



HAL
open science

BlockPerf: A hybrid blockchain emulator/simulator framework

Julien Polge, Sankalp Ghatpande, Sylvain Kubler, Jérémy Robert, Yves Le Traon

► **To cite this version:**

Julien Polge, Sankalp Ghatpande, Sylvain Kubler, Jérémy Robert, Yves Le Traon. BlockPerf: A hybrid blockchain emulator/simulator framework. IEEE Access, 2021, 9, pp.107858-107872. 10.1109/ACCESS.2021.3101044 . hal-03315798

HAL Id: hal-03315798

<https://hal.science/hal-03315798v1>

Submitted on 6 Aug 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2017.DOI

BlockPerf: A hybrid blockchain emulator/simulator framework

JULIEN POLGE¹, SANKALP GHATPANDE¹, SYLVAIN KUBLER², JÉRÉMY ROBERT³, YVES LE TRAON¹

¹University of Luxembourg - Interdisciplinary Centre for Security, Reliability and Trust, 6 rue Richard Coudenhove-Kalergi, L-1359, Luxembourg

²Université de Lorraine, CNRS, CRAN, UMR 7039, France

³Cebi Luxembourg S.A, 30, rue J.F. Kennedy, L-7327 Steinsel, Luxembourg

Corresponding author: Julien Polge (e-mail: julien.polge@uni.lu).

This work was supported by the Luxembourg National Research Fund Stability 4.0 BRIDGES19/IS/13706587.

ABSTRACT Blockchain is increasingly used for registering, authenticating and validating digital assets (financial assets, real estate, *etc.*) and transactions, governing interactions, recording data and managing identification among multiple parties in a trusted, decentralized, and secure manner. Today, a large variety of blockchain technologies is expanding in order to fulfill technical and non-technical needs and requirements. Within this context, determining and most importantly evaluating the characteristics/performance of a given blockchain platform is crucial for system designers before deploying it. A number of blockchain simulators have been proposed in the literature over the past few years, as reviewed in this paper, but are often limited in several respects (lack of extensibility, do not allow for evaluating all aspects of a blockchain. . .). This paper extends and improves a state-of-the-art simulator (BlockSim) into a new simulator called “BlockPerf” to overcome those limitations. Both simulators are compared based on a real-life (benchmarking) bitcoin scenario, whose results show that BlockPerf provides more realistic results than BlockSim, improving by around $\approx 50\%$ (in average) the outcomes.

INDEX TERMS Blockchain, Simulation, Emulation, Peer-to-peer, Consensus, Performance

I. INTRODUCTION

BLOCKCHAIN is increasingly applied to all sectors of our daily life, spanning from financial applications [1], [2] to industrial ones [3], [4]. Blockchain technology is a type of distributed ledger technology (DLT) that uses a ledger stored in a distributed manner and shared among its participants in the network [5].

Decentralization, consistency, anonymity and traceability are its intrinsic features making it an interesting technology for many applications in which such aspects must be tackled. However, it is never an easy task for researchers, developers, and practitioners to decide what blockchain technologies they should select/implement, as application requirements may significantly vary from one application to another (e.g., in terms of what data should be stored, the number of transactions to be performed, etc.), without speaking about the multiple constraints of networking, computing power and communication that the application may face [6]. Several blockchain performance assessment frameworks have been proposed in the literature to overcome this difficulty, as presented in [7]–[9], but yet they often focus on functional

aspects (e.g., type of consensus, support of smart contracts) and consider fixed performance values (collected from the literature) for dynamic criteria such as throughput, latency, block validation time, etc. The reason for this is that it is complex to formalize and model the performance of such dynamic criteria [6], [10], as they depend on a wide range of parameters and inter-dependent layers, as described in the six-layer model presented in Figure 1. Looking at significant blockchain technologies such as Bitcoin [11], Ethereum [12], or Hyperledger Fabric [13], layers differ from one framework to another, and even inside a given framework, different parameter configurations are made possible. Therefore, it is essential to evaluate the blockchain performance based on both real-life and/or simulated environments, without which the blockchain selection process cannot be optimal.

When looking at papers that experimentally evaluate blockchain performance, the majority requires the implementation of the whole system (i.e., using a large set of computers/machines) [14], which also applies to well-known practical/benchmarking tools such as BlockBench [15] or Hyperledger Caliper. Such approaches/tools incur high costs

for deployment, lack of scalability (e.g., to carry out large-scale experiments) and modularity. On the other hand, simulators can help to deploy and test blockchain technologies in large-scale infrastructure settings. To date, several blockchain simulators exist (e.g., BlockSim [16], PeerSim [17], Shadow [18], Vibes [19], etc.), but they are often limited in several respects. Recent literature reviews of existing blockchain simulation tools [20], [21] point out the fact that those tools often limit themselves to evaluate part of the blockchain system (i.e., they fail in covering all layers and associated performance metrics (a) to (o) emphasized in Figure 1). This is particularly stressed by R. Paulavicius *et al.* [20] in their systematic review and empirical survey of blockchain simulators, in which the authors conclude that “there is no ‘one-size-fits-all’ Proof-of-Work blockchain simulator that is able to accurately simulate all the layers”. To overcome such a limitation, a new hybrid blockchain emulator/simulator called *BlockPerf* is proposed, which extends the BlockSim simulator proposed by Faria *et al.* [16]. One of the main improvement lies in the fact that BlockPerf relies on real network infrastructure (at the Network layer) while simulating the upper layers, leading to more realistic results than existing simulators.

State-of-the-art simulators and the extent to which they cover the six-layer model and associated performance metrics are reviewed and discussed in section II. The architectural design of BlockPerf, and how it extends BlockSim, is presented in section III. In section IV, a performance comparison analysis between BlockPerf and BlockSim is carried out based on a real-life (benchmarking) bitcoin scenario; the conclusions and BlockPerf limitations being discussed in section V.

II. BACKGROUND AND RELATED WORK

As previously discussed, evaluating a blockchain technology or platform turns to be a complex process due to the various inter-dependent layers and associated parameters. To better explain this complexity, a slightly adapted version of the model introduced in [15] is considered in this paper, which corresponds to the six-layer model given in Figure 1. Although one may find other blockchain abstraction models in the literature, as the ones proposed by the ITU and ISO standardization bodies [22], [23], we adopted this six-layer model as it covers most aspects of a blockchain (DLT) platform, and it is straightforward to understand. Nonetheless, to allow readers to understand to what extent this model covers the ongoing ITU and ISO blockchain standard initiatives, and *vice-versa*, we provide a summary table in Table 1. Considering this model, sections II-A to II-F provide the necessary background regarding each of these layers, starting from top (Application) to bottom (Network), by detailing the key metrics – (a) to (o) – that a simulator should allow for measuring/tracking throughout a simulation run. Section II-G then discusses the extent to which existing blockchain simulators cover those layers and metrics.

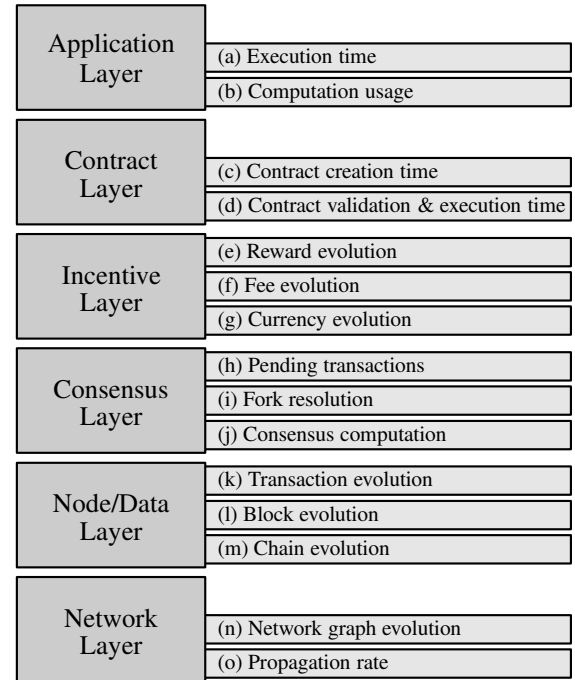


Figure 1: Blockchain abstraction layer model & associated metrics denoted by (a) to (o) in the rest of this paper.

Table 1: Extent to which the six-layer model introduced in Figure 1 covers the ITU and ISO blockchain standard initiatives, and *vice-versa*

Standards	Layer Focus					
	Application	Contract	Incentive	Consensus	Node/Data	Network
ITU-T SG16 Q22 F.751.1	✓	✓		✓	✓	✓
ITU-T SG16 Q22 F.751.2	✓	✓		✓	✓	✓
ISO/DIS 22739	✓	✓	✓	✓	✓	✓
ISO/WD TS 23258		✓		✓	✓	✓
ISO/CD 23257				✓	✓	✓
ISO/CD TR 23576					✓	✓
ISO/AWI TS 23259		✓				

A. APPLICATION LAYER

This layer manages the user interface, APIs (Application Programming Interface), and computational resources (e.g., needed for blockchain element storage, wallet creation, etc.). In many blockchains, there is the possibility to configure a node to run as a full node (storing a local copy of all transactions and blocks) or as a lightweight node (in charge of creating transactions and sending them to full nodes for validation purposes). From a simulation (or emulation¹) perspective, the following metrics should be measurable at this layer:

- (a) *Execution time*: it refers to whether the simulator keeps track of the time needed to run the simulation;

¹Simulation/Simulator and Emulation/Emulator are interchangeably used in the rest of the paper.

