LightChain: Scalable DHT-based Blockchain

Yahya Hassanzadeh-Nazarabadi, Alptekin Küpçü, and Öznur Özkasap
Department of Computer Engineering, Koç University, İstanbul, Turkey
{yhassanzadeh13, akupcu, oozkasap}@ku.edu.tr

Abstract—As an append-only distributed database, blockchain is utilized in a vast variety of applications including the cryptocurrency and Internet-of-Things (IoT). The existing blockchain solutions show downsides in communication and storage scalability, as well as decentralization. In this paper, we propose *LightChain*, which is the first blockchain architecture that operates over a Distributed Hash Table (DHT) of participating peers. *LightChain* is a permissionless blockchain that provides addressable blocks and transactions within the network, which makes them efficiently accessible by all peers. Each block and transaction is replicated within the DHT of peers and is retrieved in an on-demand manner. Hence, peers in *LightChain* are not required to retrieve or keep the entire ledger. *LightChain* is fair as all of the participating peers have a uniform chance of being involved in the consensus regardless of their influence such as hashing power or stake. We provide formal mathematical analysis and experimental results (simulations and cloud deployment) to demonstrate the security, efficiency, and fairness of *LightChain*, and show that *LightChain* is the only existing blockchain that can provide integrity under the corrupted majority power of peers. As we experimentally demonstrate, compared to the mainstream blockchains such as Bitcoin and Ethereum, *LightChain* requires around 66 times smaller per node storage, and is around 380 times faster on bootstrapping a new node to the system, and each *LightChain* node is rewarded equally likely for participating in the protocol.

Index Terms—Blockchain, Permissionless, DHT, Consensus, Storage Efficiency, Communication Efficiency, Scalability, Skip Graph.

✦

1 INTRODUCTION

Blockchain [1] is an append-only distributed database that provides a partial ordering of blocks among a set of trust-less peers. The blocks consist of transactions and are connected via immutable links from each block to its previous one on a chain that is called the *ledger*. Because they define a partial ordering of blocks without the need of a globally synchronized clock, provide a tamper-proof architecture, and establish trust over a trust-less system of independent peers, the blockchain systems are exploited in numerous decentralized applications including the cryptocurrencies [1], Internet-of-Things [2], [3], energy trading platforms [4], [5], donation systems [6], smart grids [7], and P2P cloud storage systems [8]. A blockchain system is usually modeled as a layered-architecture with at least four layers, which from bottom to top are named as *Network, Consensus, Storage, and View* layers [9]. The Network layer deals with the dissemination mechanism of the transactions and blocks among the peers of the system. The Consensus layer represents the protocols for the block generation decision-making process, which aim at providing a sequential ordering of the blocks among all the peers. The Storage layer provides the read functionality for the peers to read from the blockchain. The View layer represents the most recent state of the participating peers' data considering all the updates on the ledger from the very first to the most recent blocks.

The existing blockchain solutions show scalability problems in at least one layer of the aforementioned architecture. To the best of our knowledge, at the Network layer, all the existing blockchains operate on unstructured overlays [1], [10]–[20]. Such overlays have no deterministic, well-defined, and efficient lookup mechanism to retrieve the address of the peers, the content of the blocks, and the new transactions. Rather, the knowledge of a peer of other peers, blocks, and transactions is gained by the epidemic message dissemination among the peers (e.g., broadcasting in Bitcoin [1]) which costs a message complexity of $O(n)$ to disseminate a new block or transaction through the entire system, where n is the number of participating peers in the system.

At the Consensus layer, the existing solutions degrade the decentralization of the system by delegating the block generation decision making to a biased subset of the special peers [21], e.g., the peers with higher computational power [1], [16], [17], higher stakes [10], or longer activity history in the system [12]. The existing blockchains are also prone to the consistency problems that are caused by their probabilistic fork-resolving approach at the Consensus layer, i.e., following the longest chain of the forks as the main chain [1]. This causes the block generated by the existing consensus solutions is not considered finalized until some new blocks come after it and make that block residing on the longest chain (i.e., main chain) [22]. The main chain of existing solutions is not deterministic and can be switched abruptly once another longer chain is introduced by other peers, which compromises the consistency of blockchain.

At the Storage layer, the existing blockchains require a memory complexity of $O(b)$ by downloading and keeping the entire ledger locally at the peer's storage [9], where b is the number of blocks in the system. In other words, as peers

are not able to efficiently lookup any information within the unstructured overlay, they locally store the perceived information and gradually construct a local copy of the entire ledger, which takes $O(b)$ storage complexity. Likewise, upon joining the system, during the bootstrapping phase, a new peer needs to verify the entire state of the ledger from the very first block to the most recent one to check the integrity of the ledger [21]. This imposes a time and communication complexity of $O(b)$ at the View layer. Bootstrapping is defined as the process in which a new node constructs its view of the blockchain [9].

The best existing approach to overcome the mentioned scalability problems of the blockchains is to apply sharding [13], [14], [18], i.e., splitting the system into multiple smaller groups of peers, and each group operates in parallel on a disjoint version of the ledger. Although this reduces the required storage complexity of the nodes to $O(\frac{b}{\log n})$, it does not improve the $O(n)$ communication complexity for processing a single transaction or block [18].

In this paper, to provide a scalable blockchain architecture with fully decentralized and uniform block generation decision-making, we propose *LightChain*, which is a permissionless blockchain defined over a Skip Graph-based peer-to-peer (P2P) Distributed Hash Table (DHT) overlay [23]. *LightChain* is permissionless [24] and unmoderated [25], as it allows every peer to freely join the blockchain system and participate in the block generation decision-making without any authority. At the Network layer, *LightChain* operates on top of a Skip Graph that is a DHT-based structured P2P system with a well-defined topology and deterministic and efficient lookup strategy for data objects. We model each peer, block, and transaction by a Skip Graph node. Compared to the existing blockchains that rely on epidemic dissemination and retrieval of data with the message complexity of $O(n)$, our DHT-based *LightChain* enables participating peers to make their blocks and transactions addressable and efficiently accessible at the Network layer with the message complexity of $O(\log n)$.

At the Consensus layer, we propose Proof-of-Validation (PoV) as an efficient, immutable, and secure consensus protocol for *LightChain*. PoV is efficient as it requires only $O(\log n)$ message complexity for validating a single transaction or block. PoV is fair since each participating peer in the system has a uniform chance of being involved in the consensus regardless of its influence, e.g., processing power, available bandwidth, or stake. PoV is immutable as none of the (influential) peers in reaching a consensus can legitimately change the consensus at a later time after it is finalized. PoV is secure as the malicious peers are not able to generate and append an illegitimate transaction or block to the ledger. In PoV, the validation of each block is designated to a subset of the peers, which are chosen uniformly for each block based on its hash value (modeled as a random oracle), and are contacted efficiently using the DHT overlay. In contrast to existing consensus approaches, our PoV preserves the integrity and consistency of the blockchain in the presence of colluding adversarial peers (e.g., Sybil adversary [26]), as well as selfish miners [27], [28], as no peer can contribute to the consensus of any two consecutive blocks unless with a negligible probability in

the security parameter of the system. As we show both experimentally and analytically in the rest of the paper, there is no threshold for the number of adversarial peers in PoV, and it can be configured to remain secure and immutable even when the adversarial peers become the majority. PoV also governs a deterministic rule on resolving the forks at the Consensus layer, which makes the main chain recognized in a deterministic fashion and followed by all the peers.

At the Storage layer *LightChain* provides storage scalability and load distribution for the peers by utilizing DHT-based replication. Each peer in *LightChain* is only responsible for keeping a small subset of the randomly assigned blocks and transactions while accessing the transactions and blocks replicated on other peers of the system on an on-demand basis using the efficient Skip Graph retrievability. To preserve the data availability in presence of malicious peers, the replication in *LightChain* is done in a way that it provides at least one copy of each block and transaction accessible at any time with a very high probability in the security parameter of the system. Having b blocks in the system, compared to the existing shardless solutions that require memory complexity of $O(b)$ on each node, our *LightChain* imposes only a memory complexity of $O(\frac{b}{n})$ on each node.

At the View layer, *LightChain* provides each new node to the system with a fast bootstrapping mechanism called *Randomized Bootstrapping*, which takes $O(\log n)$ message complexity and $O(n)$ time complexity to construct its view of the entire ledger. Hence, opposed to the existing solutions that take $O(b)$ message and time complexity ($b >> n$), *LightChain* enables immediate participation on the blockchain system without the need to verify the entire blockchain for new nodes. Operating on a DHT-based overlay of peers, *LightChain* also provides a novel approach to directly obtain a particular peer's state without the need to traverse the entire ledger, which takes the message complexity of $O(\log n)$ and the time complexity of $O(1)$.

