On the Bandwidth Consumption of Blockchains

Andrei Lebedev
The University of Sydney
andrei.lebedev@sydney.edu.au

Vincent Gramoli
The University of Sydney and Redbelly Network
vincent.gramoli@sydney.edu.au

***Abstract*—With the advent of blockchain technology, the number of proposals has boomed over the past decade. The network traffic imposed by these blockchain proposals increases the cost of hosting nodes. Unfortunately, as of today, we are not aware of any comparative study of the bandwidth consumption of blockchains.**

In this paper, we propose the first empirical comparison of blockchain bandwidth consumption. To this end, we measure the network traffic of blockchain network nodes of five blockchain protocols: Algorand, Aptos, Avalanche, Redbelly and Solana. We study the variation over time, differentiate the receiving and sending traffic and analyze how this traffic varies with the number of nodes and validators.

We conclude that the transport protocol is the main factor impacting the network traffic, segregating node roles helps reduce traffic and different blockchains are differently impacted by the network size.

***Index Terms*—Network, traffic, consensus, peer-to-peer.**

I. Introduction

With the advent of blockchain technology, the number of proposals has boomed over the past decade. There exist various layer-1 blockchains offering different guarantees and performing differently at large scale. Recent studies [1], [2] have shown that the performance and robustness of some blockchains is quite far from the promises claimed by their designers. The problem stems mainly from a misunderstanding of these blockchains' underlying networking protocols. The impact of these misunderstandings is so important that Avalanche [3] and Solana [4] were even shown to stop globally when some network messages get delayed [2].

The misunderstanding of these networking protocols also presents economic drawbacks. Modern high-throughput architectures have shifted validator economics from static capital expenditures to variable operational burdens, where egress consumption acts as a primary cost driver. For instance, the prohibitively high egress tariffs of standard cloud providers can render high-performance nodes economically non-viable, potentially incurring monthly costs exceeding US$8,000 for a single validator due to misaligned billing models [5]. This structural inefficiency necessitates a strategic pivot toward bare-metal providers with unmetered allowances to avoid the "hyperscaler egress trap" [6].

As of today, there are no studies comparing the network traffic of blockchains. Even though it is well known that some blockchains (e.g., Solana [4]) favor redundancy (e.g., through erasure coding) to cope with packet losses while others (e.g., Redbelly [7]) adopt different communication patterns between validator and non validator nodes, the bandwidth consumption of these blockchains remains unclear. Understanding the network traffic is, however, crucial to improve performance and robustness of layer-1 blockchains by reducing the network congestion to make them scale or to replicate the data that are key to their robustness.

In this paper, we propose the first empirical comparison of blockchain bandwidth consumption. To this end, we build upon the series of work around the Diablo performance benchmark [1] and the STABL fault tolerance benchmark [2] to measure the network traffic of blockchain network nodes in various situations: (i) while receiving and sending messages; (ii) before, during and after the network receives transactions and (iii) as the network size grows both in terms of nodes but also validator nodes.

Using this black-box approach, we deploy five different blockchain protocols, namely Algorand [8], Aptos [9], Avalanche [3], Redbelly [7], and Solana [4]. We make the following observations:

1) The dominant factors of network traffic are the transport protocol (polling vs WebSockets) and the block propagation strategy (full block download vs. hash comparison) more than the transaction size.
2) Segregating roles between different types of nodes helps reduce the network traffic by reducing network traffic between nodes of different types.
3) Solana network traffic depends on the network size, Algorand and Redbelly network traffic increases with the validator sets and Aptos and Avalanche network traffic increases with both the number of nodes and validators.

The paper is organized as follows. In Section II, we present the background. In Section III, we present the experimental settings and our methodology. In Section IV, we present the variety of bandwidth consumption of the five blockchain protocols. In Section V, we study the distribution of network traffic over different pairs of nodes. In Section VI, we compare the bandwidth consumption before, during and after reception of transactions. In Section VII, we study the impact of the number of nodes and validators on the bandwidth consumption. In Section VIII, we study the impact of the sending rate on the bandwidth consumption. We present the related work in Section IX and we conclude in Section X.

This research is supported under Australian Research Council Discovery Project funding scheme (project number 250101739) entitled "Fair Ordering of Decentralised Access to Resources".

II. Background and Blockchain Networks

In this section, we list the characteristics of each tested blockchain network protocol.

A. The Algorand network

Algorand [8] is a blockchain protocol that shuffles participants via cryptographic sortition to enhance security. More precisely, it uses Verifiable Random Functions (VRFs) to randomly select participants for specific roles in the consensus execution. Nodes communicate with a gossip-based protocol where each node validates each message before relaying it and sends it at most once to each other node [10]. To this end, each node maintains one TCP connection per node in its neighborhood, which offers WebSockets over HTTP.

B. The Avalanche network

Avalanche, based on the Snowflake consensus protocol [3], is a probabilistic blockchain that requires, by default, a proportion of the nodes that collectively own at least 80% of the total stake to be online. In Avalanche, nodes communicate over TCP and exploit throttling to limit their resource usage. More specifically, messages and connections are rate-limited [11] to cap the amount of CPU, disk, bandwidth, and message handling that other nodes consume. Avalanche uses a dynamic proposer selection algorithm to manage network load and block production. After each parent block, it pseudo-randomly selects an ordered list of potential proposers for the next block height, weighted by stake and using a seed derived from the parent block's height and chain ID. Each proposer is assigned a minimum delay based on their position in the list before they can sign and propose a block. If no proposer acts within the cumulative delay period, any active validator may propose.

C. The Aptos network

Aptos [9] is a leader-based blockchain that uses TCP and builds upon a variant of the *Practical Byzantine Fault Tolerant (PBFT)* consensus protocol [12] called AptosBFT inheriting its cubic communication complexity. It is well-known that leaders act as bottlenecks in leader-based consensus protocols [13], like AptosBFT, and that classic blockchains suffer from redundant dissemination of the same transactions first outside and then within blocks [14], [15]. To cope with these limitations, Aptos features the Quorum Store optimization of Narwhal [16] to decouple metadata ordering from payload dissemination. In addition, Quorum Store is designed for parallel execution by all validators. As outlined in the documentation [17], validators repeat the following steps in parallel: *(1)* Pull transactions from the mempool; *(2)* Arrange transactions into batches based on gas price and select an expiration time for each batch; *(3)* Broadcast batches to all other validators; *(4)* Persist received batches, sign their digests, and send back signatures; and *(5)* Collect signatures from more than $^{2n}/_{3}$ nodes to form a *proof-of-store*. It allows validators to asynchronously broadcast transactions, offloading the leader's network interface during the consensus protocol execution [18].