- (b) *Computational resource usage*: it refers to whether the simulator keeps track of the resource usage evolution throughout the (simulation) run, which includes CPU of each node, swap space, storage capacity, etc.:

B. CONTRACT LAYER

Blockchain systems consist of scripts, also referred to as "smart contracts", that run when predetermined conditions are met (e.g., used to automate the execution of an agreement so that all participants can be immediately certain of the outcome). At this layer, the following metrics should be measurable:

- (c) *Contract creation time*: it refers to whether the simulator keeps track of the time needed to generate the different contracts (which can be of different sizes) over the simulation;
- (d) *Contract validation & execution time*: it refers to whether the simulator keeps track of the time needed to validate and execute the different contracts of the simulation scenario.

C. INCENTIVE LAYER

Participation incentive means the way how honest behavior is incentivized and dishonest discouraged. Incentives can be in the form of transaction fees and/or rewards [24]. This decision affects the implemented consensus algorithm (introduced in the next model layer) and, respectively, is affected by the selected algorithm. At this layer, the following metrics should be measurable:

- (e) *Reward evolution*: it refers to whether the simulator keeps track of the amount of cryptocurrencies distributed from the consensus process (e.g., leader election or mining) over the simulation;
- (f) *Fee evolution*: it refers to whether the simulator keeps track of the reward fees that client nodes offer to miners to incentivize them to process their transaction(s).
- (g) *Currency evolution*: it refers to whether the simulator keeps track of the currency generation rate, which evolves differently according to the implemented blockchain (e.g., in bitcoin or Ethereum, it evolves along with the hashing difficulty).

D. CONSENSUS LAYER

Consensus protocols are needed to validate the data to prevent and remove any duplicated entry and/or fraud [25], [26]. The type of blockchain to be implemented (public, private, consortium) influence profoundly the type of consensus protocol to be used, the most well-known being Proof-of-Work (PoW), Proof-of-Stake (PoA), or still Practical Byzantine fault tolerance (PBFT). At this layer, the following metrics should be measurable:

- (h) *Pending transactions*: it refers to whether the simulator keeps track of the number of transactions, over time, that are waiting to be confirmed (such a waiting area is called "Mempool" in Bitcoin, or sometimes called transaction queues);

- (i) *Fork resolution*: it refers to whether the simulator keeps track of (i_1) the number of forks that appear within the chain; (i_2) the stale rate (i.e., block discarded) throughout the simulation;
- (j) *Consensus computation*: it refers to whether the simulator keeps track of the collective (or individual) computation effort required to validate transactions and blocks.

E. NODE/DATA LAYER

The node (or data) layer is responsible for structuring the data before appending it to blocks, whose structure usually includes information such as previous block hash, Merkle root, time, bits, etc. At this layer, the following metrics should be measurable:

- (k) *Transaction evolution*: it refers to whether the simulator keeps track of the number of transactions generated per day (k_1) and whether the simulated transactions match the real-life data structure (k_2). Unlike k_1 , k_2 is not a quantifiable metric but rather a boolean metric that states whether the simulator generates transaction following the real-world blockchain specification;
- (l) *Block evolution*: it refers to whether the simulator keeps track of (l_1) the number of blocks that are validated, mined and accepted as part of the longest chain; (l_2) the (average) time taken by each block to be validated; (l_3) the block sizes, which depend on the size of the transactions they include; and finally (l_4) the number of transactions included, on average, within a block.
- (m) *Chain evolution*: it refers to whether the simulator keeps track of the length of the chain over time (i.e., the number of blocks that form the longest chain), which is a good indicator of the load a new node would have to compute if it joins the network at a given point in time.

F. NETWORK LAYER

Blockchain is a pure P2P network, which is actually an overlay network [27] for distributed object storing, searching, and sharing (e.g., Ethereum relies on the Kademlia P2P protocol [27], [28]). At this layer, the following metrics should be measurable:

- (n) *Network graph evolution* it refers to whether the simulation accurately follows the P2P protocol overlay network specifications (e.g., support for adding and discovering a node at any time) and keeps track of the network (node) evolution using network graph metrics such as Clustering Coefficient, Mean Geodesic Distance, and Diameter [29], [30];
- (o) *Throughput*: it refers to whether the simulator keeps track of the number of valid transactions per second (Tx/s) that have been incorporated as part of a valid block within the longest chain. The transactions that are considered here are the ones that reach the majority of nodes within the network depending on the underlying consensus (i.e., $\geq 50\%$ for PoW, $\geq 66\%$ for PBFT, etc.);

Table 2: Literature comparison - ● symbol indicates that the simulator makes available the performance metric without any code modification and/or post-processing, while ◐ symbol indicates that such a modification and/or post-processing is needed to obtain the desired performance metric

Framework	Application		Contract		Incentive			Consensus			Node/Data			Network	
	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)	(j)	(k)	(l)	(m)	(n)	(o)
DLSF [31]	◐	◐			●			●	◐	◐	● _{k1,k2}	● _{11,12,13,14}	◐		◐
Shadow [18]	●	●			◐				●		● _{k1}	◐	◐		●
Wang et al. [32]	◐	●			●			●			● _{k1}	● _{12,14}			
BlockSim-Net [33]	◐	●					◐		●	◐	◐ _{k1,k2}	● _{11,12,13}	◐		◐
eVibes Plasma [34]	◐		◐	●					●		◐	◐	◐		◐
Vibes [19]	●							◐			● _{k1}	● _{11,12,14}	◐		●
eVibes [35]	◐			●					●		● _{k1}	● _{11,12,14}			●
BlockSim [16], [36]	●	◐				◐			●		● _{k1}	● _{11,12}	◐		●
Goswami [37]	◐								◐		● _{k1}	● _{12,14}			
Bitcoin-Simulator [38]	◐						◐		◐		● _{k1}	● _{11,12,13,14}	◐		●
SIMBA [39]	◐								◐	◐	● _{k1}	● _{12,13,14}	◐		◐
Chin et al. [40]	◐				◐		●		●		● _{k1}	● _{11,12,14}	◐		◐
SimBlock [41], [42]	◐				◐				●		● _{k1}	● _{11,12,14}	◐		●
CIDDS [43]		◐												●	●
HIVE [44]				◐	◐	◐			◐	●	● _{k1,k2}	● _{11,12,13,14}			◐
Androulaki et al. [45]		◐							●	◐		● _{11,14}	●		
PeerSim [17]														●	●

G. RELATED WORK AND DISCUSSION

As of today, there are several blockchain simulators found in the literature, a summary of which is in Table 2. This table highlights what layers and associated metrics those simulators cover. Note that a simulator may, in some cases, cover a given layer but without necessarily providing as output a performance metric. This means that either a modification of the source code or post-processing treatments are required in order to compute the desired metric. For example, in BlockSim, the authors claim in their paper [36] that the simulator keeps track of the fee evolution, however, after analyzing it, we realized that it is not possible to retrieve it without modifying part of the code. Another example is HIVE [44] that emulates the behavior of smart contracts within the Ethereum environment, however, we found that it does not allow for analyzing the (h) metric (i.e., pending transactions). To make a distinction between simulators that make available a performance metric without requiring any code modification or post-processing, and the ones that require a modification/post-processing to obtain the metric, the following two symbols are respectively used in Table 2: ● (does not require any code modification and/or post-processing) and ◐ (does require a code modification and/or post-processing).

As a first simulator, let us mention Bitcoin-Simulator [38], which has been designed for educational purposes to help students to understand how the block generation rate (l_1) and block size (l_2) evolve over time. Although it is a well-designed pedagogical tool, it is quickly limited to carry out in-depth simulation analyzes. More research-oriented simulators have been proposed, such as Ethereum Hive [44] that have been proposed for emulating and evaluating Ethereum’s smart contracts from a validation & execution time perspective (d_1 - d_2). However, as revealed in [35], simulating a

large number of nodes becomes difficult with limited computational resources. To overcome this limitation, a series of simulators including VIBES [19], eVIBES [35], eVIBES Plasma [34] ones, as well as CIDDS [43], were proposed, but unfortunately fall short of fulfilling the promise, as reported by Lathif et al. [43]. One of the main reasons for this is that those simulators do not adequately model the transactions at the Node/Data layer (i.e., k_2), considering them as empty in the majority of the simulators. Such a consideration poses several issues when evaluating the blockchain system as a whole, knowing that transaction and block sizes may have non-negligible impacts on the overall system performance [46]. To overcome this issue, a new range of simulators, including BlockSim [16], [36], SIMBA [39], DLSF and others (cf., Table 2), have considered the transaction/block structure, algorithms for wallet creation, message signing, and so forth, thus leading to more realistic performance evaluation results. Nonetheless, the network modeling is often too simplistic, not reflecting the real behavior of the P2P overlay network (i.e., possible communication delays, network congestion, packet losses, computation resource limitations, etc.), which, in our opinion, can have significant impact on the overall system performance evaluation process. A few simulators have been designed to consider the throughput and the network graph evolution, such as CIDDS, PeerSim. However, these two simulators neglect many of the upper layers and metrics, as highlighted in Table 2.