The original contributions of this paper are as follows.

- We propose *LightChain*, the first fully decentralized scalable, permissionless, and DHT-based blockchain architecture, with asymptotic communication and storage complexities' superiority compared to the existing blockchain solutions. *LightChain* provides full decentralization and fairness in block-generation decision making by providing each peer of the system with a uniform chance of being involved in the consensus regardless of its influence, e.g., processing power, available bandwidth, or stake.
- We provide security definitions for *LightChain*, and analyze how to set its operational parameters to achieve those security features.
- We extended the Skip Graph simulator SkipSim [29] with the blockchain-based simulation scenarios, implemented and simulated the *LightChain* in extensive simulation scales of $10K$ nodes, and show its performance concerning the security features, in the presence of colluding adversarial nodes.
- We also implemented a proof-of-concept version of *LightChain* node [30], deployed it as an operational

LightChain system on Google Cloud Platform, and measured its operational overheads in practice.

2 RELATED WORKS

Network Layer: Dissemination of a new transaction or block in the existing blockchains is done via Broadcasting [1], Flooding [16], or Gossiping [17], which require $O(n)$ message complexity for a single block or transaction to be accessible by every peer of the system. On the other hand, our proposed *LightChain* applies a message complexity of $O(\log n)$ to insert a new transaction or block in the Skip Graph DHT overlay, and make it accessible by every peer of the system. Additionally, in our proposed *LightChain*, the latest state of the data objects are addressable within the network (e.g., the balance of a node in cryptocurrency applications), and retrievable with the message complexity of $O(\log n)$. Therefore, in contrast to the existing blockchains, peers in *LightChain* are not required to keep searching and retrieving the most recent blocks frequently. Rather, they can search and retrieve the latest state of their data objects of demand interest.

Consensus Layer: In *Proof-of-Work (PoW)* [1], [31], the block generation decision-making is heavily correlated with the hash power of peers, which sacrifices the fairness and decentralization of the system in favor of the nodes with higher hash power. PoW tolerates up to $\frac{1}{2}$ fraction of malicious hashing power. *Proof-of-Stake (PoS)*-based approaches, degrade the decentralization of the system by binding the block generation decision-making power of the nodes with the amount of stake they own [10], [11], [32]. PoS can tolerate less than $\frac{1}{3}$ fraction of malicious stakeholders. However, they apply a message complexity of $\Omega(n)$ to the system for block generation decision making. In the *Byzantine Fault Tolerance (BFT)*-based approaches, each node broadcasts its vote to the others, receives their votes, and follows the majority [15], [19], [20], [33]–[35]. Practical BFT (PBFT) [35] is one of the pioneers in the BFT-based class that is adaptable to the practical asynchronous setups (e.g., the Internet) where there are no known bounds on the relative execution speeds of the nodes and the message-delivery latency among them [36]. Although such BFT-based consensus protocols can tolerate up to $\frac{1}{3}$ fraction of the adversarial nodes [37], they apply a message complexity of $O(n^2)$ to the system for processing a single transaction or block [38]. Some efforts are done to improve the message overhead in such systems. For example, in Ripple [33] instead of reaching a consensus with the majority of the system, each node only reaches a consensus with a subset of nodes it has trust on. Although this slightly improves the message overhead per node depending on the size of its trusted list, it degrades the fault tolerance of the system to less than a fraction of $\frac{1}{5}$ of malicious nodes (in contrast to the classical $\frac{1}{3}$). In the sharding-based approaches the system is partitioned into disjoint subsets of peers, e.g., subsets of size $O(\log n)$ [18]. Each subset is working on an independent version of the ledger using BFT in an epoch-based manner [13], [14]. Although intra-shard transactions are processed with a lower message complexity, the inter-shard transactions should go through the main ledger, which requires the same message complexity as the case of BFT in the no-shard

scenarios. As we define in Section 1, compared to the existing PoW and PoS consensus solutions, our proposed Proof-of-Validation (PoV) is the only one that provides security, immutability, decentralization, and fairness altogether. Likewise, compared to the BFT-based mechanisms, our PoV only applies a message complexity of $O(\log n)$ to the system for reaching consensus over a block. Finally, there is no inherent adversarial fraction limitation in PoV. Rather, it receives the upper bound on the adversarial fraction of the nodes as input and tweaks its parameters to preserve the security and immutability accordingly. Considering that in PoV the block generation decision making chance is uniformly distributed among the participating nodes regardless of their influence in the system, as we experimentally show in Section 6, our proposed PoV remains secure and immutable even when the fraction of malicious nodes becomes the majority, which is in contract to the existing PoW-, PoS-, and BFT- consensus solutions.

Storage Layer: Having b blocks in the system, most of the existing blockchains apply a storage complexity of $O(b)$ on each peer [1], [9], [31]. The best existing blockchain architectures concerning storage overhead are the shard-based ones, which apply a storage complexity of $O(\frac{b}{\log n})$ [16], [18]. Current attempts on storing each transaction only on the sender and receiver nodes result in overwhelming time complexity of $O(n \times b)$ on generating new transactions [17], [39]. Compared to the existing solutions, our proposed *LightChain* requires $O(\frac{b}{n})$ storage complexity on each peer without any sharding, and with the decentralization of system fully preserved.

View Layer: To the best of our knowledge, there is no existing secure and fast bootstrapping approach as we have in our proposed *LightChain*, i.e., $O(n)$ in time and $O(\log n)$ in message complexity. Having b blocks in the system, the local self-construction of view from scratch in the existing blockchains takes the time and message complexity of $O(b)$ ($b >> n$) on each node by collecting all the blocks, reconstructing the ledger locally, and computing the state of each node by replaying all transactions.

Table 1 compares a variety of the best existing blockchain solutions to our proposed *LightChain* across different layers of blockchain protocols stack.

3 PRELIMINARIES AND SYSTEM MODEL

Skip Graphs [23] are DHT-based overlay of nodes, where each node is identified with an (IP) address and two identifiers: a numerical ID and a name ID. Each Skip Graph node can search and find the address of other nodes of Skip Graph that possess a specific numerical ID or name ID, by utilizing a search for numerical ID [23], [41], or a search for name ID [42] of those nodes, respectively. Both searches are done with the message complexity of $O(\log n)$. As the result of the searches, if the targeted identifier is available in the Skip Graph, the (IP) addresses of their corresponding nodes are returned to the search initiator. Otherwise, the (IP) addresses of the nodes with the most similar identifiers to the search target are returned. As detailed in Section 4, in *LightChain*, a Skip Graph overlay of peers is constructed by representing each peer as a Skip Graph node. We assume that each participating peer joins the Skip Graph overlay

Strategy	Network	Consensus	Storage	View
Bitcoin [1]	Broadcasting ($O(n)$)	PoW	Full ($O(b)$)	Self-Construction ($O(b)$)
Rollerchain [16]	Flooding ($O(n)$)	PoW	Distributed ($O(\frac{b}{\log n})$)	Self-Construction ($O(b)$)
Trustchain [17]	Gossiping ($O(n)$)	PoW	Distributed ($O(b)$)	Self-Construction ($O(b)$)
Snow White [10]	Broadcasting ($O(n)$)	PoS	Full ($O(b)$)	Self-Construction ($O(b)$)
PoA-Bitcoin [12]	Broadcasting ($O(n)$)	PoW-PoS	Full ($O(b)$)	Self-Construction ($O(b)$)
Elastico [13]	Broadcasting ($O(n)$)	BFT	Full ($O(b)$)	Self-Construction ($O(b)$)
Rapidchain [18]	Gossiping($o(n)$)	BFT	Distributed ($O(\frac{b}{\log n})$)	Self-Construction ($O(b)$)
BitCoin-NG [40]	Broadcasting ($O(n)$)	PoW	Full ($O(b)$)	Self-Construction ($O(b)$)
Ouroboros [11]	Broadcasting ($O(n)$)	PoS	Full ($O(b)$)	Self-Construction ($O(b)$)
Omniledger [14]	Gossiping ($O(n)$)	BFT	Full ($O(b)$)	Self-Construction ($O(b)$)
Avalanche [15]	Gossiping ($O(n)$)	Snowflake	Full ($O(b)$)	Self-Construction ($O(b)$)
LightChain	**DHT** ($O(\log n)$)	**PoV**	**Distributed** ($O(\frac{b}{n})$)	**Randomized-Bootstrapping** ($O(\log n)$)

Table 1: A comparison among the best existing blockchain solutions in a system with n nodes and b blocks. We assume an approach supports distributed storage, if the storage load of blocks and transactions is distributed among all the participating peers in a policy-based manner, e.g., replication. Also, by the self-construction at the view layer we mean a peer is required to collect all blocks, build the ledger locally, and construct its view by traversing the ledger from the tail to the head.

using the Skip Graph join protocol in a fully decentralized manner [23]. In this paper, each peer corresponds to a device connected to the Internet (e.g., a laptop, smartphone, smart TV) that executes an instance of the *LightChain* protocol. We consider the system under constant churn [43], i.e., the participating peers are dynamic between offline and online states. We assume the existence of a churn stabilization strategy [44] that preserves the connectivity of the overlay under churn.

