right panel, we plot $P_f^{(4)}(s = 30)$ as a function of the parameter $\theta^{(4)}$ for the original PEG, MC-PEG and LC-PEG algorithms using LP-sampling. From Fig. 7 right panel, we see that $\theta^{(4)}$ controls the trade-off between the probabilities of failure for the medium adversary and strong adversary. Thus, $\theta^{(4)}$ can be chosen as a hyper-parameter based on the system specifications. We also see from Fig. 7 right panel that for all the values of $\theta^{(4)}$, the LC-PEG algorithm outperforms the PEG and MC-PEG algorithm for both the medium and strong adversary.

TABLE II: Parameters used for LP-sampling and LC-PEG code construction for various CMTs in Table I. For all LDPC codes we use $d_v = 4$, $g_c^{(j)} = g_{\max}^{(j)} = g_{\min}^{(j)} + 4$, $j = 1, \dots, l$. For LC-PEG algorithm, we use $\hat{\mu}^{(j)} = \mu^{(j)} = \omega_{\min}^{(j),PEG} + \gamma$, $j = 1, \dots, l$. Under each variable that depends on the layer, we enumerate the layer numbers.

$\mathcal{T} = (n_l, R, q, l)$	$T_{th}^{(j)}$				$\hat{\theta}^{(j)}$				$g_{\min}^{(j)}$				$\omega_{\min}^{(j),PEG}$				γ	$\theta^{(j)}$			
	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4		1	2	3	4
(128, 0.5, 4, 4)	3	3	4	4	0.997	0.997	0.997	0.997	4	4	4	6	2	4	5	9	4	1	1	1	0.993
(208, 0.5, 4, 4)	3	4	4	5	0.997	0.997	0.997	0.9959	4	4	6	6	4	6	10	15	3	1	1	0.997	0.975
(200, 0.5, 4, 3)	4	4	5	-	0.997	0.997	0.998	-	4	6	6	-	6	8	13	-	3	1	0.99	0.97	-
(200, 0.4, 5, 4)	3	4	4	5	0.997	0.997	0.997	0.998	4	4	6	6	5	8	11	18	4	1	1	0.992	0.982

TABLE III: Maximum CN degree for the LDPC codes used in different layers of the CMT. Under each algorithm, we enumerate the layer numbers and specify the maximum CN degree for that layer.

$\mathcal{T} = (n_l, R, q, l)$	Ensemble	PEG				EC-PEG				MC-PEG				LC-PEG			
		1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
(128, 0.5, 4, 4)	16	8	8	9	9	10	13	12	14	8	8	9	9	8	8	9	9
(208, 0.5, 4, 4)	16	8	9	8	8	11	11	11	16	8	9	8	9	8	8	9	9
(200, 0.5, 4, 3)	16	9	9	9	-	12	11	18	-	9	9	9	-	8	9	9	-
(200, 0.4, 5, 4)	10	7	7	7	7	12	9	11	14	7	7	7	8	7	8	7	7

For completeness, we provide further examples of how our novel code constructions improve the probability of failure for different CMT parameters. In Table I, we list $P_f^{(l)}(s = 0.25n_l)$ and compare various sampling strategies and LDPC code constructions. Similar to Fig. 7 left panel, from Table I, we see that the novel co-design of the LC-PEG algorithm and LP-sampling results in the lowest probability of failure for the different CMT parameters.

In Table III, we compare the maximum CN degree for the LDPC codes used in different CMT layers for various construction techniques. We see that PEG based constructions have similar maximum CN degrees compared to the ensemble LDPC codes used in [2]. Since the IC proof size is proportional to the maximum CN degree, we conclude that the new LDPC code constructions do not significantly impact the IC proof size to improve the probability of failure.

VII. CONCLUSION

In this paper, we considered the problem of DA attacks pertinent to blockchain systems with light nodes. For various strengths of the malicious nodes, we demonstrated that, at short code lengths, a suitable co-design of specialized LDPC codes and the light node sampling strategy can result in a much lower probability of failure for the light nodes to detect DA attacks compared to schemes proposed in prior literature. As future work, we are currently investigating other code families like Reed-Muller Codes and Polar codes for use in this application.

REFERENCES

- [1] M. Al-Bassam, *et al.*, "Fraud and data availability proofs: Detecting invalid blocks in light clients," *Int. Conf. on Financial Cryptography and Data Secur.* Springer, Mar. 2021.
- [2] M. Yu, *et al.*, "Coded merkle tree: Solving data availability attacks in blockchains," *Int. Conf. on Financial Cryptography and Data Secur.*, Springer, Feb. 2020.
- [3] D. Mitra, *et al.*, "Concentrated stopping set design for coded merkle tree: Improving security against data availability attacks in blockchain systems," *2020 IEEE Inf. Theory Workshop (ITW)*, Apr. 2021, full version: <https://arxiv.org/abs/2010.07363>.