The drawbacks discussed in this section prove that designing a modular simulator that allows for testing/evaluating different blockchains, with different consensus protocols, different contract and transaction/block specifications, different incentive schemes, along with network infrastructures, turns to be a challenging task. Figure 2 provides an overview of the extent to which the reported blockchain simulators cover the




Figure 2: Overview of the extent to which state-of-the-art simulators – the 17 reported in Table 2 – cover the six-layer model introduced in Figure 1.

different layers of the six-layer model previously introduced. It can be first observed that they primarily often neglect the Contract and Incentive layers. Second, key aspects of a blockchain systems at the Network and Data/Node layers, namely the consideration of the real P2P overlay network protocol behavior (n) and transaction/block structure (k_1), result in non-optimal performance evaluation results (non-optimal compared to the reality). A new hybrid emulator/simulator tool called “BlockPerf” is proposed in this paper to overcome this limitation. BlockPerf is hybrid in the sense that it emulates the network layer to correctly address the network layer while simulating the upper layers based on statistical data modeling approaches, as presented in the next section.

III. BLOCKPERF SIMULATOR DESIGN

This section describes how BlockPerf extends BlockSim, whose main extensions are summarized in Figure 3. These layer extensions and/or adaptations are respectively discussed through sections III-A to III-E. Note that, at this stage, BlockPerf does not address/model the “Smart Contract” layer and is part of future research work.

A. APPLICATION LAYER

At this layer, BlockPerf consists of a configuration module responsible for instantiating the run’s parameters similar to the one used in BlockSim. BlockSim takes as input a list of parameters, including the size of blocks, statistical models for transaction validation, block validation, and node labels. Similarly, BlockPerf takes as input a JSON file that extends the parameters from BlockSim to include fork occurrences, the block size parameter, IP address of each node, including the ‘main’ one (corresponding to the BitcoinNode class, as summarized in Table 3) that tracks all the results using the CliStats, CpuTimeSnapshot and MemorySnapshot classes. Optionally, the input file also considers the location to store all the logs and data for each node. Note that the BitcoinNode is responsible for tracking the execution time (a) of the overall run and collecting the metrics regarding the overall computational

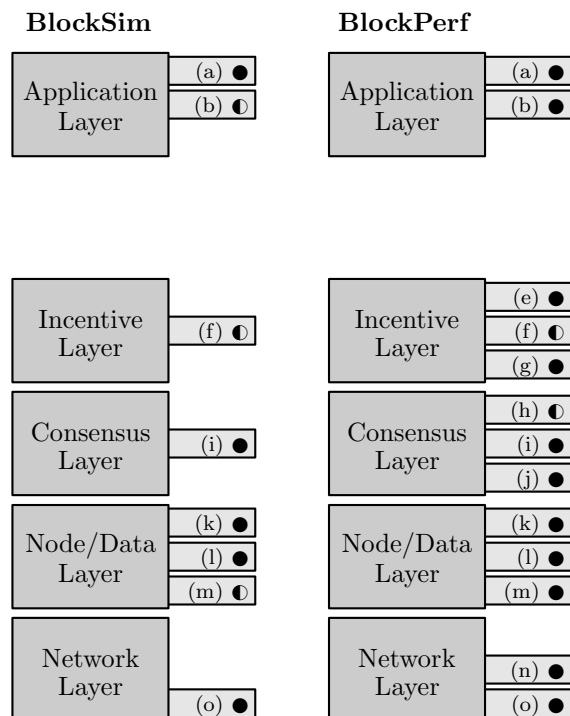


Figure 3: Illustration of how BlockPerf extends BlockSim - ● symbol indicates that the simulator makes available the performance metric without any code modification and/or post-processing, while ◐ symbol indicates that such a modification and/or post-processing is needed to obtain the desired performance metric

resource usage (b) of the nodes.

B. INCENTIVE LAYER

This layer is responsible for instantiating the models for rewarding participants, which vary from one blockchain to another. While BlockSim does not simulate the incentive layer², BlockPerf models two types of rewards: one for the generation of the valid blocks (known as *block reward*) and the second for the inclusion of a particular transaction to the block (known as *transaction fee*). These reward operations are implemented via the TxChain, BTCNode and TickEvent classes (cf., Table 3. As of today, the *block reward* is represented as a fixed amount of cryptocurrency that can be configured for the run, while the *transaction fee* is dependent on the size of the transaction, as formalized as in Eq 1, where Tx_i refers to the i^{th} transaction and $f(\text{size}(Tx_i))$ can be configured using different distributions laws (e.g., defined as a fixed fee, or following a Weibull, Log-normal, or Gamma distribution). All node wallets are continuously updated throughout the simulation run.

$$Tx_i^{fee} = f(\text{size}(Tx_i)) \quad (1)$$

²Even though the authors in [36] state that the incentive layer is partly covered, nothing from the source code provides clear evidence of this.

Table 3: Technical description of the major classes within BlockPerf, along with the layers they cover.

Class Name	Description	Application	Incentive	Consensus	Node/Data	Network
CliStats	Node state checking and logging (incl., parameters like operation time, computation)	✓		✓	✓	
CpuTimeSnapshot	Computation process and database tracking	✓			✓	✓
MemorySnapshot	Output file generation for the memory used by a given node	✓				
Proxy	Authentication management for network components and wallet information	✓				
TxChain	Chain operation handling (address generation, private-public keys, unspent Tx, etc.)		✓	✓	✓	
PublicBitcoinNode	Network stat monitoring (network graph, latency...)				✓	✓
BTCNode	Monitoring of transactions/block evolution, handshaking, etc.		✓		✓	✓
Node	Node operation handling (incl., creation of transaction, Queues)				✓	✓
TickEvent	Transactions/block timestamping		✓		✓	
TransactionQueue	Queue transaction creation (i.e., transactions to be generated by a given node)				✓	
BlockHeader	Header tracking for the chain that a node follows			✓		✓
Consensus	Consensus protocol model instantiation		✓	✓		
BitcoinNode	Connection handling and in charge of pushing model parameters to the network					✓
ZoneConfig	Node parameter consideration for the emulation layer (incl., IP addresses)					✓
Event	Event tick creation based on the models					
Runner	<i>main</i> class					

The amount of new cryptocurrencies generated over time depends on the consensus layer and the overall computational resources of the nodes in the network. BlockPerf continuously monitors such resources and allocates the newly generated cryptocurrencies to the right (mining) nodes.

C. CONSENSUS LAYER

In BlockPerf, the consensus protocol is implemented as a model that represents the consensus process and its operations. The consensus algorithm is coded within the `Consensus` class (see Table 3), which is in charge of selecting the miner/validator that builds the next block. The method of selecting the node varies from one blockchain to another, although in BlockPerf, as a first step, the Proof-Of-Work (PoW) algorithm has been implemented. As opposed to BlockSim that simulates the behavior of the consensus algorithm (i.e., the block validation and a random selection of miner), BlockPerf uses an extended approach where each mining node, similar to PoW algorithm, selects a random number, the input for all the non-confirmed transactions from its queue within the limit of block size (further discussed in Section III-D), as well as the reference of the previous known block, which allows for being closer to a real PoW process. This information is then combined and hashed recursively until a result is obtained. The mining node selects the transactions from the queue (sometimes called mempool), check whether they meet the balance requirements, verify whether the sender signatures match and that the sender's wallet has a sufficient amount of cryptocurrencies. In BlockPerf, the `BitcoinNode` class keeps track of the fork occurrences within the network.

D. NODE/DATA LAYER

Any node wishing to generate a transaction has to follow a set of operations, namely: (i) select a wallet address (from *wallet_list*) as the recipient; (ii) select a random value for the transaction; (iii) generate the transaction by sign-

ing it using its private wallet key (those operations being handled via the `TickEvent`, `Transaction Queue` and `TxChain` classes, as reported in Table 3). The set of transactions to be created per unit of time t , which is denoted by $\mathcal{X}(t) = \{Tx_1, \dots, Tx_k\} \mid k \in \mathbb{N}$, can be configured using different distributions laws (e.g., a fixed amount of transactions per unit of time, or following a Weibull-like distribution). Every new transaction created by a node calls the propagation function from the network layer to transmit this new transaction to its neighbors, who then append it to their pool. Upon reception of a transaction (Tx_i), several operations are undertaken by the recipient node, as summarized in the flow chart Figure 4 (operations that have been added compared with BlockSim being highlighted in green). Unlike BlockSim, BlockPerf models transactions as they exist in the real system (i.e., each transaction can be traced back and follow a similar process within the system).

As a second stage, the simulator has to handle the inclusion and/or removal of transactions – *from its transaction pool* – into blocks. Although the consensus layer governs operations towards deciding which block to be accepted by everyone within the network, the Data Layer is concerned with the reception of the block, being in charge of removing all the transactions which are enclosed in the last received block from its transaction pool, and finally attaching the new block to its last known state of the chain (handled by the `Node` class). These operations follow the sequence of operations given in the flow chart of Figure 5 (new operations compared with BlockSim being highlighted in green, while BlockSim's operations that have been modified are highlighted in orange).