Synchronization: We assume the system is partially synchronous [36], meaning that most of the time the process execution speeds and message-delivery times are bounded. To handle the times that the system goes asynchronous where no such bounds are assumed, processes use a timer and timeout mechanism to detect process crash failures.

A **blockchain** is a linked-list of blocks with an immutable link from each block to its previous one [9]. By an immutable link, we mean that each block points back to the collision-resistant hash value of its previous block on the linked-list. The linked-list of blocks is called *ledger*. In this paper, we call the first block of the ledger (i.e., head) as the *Genesis* block and the most recent block (i.e., tail) as the *current tail*. For each block on the ledger, the *previous* relationship is the immutable link from it to its previous block on the ledger. Blockchain defines a partial ordering of the blocks on the ledger based on the previous relationship. We say that block $blk1$ is the *immediate predecessor* of the $blk2$, if $blk2$ points back to the hash value of $blk1$ as its previous block on the ledger. In this situation, $blk2$ is the *immediate successor* of $blk1$. Due to the immutable links, the blockchain is considered as an append-only database. Updating a block of the ledger by changing its content is not allowed, and is considered an adversarial act. Re-establishing the connectivity of ledger after an update on a single block requires refreshing the hash pointers on all the subsequent blocks, which is a computationally hard problem in the existing blockchains [1].

Notations: In this paper, we say that a block is *committed* to the blockchain if it is being written by the Consensus layer protocol of the blockchain to its storage, i.e., the block passes the defined consensus verification and is being appended to the tail of the ledger. We denote the security parameter of the system by λ. We denote the hash function $H : \{0,1\}^* \rightarrow \{0,1\}^s$ as a random oracle [45], where s is the identifier size of peers in bits. We assume s is polynomial in the security parameter of the system, i.e., $s = poly(\lambda)$.

We denote the *System Capacity* by n, and define it as the maximum number of registered peers in the system, i.e., $n = O(2^s)$. Similarly, we denote the *Block Capacity* by b and define it as the maximum number of the generated blocks in the system. We assume $b >> n$.

View Layer: We assume each peer is participating in the blockchain by a set of *assets* as well as a *balance*. The assets set corresponds to the data objects that the peer initially registers on the blockchain via a transaction, and can update it later on by generating new transactions. The balance of a peer is used to cover its transaction generation fees. We consider a transaction as a state transition of the assets of the transaction's owner. View of a participating peer in our system model towards the blockchain is a table of $(numID, lastblk, state, balance)$ tuples. Each tuple represents the view of the peer concerning another peer of the system with the numerical ID of $numID$, where $lastblk$ is the hash value of the last committed block to the blockchain that contains its most recent transaction, and the current state of the assets as well as the remaining balance of $state$ and $balance$, respectively. By the current state, we mean the most recent values of the assets of the peer considering all the generated transactions by that peer from the Genesis block up to the current tail of the blockchain.

Authenticated Searches: We assume that the search queries over the Skip Graph overlay are authenticated by an authentication mechanism in the presence of an adversarial party that adaptively controls a fraction f of nodes and is aiming at conducting routing attacks to the system [46], [47]. By the authenticated searches, we mean that the validity of the search results is publicly verifiable through a search proof that is generated by the signing keys of the participating peers on the search path. The search proof also contains the attributes of the peers on the search query path (e.g., identifiers and (IP) addresses) with the last node on the search path considered as the search result. We assume the success chance of the adversary on breaking the authenticated search mechanism and forging a search proof is limited to some negligible function $\epsilon(\lambda)$ (e.g., $\epsilon(\lambda) = 2^{-\lambda}$).

4 LightChain System Architecture [1]

4.1 Overview:

Overlay: In *LightChain* the peers, as well as the transactions and blocks, are represented by Skip Graph nodes. Each peer invokes the join protocol of Skip Graph [23] using its identifiers and (IP) address and joins the system. Both identifiers of a peer (i.e., its name ID and numerical ID) are the hash value of its public key (i.e., verification key). In this paper, we assume the peers utilize a digital signature scheme that is existentially unforgeable under adaptive chosen message attack [45]. Once the peer joins the Skip Graph overlay, it can efficiently search for any other peer of the system with the message complexity of $O(\log n)$. Upon joining the Skip Graph overlay, the peer creates its view of the blockchain using *LightChain*'s Randomized Bootstrapping feature without the need to download and process the entire ledger.

Transaction Generation: In *LightChain*, a transaction represents a state transition of the assets of a peer, which is denoted by the *owner* peer of that transaction. For example, in cryptocurrency applications, the asset of a peer is its monetary wealth, and a transaction represents the state transition of the monetary wealth of it. The owner peer casts the state transition into a transaction, computes the identifiers of validators, searches for the validators over Skip Graph overlay, and asks them to validate its transaction. To be validated, each transaction needs to be signed by a system-wide constant number of validators, where their identifiers are chosen randomly for each transaction to ensure security. In addition to security, the idea of validating transactions makes participating nodes in the block generation needless of going through the validation of individual transactions.

Block Generation: Once the transaction gets validated, the owner inserts it as a node into the Skip Graph overlay, which makes it searchable and accessible by any other peer. The insertion of the transaction is done by invoking the join protocol of Skip Graph using the transaction's identifiers (i.e., its hash value) but the (IP) address of the owner peer itself. The Skip Graph peers route the messages on behalf of the transaction nodes they hold. This idea is similar to the other existing DHTs like Chord [49], and enables *LightChain* peers to search and find the new transactions. Upon finding new validated transactions, each peer is able to cast them into blocks, go through the validation procedure (similar to the transactions' case), and insert the validated block into the Skip Graph overlay. Each transaction's owner then removes its transaction node from the overlay once it is successfully included in a validated block. The idea of representing each transaction and block by a Skip Graph node results in any search for the peer or the transactions and blocks that it holds to be routed to the peer's (IP) address, rendering them accessible by every other peer in a fully decentralized manner with the message complexity of $O(\log n)$. Hence, in *LightChain*'s Skip Graph overlay, there exist three types of nodes: peers, transactions, and blocks. In

other words, the Skip Graph overlay acts as a distributed database of the transactions and blocks that are owned by the peers. As elaborated later, the *previous* relationship of blocks stored in a distributed manner on distinct peers defines a blockchain.

Distributed Storage Layer: In *LightChain*, each block and transaction is replicated on its owner and validators to support availability, accessibility, and fault tolerance. Using searchable blocks and transactions as well as replication, in *LightChain* we introduce the idea of a distributed storage layer for the blockchain where participating peers in the consensus only need to keep and maintain a subset of the blocks, and not the ledger entirely. In the rest of this section, unless stated otherwise, by the term node, we mean a peer.

Incentive Mechanism: As an incentive mechanism, *LightChain* employs a monetary balance for each participating peer to exchange with other peers and cover the operational fees of appending data to the blockchain [1]. *LightChain* rewards the peers' honest contribution to maintain the connectivity of the system, provide validation service, and generate blocks. Moreover, *LightChain* encourages honest peers to audit other peers, by rewarding the detection and report of adversarial acts. Malicious behavior is penalized by *LightChain* upon detection, and the adversarial peers are blacklisted and gradually isolated from the system.