- [4] S. Nakamoto, "Bitcoin: A peer to peer electronic cash system," 2008. [Online] Available: <https://bitcoin.org/bitcoin.pdf>.
- [5] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, Apr. 2014.
- [6] M. J. Casey and P. Wong, "Global supply chains are about to get better, thanks to blockchain," *Harvard Business Review*, Mar. 2017. [Online] Available: <https://hbr.org/2017/03/global-supply-chains-are-about-to-get-better-thanks-to-blockchain>.
- [7] X. Wang, *et al.*, "Survey on blockchain for internet of things," *Comput. Commun.*, vol. 136, pp. 10-29, 2019.
- [8] M. Mettler, "Blockchain technology in healthcare: The revolution starts here," *IEEE 18th Int. Conf. on e-Health Network., Apps., and Services (Healthcom)*, Sept. 2016.
- [9] T. Rocket, *et al.*, "Scalable and probabilistic leaderless BFT consensus through metastability," *arXiv:1906.08936*, Jun. 2019.
- [10] Online: <https://www.blockchain.com/charts/blocks-size>, accessed: Jan. 10, 2022.
- [11] Online: <https://etherscan.io/chartsync/chaindefault>, accessed: Jan. 10, 2022.
- [12] Light ethereum subprotocol, <https://github.com/ethereum/devp2p/blob/master/caps/les.md>, accessed: Jan. 10, 2022.
- [13] Online: <https://github.com/ethereum/research/wiki/A-note-on-data-availability-and-erasure-coding>
- [14] T. Richardson, and R. Urbanke, "Modern coding theory," *Cambridge: Cambridge University Press*, 2008.
- [15] X. Jiao, *et al.*, "Eliminating small stopping sets in irregular low-density parity-check codes," *IEEE Commun. Lett.*, vol. 13, no. 6, pp. 435-437, Jun. 2009.
- [16] K. M. Krishnan, and P. Shankar, "Computing the stopping distance of a Tanner graph is NP-hard," *IEEE Trans. on Inf. Theory*, vol. 53, no. 6, pp. 2278-2280, Jun. 2007.
- [17] X.Y. Hu, *et al.*, "Regular and irregular progressive edge-growth tanner graphs," *IEEE Trans. on Inf. Theory*, vol. 51, no. 1, pp. 386-398, Jan. 2005.
- [18] A. Orlitsky, *et al.*, "Stopping set distribution of LDPC code ensembles," *IEEE Trans. on Inf. Theory*, vol. 51, no. 3, pp. 929-953, Mar. 2005.
- [19] S. Cao, *et al.*, "CoVer: Collaborative light-node-only verification and data availability for blockchains," *IEEE Int. Conf. on Blockchain*, Nov. 2, 2020.
- [20] Trifecta Blockchain Team, "Trifecta: The blockchain trilemma solved," <http://pramodv.ece.illinois.edu/pubs/Whitepaper2019-9.pdf>
- [21] P. Sheng, *et al.*, "ACeD: Scalable data availability oracle," *arXiv:2011.00102*, Oct. 2020.
- [22] D. Mitra, *et al.*, "Communication-efficient LDPC code design for data availability oracle in side blockchains," *IEEE Inf. Theory Workshop (ITW)*, Oct. 2021.
- [23] T. Tian, *et al.*, "Construction of irregular LDPC codes with low error floors," *IEEE Int. Conf. on Commun.*, May 2003.
- [24] S. Kim, *et al.*, "LDPC code construction with low error floor based on the IPEG algorithm," *IEEE Commun. Lett.*, vol. 11, no. 7, pp. 607-609, Jul. 2007.
- [25] M. Dai, *et al.*, "A low storage requirement framework for distributed ledger in blockchain," *IEEE Access*, vol. 6, pp. 22970-22975, Mar. 2018.
- [26] Q. Huang, *et al.*, "Downsampling and transparent coding for blockchain" *arXiv:1911.01778*, Nov. 2019.
- [27] S. Kadhe, *et al.*, "SeF: A secure fountain architecture for slashing storage costs in blockchains," *arXiv:1906.12140*, Jun. 2019.
- [28] S. Li, *et al.*, "PolyShard: coded sharding achieves linearly scaling efficiency and security simultaneously," *IEEE Trans. on Inf. Forensics and Secur.*, vol. 16, Jul. 2020.
- [29] D. Perard, *et al.*, "Erasure code-based low storage blockchain node," *IEEE Int. Conf. on Internet of Things (iThings) and IEEE Green Comput. and Commun. (GreenCom) and IEEE Cyber, Physical and Social Comput. (CPSCom) and IEEE Smart Data (SmartData)*, Jul. 2018.