The linked blocks form a chain, but nodes may have a different state of the chain (or ledger) due to network communication delays (this is termed forks). As opposed to many existing simulators, including BlockSim, each node in BlockPerf has its local ledger implemented, which means that if a node discovers that the state of the chain is different from

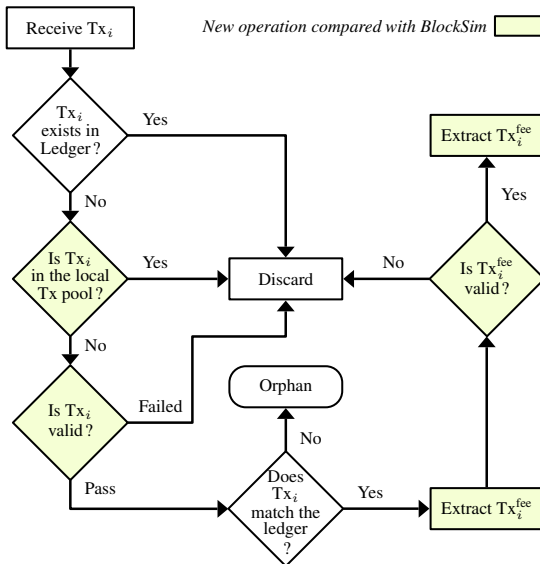


Figure 4: Flowchart of transaction operation sequence upon reception of Tx_i in BlockPerf

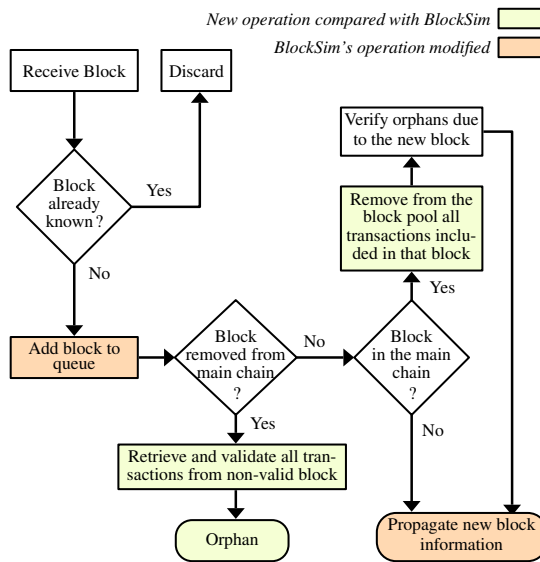


Figure 5: Flowchart of transaction inclusion/removal from a block in BlockPerf

its own, it sends a request to neighboring nodes for getting the full state of the chain. This allows being closer to reality, which should result in more realistic simulation performance. This improvement, compared with BlockSim, is highlighted in the flow chart given in Figure 6 (cf. green boxes/steps).

E. NETWORK LAYER

The network layer of BlockPerf differs significantly from BlockSim, and many other simulators reported in Table 3, since, in BlockPerf, the blockchain network is emulated over distributed (physical) nodes. The benefit of doing so is that the simulation should lead to results closer to reality, as

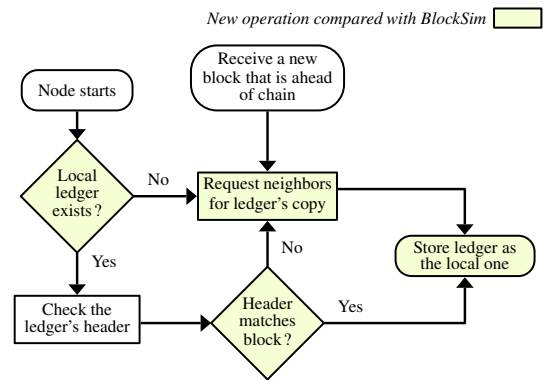


Figure 6: Flowchart of local ledger in BlockPerf.

BlockPerf implements a P2P protocol similar to the ones used by real blockchain systems using advertisement messages (e.g., GETADDR and ADDR) to discover neighboring nodes. BlockPerf also integrates:

- *a wallet management module*: in charge of generating wallet addresses to uniquely identify all nodes. This module also integrates the protocol to broadcast wallet-related information in the network (via WALLADD messages) to make all nodes aware of existing addresses;
- *a reputation score module*: in charge of assigning negative scores from the connections where malformed messages were received (e.g., if the transaction correctness fails, nodes keep receiving outdated block information). This module is fulfilled by the `PublicBitcoinNode` class, which is responsible for keeping a log of the neighboring nodes and associated scores. Such logs allow BlockPerf to keep track of the network graph evolution (n) over the simulation run.

F. BLOCKPERF IMPLEMENTATION

Table 3 provides an overview of the important programming classes underpinning BlockPerf, whose implementation is combination of C/C++ programming languages for the network layer, and Python for the upper layers. The source code of the BlockPerf simulator is available on GitHub³.

IV. VALIDATION AND EVALUATION

This section aims to compare the extent to which BlockPerf provides realistic results, but also to which it outperforms BlockSim. Figure 7 provides an overview of the experimental process that has been conducted in this respect, which consists of five main stages denoted by ① to ⑤ in Figure 7. First (①), a real-life blockchain network (`bitcoind`) consisting of 40 nodes spread over the world have been implemented, which corresponds to the benchmarking Testbed in this study whose metrics (a) to (o) have been collected/measured whenever possible (see ②). Based on those metrics, several parameters are specified as input parameters of BlockPerf

³BlockPerf repo: <https://github.com/Deadlyelder/BlockPerf>

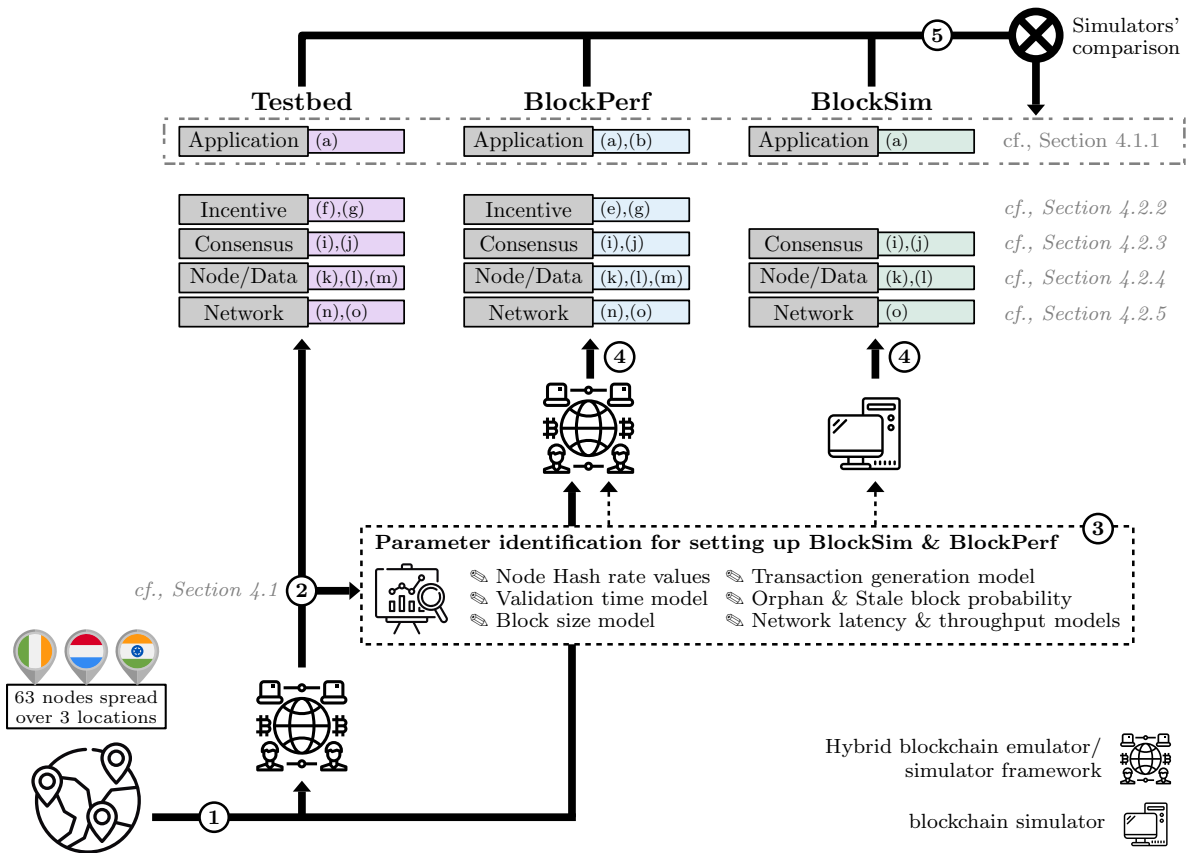


Figure 7: Experimental process set up and conducted in this paper to compare BlockSim and BlockPerf simulators

and BlockSim (see ③ and ④), as will be presented in section IV-A. BlockPerf and BlockSim are then compared against the results obtained for the (benchmarking) Testbed in section IV-C (see ⑤). A discussion about our experiments is finally given in section IV-C.