4.2 Structure of Transactions and Blocks

A *LightChain* transaction, tx, is represented by a $(prev, owner, cont, search\_proof, h, \sigma)$ tuple, where $prev$ is the hash value of a committed block to the blockchain. We use the $prev$ pointer for each transaction tx to define an order of precedence between tx and all the blocks and transactions in the blockchain without the need for any synchronized clock. The block that is referred by $prev$ takes precedence over tx. All the transactions included in the $prev$ block are assumed to be committed before tx in the essence of time. Following the same convention, all the blocks and transactions that precede $prev$, also precede tx. The *owner* represents the identifier of the owner node in the Skip Graph overlay that generates the transaction tx. The contribution field (i.e., $cont$) of a transaction denotes the state transition of the assets of the owner node, e.g., a monetary remittance between two peers in cryptocurrency applications. The $search\_proof$ field of a transaction is the authenticated proof of searches over the peers of the Skip Graph overlay to find the validators of the transaction tx, as explained before. The h field of the transaction tx is the hash value of the transaction, which is computed as shown by Equation 1. The σ field of the transaction tx contains the signatures of both the owner as well the validators on its hash value h. The validators' signature is a part of *LightChain*'s consensus strategy and is explained later.

$$h = H(prev||owner||cont||search\_proof) \qquad (1)$$

Similarly, a *LightChain* block blk is defined by a $(prev, owner, \mathcal{S}, search\_proof, h, \sigma)$ tuple. $\mathcal{S}$ represents the set of all the transactions that are included in the block blk. The h field of block blk is its hash value, which is computed as shown by Equation 2. The σ field contains the

---

signatures of both the block's owner as well as the block's validators on its hash value h.

$$h = H(prev||owner||\mathcal{S}||search\_proof) \quad (2)$$

4.3 Network Layer: Skip Graph overlay

In our proposed *LightChain*, we represent each peer, transaction, and block by a Skip Graph node. This way, all the peers, transactions, and blocks are addressable within the network. In other words, participating nodes (i.e., peers) in *LightChain* exploit the Skip Graph overlay to search for each other, as well as each others' blocks and transactions. Both the numerical ID and name ID of the peers are the hash value of their public key using a collision-resistant hash function. As in a Skip Graph, nodes' identifiers define the topology [23]; hence, considering the hash function as a random oracle results in the uniform placement of peers in Skip Graph overlay topology, which limits the adversarial power on tweaking the Skip Graph topology for its advantage.

The numerical ID and name ID of a transaction (or a block) in the Skip Graph overlay is its hash value (i.e., h) and its corresponding *prev* field value, respectively. This regulation enables peers to traverse the *LightChain*'s ledger in both forward and backward directions. Following this convention, in *LightChain*, having a block with numerical ID (i.e., the hash value) of h and previous pointer value of *prev*, the (IP) address of the peers that hold the immediate predecessor block are obtained by performing a search for the numerical ID of *prev* in the Skip Graph overlay [23]. Similarly, the (IP) address of the peers holding the immediate successor transaction(s) or block(s) in the blockchain are obtainable by performing a search for name ID of h over the Skip Graph overlay. This follows the fact that all the immediate successors of the block have its hash value h as their name ID. This feature of the *LightChain* enables the peers to efficiently update their view towards the tail of the blockchain by performing a search for the name ID of their local tail [42]. The search returns all the blocks that are appended subsequently to their local tail, as well as all the new validated transactions that are waiting to be included in blocks. Additionally, using this feature, a peer does not need to store the entire blockchain locally. Rather, having only a single block of the ledger enables the peer to efficiently retrieve the predecessor and successor blocks to it with the message complexity of $O(\log n)$.

Figure 1 illustrates this convention of *LightChain*, where a peer that only has $blk2$ is able to efficiently retrieve its immediate predecessor (i.e., $blk1$) by searching for the numerical ID [23] of its *prev* value (i.e., $blk2.prev = blk1.h$) in a fully decentralized manner. The search is responded by the replicas of $blk1$ with their (IP) addresses, and hence the predecessor of $blk2$ (i.e., $blk1$) is retrievable efficiently by directly contacting those replicas. Similarly, the peer that only possesses $blk2$ is able to perform a search for name ID [42] over its hash value (i.e., $blk2.h$) to retrieve the immediate successor block that comes after $blk2$. As the result of the search for the name ID of $blk2.h$, replicas of $blk3$ respond to the search initiator peer with their (IP) addresses, and $blk3$ is retrievable efficiently by directly contacting those replicas. In the case where a single block has several
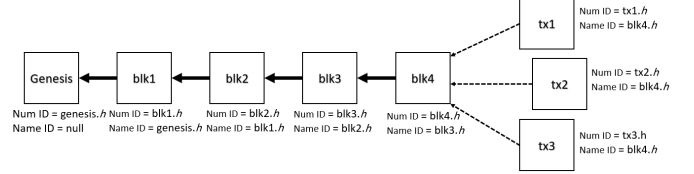

Figure 1: The *LightChain* regulation on name IDs and numerical IDs. Numerical ID (i.e., Num ID) of a block or transaction is its hash value, and name ID is its corresponding *prev* value.

successor blocks, the search initiator receives a response from each of the immediate successor block replicas. In the example of Figure 1, considering $blk4$ as the current tail of the blockchain, as discussed later in this section, the newly generated transactions that succeed $blk4$ (i.e., $tx1$, $tx2$, and $tx3$) are efficiently retrievable by performing a search for the name ID using $blk4.h$.

4.4 Consensus Layer: Proof-of-Validation (PoV)

The consensus layer of *LightChain* is based on our proposed PoV, which is a fair, immutable, secure, and efficient consensus protocol (see Section 1 for more details on these features). A transaction or block is considered as validated once it successfully passes the PoV consensus. Note that a validated transaction's contribution is not considered effective and authoritative unless it is included in a validated block that is committed to the blockchain. As detailed in the following, PoV provides a set of α uniformly chosen validators for each transaction or block. PoV then considers a transaction or block as valid if its hash value h is signed by t validators. Both t and α are constant protocol parameters, which are called the *Signatures Threshold* and *Validators Threshold*, respectively. We formally discuss this in Section 5, and develop a formulation for deciding on the proper values of the *Signatures Threshold* and *Validators Threshold* considering the security of the system.

4.4.1 Transaction Generation and Validation

For a transaction tx, the numerical ID of each validator is chosen uniformly as shown by Equation 3, where v_i is the numerical ID of the i^{th} validator in the Skip Graph overlay.

$$v_i = H(tx.prev||tx.owner||tx.cont||i) \quad (3)$$

The transaction's owner then conducts a search for the numerical ID of the validator (i.e., v_i) within the Skip Graph overlay. If there exists a peer with the numerical ID of v_i in the overlay, the owner receives its (IP) address. Otherwise, it receives the (IP) address of the peer with the largest available numerical ID that is less than v_i. Both cases are supported by an authenticated search proof that is generated by the Skip Graph peers on the search path and is delivered to the owner. The authenticated proof of the search for the numerical ID of the i^{th} validator is denoted by $search\_proof_i$, which also contains all the (IP) addresses and identifiers of the Skip Graph peers on the search path. The last peer on the search path of v_i is designated as the i^{th} validator. The transaction's owner then adds the authenticated search proof for all the validators to the transaction, computes its hash value h as specified by Equation 1, signs the hash value, and appends her signature to σ. The transaction's owner then contacts the validators asking for the validation of the tx. Each validator evaluates

the soundness, correctness, and authenticity of tx, as well as the balance compliance of its owner to cover the fees. As the validation result for tx, the transaction owner either receives a signature over h or $\perp$ from a contacted validator.

Soundness: A transaction tx is sound if it does not precede the latest transaction of its owner on the blockchain, i.e., its $prev$ should point to the hash value of a validated and committed block on the ledger with no transaction of its owner in any of the subsequent blocks. In other words, soundness requires a causal ordering among the committed transactions of each owner. This is both to counter double-spending from the same set of assets, as well as to make the validation of a transaction a one-time operation, i.e., the owner of a validated tx transaction can append it to the blockchain as long as it does not generate any new transaction on the blockchain that precedes tx based on $prev$. Considering the soundness, at most one of the concurrently generated and validated transactions of a peer has the chance to be included in a new block. Once one of its transactions is included in a block, the others become unsound, cannot be included in the same block or further blocks, and should go over re-validation. Therefore, besides preventing the double-spending, soundness provides a uniform chance for the transaction generators to include their transaction into each new block.

Correctness: For a transaction tx to be correct, its contribution field (i.e., $cont$) should represent a valid state transition of the owner's assets. The compliance metric is application dependent. For example, in cryptocurrency applications, for a transaction to be correct, the owner's account should have enough balance to cover the remittance fee (i.e., the contribution).

Authenticity: The evaluation of authenticity is done by checking the correctness of h based on Equation 1, verifying σ for the inclusion of a valid signature of the transaction's owner over h, and verifying $search\_proof$ for all the validators of tx. A validator rejects the validation of tx as unauthenticated if any of these conditions is not satisfied.

Balance Compliance: As an incentive mechanism to participate in the validation, *LightChain* considers a validation fee in the favor of the t validators of the transaction tx that sign its hash value and make it validated. Also, *LightChain* considers a routing fee in the favor of all the Skip Graph peers that participate in finding the transaction's validators, i.e., the peers in $search\_proof$. A transaction tx passes the balance compliance part of validation if its owner has enough balance to cover the validation and routing fees. The balance compliance validation is done based on the view of the validator towards the blockchain. Both the routing and validation fees are fixed-value protocol parameters and are the incentive mechanism for the peers to perform the routing and validation honestly [1], [40]. The fees also prevent Sybil adversarial peers from indefinitely generating transactions by circulating the adversarial balance among themselves and continuously congesting the system with the validation of adversarial transactions.

