



**HAL**  
open science

## SoK: DAG-based Consensus Protocols

Mayank Raikwar, Nikita Polyanskii, Sebastian Müller

► **To cite this version:**

Mayank Raikwar, Nikita Polyanskii, Sebastian Müller. SoK: DAG-based Consensus Protocols. 2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), May 2024, Dublin, France. pp.1-18, 10.1109/ICBC59979.2024.10634358 . hal-04943549

**HAL Id: hal-04943549**

**<https://hal.science/hal-04943549v1>**

Submitted on 12 Feb 2025

**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.

# SoK: DAG-based Consensus Protocols

Mayank Raikwar,\* Nikita Polyanskii,† Sebastian Müller‡

\*University of Oslo, Norway

Email: mayankr@ifi.uio.no

†IOTA Foundation, Berlin, Germany

Email: nikita.polyansky@gmail.com

‡Aix Marseille Université, CNRS, Centrale Marseille, France

Email: sebastian.muller@univ-amu.fr

**Abstract**—This paper is a Systematization of Knowledge (SoK) on Directed Acyclic Graph (DAG)-based consensus protocols, analyzing their performance and trade-offs within the framework of consistency, availability, and partition tolerance inspired by the CAP theorem.

We classify DAG-based consensus protocols into availability-focused and consistency-focused categories, exploring their design principles, core functionalities, and associated trade-offs. Furthermore, we examine key properties, attack vectors, and recent developments, providing insights into security, scalability, and fairness challenges. Finally, we identify research gaps and outline directions for advancing DAG-based consensus mechanisms.

## I. INTRODUCTION

Distributed Ledger Technology (DLT) has become fundamental in supporting secure and transparent transaction systems across decentralized networks. Maintaining an immutable and append-only ledger, DLT enables a trustless environment where participants can transact directly without intermediaries. This technology, particularly its application in cryptocurrencies like Bitcoin and Ethereum, has attracted wide interest due to its potential for enhanced transparency, operational efficiency, and decentralization. However, standard blockchain architectures face key performance limitations, including low throughput and high confirmation latency. Studies have further highlighted trade-offs in speed, security [1], and performance [2], as well as the inherent challenge in achieving decentralization, consistency, and scalability simultaneously (the DCS-satisfiability theorem) [3], [4]. In response, DAG-based DLTs have emerged, using alternative consensus structures to improve scalability and support more efficient consensus mechanisms.

### A. What is a DAG-based Consensus Protocol

Traditional blockchain protocols, such as Bitcoin’s Nakamoto consensus [5], arrange blocks of transactions sequentially in a chain, where each block references its predecessor, extending back to the genesis block. Consensus in these systems often proceeds in rounds, with new blocks added to the longest chain, which participants recognize as the valid ledger. This linear structure, however, imposes limitations on scalability and throughput.

In contrast, DAG-based consensus protocols allow blocks to reference multiple predecessors, creating a Directed Acyclic Graph (DAG) structure. This referencing mechanism enables a




Fig. 1. Blockchain and blockDAG

framework where multiple blocks can be added concurrently, supporting parallel processing of transactions. By diverging from the sequential constraints of traditional models, DAG-based protocols offer improved scalability and flexibility in transaction processing.

### B. Why a DAG-based Consensus Protocol

In the pursuit of transcending the inherent trade-off between security and performance and in response to the performance bottlenecks, DAG-based consensus protocols were proposed as a solution. These protocols promise high scalability and fast confirmation of transactions, hence effectively addressing the intricate balance between security and performance. Following, we describe the promises and challenges of DAG-based consensus protocols compared to linear chain protocols.

*Advantages and promises of DAG-based consensus protocols compared to blockchains*

- 1) *Scalability*: DAG-based protocols can process transactions in parallel rather than sequentially, as in blockchains.

This capability offers improved scalability and throughput compared to sequential blockchains.