A. BENCHMARKING TESTBED

The (benchmarking) Testbed is made of a modified version of the reference bitcoin implementation, namely `bitcoind`⁴. The evaluation is performed during 11 days, whose blockchain network consists of 23 full nodes and 40 light nodes spread over three locations, namely Ireland, Luxembourg and India. During the course of the experiment, transactions are randomly created. Based on this experimental configuration setting, metrics (a) to (o) are measured, whenever possible, and then reused for two purposes:

- 1) to serve as benchmarking metrics to evaluate the extent to which BlockPerf and BlockSim deviate from the reality (see ⑤ in Figure 7);
- 2) to identify, from the testbed run, several parameters that need to be configured as inputs of BlockPerf and BlockSim such as node hash rate values, block size distribution model, etc. (see ③-④). To this end, a similar process as the one defined in [16] has been applied to

⁴A slight difference can be noted, as our testbed is a private network where all nodes are controlled and maintained throughout the experiment.

extrapolate the probability distribution model for each parameter.

Table 4 summarizes all the metric values obtained from the Testbed, and, for some of them refer to specific figures/graphs. For example, looking at metric k1, Table 4 reports the average number of transactions/day, while Figure 11 reports the exact number transactions over the 11 days of the experiment).

B. EXPERIMENT WITH BLOCKPERF AND BLOCKSIM & COMPARISON

In this section, the performance of BlockPerf and BlockSim is compared against the benchmarking testbed results. BlockSim experiments have been conducted on a single machine with a 2.40 GHz Intel Xeon E5 CPU and 4GB of RAM. In contrast, BlockPerf has been deployed over distributed nodes – located in Ireland, Luxembourg and India – with the same computational resource (i.e., 2.40 GHz Intel Xeon E5 CPU and 4GB of RAM). Sections IV-B1 to IV-B5 discuss the comparison analyses regarding the different layers, as emphasized in Figure 7 (see ⑤). Note that for this comparative analysis, only the metrics that are covered and for which the results are available (i.e., metrics with ● in Figure 3 and listed in Figure 7) are considered.

Table 4: Overview of the results

	Metrics	Evolution (cf.,)	Testbed	BlockPerf	BlockSim
Application	(a) Execution time (in hours)	-	264h	8h	0.5h
	(b) Computation usage (Average CPU)	-	80%	90%	-
	(Average memory)	-	40%	86%	-
	(Average storage)	-	2.6GB	2.9 GB	-
Incentive	(e) Average reward per day	Figure 8	-	0.043991	-
	(f) Fee evolution	-	<i>by default</i>	0.00001	-
	(g) Currency generat. rate per time unit	Figure 9	720 BTC	384 BTC	-
Consensus	(i_1) Total number of forks	-	2	5	2
	(i_2) Total number of blocks discarded	Figure 10	4	13	5
	(j) Min-Max consensus computation effort	-	10-45 MH/s	9-42 MH/s	1-2 MH/s
Node/Data	(k_1) Average Tx/day	Figure 11	7341	4504	14184
	(l_1) Total number of validated blocks	Figure 12(a)	1979	1118	4132
	(l_2) Average time to validate blocks	Figure 12(b)	10.7min	10.9min	10.4min
	(l_3) Average block size	Figure 12(c)	1.6 MB	1.1 MB	2 MB
	(l_4) Average number of Tx/block	Figure 12(d)	41	44	38
Network	(m) Chain's length (total number of blocks)	Figure 13	1980	1119	4133
	(n) Network graph evolution (Cluster Coefficient)	-	0.473217	0.469478	1
	(Mean Geodesic Distance)	-	2.08631	2.11437	1
	(Diameter)	-	4	4	1
	(o) Average throughput (Tx/s)	Figure 17	1.5	1.6	2

1) Application Layer

From an *Execution time* (a) perspective, BlockSim and BlockPerf allow for simulating the experiment at a faster pace compared to the testbed, the former being 528 times faster while the later is 33 times faster. This substantial difference is explained by the fact that BlockSim does not rely on a real-life network layer. In terms of *Computational resource usage* (b), the average CPU, memory and storage metrics have been reported in Table 4, although it cannot be concluded that one simulator is better than another for those metrics.

2) Incentive Layer

Unlike BlockSim and the Testbed, BlockPerf stores log about the *Reward evolution* (e), as they are generated within the run. Figure 8 provides insight into the evolution of the reward over the 11 days of simulation. On average, the reward is 0.043991 BTC per day, equally distributed to the miners with a standard deviation of 0.0002 throughout the run. Although no comparison for this metric could be made with the Testbed and BlockSim, we nonetheless report those results as they could serve as benchmarks for other researchers.

The *Currency generation rate* (g) can be compared with the Testbed, but not with BlockSim (metric not available). Figure 9 shows the evolution of the currency generated over the course of the simulation, 384 BTC being generated on average by BlockPerf, against 720 BTC generated by the Testbed. Despite this difference, what is interesting to note is that the evolution in BlockPerf and Testbed follows a similar trend.

3) Consensus Layer

During the simulation, a number of *Fork resolutions* (i_1) are handled by the Testbed (2 in total, one at day 4 and one at day 5), by BlockSim (2 in total, one at day 4 and one at day 5), and by BlockPerf (4 in total, three at day 5 and two at day 6). These forks led to the appearance of stale blocks (i_2), whose




Figure 8: Reward evolution over days (e)




Figure 9: Currency generated over days (g)

occurrences over the 11 days of the experiment have been reported in Figure 10. Two observations can be made: (i) stale blocks mostly appear in the same timeframe (between day 4 and 6); (ii) BlockPerf provides more realistic results than BlockSim, as the number of stale blocks in BlockPerf (4 in total) is closer to the Testbed (6) than BlockSim (13).

Regarding now the *Consensus computation effort* (j), BlockPerf also provides more realistic results, ranging between 9-42 MH/s, while the Testbed ranges between 10-45 MH/s and BlockSim between 1-2 MH/s. The reason for such a difference is that, in BlockSim, even when specifying as input parameters the min and max hash rate values (i.e., 10-45 MH/s identified through the benchmarking phase, cf.




Figure 10: Stale blocks per days (i_2)




Figure 11: Transaction evolution (k_1) & relative absolute error with respect to (w.r.t) the Testbed

section IV-A), it seems that BlockSim does not consider that parameters in the simulation run, always applying 1-2 MH/s as hash rate range.

4) Node/Data Layer

As the first metric of that layer, let us look at the *Transaction evolution* (k_1) in Figure 11, which shows the number of transactions generated daily by the Testbed and the two simulators. It can be observed that BlockPerf closely follows the Testbed’s transaction evolution, thus providing more realistic results than BlockSim. Indeed, BlockPerf and BlockSim respectively generate 4504 and 14184 Tx/day on average, while the Testbed generates 7341 Tx/day. For a more accurate view on the extent that BlockPerf leads towards realistic results, we show the relative absolute errors of BlockPerf and BlockSim with respect to the results obtained from the testbed is represented in the form of a boxplot in Figure 11. For example, the max value of the BS (BlockSim)-related boxplot (i.e., ≈ 12000) indicates that, over the 11 days of experiments, the maximal difference in terms of number of transactions generated by the Testbed and BlockSim is 12000, which – if we look at the transaction evolution in Figure 11 – corresponds to day 5. Looking at the error median values, it can be noted that the difference/error is quite significant in BlockSim, being approximately twofold higher compared to BlockPerf. In addition of k_1 , let us also note that BlockPerf follows the *real-life transaction specification* (k_2), while BlockSim adopts a more simplistic model, which may explain part of the difference in results of the transaction evolution above-discussed (k_1).

Let us now discuss about the *Block evolution* (l), which consists of four sub-metrics (l_1 to l_4). Figure 12(a) shows

the evolution of the number of valid blocks (l_1) over the 11 days of experiment. It can be observed that BlockPerf produces fewer blocks than the testbed, which is likely due to the unpredictable nature of the network connections and the occurrence of forks. On the other hand, BlockSim produces a number of blocks ($\approx 400/\text{day}$) that is twice higher than the Testbed ($\approx 200/\text{day}$). Overall, as evidenced through the boxplot in Figure 12(a), BlockPerf provides more realistic results than BlockSim, whether in terms of min, quartiles and max values. Looking now at the second sub-metric (l_2 : block validation times), both BlockPerf and BlockSim provide similar results to the testbed, as shown in Figure 12(b). Indeed, the three curves followed the same trend and the relative absolute error was ≤ 1 min (i.e., one-tenth of the total time required by the Testbed). Figure 12(c) then provides insight into the evolution throughout the run of the block validation times (l_3), along with the difference – *relative absolute error to be precise* – between BlockPerf/BlockSim and the Testbed. It is interesting to note that, when looking at the boxplot, BlockSim provides closer results with the Testbed than BlockPerf; however, when looking at the day-by-day block validation times, BlockPerf evolves in a similar way to the Testbed, while BlockSim does not. Finally, regarding l_4 (i.e., the average number of transactions mined within blocks per day), the two simulators provide similar results to the Testbed, following a similar trend over the course of the run and having similar differences/errors compared to the Testbed (the min and max errors being the same for BlockPerf and BlockSim).