Once a transaction tx receives t valid signatures issued by its uniformly designated validators, it is considered valid and is inserted as a Skip Graph node by its owner, which makes it accessible by other participating peers of the system to be included in a block. The numerical ID of tx is

$tx.h$, and the name ID of tx is $tx.prev$ (block). This enables any Skip Graph peer to conduct a search for name ID on the hash value of any ledger's block within the Skip Graph overlay and find all the new transactions that are pointing back to that block. Setting $tx.prev$ to point to the tail of the ledger hence increases the chance of tx for being discovered by other peers.

4.4.2 Block Generation and Validation

In *LightChain*, a peer that generates blocks is called a block owner. Once a block owner collects at least min_tx newly generated transactions that have not been included in any committed block to the ledger, it casts them into a new block blk, and sends the block for validation. By casting transactions into blk we mean including the collected transactions into the set $\mathcal{S}$ as discussed earlier (i.e., Section 4.2). min_tx is an application-dependent fixed-value parameter of *LightChain* denoting the minimum number of transactions that should be included in a block. To have blk validated, the block owner computes the numerical ID of the i^{th} validator as shown by Equation 4, where $prev$ is the hash value of the current tail of the ledger. The rest of the finding validators procedure is similar to the transaction case. On receiving a validation request for a block blk, each of its PoV validators checks the authenticity and consistency of blk itself, as well as the authenticity and soundness of all transactions included in $\mathcal{S}$ (as discussed earlier). The authenticity evaluation of blocks is done similar to the transactions.

$$v_i = H(prev||owner||\mathcal{S}||i) \tag{4}$$

Consistency: A block blk is said to be consistent if its $prev$ pointer points to the current tail of the blockchain; otherwise it is inconsistent. By the current tail of the blockchain, we mean the most recent view of the validators towards the tail of the chain. However, it is likely for the current tail of the blockchain to be updated during the validation of a newly generated block and make the block under validation inconsistent. To avoid forking the ledger, upon the start of validation, all randomly chosen PoV validators continuously monitor the updates on the tail of the ledger and terminate the validation with rejection (at any step) upon detecting a potential fork that is caused by the current block under validation.

After the blk gets validated, its owner inserts it in the Skip Graph overlay as a node. As the incentive mechanism of *LightChain*, the owner of a block receives a block generation reward once its block gets validated and committed to the blockchain. The block generation reward is a fixed-value parameter of *LightChain* that acts both as an incentive mechanism for encouraging the peers to participate progressively in generating blocks, as well as a mean for wealth creation. In this paper, we assume that the generation reward for a block is larger than its validation and routing fees.

4.4.3 Fork-free mechanism:

To resolve the forks caused by the simultaneously validated blocks, *LightChain* governs a fork-free mechanism, which is a deterministic approach that instructs all the peers to solely follow the block with the lowest hash value upon a fork. Upon a fork, we call the block with the lowest hash value as the winner block and the other participating blocks of the

fork as the knocked-out ones. The knocked-out block owners remove their block from the Skip Graph overlay, update their set of transactions by dropping the transactions that are included in the winner block, adding the new transactions to reach the min_tx threshold, and restart the validation procedure. The knocked-out block owners neither gain any block generation reward nor lose any balance because of the fees, as these fees and rewards are not effective unless the block is successfully committed to the ledger. To ensure that a newly appended validated block blk to the ledger does not undergo any further fork rivalry, and is considered committed, effective, and finalized, *LightChain* waits for only one further block to be appended subsequently to blk. In this way, as explained earlier, all the forks at the height of the blk are considered as potential forks and are rejected by the consistency checking mechanism of PoV validators.

4.5 Storage Layer: Replication

To provide efficient retrievability and data availability under churn [43] in *LightChain*, each transaction or block is replicated in the local storage of its owner as well as its uniformly chosen PoV validators, and presented by them as Skip Graph nodes. This makes the transactions and blocks efficiently searchable by all the participating peers in the system. As we show in Section 6, the parameters of *LightChain* are chosen in a way that at least one (honest) replica for each transaction and block is always available in expectation. Hence, in *LightChain*, peers do not need to store or download the entire ledger. Rather, they access the transactions and blocks in an on-demand manner, i.e., a peer searches for a transaction or block upon a need and retrieves it efficiently from the overlay.

4.6 View Layer: Randomized Bootstrapping and Fast Retrieval

Randomized Bootstrapping: We define the *view introducers* of a new peer as the set of randomly chosen peers that share their view of the blockchain with the newly joined peer. Upon joining the overlay, a new peer computes the numerical IDs of its view introducers based on Equation 5, where $new\_peer.numID$ is the numerical ID of the new peer and $view\_intro_i$ is the numerical ID of the i^{th} view introducer of it. The new peer then conducts an authenticated search for the numerical ID of $view\_intro_i$ within the overlay [46], contacts the peer in the search result, and obtains its view of the blockchain. The new peer continues iterating over i until it obtains t consistent views. As we show later, we determine t and α in such a way that a new peer obtains t consistent views of the honest peers by iterating i over $[1, \alpha]$.

$$view\_intro_i = H(new\_peer.numID||i) \qquad (5)$$

Fast Retrieval: Tracking the updates on the entire view of other peers' assets requires a peer to keep its local view updated with the new blocks, which is a plausible assumption in the majority of the existing solutions [9]. However, in addition to this traditional approach, *LightChain* enables each peer to directly retrieve the latest assets' state of another peer of interest without the need to keep track of the new blocks on the ledger and sequentially applying the updates. This is done by the additional representation
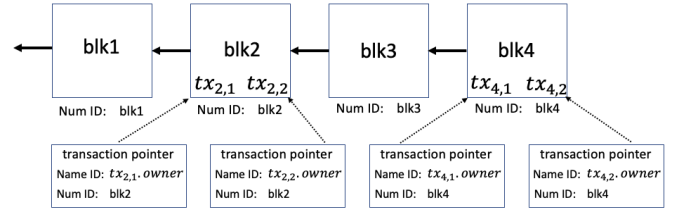

Figure 2: Transaction pointers for the transactions of blocks $blk2$ and $blk4$. Transactions and pointers of other blocks are not shown for sake of simplicity.

of each block with multiple Skip Graph nodes, one per included transaction, which we call them the associated *transaction pointers* of that block. In this approach, each transaction tx that is included in a committed block blk is represented by a transaction pointer node (i.e., $pointer$) with $pointer.nameID = tx.owner$ and $pointer.numID = blk.h$. The transaction pointer nodes associated with each block are inserted by the block's owner and replicated on the block's PoV validators. Figure 2 depicts an example of transaction pointers for transactions of block $blk2$ and $blk4$ on the ledger. Assume that by the time $blk2$ is committed to the ledger, it contains the transaction $tx_{2,2}$, which is the most recent transaction of $tx_{2,2}.owner$ at the height of $blk2$. Hence, any peer that is solely interested in knowing the latest state of the owner of $tx_{2,2}$ performs a search for the transaction pointer with the name ID of the owner node's identifier, i.e., $tx_{2,2}.owner$ in this example. As result of the search, the querying peer obtains a copy of the transaction pointer. Since the numerical ID of transaction pointer corresponds to the hash of the block that the transaction is included in, by having the transaction pointer, the querying peer finds that the latest transaction of $tx_{2,2}.owner$ is included in block $blk2$. By performing a search for the block numerical ID of $blk2$, the querying peer obtains the content of $blk2$, which includes $tx_{2,2}$ that also contains the latest state of $tx_{2,2}.owner$ affected by the transaction. To keep track of the latest updates over the state of assets, both the owner and validators of a block should drop each of their associated transaction pointers once an update on the corresponding assets appears on a newer committed block to the ledger. In this way, the transaction pointers act as moving flags within the ledger each pointing to the latest transaction of a peer on the ledger. For example, in Figure 2 assume that once $blk4$ is committed to the ledger, it contains a new transaction $tx_{4,2}$ from the same owner as $tx_{2,2}$ (i.e., $tx_{2,2}.owner = tx_{4,2}.owner$). In this situation, $blk2$ no longer holds the latest transaction of $tx_{4,2}.owner$. Hence, the owner and PoV validators of $blk2$ drop the transaction pointer for $tx_{2,2}$ from the Skip Graph overlay. Taking down a pointer node from the overlay is simply done by performing the Skip Graph node deletion operation [23] by the owner and each of the validators in a fully decentralized manner. This is for the sake of better efficiency of the search, and to make sure that the transaction pointers always point to the most recent states. Not dropping the (old) pointers after a new update is counted as misbehavior, which we address by the misbehavior detection strategy of *LightChain*. To address the network asynchrony, however, the block owner and PoV validators are allowed to take down the pointers within a certain block interval upon a new transaction on the associated set of assets that they maintain the pointers.