- [30] P. Daian, *et al.*, "Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake," *Financial Cryptography*, Sept. 2019.
- [31] S. Park, *et al.*, "Spacemint: A cryptocurrency based on proofs of space." *Financial Cryptography*, Springer, Feb. 2018.
- [32] A. Sarıduman, *et al.*, "An integer programming-based search technique for error-prone structures of LDPC codes," *AEU-Int. Journal of Electronics and Commun.*, vol. 8, no. 11, pp. 1097-1105, Nov. 2014.
- [33] Y. T. Lee, and A. Sidford, "Efficient inverse maintenance and faster algorithms for linear programming," *IEEE Annual Symp. on Foundations of Computer Science*, Oct. 2015.
- [34] E. B. Sasson, *et al.*, "Zerocash: decentralized anonymous payments from bitcoin," *IEEE Symp. on Secur. and Privacy*, May 2014.
- [35] K. Nazirkhanova, *et al.*, "Information dispersal with provable retrievability for rollups," *arXiv:2111.12323*, Nov. 2021.
- [36] E. B. Sasson, *et al.*, "Scalable, transparent, and post-quantum secure computational integrity," *IACR Cryptol. ePrint Arch*, 2018.

APPENDIX

1) *Proof of Lemma 2:* $P_{(f,med)}^{(l)}(s) = \max_{k \in \{1,2,\dots,|\Psi_l|\}} P_f^{(l)}(s; \psi_k^{(l)}) = \max_{k \in \{1,2,\dots,|\Psi_l|\}} \left(1 - \sum_{i: v_i^{(l)} \in \psi_k^{(l)}} x_i\right)^s$
 $= \left[\max(1 - \Pi^{(l)}x)\right]^s$. Recall that Ψ_j^∞ is set of all stopping sets of H_j . We have the following:
 $P_{(f,str)}^{(l)}(s) = \max_{\psi \in \Psi_l^\infty} P_f^{(l)}(s; \psi) = \max \left(\left[\max(1 - \Pi^{(l)}x)\right]^s, \max_{\psi \in \Psi_l^\infty, \text{size}(\psi) \geq \mu_l} P_f^{(l)}(s; \psi) \right)$
 $\leq \max \left(\left[\max(1 - \Pi^{(l)}x)\right]^s, \left[1 - \beta^{(l)}\mu_l\right]^s \right) = \left(\max \left(\max(1 - \Pi^{(l)}x), 1 - \beta^{(l)}\mu_l \right) \right)^s$. The second term in the maximum of $P_{(f,str)}^{(l)}(s)$ is because $\max_{\psi \in \Psi_l^\infty, \text{size}(\psi) \geq \mu_l} P_f^{(l)}(s; \psi) \leq \left(1 - \beta^{(l)}\mu_l\right)^s$.

2) *Proof of Lemma 3:* For $1 \leq j \leq l-1$, the i^{th} column of $A^{(j)}$ (see Section IV-A) corresponds to VN $v_i^{(l)}$ of the base layer and the non-zero positions in the i^{th} column (two per column) correspond to the symbols of layer j which are part of the Merkle proof of $v_i^{(l)}$. Thus, for a sampling strategy $(x, \beta^{(l)})$ and $x^{(j)} = A^{(j)}x$, $1 \leq j \leq l$, it is easy to see that $x_k^{(j)}$ is the probability that $v_k^{(j)}$ is sampled. Now, consider a stopping set ψ that belongs to an intermediate layer j . Note that the Merkle proof for a base layer sample contains a single data and a single parity symbol from layer j and is deterministic given the base layer sample. If both the symbols (VNs) exist in ψ , it is possible for a single base layer symbol to sample ψ at two VNs. To avoid over-counting, we have defined the matrices $\Delta^{(j)}$ in Section IV-A. $\Delta^{(j)}$ has the property that $\Delta_{ki}^{(j)}$ is 1 if the i^{th} base layer symbol (i.e., $v_i^{(l)}$) samples, via its Merkle proof from layer j , the k^{th} stopping set of Ψ_j and zero otherwise. Thus, for a sampling strategy $(x, \beta^{(1)}, \dots, \beta^{(l)})$, it is not difficult to see that $P_{(f,med)}^{(j)}(s) = \left[\max(1 - \Delta^{(j)}x)\right]^s$, $1 \leq j \leq l$.