The last metric of the Node/Data layer refers to the length of the chain (m), whose evolution over the 11 days of the experiment for the Testbed and the two simulators are plotted in Figure 13. It can be observed that BlockPerf provides more realistic results than BlockSim, which becomes increasingly significant over time. Indeed, the longer the simulation run, the higher the difference between BlockSim and the Testbed. This is also confirmed by the boxplot given in Figure 13, the minimal error values of the two simulators being approximately the same, while the quartiles, median and max values become increasingly higher for BlockSim.

5) Network Layer

Two distinct metrics are analyzed at this layer, namely how the network graph evolves over time (n), and what throughput performance is possible with the implemented blockchain and scenario (o).

Regarding the first metric (n), Figures 14, 15 and 16 provide the network graph evolution over two consecutive days (days 1 and 2) respectively regarding the Testbed, BlockPerf and BlockSim. Pink nodes correspond to light nodes and yellow ones to full ones, each node being denoted by the country’s initial (L: Luxembourg, I: Ireland; In: India) and the node’s numbering. A twofold observation can be drawn from those graphs. First, unlike BlockPerf and the Testbed, BlockSim assumes that all nodes are interconnected in a fully connected mesh throughout the experimental run, adding

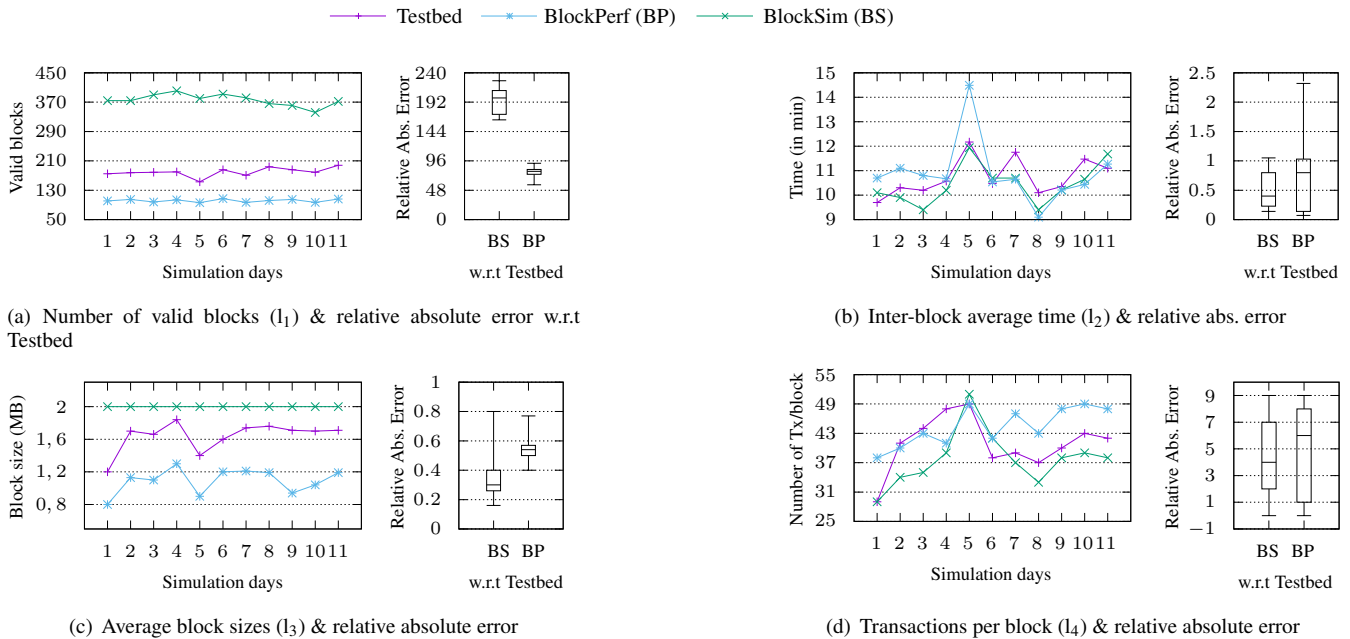


Figure 12: Block evolution (I_1 - I_2)




Figure 13: Chain's length (m)

that it does not distinguish lightweight and full blockchain nodes. A second observation is that the P2P logical network infrastructure evolves quite substantially from one day to another, whether regarding the Testbed or BlockPerf, which are due, among other things, to eventual connection losses, ongoing traffic, . . . The Diameter⁵, mean geodesic distance⁶, and clustering coefficient⁷ reported in Table 4 show that the graphs evolve in a similar manner for BlockPerf and BlockSim, which is not the case for BlockSim.

Throughput is a key metric that provides a good performance indicator of a given blockchain system. It is nonetheless tricky to analytically formalize it, as it depends on

⁵Diameter: $\max(\text{dist}(n_i, n_j)) \forall n_i, n_j \in \mathcal{N}$, with \mathcal{N} the set of nodes (full and lightweight) from the network graph.

⁶Mean geodesic distance: $\text{mean}(\text{dist}(n_i, n_j)) \forall n_i, n_j \in \mathcal{N}$, with $\mathcal{N} = \{n_1, \dots, n_N\}$ the set of nodes (full and lightweight) from the network graph.

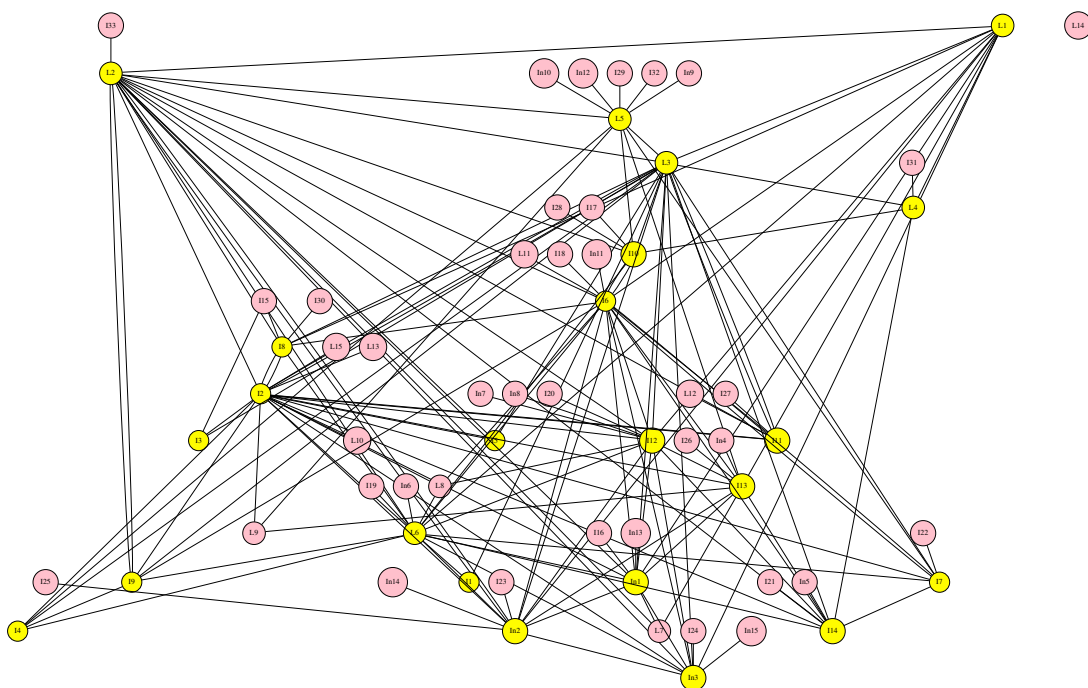
⁷Clustering coefficient of entire graph: $C = \frac{1}{N} \sum_{i=1}^N C_i$, where C_i refers to the clustering coefficient of a node n_i that is calculated as follows: $C_i = \frac{2 \cdot L_i}{k_i(k_i - 1)}$, k_i referring to the degree of node n_i and L_i to the number of edges between the k_i neighbors of n_i .

multiple processes across the different layers. Indeed Application (for generation), Node/Data (for accessing rewards), Consensus (mining), Incentive (for rewards), and Network layers (for final throughput), all combined, influence the throughput of a given blockchain system. Figure 17 shows the evolution of the average throughput for the considered scenario, along with the boxplot showing the relative absolute errors of BlockSim and BlockPerf when compared to the Testbed. The results show that throughput ranges between 1 and 2.5 Tx/s for both the Testbed and the two simulators (BlockPerf having slightly closer results to the Testbed than BlockSim).


C. DISCUSSION

The previous section has shown that BlockPerf outperforms BlockSim for most of the metrics (a)-(o), or provide more realistic results to be more precise. We firmly believe that this is mainly due to the assumption made by BlockSim (i.e., non-consideration of the network-level change within the system). Furthermore, our approach presents several advantages. First, it covers all the layers (except the contract one) and allows to obtain a higher number of performance metrics. Also, by emulating the network layer, our approach replicates a more realistic behavior, thus leading to more realistic results than BlockSim. Finally, even if it is not yet the case, our approach is aimed at making the plugging of any blockchain technology easier, as our code is structured following the six-layer model described in Figure 3.