Operation	Type	Complexity
Joining Skip Graph Overlay	Message	$O(\log n)$
Randomized Bootstrapping	Message	$O(\log n)$
TxB Generation/Validation	Time	$O(\log n)$
TxB Generation/Validation	Message	$O(\log n)$
TxB Storage	Storage	$O(\frac{b}{n})$
Direct State Retrieval	Message	$O(\log n)$

Table 2: Summary of asymptotic operational complexities of *LightChain* in a system with n nodes and b blocks. TxB stands for "Transaction and Block".

This allows them to have enough time to discover the new updates without being subject to misbehavior. The length of the block interval (i.e., number of blocks between two transaction pointers over the same set of assets) is a constant protocol parameter that is application dependent.

4.7 Incentive Mechanism

As discussed earlier, the block generation reward and the routing and validation fees constitute the incentive mechanism of *LightChain* for the peers to retain their honest behavior. The counterpart of honest behavior is the *misbehavior*, which we define as any sort of deviation from the described *LightChain*'s protocol and architecture. As detailed earlier, for the transactions and blocks that are gone through the consensus layer, we consider the randomly chosen PoV validators to check the submitted transaction or block against the misbehavior. As we show in Section 5, we set the operational parameters of *LightChain* in a way that an adversarial peer cannot convince the PoV validators on misbehavior unless with a negligible probability in the security parameter. We also introduce the *misbehavior detection* as an extra level of adversarial countermeasure, especially for the adversarial actions that are not gone through the PoV, e.g., direct submission of an invalid block to the ledger. Each peer of *LightChain* also acts as an auditor for other peers' behavior and gains a *misbehavior audition reward* by reporting their misbehavior. As an auditor, any peer should be able to evaluate a block or transaction in the same way that its PoV validators do during the validation.

Upon a misbehavior detection, the auditor generates a transaction with the evidence of the misbehavior in the contribution field (e.g., a soundness violating transaction in a committed block). The transaction then goes through the same PoV validation process as described earlier, except that the validators verify the correctness of the transaction as the correctness of the reported evidence. Once the transaction is validated and placed into a committed block to the ledger, the misbehaving peer is penalized by a misbehavior penalty fee. The misbehavior penalty fee is another constant parameter of *LightChain* that is application dependent, and is paid in the favor of the auditor peer who reports the misbehavior. Once a misbehavior is recorded for a peer on a committed block, its identifier is blacklisted. The blacklisted peers are isolated by the honest peers and hence can no longer participate in *LightChain*, i.e., any incoming message from the blacklisted peers is discarded by the honest peers. Table 2 summarizes the asymptotic operational complexities of *LightChain*.

5 SECURITY ANALYSIS

This section provides a summary of the security analysis of *LightChain* and skips details and intermediate steps due to the page limit. The interested readers are referred to the full version of our work [48] for the detailed security analysis.

Definitions: We analyze the security of *LightChain* from the lens of integrity, data availability, and service availability [50]. We define the *integrity* of *LightChain* as the property that the views of the peers towards the ledger are not being changed, except by committing a new block of validated transactions to the current tail of the ledger solely by the designated PoV validators of that block. We define the *data availability* as the property that every validated transaction or committed block is being accessible in a timely fashion [50]. We define the *service availability* as the availability of the Consensus and View layer protocols. The Consensus layer service availability means that the honest nodes who follow the *LightChain* protocols should be able to find enough honest PoV validators, in expectation, to validate their transaction and block, as specified in Section 4.4. The availability of the View layer means that the honest nodes should be able to find enough honest view introducers, in expectation, to bootstrap to the system, as specified in Section 4.6. Defining the service availability in expectation means that the honest nodes should be able to eventually bootstrap to the system and have their transactions and blocks eventually validated.

Adversarial and Honest Behavior: We model all participating peers as probabilistic Turing machines, whose running time is polynomial in the security parameter of the system (i.e., λ). We define honest behavior as the one that follows *LightChain* protocols specified in Section 4, and maintains its availability and accessibility in the network. On the contrary, the adversarial behavior is the one that deliberately deviates from the *LightChain* protocols at arbitrary points. Similar to BFT-based approaches [37], we assume the existence of a Sybil adversarial party [26] that adaptively takes control over at most a fraction f of peers in the system. The adversarial party aims at compromising the integrity, data availability, or service availability of *LightChain* by orchestrating attacks through its controlled set of corrupted peers. In our analysis, to presume the worst-case scenario, we assume no churn and failure for corrupted peers under the control of the adversary. For the honest peers, on the other hand, we model the churn with a uniform failure probability of q, under the crash-recover model [36], i.e., an honest node fails with a probability of q and recovers back online after a while.

Assumptions: Based on the above definitions we summarize the assumptions in the following corollaries. We assume a *LightChain* system with n peers, a partially synchronous network, an authenticated routing mechanism (as defined in Section 3), where honest nodes follow a uniform failure probability of q under the crash-recover model. We assume the existence of the negligible probability $\epsilon(\lambda)$ and a probabilistic polynomial-time adversary that adaptively controls a fraction f of all peers in the system.

Corollary 5.1. **(Integrity and Data Availability):** Under the specified assumptions above, selecting the *Validators Threshold* α and the *Signatures Threshold* t in such a way that it satisfies the inequalities shown by Equations 6 and 7, respectively, yields in the success probability of the adversary on breaking the integrity or data availability

of *LightChain* bounded by $\epsilon(\lambda)$.

$$\alpha \geq \frac{(\sqrt{f} + \sqrt{f \times (\psi^{-1}(1 - \epsilon(\lambda)))^2 + 4))^2}}{4(1 - f)} \quad (6)$$

$$t \geq (\sqrt{\alpha f(1 - f)} \times \psi^{-1}(1 - \epsilon(\lambda))) + \alpha f + 1 \quad (7)$$

Corollary 5.2. **(Service Availability):** Under the specified assumptions above, setting values of the *Validators Threshold* α and the *Signatures Threshold* t in such a way that it satisfies the inequality shown by Equation 8, yields in *LightChain* providing its service availability (where one replica per block is available) in expectation.

$$t \leq \frac{\alpha(1 - f)(1 - q)}{f + (1 - f)(1 - q)} \quad (8)$$

6 EXPERIMENTAL RESULTS

6.1 Simulation Results

Setup: We implemented *LightChain* over the SkipSim simulator [29], where each node follows the Bitcoin churn trace with an expected online and offline duration of 10.6 and 2.8 hours, respectively [43], and generates a transaction per hour [51], [52]. We simulated for 100 *LightChain* systems each $n = 10,000$ peers. Each system was simulated for 48 hours.

Integrity: Figure 3.a shows the success probability of the adversary that controls fraction f of corrupted peers on compromising the integrity of the system as defined in Section 5. In this figure, $f = 0.16$ corresponds to the largest fraction of colluding hash power in the Bitcoin network [10]. Considering that each node in *LightChain* has a uniform chance of involvement in consensus, $f = 0.33$ is beyond the adversarial fraction threshold of BFT and PoS-based blockchains, e.g., Hyperledger [19]. Likewise, $f = 0.51$ is beyond the adversarial power threshold of the PoW, e.g., Bitcoin [1]. Supported by Corollary 5.1, as shown by Figure 3.a, for each simulated adversarial fraction, there exists a certain value of α in which the success probability of adversarial peers in compromising the integrity of *LightChain* exponentially converges to zero with the growth of t. **This implies that the integrity of *LightChain* is preserved even when corrupted peers become the majority.**

Service Availability: Figure 3.b presents the integrity aspect of *LightChain* when its service availability is preserved. The validator thresholds of Figure 3.b are obtained by applying both Corollaries 5.1 and 5.2. These corollaries are infeasible to satisfy together for $f > 0.5$. Hence, although there is no adversarial bound for the integrity of *LightChain* alone, its integrity under service availability is preserved for the adversarial fractions less than 0.5. This means that in contrast to the state-of-the-art blockchains such as Bitcoin [1], Ethereum [31], and Hyperledger [19], that are fully compromised once the adversarial fraction of nodes goes beyond their inherent threshold, our proposed *LightChain* system halts when the adversarial fraction of nodes goes beyond its configured operational parameters. This allows the blockchain community to notice the security risks and perform a decentralized bootstrapping of the system by discarding the nodes with suspicious behavior, which is left as future work.