D. The Redbelly network

Redbelly Blockchain [7] is a scalable blockchain built on the Democratic Byzantine Fault Tolerant (DBFT) consensus algorithm [19] that is *leaderless* (i.e., non leader-based) and deterministic, and works in a partially synchronous environment. To enhance scalability further, Redbelly uses a collaborative approach, appending a superblock with as many valid proposed blocks as possible. This way the number of transactions per appended block can grow linearly with the number of nodes [7]. Redbelly's nodes communicate using TCP and features a Scalable variant of the EVM, called SEVM. It was shown to perform well under realistic dApps particularly in a large geo-distributed environment when compared to other modern blockchains [14].

E. The Solana network

Solana [4] is a leader-based blockchain that may fork. In order to determine whether a transaction is committed, Solana requires 30 additional blocks to be appended after the transaction's block. Nodes communicate over the QUIC network protocol [20] to exchange transactions. Nodes split blocks into chunks that they disseminate in a hierarchical structure, called Turbine [21], through UDP.

III. Experimental Settings

In this section, we explain how we deploy blockchain protocols and measure their bandwidth consumption.

a) Distributed system setup: Our experimental setup consists of a distributed system of 25 VMs running Ubuntu 24.04.1 LTS on a Proxmox cluster of physical servers, each equipped with 4x AMD Opteron 6378 16-core CPUs at 2.40 GHz, 256 GB of RAM, and 10 GbE NICs. Each experiment runs a blockchain protocol with 5 client VMs and up to $N = 20$ blockchain node VMs. We also refer to V as the number of validators such that $N, V \in \{5, 10, 15, 20\}$ and $V \leq N$. This 25-node setup is sufficient to reproduce trends observed in geo-distributed settings as was demonstrated recently [22].

For each blockchain, we select the following versions: Algorand v3.27.0, Aptos v1.25.1, Avalanche C-Chain v1.12.1, Redbelly v0.36.2 and Solana Agave v2.0.20. All experiments follow a fixed timeline: a 100 s "Setup" phase with no transactions, a 100 s "Workload" phase where transactions are sent, and a 100 s "Cooldown" phase with no transactions, allowing remaining transactions to commit. During the Workload phase, the 5 clients send transactions to the first 5 validators. The target load is distributed equally (e.g., 40 TPS each for a total of 200 TPS). Each transaction is sent to a single node, which is then queried for finality using block streaming or polling.

The resources of each VM, 4 vCPUs and 8 GB of memory, intentionally mimic the resources of a commodity computer run by an individual in a blockchain network. Note that this specification is lower than what some blockchains typically recommend, including Aptos [23], Avalanche [24] or Redbelly [25], however, strict hardware requirements on remote nodes remain hard to enforce and a unique configuration is necessary for our comparison.

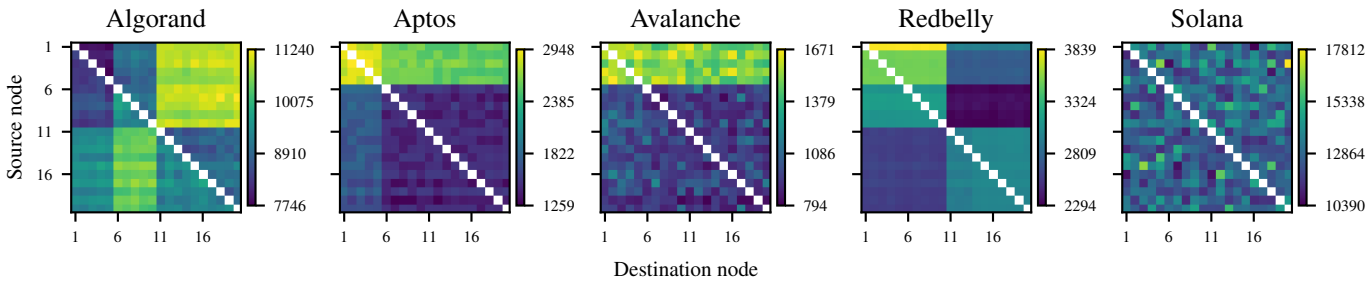

Fig. 1. Heatmaps $M_{i,j}$ of the bandwidth used in KiB between sending node n_i and receiving node n_j for each blockchain, 20 nodes, 20 validators.

b) Measuring peer-to-peer bandwidth: We implemented a fine-grained bandwidth monitoring system to capture traffic usage between blockchain nodes and clients. Our approach provides pairwise measurements of the traffic exchanged between each node in the network. The system utilizes STABL observer processes running on blockchain VMs and relies on the Linux iptables firewall infrastructure to perform non-intrusive packet accounting.

For each node under observation, we programmatically install a set of iptables rules. These rules create custom accounting chains that contain a specific rule for every other peer in the experiment. Each rule is configured to match packets based on their source (for incoming traffic) or destination (for outgoing traffic) IP address.

A monitoring script then periodically queries the byte counters associated with each of these per-peer rules and immediately resets them to zero. This process yields a time series where each data point represents the average transmission (TX) and reception (RX) rate over the preceding interval, allowing us to precisely analyze network behavior.

IV. Varying Consumption

To illustrate how blockchains consume bandwidth, we compare their bandwidth usage heatmaps and transaction size.

A. Bandwidth consumptions as heatmaps

Figure 1 depicts one heatmap per blockchain protocol for nodes that all act as validators where the color of cell $M_{i,j}$ represents the bandwidth consumed by node n_i when sending to node n_j. A warmer color (e.g., yellow) thus represents a higher bandwidth usage than a colder color (e.g., dark blue) while the white color indicates zero or negligible traffic. Note that nodes $n_1, \ldots, n_5$ are the nodes receiving transaction requests from the clients, while nodes $n_6, \ldots, n_{20}$ may receive messages from other nodes but not transactions directly from clients. With a maximum of 17,812 KiB, Solana consumes more bandwidth than the other tested blockchains, namely Algorand, Aptos, Avalanche and Redbelly. In particular, Algorand, the second most bandwidth consuming blockchain uses a maximum of 11,240 KiB.

The fact that Solana consumes 58% more than Algorand can be due to more metadata sent per transaction or higher duplication of the same information. This can be explained by Solana maximizing dissemination of information despite failures [26] with erasure coding. Its lack of fault tolerance [2] probably motivated this design decision. Erasure coding relies on sending additional data in order to reconstruct the relevant information in case of partial loss. Other blockchains consume less traffic likely because they do not use erasure coding.

B. Transaction sizes

In order to exclude other factors that could have invalidated our hypothesis that Solana uses more bandwidth due to erasure coding, we measured empirically the size of transactions sent by each blockchain. The transaction sizes of some distributed ledger (e.g., Corda) were already seen as excessive (e.g., 8 KiB), which could bloat bandwidth consumption [27].