- 2) *Latency*: Latency in DAG-based protocols, particularly regarding transaction confirmation time, varies with the underlying consensus mechanism. In PoW-based protocols, a DAG structure enables shorter block times and faster transaction confirmation. In some Proof of Stake (PoS)-based protocols, using additional reliable broadcast primitives increases the latency. Therefore, the extent to which DAG-based protocols reduce latency is conditional on the specific consensus approach and the architectural decisions around synchrony and security.
- 3) *Flexibility*: The DAG architecture allows for more flexible consensus mechanisms and can adapt to various network conditions, potentially making it more versatile than a blockchain. For example, in traditional blockchains, every block serves three roles: acting as a leader that validates transactions, providing content in the form of transaction data, and voting on the causal history. However, blocks in a DAG structure can have differentiated roles, e.g., see Section III-A2, enabling a more distributed approach to consensus and transaction validation.
- 4) *Parallel writing*: DAG-based protocols enable concurrent block production, diverging from the leader-based block generation model of traditional blockchains. This architecture permits multiple participants to simultaneously append transactions or blocks to the ledger, effectively broadening writing access. In its most extreme situation, every participant may independently produce blocks, eliminating reliance on a single or group of block producers. This property may remove the need for additional mempools of transactions, reduce the latency, and remove bottlenecks associated with sequential block generation.

#### *Challenges in DAG-based consensus protocols compared to blockchains*

- 1) *Security risks*: The increased complexity of DAG structures introduces additional potential attack vectors, as more intricate systems often expose a broader range of vulnerabilities compared to simpler protocols.
- 2) *Understanding and adoption*: The more complex nature of DAG-based protocols can make them harder to understand and adopt, particularly for developers accustomed to traditional blockchain technology. Protocols that only provide partial ordering may encounter more difficulties adapting to existing infrastructures.
- 3) *Consensus mechanism maturity*: Consensus mechanisms used within DAG-based DLTs are often newer and less tested than those used with blockchains, creating uncertainty around their long-term stability and resilience against evolving network threats.
- 4) *Reachability challenge*: The concept of reachability within a DAG refers to the ability of a new block to reference earlier transactions or blocks, which is central to the structure's integrity and efficient functioning. This ability has ramifications for functionalities, such as pruning (removing

old data to save space) and the operation of light clients (nodes that do not store the full ledger data).

The number of DAG-based consensus protocols has been increasing. However, there have been few attempts to consolidate these protocols and provide a comprehensive analysis of them. As the DAG-based DLT community grows, it is becoming increasingly important to have a holistic understanding of the various architectures and trade-offs involved.

#### *C. Contribution*

This SoK offers:

- A structured classification of DAG-based consensus protocols into Availability-focused and Consistency-focused categories, analyzing their attributes, trade-offs, and consensus approaches (Sec. III).
- A systematic overview of attack vectors relevant to each protocol category, detailing potential vulnerabilities and countermeasures (Sec. IV).
- An examination of desirable properties currently emphasized in DAG-based DLT research (Sec. V).
- A discussion of recent advancements in DAG-based DLTs that relate to consensus mechanisms but extend beyond the primary classification (Sec. VI).
- An outline of research gaps and future directions to inform and guide ongoing work in the field (Sec. VII).

*Note*: This paper focuses on the conceptual/architectural design of the DAG-based consensus protocols. However, the paper does not discuss or present any performance evaluation of the included protocols.

#### *D. Related Work*

Recently, there has been a growing interest in DAG-based protocols in both industrial and academic circles. Numerous DAG-based consensus protocols have been developed to optimize the complex interplay between security and performance. This paper seeks to systematize a wide variety of research on these DAG-based consensus protocols. An overview of the development of these protocols is depicted in Figure 2.

Wang et al. [8] provides a review of DAG-based systems, exploring essential aspects such as consensus mechanisms, security, and performance. However, the study does not address emerging insights and critical open research questions in DAG-based consensus. Furthermore, recent advancements in consistency-focused DAG protocols necessitate an updated, synthesized reference for researchers.

In contrast, [9] organizes DAG-based ledgers into structural categories—main chain, parallel chains, natural topology, and layered DAGs—providing a taxonomy based on the consensus structure. Our SoK introduces a different framework by classifying DAG protocols through an availability and consistency lens. This approach enables a more conceptual understanding of the trade-offs and challenges relevant to DAG-based consensus protocols.

Other studies have examined distributed ledger technologies (DLTs) across a broader spectrum. Bellaj et al. [10] survey




Fig. 2. The evolution of DAG-based consensus protocols over time. Our estimation of the total number of protocols is based on a manual review of papers citing previous work starting with Hashgraph [6] and Tangle [7] - in Google Scholar and their related work sections. This count, intended as a rough approximation, highlights the growing diversity and interest in DAG protocols.

DLTs using a layered model, categorizing them into chained, chainless, and hybrid types, providing a high-level overview that includes both blockchain and non-blockchain structures. Kannengießer et al. [11] analyze various DLT characteristics and the inherent trade-offs between them, highlighting the diverse operational and structural choices within DLT systems. Wu et al. [12] review both chain-based and DAG-based ledgers, offering a taxonomy that differentiates consensus mechanisms and structural variations. They also identify open research questions relevant to each category, giving insights into potential future developments.