One may wonder whether the number of nodes used in our experiments is not too small. This parameter is not critical in this study as the objective is not about showing how efficient a given blockchain is (e.g., from a scalability standpoint), but



(a) P2P network topology underlying the Testbed at day 1 of the experiment




(b) P2P network topology underlying the Testbed at day 2 of the experiment

Figure 14: Testbed’s network graph evolution over two days of experiment

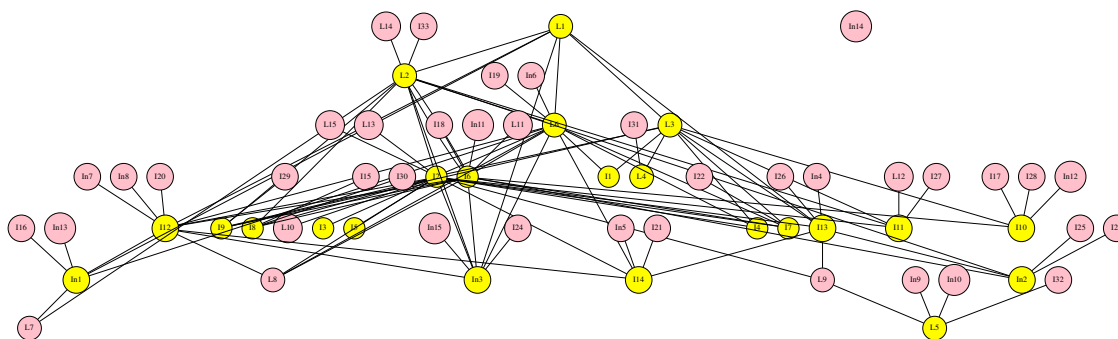
rather on showing that the proposed simulator provides more realistic evaluation results than an existing one (BlockSim in this case). As a consequence, we limited the experiment to be large enough to capture a realistic scenario, and short enough to be manageable (in terms of costs).

One may also wonder whether the size of the datasets is

large enough to draw relevant conclusions about the comparison analysis. In this regard, let us note that all the graphs given in Figures 8 to 17 only report the average values of all the measurements obtained on a daily basis for the corresponding metrics. Just to give an indication about the number of measurements that have been collected every day, around



(a) P2P network topology underlying BlockPerf at day 1 of the experiment



(b) P2P network topology underlying BlockPerf at day 2 of the experiment

Figure 15: BlockPerf’s network graph evolution over two days of experiment




Figure 16: BlockSim’s network graph remaining unchanged throughout the simulation run

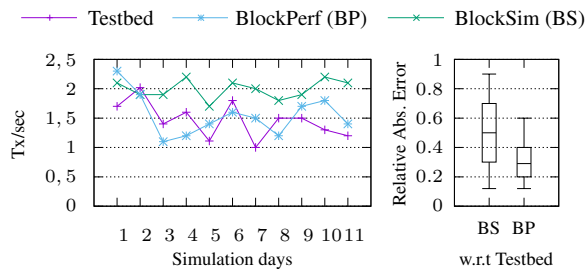


Figure 17: Throughput (o) & relative absolute error w.r.t Testbed

10000 measurements/day were collected regarding metrics (e) and (k_1), 200 measurements/day regarding (l_1)-(l_3), and about 50 regarding l_4 . Given this amount of measurements, we confidently state that the conclusions drawn from our comparison analysis are relevant (statistically speaking).

V. CONCLUSION, LIMITATIONS & PERSPECTIVES

A. CONCLUSION

Today, a large variety of blockchain technologies is expanding in order to fulfill technical and non-technical needs and requirements. Within this context, determining and, most importantly, evaluating the characteristics/performance of a given blockchain platform is crucial for system designers before deploying it. In this respect, several blockchain simulators have been proposed in the literature over the past few years. Still, they are often limited in several respects (lack of extensibility, failure in covering all aspects/metrics underpinning a blockchain system, etc.). In this paper, the six-layer model introduced by [15] (Application, Contract, Incentive, Consensus, Node/Data, Network) is considered, against which 15 performance metrics have been mapped.

To overcome the limitations of state-of-the-art blockchain simulators, a new one called BlockPerf is proposed in this paper, an extension to the existing BlockSim simulator. BlockPerf tries to cover as much as possible the layers mentioned above along with its metrics. A comparative analysis with BlockSim is carried, which has demonstrated that BlockPerf provides more realistic results than BlockSim (e.g., at the Node/Data and Network layers) respectively improved by 39% and 55% in average. However, several limitations remain to be addressed in the future, as discussed in the next section. The comparison analysis conducted in this paper is built upon a benchmarking bitcoin scenario.

B. LIMITATIONS & PERSPECTIVES

Several limitations to BlockPerf are to be addressed for it to be a more exhaustive simulator. First, the deployment cost of deploying nodes within BlockPerf on different geographical locations is not a simple and straightforward task, as it requires careful planning and deployment of nodes, connecting them with the main interface within BlockPerf. However, such a layer is, in our opinion, crucial to obtain simulation results closer to reality.

Second, BlockPerf, in its current form, only supports Bitcoin-related experiments, and some efforts still remain to be done to allow for simulating different blockchain platforms. For example, a blockchain-based system such as the IOTA, which uses graph-based transaction chains, would require more extensive tuning of BlockPerf to replicate the effects of its real-world deployment.

Third, the *Contract layer* still remains to be developed within BlockPerf. The ability of deploying contracts has been a key element and argument within all the blockchain technologies that emerged over the last years, and being able to analyse them within the simulation environment would likely provide further insights on its usage. However, covering/integrating such a layer in a simulator is particularly challenging because each smart contract execution happens within a virtual computing environment (e.g., EVM in Ethereum), which is called by every node execute the instructions contained by that contract. This environment itself relies on other layers for execution, including the execution of certain transactions or changing the state of nodes, which makes it difficult to obtain realistic execution effects on the overall system, as discussed in [?]. Given this complexity, we believe that the adoption of a hybrid approach that uses the virtual computing environment as an emulation layer – *in a similar manner as done in BlockPerf with the network layer* – is an efficient way to cover/integrate the smart contract layer in a simulator.

ACKNOWLEDGEMENT

Supported by the Luxembourg National Research Fund Stability 4.0 BRIDGES19/IS/13706587.

References

- [1] A. A. Monrat, O. Schelén, and K. Andersson, "A survey of blockchain from the perspectives of applications, challenges, and opportunities," *IEEE Access*, vol. 7, pp. 117 134–117 151, 2019.
- [2] J. Abou Jaoude and R. G. Saade, "Blockchain applications—usage in different domains," *IEEE Access*, vol. 7, pp. 45 360–45 381, 2019.
- [3] J. Al-Jaroodi and N. Mohamed, "Blockchain in industries: A survey," *IEEE Access*, vol. 7, pp. 36 500–36 515, 2019.
- [4] U. Bodkhe, S. Tanwar, K. Parekh, P. Khanpara, S. Tyagi, N. Kumar, and M. Alazab, "Blockchain for industry 4.0: A comprehensive review," *IEEE Access*, vol. 8, pp. 79 764–79 800, 2020.
- [5] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey," *International Journal of Web and Grid Services*, vol. 14, no. 4, pp. 352–375, 2018.
- [6] C. Fan, S. Ghaemi, H. Khazaei, and P. Musilek, "Performance evaluation of blockchain systems: A systematic survey," *IEEE Access*, vol. 8, pp. 126 927–126 950, 2020.
- [7] S. Maranhão, J.-. Seigneur, and R. Hu, "Towards a standard to assess blockchain & other dlt platforms," in *ITU*, 2019.
- [8] F. Gräbe, N. Kannengiesser, S. Lins, and A. Sunyaev, "Do not be fooled: Toward a holistic comparison of distributed ledger technology designs," in *Proceedings of the 53rd Hawaii International Conference on System Sciences*, 2020.
- [9] J. Polge, J. Robert, and Y. Le Traon, "Permissioned blockchain frameworks in the industry: A comparison," *ICT Express*, 2020.
- [10] W.-T. Tsai, R. Wang, S. Liu, E. Deng, and D. Yang, "Compass: A data-driven blockchain evaluation framework," in *IEEE International Conference on Service Oriented Systems Engineering*, 2020, pp. 17–30.
- [11] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