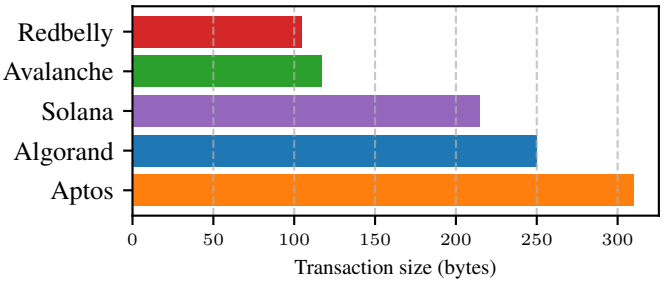

Fig. 2. Transaction size per blockchain.

Fig. 2 compares the transaction sizes of the different blockchains. It shows in particular that, even though Solana uses more bandwidth as discussed in Section IV, it does not use the largest transactions. Actually, Aptos' transaction size is 310 B or 44% larger than Solana's (215 B). Finally, note that even Algorand's transactions, with a size of 250 B, are larger than Solana's transactions. To conclude, Solana transaction size is not the main reason for higher bandwidth consumption.

V. Consumption Skewness over Routes

A high bandwidth usage does not necessarily induce a detrimental performance, especially when the bandwidth consumption is well-balanced over multiple routes. In fact, previous works have shown that balancing an amount of information that is quadratic in the number of network nodes over a quadratic number of routes of this network could be more efficient than concentrating a linear amount of information at a bottleneck [28]. It is thus crucial to understand how bandwidth consumption is balanced over the network.

A. Some nodes exchange more data than others

In the Aptos heatmap of Fig. 1, we can see that some nodes of the Aptos network seem to consume more bandwidth than other nodes of the same network, and traffic is heavily unbalanced across nodes. Five nodes, the ones that receive the transactions sent by the clients, send more messages than all the other nodes as indicated by the top five rows in light colors. They also send more messages to themselves than to the rest of the nodes, as indicated by the yellow 5-by-5 submatrix of the top left corner of the heatmap. This is due to its Quorum Store optimization [17] that puts more load on the receivers of transactions than on the rest of the network by requiring them to collect signatures. More precisely, these validator nodes have to sign transaction batch digests and collect the produced signatures from a quorum of blockchain nodes to form a "proof-of-store". This signature collection puts more bandwidth pressure on the nodes responsible for signing.

In the Solana heatmap of Fig. 1, the color shows that there is no clear per-node distinction as no column or row stands out. As a result, the bandwidth consumption of Solana appears generally more balanced than the one of Aptos. It is interesting to note that Solana, which consumes a lot of bandwidth as we showed in Section IV, manages to balance the load effectively.

Finally, the Algorand heatmap of Fig. 1 shows some imbalance in that the second half of the nodes $n_{11}, \ldots, n_{20}$ receive and send more messages than others.

B. Some nodes send more than they receive

Another interesting dimension to consider is the level of unbalance between sending and receiving traffic. The nodes receiving transactions could either consume more bandwidth by propagating the information or, instead, consume less than the nodes agreeing on which block to append. For example, the Avalanche heatmap of Fig. 1 shows that the nodes of Avalanche that receive transactions generate more traffic than what they receive.

The Avalanche heatmap of Fig. 1 shows the five top rows in lighter color than the five left columns. This indicates that the five Avalanche nodes that receive transactions send more data than they receive. By contrast with Aptos, they do not need to send more data to themselves than to the rest of the network. This is explained by the fact that they have to propagate the transactions they receive to the rest of the network without needing to collect signatures from each other.

The Redbelly heatmap of Fig. 1 shows one particular node sending more messages than others. This is explained by having a weak coordinator that sends a particular message in the first round of the DBFT binary consensus protocol [19]. Note that this coordinator is weaker than a leader in that the consensus terminates even when it is faulty. Generally, the first five nodes send more messages than others because they forward the received transactions. Finally, the remaining differences could be due to the node placement and the reordering of messages, requiring some nodes to request the batch of proposed transactions from other nodes.

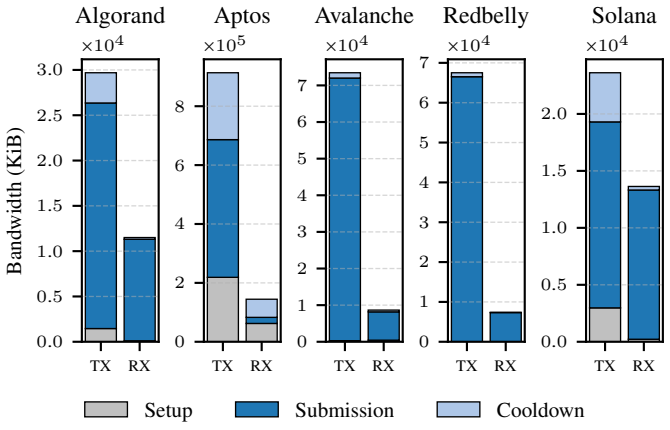

Fig. 3. Client-side network traffic (TX and RX) stacked by experimental phase. Y-axis scales differ across subplots to accommodate variations in magnitude.

C. The dissemination reliability of Solana requires more traffic

Bandwidth consumption can have some advantages, for example, when redundant information copes with packet losses. The trade-off between bandwidth efficiency and dissemination reliability is particularly apparent in Solana. Previous research has highlighted Solana's ability to tolerate packet losses better than other blockchains [2]. Our experiments quantify the cost of this dissemination reliability.

The Solana heatmap of Fig. 1 shows the bandwidth consumption is well balanced and very high. Unlike other protocols that attempt to minimize redundant transmissions (resulting in the dark blue or empty regions seen in Algorand or Redbelly), Solana appears to utilize a "flood" or highly redundant propagation mechanism, likely related to its Turbine block propagation protocol and erasure coding schemes. While this results in significantly higher total bandwidth consumption, it ensures that data is recoverable and available to all nodes, even in the presence of network failures.

VI. Consumption Skewness over Time

To understand when bandwidth is consumed, we differentiate the bandwidth used at the three stages, Setup, Submission, and Cooldown, of our experiments mentioned in Section III.

Fig. 3 differentiates the Transmitted (TX) data, which are sent from nodes to clients, from the Received (RX) data, which are sent from clients to nodes, when 5 clients submit 19,995 transactions (3,999 each) to 5 nodes.

A. Impact of communication protocols on idle traffic

A distinct disparity is visible in the overhead traffic during the non-submission phases (Setup and Cooldown). Aptos exhibits significantly higher bandwidth consumption during these idle periods compared to other blockchains. For instance, during Phase 0 (Setup), Aptos nodes transmitted 218,856 KiB of data, whereas Avalanche and Redbelly transmitted only 298 KiB and 11 KiB, respectively.