## II. DAG-BASED CONSENSUS IN A NUTSHELL

A DAG-based consensus protocol arranges blocks in a DAG instead of a linear chain. The directed edges in the DAG establish a causal order linked by cryptographic means, providing evidence of node communication. Nakamoto’s proposal already presents this concept of causal dependency, but DAG-based consensus protocols use it more explicitly. Crucially, the architecture allows each block to reference multiple predecessors, thereby facilitating the inclusion of more blocks. This contrasts with the traditional blockchain approach, which privileges the blocks along the longest chain.

More precisely, the DAG consists of vertices and directed edges between them. Each vertex represents a block, while the directed edges represent the relationships between these vertices. Specifically, an edge between two vertices represents a partial order relationship, where one vertex verifies, confirms, or witnesses the other vertex. An edge represents a hash reference from one vertex to another. In this way, a reference gives a causal order of the vertices (blocks). The fact that the DAG is acyclic and thus does not contain any cycle guarantees

the absence of circular dependencies in our DAG structure. We also refer to Figure 1 for an illustrative comparison between a blockchain and a (block)DAG.

In a traditional linear blockchain, the longest chain rule serves two purposes: 1) agreeing on the included blocks and 2) establishing their order. Transitioning to DAG-based structures necessitates redefining these mechanisms due to the non-linear nature of DAGs. Two principal categories of protocols have emerged: those that perform both tasks directly on the DAG itself and Byzantine Fault Tolerant (BFT) protocols that utilize supplementary broadcast primitives for the agreement on the included blocks.

Within this SoK, the analysis of principal design decisions is essential. These decisions determine the structural and operational attributes of the underlying DLT systems and impact scalability, security, and efficiency. Such scrutiny is crucial for comparing the attributes and constraints of different DLT architectures. Therefore, we will present these design choices for DAG-based DLTs.

### A. Ordering

In distributed ledgers, consensus is often associated with establishing a total ordering of blocks. However, total ordering is not always essential, particularly for payment systems (e.g., [13]–[15]). In these cases, a partial ordering—provided naturally by the underlying DAG structure—can be sufficient to confirm transactions. This approach can significantly improve throughput and efficiency for applications that do not rely on strict sequential processing of transactions.

### B. Ledger Model

Most consensus protocols are agnostic to the underlying ledger models as they rely on a total ordering of transactions. However, some protocols, like those in SPECTRE [16], IOTA [17], and Avalanche [18], deviate by adopting a partial order, challenging the idea that total ordering is indispensable.

Two prevalent ledger models serve as the backbone for transaction management:

- 1) *UTXO (Unspent Transaction Output)*: In a UTXO model, transactions are transfers of value from previous transaction outputs to new unspent outputs in an inductive way. New unspent transaction outputs are called UTXO and are used as inputs for new transactions. Bitcoin [5], and Cardano [19], epitomises this approach. It was adapted to the DAG-setting by IOTA [17] and Avalanche [18]<sup>1</sup> since it sidesteps the need for total ordering by structuring the UTXOs as a DAG, promising higher parallelism in transaction processing.
- 2) *Account-based*: In an account-based model, each public address is considered an account. Each account has a balance associated with it. Transactions signify direct value transfers between these accounts, enabling a straightforward balance update mechanism exemplified by Ethereum [20].

<sup>1</sup>The Avalanche crypto project is no longer pursuing the case of UTXOs and their Avalanche consensus protocol on its main net and currently uses an account-based ledger state with a consensus protocol named Snowman.

Object-based and message-based models offer a more general description of possible ledger models. In object-based ledgers, such as the UTXO model, transactions modify objects and thus only affect local states. Due to its localized change impact, this model naturally lends itself to sharding solutions. The causality in object-based models forms a DAG, suggesting that many scenarios do not require a total ordering of transactions thanks to their innate parallelism, as explored in the reality-based ledger model [21]. Conversely, message-based ledgers conceptualize transactions as messages that induce changes across a global state, highlighting a fundamental distinction in how transactions are processed and effectuated.

A mention of owned and shared objects and how they interplay in contemporary platforms like SUI [22] further refines this classification, showcasing a transition towards more granular and flexible state management within DLTs. Such differentiation could enhance understanding and stimulate the development of more scalable ledger solutions, as done in [23].