Data Availability: Figure 3.c illustrates the average number of available replicas for each block in the system at each time slot over 48 hours of the simulation as t grows. Each validated block is replicated $t + 1$ times in the system, i.e., t validators as well as the owner itself. As shown by Figure 3.c, the average availability of the blocks increases linearly with respect to t. With the uniform failure probability of q, having $t + 1$ replicas for a block results in $(t + 1) \times (1 - q)$ available replicas in expectation. Hence, as shown by Figure 3.c choosing $t \geq \frac{1}{1-q} - 1$ results in an expected availability of at least one replica at each time slot. In Figure 3.c, the number of replicas for $t = 1$ is about 1.6 replicas on the average, and grows linearly as t increases. An identical behavior is observed for the transactions' availability.

6.2 Cloud Platform Deployment

We implemented and deployed a proof-of-concept version of the *LightChain* [30], [53] with 1000 nodes for the adversarial fractions of 0.16 and 0.33, and block sizes of 10, 20, and 40 transactions on the Google Cloud Platform, where each node generates a transaction per second for a total of 1000 transactions per node. Due to the quota limitations, we run each node on a low-power hardware setup with 2.2 GHz of processing power and 3.9 GB of memory. To avoid these low-power nodes crashing of a high degree of concurrency, and memory limitation, we serialized those operations of the nodes that require synchronized communication, i.e., resources getting blocked for a response or timeout. Figure 4.a shows the block formation time, which is the time that it takes for a node to collect a certain number of newly validated transactions into a block, generate the block metadata as specified in Section 4, and request the validation of the block. Figure 4.b shows the average local block validation time at each validator. Finally, Figure 4.c presents the block size in KB as the number of transactions grows. The difference between $f = 0.16$ and $f = 0.33$ in the slope of growth of the block formation time (i.e., Figure 4.a) and block size (i.e., Figure 4.c) is due to the 10 times growth in the number of validators per block. A similar pattern is also observed in the transaction time (i.e., generation, validation, and insertion in the Skip Graph overlay), which is about 7 and 82 seconds for $f = 0.16$ and $f = 0.33$ on average, respectively. As mentioned earlier due to hardware limitations the operations like transaction and block generation are serialized in our implementation on the Google Cloud Platform. This difference is however not evident in the block validation process (i.e., Figure 4.b) since verifying the validators signature is a lighter operation than the network-based operations (i.e., transactions and block generation) and is done with some degree of concurrency.

6.3 Comparison With Mainstreams

Advantages: As shown by our experimental results in Figure 3.a, *LightChain* preserves its integrity under the corrupted majority power of the system (e.g., $f = 0.51$), which is in contrast to all mainstream blockchains including Bitcoin [1] and Ethereum [31] that are fully compromised on integrity under this adversarial fraction. Also, as shown in Figure 3.b, *LightChain* preserves both its integrity and service availability under the corrupted power fraction of
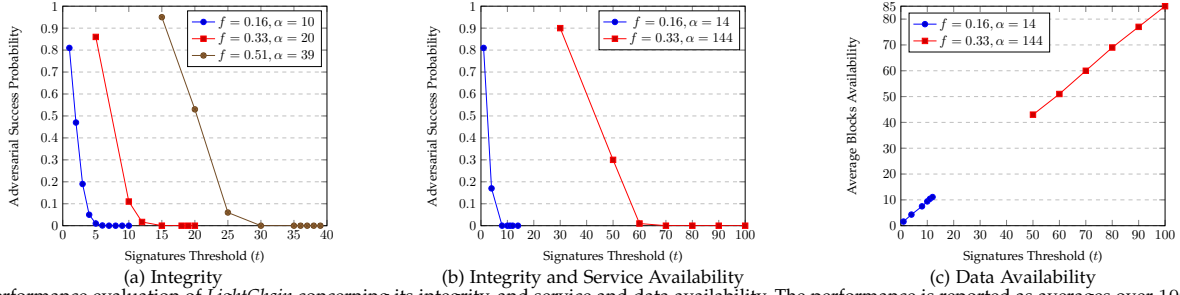
Figure 3: Performance evaluation of *LightChain* concerning its integrity, and service and data availability. The performance is reported as averages over 100 *LightChain* systems each with 10, 000 nodes. Each system is simulated for 48 hours under the real churn traces of Bitcoin [43].
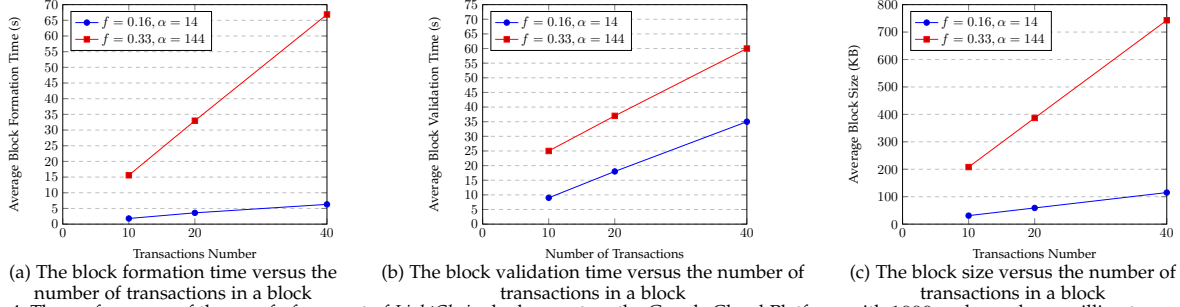

Figure 4: The performance of the proof-of-concept of *LightChain* deployment on the Google Cloud Platform with 1000 nodes and one million transactions.

$f = 0.33$ of system, which is in contrast to the mainstream PoS- and BFT-based blockchains (e.g., Hyperledger [19]). Considering our Google Cloud Platform deployment with 1, 000 nodes, one million transactions, 25, 000 blocks, and $f = 0.16$, *LightChain* only has a storage overhead of about 30 MB per node to maintain the transactions and blocks and keep up with the protocol. This is around 66 times less storage compared to the mainstream blockchains such as Bitcoin and Ethereum that result in a storage overhead of about 2 GB per node under the same setup. Also, in our cloud deployment, bootstrapping a new *LightChain* node to the system at the block height of 25, 000 via the randomized bootstrapping protocol takes around 28 seconds on average. Considering the average block retrieval time of around 400 ms in our deployment, *LightChain* is around 380 times faster on bootstrapping a new node than the mainstream blockchains such as Bitcoin and Ethereum, which take around 3 hours by imposing a new node to download and process all 25, 000 generated blocks. Finally, under the mentioned setup, each *LightChain* node is involved in and rewarded for an average of 348 block generation decision makings with a standard deviation of around 1.3 blocks, which denotes a uniform distribution of the consensus involvement among the nodes. This is in contrast to the mainstream blockchains like Bitcoin where nodes are involved in the block generation decision making based on their influence in the system, i.e., hash power.

Shortcomings: *LightChain*'s implementation is currently at the prototype level, which hinders back its transaction throughput to around a single transaction per second on average. This makes the transaction throughput of *LightChain* lower than the mainstream blockchains like the Bitcoin and Ethereum that are capable of processing around 4 and 12 transactions per second under their production-level implementation [51]. As our future work, we plan to optimize the throughput of *LightChain* through a production-level imple-

mentation and software engineering-level optimizations.

7 CONCLUSION

To improve the communication and storage scalability, as well as the decentralization of blockchain architectures, we proposed *LightChain*, which is a novel blockchain architecture that operates over a DHT overlay. *LightChain* provides addressable peers, blocks, and transactions within the network, which makes them efficiently accessible in an on-demand manner. Using *LightChain*, no peer is required to store the entire ledger. Rather, each peer replicates a random subset of the blocks and transactions, and answers other peer's queries on those. *LightChain* is a fair blockchain as it considers a uniform chance for all the participating peers to be involved in the consensus protocol regardless of their influence in the system (e.g., hashing power or stake). We analyzed *LightChain* both mathematically and experimentally regarding its security and performance.

ACKNOWLEDGEMENT

REFERENCES

[1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

[2] S. Aggarwal, R. Chaudhary, G. S. Aujla, N. Kumar, K.-K. R. Choo, and A. Y. Zomaya, "Blockchain for smart communities: Applications, challenges and opportunities," *Journal of Network and Computer Applications*, 2019.

[3] N. Kumar and A. Miglani, *Probabilistic Data Structures for Blockchain-Based Internet of Things Applications*. CRC Press, 2021.

[4] A. Jindal, G. S. Aujla, and N. Kumar, "Survivor: A blockchain based edge-as-a-service framework for secure energy trading in sdn-enabled vehicle-to-grid environment," *Computer Networks*, 2019.

[5] R. Chaudhary, A. Jindal, G. S. Aujla, S. Aggarwal, N. Kumar, and K.-K. R. Choo, "Best: Blockchain-based secure energy trading in sdn-enabled intelligent transportation system," *Computers & Security*, 2019.