This massive overhead in Aptos is attributed to its client communication architecture. While Avalanche, Redbelly, and

Solana utilize WebSocket streaming to push updates to clients efficiently, Aptos clients rely on polling. Consequently, Aptos clients repeatedly request data even when blocks are empty, resulting in substantial bandwidth usage (≈227,000 KiB in Phase 2) despite the absence of new transaction submissions.

B. Block verification and data efficiency

During the Submission phase (Phase 1), a clear divergence in data efficiency emerges between Solana and the chains compatible with the Ethereum Web3 WebSocket API (Avalanche, Redbelly) that use the same methods to send and listen to blocks as Ethereum.

a) Full block transmission: These protocols exhibit a high TX-to-RX traffic ratio. For example, Redbelly received 7,314 KiB of transaction data (RX) but transmitted 66,511 KiB (TX) back to clients. This amplification occurs because these protocols require nodes to broadcast the entire block (containing transaction bodies and metadata) to every client for verification. Although a client only needs transaction hashes to confirm finality, it must download the full block payload.

b) Signature-based verification: Solana demonstrates a more balanced traffic profile during submission (RX: 13,099 KiB vs. TX: 16,325 KiB). Despite Solana having a relatively large raw transaction size (215 B) compared to Redbelly (105 B) or Avalanche (117 B), its outgoing traffic is significantly lower. This efficiency stems from Solana's verification mechanism: clients requesting block confirmation do not need to download the full block body. Instead, they request only the hashes or signatures of committed transactions and compare them against their locally stored payloads. This selective data retrieval significantly reduces the egress bandwidth required by the nodes.

In summary, while the raw transaction sizes (ranging from 105 B on Redbelly to 310 B on Aptos, as displayed in Fig. 2) play a role in bandwidth usage, the dominant factors influencing network traffic are the choice of transport protocol (polling vs. WebSockets) and the block verification strategy (full block download vs. hash comparison).

VII. Consumption Skewness over Scale

We now measure the increase in bandwidth consumption, denoted S_b, as a function of the number N of blockchain nodes and the number V of validators in blockchain b. Interestingly, we identify that the bandwidth of Aptos, Avalanche and Solana increases with N and the bandwidth of Algorand, Aptos, Avalanche and Redbelly increases with V.

A. Large consumption variations depending on network scale

Fig. 4 depicts a comprehensive grid of heatmaps for every blockchain across different combinations of N nodes and V validators. As in Section IV-A, the bandwidth consumed by node n_i sending to node n_j is represented by the color in cell $M_{i,j}$ of the heatmap where a brighter color means higher.

The raw data reveals significant disparities in data exchange across the protocols. In a fully interconnected small network (5 nodes, all validators), the difference in per-link bandwidth is striking: while Redbelly and Avalanche maintain a lean footprint with approximately 1,700 KiB and 2,800 KiB transmitted per pair respectively, Algorand consumes considerably more, averaging around 6,500 KiB per link. Solana, however, operates at a completely different order of magnitude, with cells in the 5-node configuration consistently showing over 42,000 KiB of traffic—nearly 25 times the bandwidth usage of Redbelly for the same workload.

Furthermore, the data highlights a clear hierarchy of network load based on node roles. In mixed configurations (e.g., 20 nodes with 5 validators), the bandwidth intensity within the validator group (the top-left 5×5 sub-matrix) is drastically higher than the traffic involving non-validator nodes. For instance, in the 10-5 setup, for Aptos and Algorand, the traffic between two validators remains in the thousands of KiB (e.g., ≈3,600 KiB for Aptos), whereas traffic originating from non-validator nodes often drops to the low hundreds (e.g., ≈200 KiB), illustrating a highly centralized bandwidth burden on the consensus committee.

B. Validators communicate more with themselves

A distinct segmentation of the network topology is visible in most protocols. For Algorand, Aptos, Avalanche, and Redbelly, the top-left $V \times V$ sub-matrix is consistently dense and bright, confirming that validators communicate heavily among themselves to achieve consensus. However, the behavior regarding non-validator nodes varies significantly.

As observed, Solana is the unique outlier. It is the only blockchain where the non-validator nodes communicate directly with other non-validator nodes (the bottom-right quadrant of the heatmaps). For instance, in the 10-5 configuration, the heatmap is uniformly populated, indicating a full mesh topology where every node, regardless of its role, exchanges data with every other node.

In contrast, blockchains like Redbelly and Avalanche show a clear separation. While validators send data to non-validators (top-right quadrant) and receive data from them (bottom-left quadrant), non-validator nodes do not communicate with each other. This is evident in the 10-5 matrices for both systems, where the bottom-right 5×5 sub-matrix is largely empty.

Finally, Algorand and Aptos exhibit the strictest separation in certain configurations. Non-validator nodes act almost exclusively as passive receivers or pull-based clients. In Algorand's 10-5 configuration, while validators send data to non-validators (rows 0-4 to columns 5-9), the traffic from non-validators back to validators is negligible, and traffic between non-validators is non-existent.

C. Aptos pattern change with the number of validators

Fig. 4 reveals a dynamic dissemination strategy in Aptos that depends on the ratio of validators to the total network size. In configurations where the number of validators is small relative to the network size, validators appear to broadcast to all non-validators. For example, in the 10-5 configuration and 20-5 configuration, the top-right quadrants are dense, showing