### C. Consensus Participation and Writing Access

The model of consensus participation determines who can validate and contribute to the DAG. This model directly impacts the system’s openness, security, and decentralization. Consensus participation generally falls into four categories:

- 1) *Lottery-based*: Participation is determined by a Proof-of-Work (PoW) or Proof-of-Stake (PoS) lottery mechanism, where nodes are selected based on computational work or stake holdings.
- 2) *Permissioned*: Participation is restricted to a predefined set of validators granted access by the network initiator or governed by protocol-defined rules. All participants are typically aware of one another.
- 3) *Committee-based*: A rotating subset of validators is selected to participate in consensus for a designated period, allowing for flexibility while maintaining control over participation.
- 4) *Open Participation*: Any node can join as a validator without permission, allowing fully open participation in the consensus process.

Recent work by Lewis-Pye and Roughgarden [24] outlines a hierarchy of “degrees of permissionlessness” describing how participants’ knowledge affects the consensus mechanism.

While many permissioned protocols could theoretically incorporate committee-based selection, we categorize them as permissioned unless committee selection is explicitly addressed in their theoretical design.

Writing access, or the ability to propose new blocks, typically aligns with consensus participation. However, specific protocols like [17] and [25] may allow broader access to block proposals while limiting consensus participation to a subset of nodes, balancing decentralization and efficiency.

### D. Network Model

Consensus protocols are designed based on different network models that describe how communication occurs between the participants. These models are often defined by the potential power of an adversary to control message delays. Three

commonly used communication models are synchronous, asynchronous, and partially synchronous. To comprehend these models, it is essential to have a formal understanding of the following key concepts in network systems.

The *delay parameter*  $\Delta$  defines the maximum delay (in time steps) a message suffers in a synchronous phase. The parameter  $\Delta$  is known upfront to the nodes in the protocol. The *Global Stabilization Time* (GST) dictates when the underlying communication network switches from asynchronous to synchronous. Unlike  $\Delta$ , GST is not known upfront in the protocol and can be chosen by an adversary. More precisely:

- 1) *Synchronous*: In a synchronous model, the messages between nodes are delivered within known and predictable time frames. Nodes follow a common global clock or synchronized time intervals to operate in lockstep. In short, for any message sent, an adversary can delay its delivery by at most  $\Delta$ . This model is used in time-sensitive systems, making it suitable for some DLT-based real-time systems.
- 2) *Asynchronous*: In an asynchronous model, the messages between nodes experience varying delays. An adversary can delay the delivery of any message sent by any finite amount of time, but eventually, the message gets delivered. This offers the most flexibility, as message delivery times are unpredictable. This model is standard and used in scenarios where global synchronization is impractical.
- 3) *Partially Synchronous*: In a partially synchronous model, messages are delivered within an unknown finite time. In the equivalent *eventually synchrony* model, the GST event will occur after a certain, unspecified time has passed. Moreover, any message sent at time  $t$  must be delivered by time  $\Delta + \max(t, \text{GST})$ . This model offers a balance between flexibility and accuracy. Nodes in this model aim to operate with timers that can measure time  $\Delta$  after an event, and there is no guarantee about the exact timing of message deliveries. These models are common in network environments where transient conditions may cause disruptions. After GST, the network guarantees that messages are delivered within the delay parameter  $\Delta$ , providing a window of predictability that can be very useful for consensus algorithms, especially in fault-tolerant systems that must operate under the assumption of periodic network instability. This model is particularly relevant for DLT systems that aim to achieve consensus despite unpredictable network conditions.

### E. Dynamic Availability

In certain consensus protocols, as highlighted by Neu et al. (2022) [26], nodes operate within a fluid participation framework that aligns well with consensus algorithms built on the “sleepy” model, put forth by Pass and Shi (2017) [27]. In these consensus protocols, nodes can exhibit dynamic availability, seamlessly transitioning between periods of activity and dormancy up to a specified *Global Awake Time* (GAT). After this moment, an assumption is made that all honest nodes will be consistently online, facilitating consensus in an environment

reflective of real-world participation and connectivity patterns. Thus, dynamic availability ensures transactions are processed and finalized despite any transient fault.

#### F. Additional Broadcast Primitive

DAG-based consensus protocols adopt different strategies to ensure the reliable propagation of blocks. These strategies can be broadly classified based on their reliance on additional reliable broadcast primitives.