[6] O. Biçer and A. Küpçü, "Anonymous, attribute based, decentralized, secure, and fair e-donation," *Proceedings on Privacy Enhancing Technologies*, vol. 2020, no. 4, pp. 196–219, 2020.

[7] N. Kumar, S. Aggarwal, and P. Raj, "The blockchain technology for secure and smart applications across industry verticals," *Academic Press*, 2021.

[8] Y. Hassanzadeh-Nazarabadi, A. Küpçü, and O. Ozkasap, "Decentralized utility-and locality-aware replication for heterogeneous dht-based p2p cloud storage systems," *TPDS*, 2019.

[9] K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. G. Sirer *et al.*, "On scaling decentralized blockchains," in *FC*.   Springer, 2016.

[10] P. Daian, R. Pass, and E. Shi, "Snow white: Provably secure proofs of stake," Cryptology ePrint Archive, 2016, https://eprint.iacr.org/2016/919.

[11] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *CRYPTO*. Springer, 2017.

[12] I. Bentov, C. Lee, A. Mizrahi, and M. Rosenfeld, "proof of activity: Extending bitcoin's proof of work via proof of stake [extended abstract]," *Performance Evaluation Review, ACM*, 2014.

[13] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *SIGSAC*.   ACM, 2016.

[14] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, "Omniledger: A secure, scale-out, decentralized ledger via sharding," in *IEEE SP*, 2018.

[15] T. Rocket, "Snowflake to avalanche: A novel metastable consensus protocol family for cryptocurrencies," 2018.

[16] A. Chepurnoy, M. Larangeira, and A. Ojiganov, "A prunable blockchain consensus protocol based on non-interactive proofs of past states retrievability," *arXiv preprint arXiv:1603.07926*, 2016.

[17] P. Otte, M. de Vos, and J. Pouwelse, "Trustchain: A sybil-resistant scalable blockchain," *FGCS*, 2017.

[18] M. Zamani, M. Movahedi, and M. Raykova, "Rapidchain: Scaling blockchain via full sharding," in *CCS*.   ACM, 2018.

[19] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich *et al.*, "Hyperledger fabric: a distributed operating system for permissioned blockchains," in *EuroSys*.   ACM, 2018.

[20] C. Cachin, "Architecture of the hyperledger blockchain fabric," in *DCCL*, 2016.

[21] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, "An overview of blockchain technology: Architecture, consensus, and future trends," in *BigData congress*.   IEEE, 2017.

[22] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey," *IJWGS, IEL*, 2018.

[23] J. Aspnes and G. Shah, "Skip graphs," *ACM TALG*, 2007.

[24] K. Wüst and A. Gervais, "Do you need a blockchain?" in *CVCBT*. IEEE, 2018.

[25] M. S. Dousti and A. Küpçü, "Moderated redactable blockchains: A definitional framework with an efficient construct," 2020.

[26] J. R. Douceur, "The sybil attack," in *International workshop on peer-to-peer systems*.   Springer, 2002.

[27] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," *Communications of the ACM*, 2018.

[28] O. Biçer and A. Küpçü, "Fortis: Forgeable timestamps thwart selfish mining," *IACR ePrint Cryptology Archive, Report 2020/1290*, 2020.

[29] Y. Hassanzadeh-Nazarabadi, A. U. Şahin, Ö. Özkasap, and A. Küpçü, "Skipsim: Scalable skip graph simulator," in *ICBC*. IEEE, 2020.

[30] Y. Hassanzadeh-Nazarabadi, N. Nayal, S. S. Hamdan, Ö. Özkasap, and A. Küpçü, "A containerized proof-of-concept implementation of lightchain system," in *ICBC*.   IEEE, 2020.

[31] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, 2014.

[32] V. Buterin and V. Griffith, "Casper the friendly finality gadget," *arXiv preprint arXiv:1710.09437*, 2017.

[33] D. Schwartz, N. Youngs, A. Britto *et al.*, "The ripple protocol consensus algorithm," *Ripple Labs Inc White Paper*, 2014.

[34] J. Kwon, "Tendermint: Consensus without mining," 2014.

[35] M. Castro, B. Liskov *et al.*, "Practical byzantine fault tolerance," in *OSDI*, 1999.

[36] A. S. Tanenbaum and M. Van Steen, *Distributed systems: principles and paradigms*.   Prentice-Hall, 2007.

[37] M. Pease, R. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *JACM*, 1980.

[38] Y. Xiao, N. Zhang, W. Lou, and Y. T. Hou, "A survey of distributed consensus protocols for blockchain networks," *IEEE Communications Surveys & Tutorials*, 2020.

[39] E. Harris-Braun, N. Luck, and A. Brock, "Holochain-scalable agentcentric distributed computing," *Alpha*, 2018.

[40] I. Eyal, A. E. Gencer, E. G. Sirer, and R. Van Renesse, "Bitcoin-ng: A scalable blockchain protocol." in *NSDI*, 2016.

[41] Y. Hassanzadeh-Nazarabadi, A. Küpçü, and Ö. Özkasap, "Locality aware skip graph," in *IEEE ICDCSW, 2015*.

[42] Y. Hassanzadeh-Nazarabadi, A. Küpçü, and O. Ozkasap, "Laras: Locality aware replication algorithm for the skip graph," in *IEEE NOMS*, 2016.

[43] M. A. Imtiaz, D. Starobinski, A. Trachtenberg, and N. Younis, "Churn in the bitcoin network: Characterization and impact," in *ICBC*.   IEEE, 2019.

[44] Y. Hassanzadeh-Nazarabadi, A. Küpçü, and Ö. Özkasap, "Interlaced: Fully decentralized churn stabilization for skip graph-based dhts," *Journal of Parallel and Distributed Computing, Elsevier*, 2020.

[45] J. Katz and Y. Lindell, *Introduction to modern cryptography*.   CRC press, 2014.

[46] S. T. Boshrooyeh and O. Ozkasap, "Guard: Secure routing in skip graph," in *IFIP Networking*.   IEEE, 2017.

[47] S. Taheri-Boshrooyeh, A. U. Şahin, Y. Hassanzadeh-Nazarabadi, and Ö. Özkasap, "A proof-of-concept implementation of guard secure routing protocol," in *SRDS*.   IEEE, 2020.

[48] Y. Hassanzadeh-Nazarabadi, A. Küpçü, and Ö. Özkasap, "Lightchain: A dht-based blockchain for resource constrained environments," *arXiv preprint arXiv:1904.00375*, 2019.

[49] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," *ACM SIGCOMM Computer Communication Review*, 2001.

[50] M. T. Goodrich and R. Tamassia, *Introduction to computer security*. Pearson, 2011.

[51] "Blockchain charts," https://www.blockchain.com/en/charts.

[52] "Global bitcoin nodes distribution," https://bitnodes.earn.com/.

[53] Y. Hassanzadeh-Nazarabadi, N. Nayal, S. S. Hamdan, A. U. Şahin, Ö. Özkasap, and A. Küpçü, "Skip graph middleware implementation," in *SRDS*.   IEEE, 2020.

Yahya Hassanzadeh-Nazarabadi received his Ph.D. degree in Computer Science and Engineering from Koç University, Istanbul, Turkey, in 2019. His research interests are distributed cloud storage systems, blockchains, and security. Dr. Hassanzadeh is currently a Senior Distributed Systems Engineer at Dapper Labs, Vancouver, Canada.

Alptekin Küpçü received his Ph.D. degree from Brown University Computer Science Department in 2010. Since then, he has been working as a faculty at Koç University, and leading the Cryptography, Cyber Security & Privacy Research Group he has founded. Dr. Küpçü has various accomplishments including 6 international patents granted, 11 funded research projects (for 9 of which he was the principal investigator), 2 European Union COST Action management committee memberships, 4 Koç University Teaching Awards, 4 Outstanding Young Scientist Awards (BAGEP, GEBİP, METU Parlar Foundation, IEEE Turkey), ACM and IEEE Senior Member Awards, and the Royal Society of UK Newton Advanced Fellowship. For more information, visit https://crypto.ku.edu.tr

Öznur Özkasap received her Ph.D. degree in Computer Engineering from Ege University in 2000 and was a Graduate Research Assistant with the Department of Computer Science, Cornell University where she completed her Ph.D. dissertation. She is currently a Professor with the Department of Computer Engineering, Koç University, and leading the Distributed Systems and Reliable Networks (DISNET) Research Lab. She is the General Co-Chair of BCCA 2020 (International Conference on Blockchain Computing and Applications), and the Area Editor for Future Generation Computer Systems, Elsevier and Cluster Computing, Springer. She is a recipient of The Informatics Association of Turkey 2019 Prof. Aydın Köksal Computer Engineering Science Award.