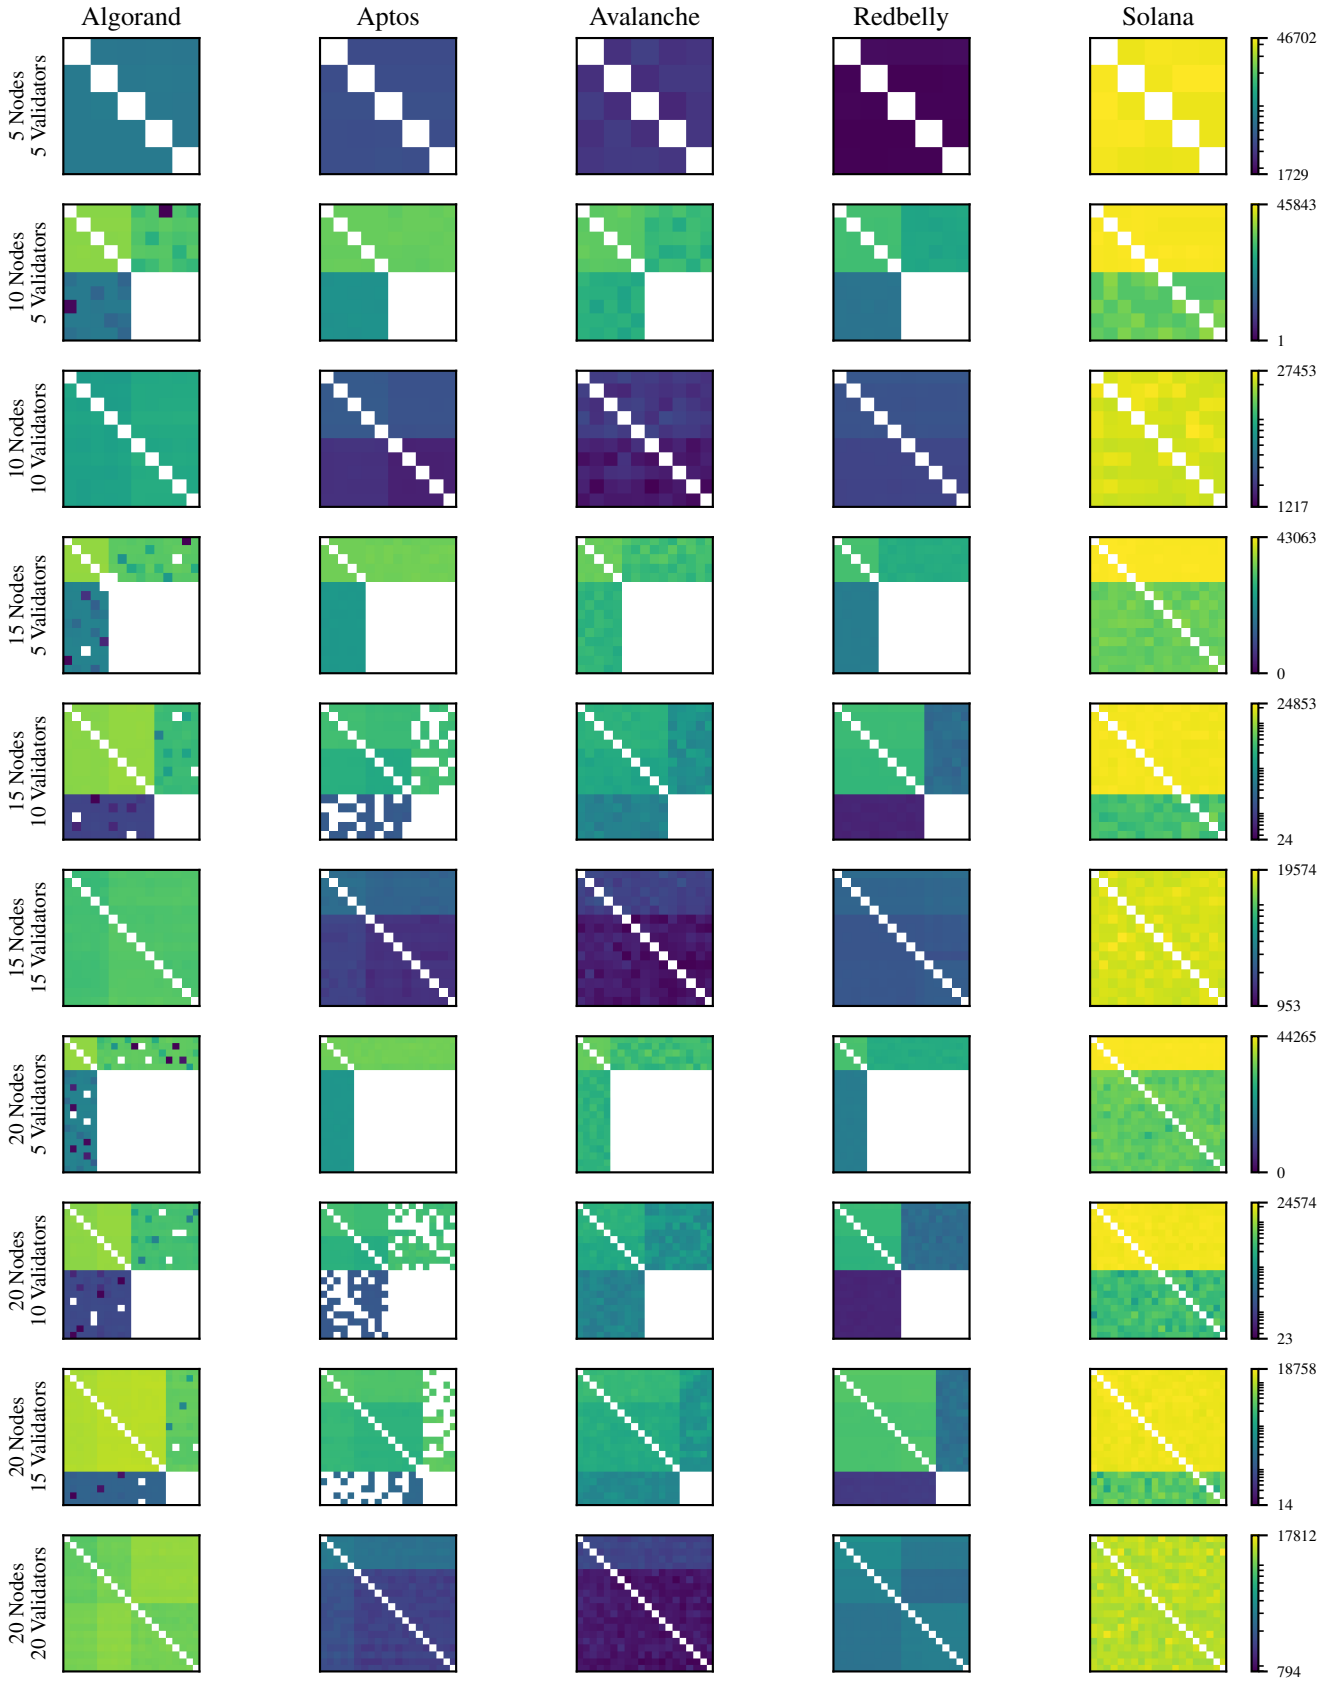

Fig. 4. Heatmaps $M_{i,j}$ of the bandwidth used in KiB between sending node n_i and receiving node n_j for each blockchain.