Now, let us consider the strong adversary. Since a Merkle proof contains one data and one parity symbol from every intermediate layer, all data (parity) symbols are sampled jointly. As such, we can bound the probability of sampling a stopping set ψ of size $\geq u_j$, $1 \leq j < l$, by $P_f^{(j)}(s = 1; \psi) \leq 1 - \sum_{x_i^{(j)}: v_i^{(j)} \in \psi, v_i^{(j)} \text{ is a data symbol}} x_i^{(j)}$ and $P_f^{(j)}(s = 1; \psi) \leq 1 - \sum_{x_i^{(j)}: v_i^{(j)} \in \psi, v_i^{(j)} \text{ is a parity symbol}} x_i^{(j)}$. Summing the two inequalities and dividing over 2 yields

$P_f^{(j)}(s = 1; \psi) \leq 1 - \frac{1}{2} \sum_{x_i^{(j)}: v_i^{(j)} \in \psi} x_i^{(j)} \leq 1 - \frac{1}{2} \beta^{(j)} \mu_j$. Finally, use $P_f^{(j)}(s; \psi) = (P_f^{(j)}(s = 1; \psi))^s$.

3) *Proof of Lemma 4:* Soundness fails if the light nodes get back all the requested samples but no honest full node is able to fully decode the entire CMT. We consider two cases:

i) There is a DA attack at layer j : In this case, no honest full node will be able to decode layer j of the CMT. Light nodes fail to detect this DA attack using the overall greedy sampling strategy described in Remark 3 with probability $P_f^{(1)}(s) = \max_{\omega^{(j)} < \mu_j} \left[[1 - \tau(S_{greedy}^{(\rho s, j)}, \omega^{(j)})] \left(1 - \frac{\omega^{(j)}}{n_j}\right)^{s - \rho s} \right]$. The term inside the maximum is the probability of failure using the overall greedy sampling strategy when the weak adversary hides a stopping set of size $\omega^{(j)}$.

ii) There is no DA attack: In this case, light nodes will accept the block. Soundness failure occurs when honest full nodes are not able to decode the entire CMT from the samples broadcasted by the light nodes. Let $P_f^{(2)}(s)$ be the probability of this event. To bound $P_f^{(2)}(s)$, we use the following property of the CMT which was proved in [21]: the Merkle proof of η fraction of distinct base layer coded symbols have at least η fraction of distinct coded symbols from each layer of the CMT. Thus for $\eta_{rec} = \left(\max_{1 \leq j \leq l} \frac{n_j - \omega_{\min}^{(j)} + 1}{n_j} \right)$, if a full node has η_{rec} fraction of distinct coded symbols from the base layer of the CMT, then it has at least η_{rec} fraction or at least $\eta_{rec} n_j$ distinct coded symbols from layer j of the CMT. Since $\eta_{rec} n_j \geq n_j - \omega_{\min}^{(j)} + 1$, using these $\eta_{rec} n_j$ distinct coded symbols, the full node will be able to successfully decode layer j , $\forall 1 \leq j \leq l$. Let Z be the total number of distinct base layer coded symbols collected by a honest full node from the random portion of the light node's overall greedy sampling strategy. Then, we have $P_f^{(2)}(s) \leq P(Z \leq \eta_{rec} n_l) \leq \binom{n_l}{\eta_{rec} n_l} \frac{(\eta_{rec} n_l)^{Ms(1-\rho)}}{n_l^{Ms(1-\rho)}} \leq 2^{[\mathcal{H}(\eta_{rec}, 1-\eta_{rec}) n_l - Ms(1-\rho) \log(\frac{1}{\eta_{rec}})]}$. The probability of soundness failure is smaller than the maximum of the above two cases. Moreover, in our system, for the same reasons as [2], soundness implies agreement (since each light node is connected to at least one honest full node and honest full nodes form a fully connected graph; see network model in Section II-C). Thus, $P_f^{S,A} \leq \max(P_f^{(1)}(s), P_f^{(2)}(s))$ completing the proof.

4) *Proof of Lemma 5:* Again we consider the two cases described in the proof of Lemma 4. For the first case, light nodes fail to detect the DA attack at layer j using LP-sampling with probability $P_f^{(1)}(s) = \max_{1 \leq j \leq l} P_{f,med}^{(j)}(s)$ and $P_f^{(1)}(s) = \max_{1 \leq j \leq l} P_{f,str}^{(j)}(s)$ for the medium and the strong adversary, respectively. For the second case, let Z be the total number of distinct base layer coded symbols collected by a honest full node when light nodes use LP-sampling. We have $P_f^{(2)}(s) \leq P(Z \leq \eta_{rec} n_l) \leq \binom{n_l}{\eta_{rec} n_l} \left(\sum_{i=1}^{\eta_{rec} n_l} x_{[i]} \right)^{Ms} \leq 2^{[\mathcal{H}(\eta_{rec}, 1-\eta_{rec}) n_l - Ms \log \left(\frac{1}{\sum_{i=1}^{\eta_{rec} n_l} x_{[i]}} \right)]}$. Similar to the proof of Lemma 4, soundness implies agreement and we have $P_f^{S,A} \leq \max(P_f^{(1)}(s), P_f^{(2)}(s))$.