- [12] "A next-generation smart contract and decentralized application platform," <https://github.com/ethereum/wiki/wiki/White-Paper>, last accessed 13/04/2021.
- [13] 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 *Proceedings of the Thirteenth EuroSys Conference*. ACM, 2018, p. 30.
- [14] S. Smetanin, A. Ometov, M. Komarov, P. Masek, and Y. Koucheryavy, "Blockchain evaluation approaches: State-of-the-art and future perspective," *Sensors*, vol. 20, no. 12, p. 3358, 2020.
- [15] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "Block-bench: A framework for analyzing private blockchains," in *Proceedings of the 2017 ACM International Conference on Management of Data*, 2017, pp. 1085–1100.
- [16] C. Faria and M. Correia, "Blocksim: Blockchain simulator," in *2019 IEEE International Conference on Blockchain (Blockchain)*. IEEE, 2019, pp. 439–446.
- [17] A. Montresor and M. Jelasity, "Peersim: A scalable p2p simulator," in *2009 IEEE Ninth International Conference on Peer-to-Peer Computing*. IEEE, 2009, pp. 99–100.
- [18] A. Miller and R. Jansen, "Shadow-bitcoin: Scalable simulation via direct execution of multi-threaded applications," in *8th Workshop on Cyber Security Experimentation and Test (CSET'15)*, 2015.
- [19] L. Stoykov, K. Zhang, and H.-A. Jacobsen, "Vibes: fast blockchain simulations for large-scale peer-to-peer networks," in *Proceedings of the 18th ACM/IFIP/USENIX Middleware Conference: Posters and Demos*, 2017, pp. 19–20.
- [20] R. Paulavičius, S. Grigaitis, and E. Filatovas, "A systematic review and empirical analysis of blockchain simulators," *IEEE access*, vol. 9, pp. 38 010–38 028, 2021.
- [21] H. Huang, W. Kong, S. Zhou, Z. Zheng, and S. Guo, "A survey of state-of-the-art on blockchains: Theories, modelings, and tools," *ACM Computing Surveys (CSUR)*, vol. 54, no. 2, pp. 1–42, 2021.
- [22] "Technical paper hstp.dlt-rtf: Distributed ledger technology: Regulatory framework," https://www.itu.int/dms_pub/itu-t/opb/tut/T-TUT-DLT-2019-RF-PDF-E.pdf, accessed: 2021-04-12.
- [23] "Blockchain and distributed ledger technologies - Reference architecture," International Organization for Standardization, Geneva, CH, Standard (Under Development), Apr. 2020.
- [24] A. A. Monrat, O. Schelén, and K. Andersson, "A survey of blockchain from the perspectives of applications, challenges, and opportunities," *IEEE Access*, vol. 7, pp. 117 134–117 151, 2019.
- [25] L. M. Bach, B. Mihaljevic, and M. Zagar, "Comparative analysis of blockchain consensus algorithms," in *2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, 2018, pp. 1545–1550.
- [26] D. Mingxiao, M. Xiaofeng, Z. Zhe, W. Xiangwei, and C. Qijun, "A review on consensus algorithm of blockchain," in *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2017, pp. 2567–2572.
- [27] W. Wang, D. T. Hoang, P. Hu, Z. Xiong, D. Niyato, P. Wang, Y. Wen, and D. I. Kim, "A survey on consensus mechanisms and mining strategy management in blockchain networks," *IEEE Access*, vol. 7, pp. 22 328–22 370, 2019.
- [28] L. Zheng, X. Helu, M. Li, and H. Lu, "Automatic discovery mechanism of blockchain nodes based on the kademia algorithm," in *International Conference on Artificial Intelligence and Security*. Springer, 2019, pp. 605–616.
- [29] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of modern physics*, vol. 74, no. 1, p. 47, 2002.
- [30] M. A. Javarone and C. S. Wright, "From bitcoin to bitcoin cash: a network analysis," in *Proceedings of the 1st Workshop on Cryptocurrencies and Blockchains for Distributed Systems*, 2018, pp. 77–81.
- [31] "Dslf: Distributed ledger simulation framework," <https://github.com/i13-msrg/dslf>, accessed: 2021-01-25.
- [32] B. Wang, S. Chen, L. Yao, B. Liu, X. Xu, and L. Zhu, "A simulation approach for studying behavior and quality of blockchain networks," in *International Conference on Blockchain*. Springer, 2018, pp. 18–31.
- [33] N. Agrawal, O. Biçer, A. Kıpçü et al., "Blocksim-net: A network based blockchain simulator," *arXiv preprint arXiv:2011.03241*, 2020.
- [34] "evibes plasma simulator," <https://github.com/i13-msrg/evibes-plasma>, accessed: 2021-01-25.
- [35] A. Deshpande, P. Nasirifard, and H.-A. Jacobsen, "evibes: Configurable and interactive ethereum blockchain simulation framework," in *Proceedings of the 19th International Middleware Conference (Posters)*, 2018, pp. 11–12.
- [36] M. Alharby and A. van Moorsel, "Blocksim: A simulation framework for blockchain systems," *SIGMETRICS Perform. Eval. Rev.*, vol. 46, no. 3, pp. 135–138, Jan. 2019. [Online]. Available: <https://doi.org/10.1145/3308897.3308956>
- [37] S. Goswami, "Scalability analysis of blockchains through blockchain simulation," 2017.
- [38] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, "On the security and performance of proof of work blockchains," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 3–16.
- [39] S. M. Fattahi, A. Makanju, and A. M. Fard, "Simba: An efficient simulator for blockchain applications," in *2020 50th Annual IEEE-IFIP International Conference on Dependable Systems and Networks-Supplemental Volume (DSN-S)*. IEEE, 2020, pp. 51–52.
- [40] Z. H. Chin, T. T. V. Yap, and I. K. Tan, "Simulating difficulty adjustment in blockchain with simblock," in *Proceedings of the 2nd ACM International Symposium on Blockchain and Secure Critical Infrastructure*, 2020, pp. 192–197.
- [41] Y. Aoki, K. Otsuki, T. Kaneko, R. Banno, and K. Shudo, "Simblock: A blockchain network simulator," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2019, pp. 325–329.
- [42] R. Banno and K. Shudo, "Simulating a blockchain network with simblock," in *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, 2019, pp. 3–4.
- [43] M. R. A. Lathif, P. Nasirifard, and H.-A. Jacobsen, "Cidds: A configurable and distributed dag-based distributed ledger simulation framework," in *Proceedings of the 19th International Middleware Conference (Posters)*, 2018, pp. 7–8.
- [44] "Hive simulation tool," <https://github.com/ethereum/hive>, accessed: 2021-01-25.
- [45] E. Androulaki, G. O. Karame, M. Roeschlin, T. Scherer, and S. Capkun, "Evaluating user privacy in bitcoin," in *International Conference on Financial Cryptography and Data Security*. Springer, 2013, pp. 34–51.
- [46] G. Leduc, S. Kubler, and J.-P. Georges, "Innovative blockchain-based farming marketplace and smart contract performance evaluation," *Journal of Cleaner Production*, vol. 306, p. 127055, 2021.



JULIEN POLGE is a Ph.D. student at the Interdisciplinary Centre for Security, Reliability and Trust (SnT) in the University of Luxembourg. He received his M.Sc. in Complex Systems Engineering from the University of Lorraine (France) in 2017. He is part of the SerVal group (SEcurity, Reasoning and VALidation) headed by Pr. Yves Le Traon. His research work is conducted in collaboration with the industry and focuses on Industrial Internet of Things, Machine Learning and Distributed Ledger Technologies.



SANKALP GHATPANDE is R&D Engineer with position of Research Associate at the Interdisciplinary Centre for Security, Reliability and Trust (SnT) in the University of Luxembourg. He received his M.S degree from University of Luxembourg in 2016. Since 2016, he has worked in industry and academic R&D projects, national and international, as Engineer focused on Blockchain, IoT and Cryptography. He has also been active in multiple volunteering projects involving implementations of specific software(s), code-analysis and multiple open source projects. Recently, he has been involved in data analysis and machine learning project within the financial domain with industrial partner.



SYLVAIN KUBLER is Associate Professor at the Research Center for Automatic Control (CRAN) in Université de Lorraine, and before Research Associate at the Interdisciplinary Centre for Security, Reliability and Trust (SnT) in the University of Luxembourg (2015-2017) and Postdoctoral Researcher at Aalto University (2013-2015). He received his M.Sc. degree and Ph.D. degree in Computer Science and Engineering from the Université de Lorraine (France) respectively in 2009 and 2012. He was awarded the best Thesis in Automatic Control from the IFAC French Workgroup GdR MACS/Club EEA. He has broad expertise in Internet of Things, Semantic Web, Blockchain, and Multicriteria Decision Support Systems.



JÉRÉMY ROBERT is a Digital Technology Specialist at the Cebi Luxembourg S.A., an automotive and household appliances industry. He received his M.Sc. and Ph.D. degree in Computer Science and Engineering from the University of Lorraine (France), respectively, in 2009 and 2012. He has broad expertise in industrial and embedded networks since his PhD research focused on the use of switched Ethernet embedded in the future space launchers. Since 2015, his work was more about heterogeneous data communication challenges in the Industrial Internet of Things and the implementation of messaging services and high-level data formats. As the major research work was conducted in collaboration with the industry, he naturally joined Cebi Luxembourg S.A. for supporting the deployment of the Industry 4.0 project at the group level.



YVES LE TRAON is full professor of Computer Science at University of Luxembourg, in the domain of software engineering, with a focus on software testing, software security, and data-intensive systems. He is currently vice-director of the Interdisciplinary Centre for Security, Reliability and Trust (SnT center) and head of the SerVal group (SEcurity, Reasoning and VALidation), which is composed of around 25 researchers. He was head of the CSC Research Unit (Dept. of Computer Science) for the period of 2013-2016.

...