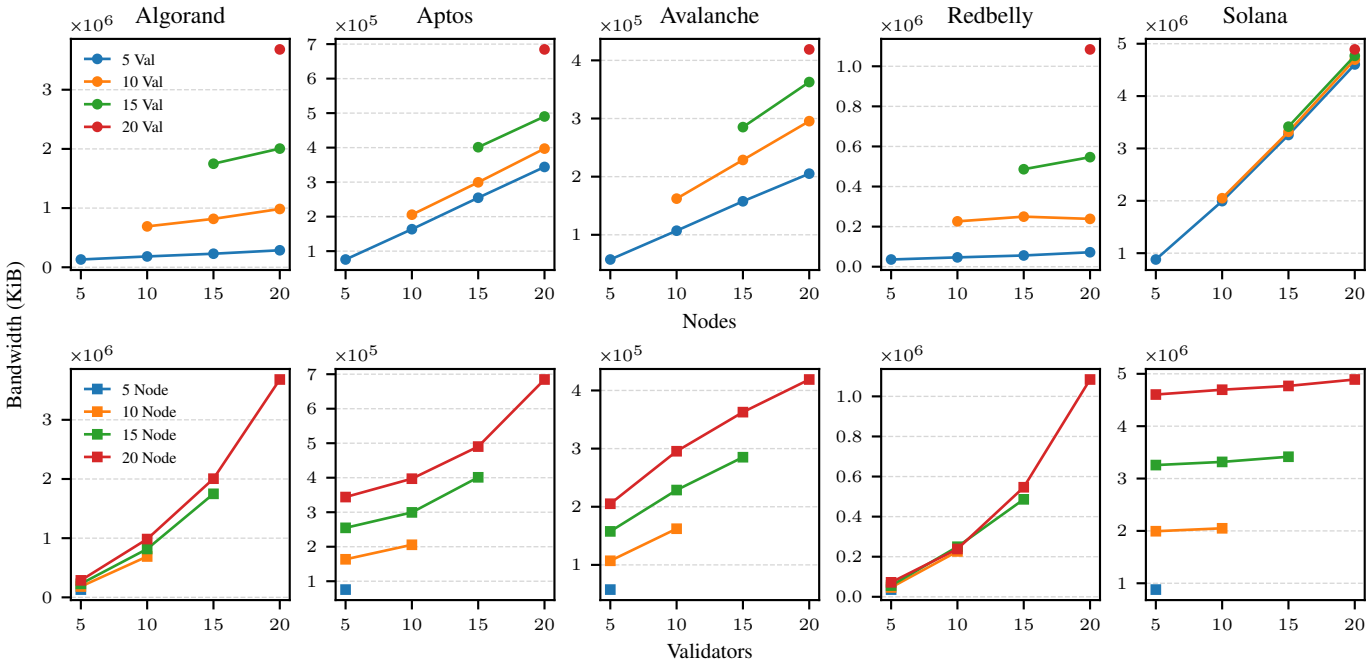

Fig. 5. Total bandwidth consumption (KiB) for blockchains under test. The top row plots bandwidth against the number of nodes (grouped by validator count), while the bottom row plots bandwidth against the number of validators (grouped by node count).

that the 5 validators are sending data to all 5 (or 15) non-validator nodes.

However, as the number of validators increases, Aptos shifts strategy to reduce bandwidth overhead. In the 15-node, 10-validator configuration, the top-right quadrant becomes sparse. Specific validators only communicate with specific non-validator nodes rather than broadcasting to the entire set. This suggests a sharding or randomized gossip approach to dissemination when the validator set is large, likely to prevent bandwidth saturation on individual nodes.

D. Blockchains vary with regards to bandwidth scalability

Fig. 5 depicts the trends of bandwidth consumption depending on the scale, where each figure provides: line plots showing S_b as a function of the total number N of nodes, and line plots showing S_b as a function of the number V of validators. Here, S_b represents the sum of bandwidth measured in KiB for blockchain b during the experiment duration.

We can clearly see from Fig. 5 that S_{Solana} bandwidth consumption increases primarily with the number of nodes, regardless of the number of validators. For instance, with 5 validators, increasing the total node count from 5 to 20 causes bandwidth to surge from 879,268 KiB to 4,604,375 KiB—a five-fold increase. However, keeping the node count fixed at 20 and increasing validators from 5 to 20 results in a negligible increase from 4,604,375 KiB to 4,891,524 KiB. This confirms that Solana requires all nodes to communicate with each other, creating a high-bandwidth mesh topology dependent on the network size rather than the validator set size.

In contrast, the bandwidth consumed $S_{b'}$ for blockchains $b' \in \{Algorand, Redbelly\}$ depends almost entirely on the number of validators. For Algorand, with 5 validators, increasing nodes from 10 to 20 only increases usage from 182,790 KiB to 287,032 KiB. However, increasing validators has a significant impact: with 20 nodes, moving from 5 to 20 validators causes usage to surge from 287,032 KiB to 3,681,191 KiB—an increase of over 12×. Redbelly exhibits an even higher ratio, jumping from 72,098 KiB (20 nodes, 5 validators) to 1,084,185 KiB (20 nodes, 20 validators).

Finally, for blockchains $b'' \in \{Aptos, Avalanche\}$ the bandwidth consumed $S_{b''}$ grows with both V and N. Note that the bandwidth increase appears superlinear with the number of validators in Aptos but sublinear in Avalanche. In Aptos, with 20 nodes, the bandwidth grows moderately between 5 and 15 validators (343,968 KiB to 490,237 KiB) but jumps sharply when reaching 20 validators (684,649 KiB). Conversely, Avalanche shows diminishing returns in bandwidth growth as validators are added to a 20-node network, rising from 205,154 KiB (5 validators) to 418,852 KiB (20 validators), indicating a more consistent propagation overhead.

Based on these trends, we can classify blockchains in three categories:

1) **Node-dependent (Solana).** Bandwidth consumption increases with the total network size (N), indicating a topology where all nodes propagate data heavily.
2) **Validator-dependent (Algorand and Redbelly).** Bandwidth consumption increases primarily with the size of the consensus committee (V). Non-validator nodes are passive consumers.
3) **Hybrid (Avalanche and Aptos).** Bandwidth consumption increases with both N and V, suggesting a topology where non-validator nodes participate in propagation.

VIII. Consumption Variation with TPS

The bandwidth consumption could simply be a consequence of the rate at which our experiments send transactions. Below, we vary this sending rate to observe how it impacts bandwidth consumption. To this end, we conducted a set of experiments where the total workload size was kept constant at approximately 20,000 transactions. We varied the target TPS from 100 to 500, inversely adjusting the duration (from 200 s down to 40 s) to maintain the fixed transaction count. Fig. 6 presents these results from the perspective of total absolute bandwidth.

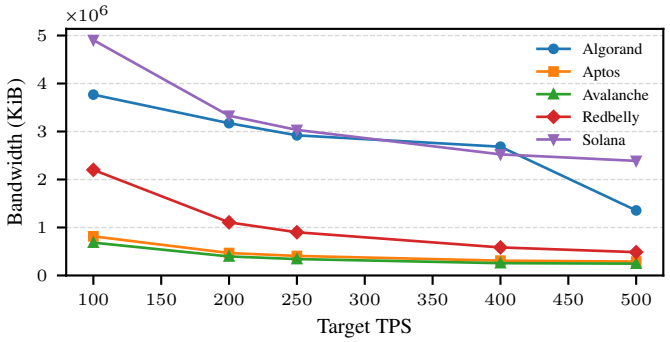

Fig. 6. Total bandwidth (KiB) vs. Target TPS for blockchains under test under varying load rates. The workload was fixed at 20,000 total transactions per run, with experiment duration scaling inversely to TPS (200 s down to 40 s).

Fig. 6 plots the total absolute bandwidth consumed by each blockchain. A consistent trend is visible across all five protocols: total bandwidth consumption decreases as TPS increases. Since the number of transactions is fixed, this decrease indicates that a significant portion of bandwidth usage is time-dependent rather than transaction-dependent. At lower TPS (longer duration), "background" traffic, such as heartbeats, empty block proposals, and consensus maintenance, accumulates, inflating the total footprint.

Solana consistently consumes the highest absolute bandwidth, ranging from approximately 4.7 GiB (4,905,820 KiB) at 100 TPS down to 2.3 GiB (2,387,814 KiB) at 500 TPS. Algorand follows as the second most bandwidth-intensive chain (3.6 GiB to 1.3 GiB), while Avalanche and Aptos remain the most efficient, with Aptos consuming as little as 283 MiB (289,834 KiB) at 500 TPS. Notably, in terms of absolute bandwidth, Algorand consistently consumes more than Redbelly across all data points (e.g., at 100 TPS: Algorand ≈3.6 GiB vs. Redbelly ≈2.1 GiB).

IX. Related Work

While existing literature extensively covers execution metrics or theoretical network propagation, empirical, multi-protocol bandwidth comparisons remain overlooked. Table I provides a synthesis of the related work.

a) Blockchain benchmarking: Frameworks such as Blockbench [29] and Hyperledger Caliper [30] standardize the evaluation of throughput, latency, and fault tolerance, while Gromit [34] addresses ad-hoc testing limitations. However, these frameworks prioritize execution capacity (TPS) and treat the network layer as secondary. Our work builds upon Diablo [1] and Stabl [2] to isolate bandwidth consumption, a critical metric overlooked by execution-focused benchmarks.

TABLE I
Comparison of related work and our contribution

Work	Primary Focus	Empirical BW Measurement	Multi-Protocol Comparison
[29], [30]	Execution Capacity (TPS/Latency)	No	Yes
[31], [32]	Network & Block Propagation	Partial	No (Single Chain)
[1], [2]	Execution & Fault Tolerance	No	Yes
[28], [33]	Topology & Comm. Complexity	Partial	Yes
This Paper	**Bandwidth Consumption**	**Yes**	**Yes**

b) Network traffic and block propagation: Early studies analyzed propagation delays in Bitcoin [31], leading to bandwidth optimization protocols like Graphene [32], Compact Blocks [35], and the BloXroute [36] Layer-0 CDN. While these works propose techniques to minimize traffic, our paper provides a comparative measurement of modern Layer-1 protocols. We reveal that contemporary systems like Solana often prioritize data redundancy over the bandwidth efficiency emphasized in earlier Bitcoin and Ethereum research.

c) Communication complexity vs. empirical reality: While theoretical literature favors linear communication complexity ($O(n)$) protocols, such as HotStuff [37], over quadratic BFT implementations [12], theoretical bounds often fail to predict real-world behavior. By contrast, empirical evidence demonstrates that quadratic complexity can achieve significantly better performance when well balanced across a quadratic number of routes of a WAN [28]. Furthermore, dense network topologies have been shown to improve performance under load [33], a finding correlated with higher energy usage [38]. Our empirical results confirm this divergence between redundancy-heavy architectures like Solana and the lean traffic profiles of Redbelly.

X. Conclusion

We compared the bandwidth consumption of blockchains. Our results show that Solana consumes significant bandwidth due to redundancy, while Aptos suffers from high idle overhead. Crucially, Solana, Aptos and Avalanche network traffic increases with network size, whereas that of Algorand and Redbelly grows with the validator count.

As a result, we offer three recommendations: (i) a push-based WebSockets design over a polling one to reduce idle overhead; (ii) a hash-based or signature-based selective data retrieval rather than full-block payload transmission for verification; and (iii) to balance erasure-coding redundancy against the risk of rapid network congestion at scale.

As future work, we plan to extend our experimental environments to evaluate the impact of highly dynamic networks.

REFERENCES

[1] V. Gramoli, R. Guerraoui, A. Lebedev, C. Natoli, and G. Voron, "DIABLO: A benchmark suite for blockchains," in *Proceedings of the Eighteenth European Conference on Computer Systems (EuroSys)*, G. A. D. Luna, L. Querzoni, A. Fedorova, and D. Narayanan, Eds. ACM, 2023, pp. 540–556.

[2] V. Gramoli, R. Guerraoui, A. Lebedev, and G. Voron, "STABL: The sensitivity of blockchains to failures," in *Proceedings of the 26th International Middleware Conference*, ser. Middleware '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 202–214. [Online]. Available: https://doi.org/10.1145/3721462.3730952

[3] Team Rocket, M. Yin, K. Sekniqi, R. van Renesse, and E. G. Sirer, "Scalable and probabilistic leaderless BFT consensus through metastability," arXiv, Tech. Rep., Aug. 2020. [Online]. Available: http://arxiv.org/abs/1906.08936

[4] A. Yakovenko. (2021) Solana: A new architecture for a high performance blockchain v0.8.13. Accessed: Jan. 12, 2026. [Online]. Available: https://solana.com/solana-whitepaper.pdf

[5] Hivelocity. (2026) The ultimate guide to solana validator infrastructure. Accessed: Jan. 12, 2026. [Online]. Available: https://www.hivelocity.net/kb/solana-validator-infrastructure/

[6] T. Wevelsiep. (2025) Public internet egress costs: AWS vs Azure vs GCP vs Hetzner comparison. Accessed: Jan. 12, 2026. [Online]. Available: https://wz-it.com/en/blog/public-internet-egress-costs-comparison-aws-azure-gcp-hetzner/

[7] T. Crain, C. Natoli, and V. Gramoli, "Red Belly: a secure, fair and scalable open blockchain," in *Proceedings of the 42nd IEEE Symposium on Security and Privacy (S&P)*. IEEE, May 2021.

[8] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling Byzantine agreements for cryptocurrencies," in *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP)*. ACM, 2017, pp. 51–68.

[9] Aptos Foundation. (2022, Aug.) The Aptos blockchain: Safe, scalable, and upgradeable Web3 infrastructure. [Online]. Available: https://aptosfoundation.org/whitepaper/aptos-whitepaper_en.pdf

[10] M. Conti, A. Gangwal, and M. Todero, "Blockchain trilemma solver Algorand has dilemma over undecidable messages," in *Proceedings of the 14th International Conference on Availability, Reliability and Security*. New York, NY, USA: ACM, 2019. [Online]. Available: https://doi.org/10.1145/3339252.3339255

[11] Avalanche. (2024) AvalancheGo configs and flags. Accessed: Aug. 31, 2024. [Online]. Available: https://docs.avax.network/nodes/configure/configs-flags

[12] M. Castro and B. Liskov, "Practical Byzantine fault tolerance," in *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, ser. OSDI '99. USA: USENIX Assoc., 1999, pp. 173–186.

[13] G. Voron and V. Gramoli, "Dispel: Byzantine SMR with distributed pipelining," arXiv, Tech. Rep. 1912.10367, 2020.

[14] D. Tennakoon, Y. Hua, and V. Gramoli, "Smart Redbelly blockchain: Reducing congestion for Web3," in *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. St. Petersburg, FL, USA: IEEE, May 2023, pp. 940–950. [Online]. Available: https://ieeexplore.ieee.org/document/10177397/

[15] D. Tennakoon and V. Gramoli, "Deconstructing the Smart Redbelly blockchain," *IEEE Trans. Comput.*, 2024.

[16] G. Danezis, L. Kokoris-Kogias, A. Sonnino, and A. Spiegelman, "Narwhal and Tusk: a DAG-based mempool and efficient BFT consensus," in *Proceedings of the Seventeenth European Conference on Computer Systems*. Rennes France: ACM, Mar. 2022, pp. 34–50. [Online]. Available: https://dl.acm.org/doi/10.1145/3492321.3519594

[17] B. Cho. (2023) AIP-26 – Quorum Store. Accessed: Jan. 9, 2026. [Online]. Available: https://github.com/aptos-foundation/AIPs/blob/main/aips/aip-26.md

[18] B. Cho and A. Spiegelman. (2023, May) Quorum Store: How consensus horizontally scales on the Aptos blockchain. Accessed: Mar. 7, 2025. [Online]. Available: https://medium.com/aptoslabs/quorum-store-how-consensus-horizontally-scales-on-the-aptos-blockchain-988866f6d5b0

[19] T. Crain, V. Gramoli, M. Larrea, and M. Raynal, "DBFT: efficient leaderless Byzantine consensus and its application to blockchains," in *17th IEEE International Symposium on Network Computing and Applications NCA*. IEEE, 2018, pp. 1–8.

[20] IETF, "RFC 9000 – QUIC: A UDP-based multiplexed and secure transport," 2021.

[21] Anza. Turbine block propagation. Accessed: Mar. 7, 2025. [Online]. Available: https://docs.anza.xyz/consensus/turbine-block-propagation

[22] A. Lebedev and V. Gramoli, "On the relevance of blockchain evaluations on bare metal," in *7th Symposium on Distributed Ledger Technologies (SDLT)*. Springer Nature Singapore, 2023.

[23] Aptos. (2025, Feb.) Node requirements. Accessed: Mar. 7, 2025. [Online]. Available: https://aptos.dev/en/network/nodes/validator-node/node-requirements

[24] M. Nadeau. (2023, Nov.) The fundamentals of Avalanche. Accessed: Mar. 7, 2025. [Online]. Available: https://tokenterminal.com/crypto-research/avalanche#decentralization

[25] Redbelly. Node hardware specification requirements — Vine (Redbelly developer portal). Accessed: Mar. 7, 2025. [Online]. Available: https://vine.redbelly.network/nds-node-hardware-specification-requirements

[26] R. Chern. (2023, Oct.) Turbine: Block propagation on Solana. Accessed: Mar. 7, 2025. [Online]. Available: https://www.helius.dev/blog/turbine-block-propagation-on-solana

[27] D. Hyland, J. Sousa, G. Voron, A. Bessani, and V. Gramoli, "Ten myths about blockchain consensus," in *Blockchains*, S. Ruj, S. S. Kanhere, and M. Conti, Eds. Cham: Springer International Publishing, 2024, vol. 105, pp. 3–24. [Online]. Available: https://link.springer.com/10.1007/978-3-031-32146-7_1

[28] G. Voron and V. Gramoli, "Invited paper: Planetary scale byzantine consensus," in *Proceedings of the 5th Workshop on Advanced Tools, Programming Languages, and PLatforms for Implementing and Evaluating Algorithms for Distributed Systems*, ser. ApPLIED 2023. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: https://doi.org/10.1145/3584684.3597270

[29] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "Blockbench: A framework for analyzing private blockchains," in *Proceedings of the 2017 ACM International Conference on Management of Data*, 2017, p. 1085–1100.

[30] D. Kelsey. (2024) Hyperledger Caliper. Accessed: Jan. 12, 2026. [Online]. Available: https://hyperledger.github.io/caliper/

[31] C. Decker and R. Wattenhofer, "Information propagation in the bitcoin network," in *IEEE P2P 2013 Proceedings*, 2013, pp. 1–10.

[32] A. P. Ozisik, G. Andresen, B. N. Levine, D. Tapp, G. Bissias, and S. Katkuri, "Graphene: efficient interactive set reconciliation applied to blockchain propagation," in *Proceedings of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 303–317. [Online]. Available: https://doi.org/10.1145/3341302.3342082

[33] V. Di Perna, M. Bernardo, F. Fabris, S. Amaro, M. Matos, and V. Schiavoni, "Impact of network topologies on blockchain performance," in *Proc. 19th ACM International Conference on Distributed and Event-Based Systems (DEBS)*, Jun. 2025.

[34] B. Nasrulin, M. De Vos, G. Ishmaev, and J. Pouwelse, "Gromit: Benchmarking the performance and scalability of blockchain systems," in *2022 IEEE International Conference on Decentralized Applications and Infrastructures (DAPPS)*. Newark, CA, USA: IEEE, Aug. 2022, pp. 56–63. [Online]. Available: https://ieeexplore.ieee.org/document/9899852/

[35] M. Corallo. (2016) Compact block relay. bip 152. Accessed: Jan. 12, 2026. [Online]. Available: https://github.com/bitcoin/bips/blob/master/bip-0152.mediawiki

[36] U. Klarman, S. Basu, A. Kuzmanovic, and E. G. Sirer. (2019) bloxroute: A scalable trustless blockchain distribution network. Accessed: Jan. 12, 2026. [Online]. Available: https://bloxroute.com/wp-content/uploads/2019/11/bloXrouteWhitepaper.pdf

[37] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "HotStuff: BFT consensus with linearity and responsiveness," in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, 2019.

[38] V. P. Di Perna, V. Schiavoni, F. Fabris, and M. Bernardo, "Blockchain energy consumption: Unveiling the impact of network topologies," in *2025 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2025, pp. 1–10.