Many DAG-based systems use established Reliable Broadcast (RB) protocols [28] to distribute blocks. These systems implement an RB protocol as an additional layer, ensuring consistent and reliably disseminated information among nodes, even in the face of network challenges or adversarial behaviour.

Conversely, some protocols employ the blockDAG structure itself to replicate the functionality of a reliable broadcast mechanism. The DAG's architecture inherently distributes data across the network so that blocks are organically verified and propagated without needing a specialized RB protocol.

Both approaches aim to maintain a secure and coherent state across distributed participants. However, they differ fundamentally in whether they view the blockDAG as sufficient for reliable broadcasting or require an additional overlay of RB protocol to strengthen information dissemination and security. The additional use of RB typically influences latency by increasing it for several network rounds.

#### G. Adversarial Model

The adversary model defines attacker capabilities essential for setting security parameters and building resilient consensus mechanisms. Typically, these models assume adversaries with bounded computational power, limited to polynomial-time algorithms or unbounded power, capable of breaking cryptographic primitives. This model is intrinsically connected to the network's communication model. The FLP impossibility result, [29], highlights the challenges of achieving consensus in a purely asynchronous network with just one faulty node. To address this, DLTs may assume partial synchrony or incorporate elements of randomness to secure consensus against adversarial actions within probabilistic bounds. In systems that use PoW and PoS offering probabilistic finality, the critical threshold is 50%. This means that if more than 50% of the participants are honest and not faulty, the network is operational, guaranteeing safety and liveness. On the other hand, BFT mechanisms adopt a 2/3 threshold, requiring more than 2/3 of honest and non-faulty nodes, and offer a deterministic finality. These two thresholds are fundamental to the security models of DLTs and reflect a balance between resilience to adversarial control and consensus outcome under varying assumptions of network synchrony.

#### H. Leader-based Vs Leaderless Consensus

Consensus mechanisms within DLTs typically adopt either a leader-based or a leaderless structure. Leader-based consensus protocols designate a leader to propose an orderly view of the ledger transactions. This view becomes committed when a

sufficient consensus, or votes, affirming the view is reached. Leaders may be chosen through various mechanisms, such as PoW/PoS lottery or a round-robin selection.

In contrast, leaderless consensus protocols decentralize the proposal process, with decisions emerging from the collective input of all nodes.

#### I. Mempool

For protocols where writing access is probabilistic or requires selection from a group, such as lottery-based or committee-based systems, a mempool is essential.

A mempool, short for memory pool, is a temporary storage for transactions broadcasted to the network but not yet included in a block. The mempool serves several important roles:

- 1) *Transaction Queue*: It acts as a queue for pending transactions, maintaining them in an accessible state until they are selected and confirmed by a node with writing access.
- 2) *Prioritization*: The mempool can prioritize transactions based on specific criteria, such as fee rates, ensuring that transactions with higher fees are confirmed more quickly.
- 3) *Availability*: It ensures that transactions remain available for selection and confirmation, even when uncertain which node will write the next block into the ledger.
- 4) *Synchronization*: In systems with multiple potential writers or validators, the mempool is crucial for synchronizing the state across these entities, allowing them to view and validate the same set of unconfirmed transactions.

Some protocols use a blockDAG as a distributed mempool for pending transactions, eliminating the need for a separate mempool structure. The blockDAG then maintains a real-time "ledger" of unconfirmed transactions, making the system more transparent and accessible. This more structured mempool can mitigate transaction duplicates; see also discussion in Section VII.

#### J. Finality

Finality refers to the point at which a transaction or a block of transactions is considered irreversible, permanently part of the ledger, and its execution is settled. Finality comes in two forms: probabilistic and deterministic (or absolute).

- 1) *Probabilistic Finality*: It means the order of a block can be reverted, but the risk decreases as it gets more embedded in the blockchain or blockDAG. In PoW systems, finality is probabilistic, and in most cases, the probability of a block being reverted decreases exponentially over time.
- 2) *Deterministic Finality*: It ensures that the order of a block or the presence of a transaction in the ledger is permanent once it is recorded. Protocols that use BFT ideas usually offer deterministic finality. They require an agreement

<sup>2</sup>The finality results included here represent claims or descriptions by the respective protocol authors. Many protocols listed lack complete formal proofs within their stated models, and results are often based on suggested behaviour rather than rigorous validation.

<sup>3</sup>The network model is not always given explicitly in the corresponding papers. In this case, we attempt to classify it based on the information in the paper. These cases are marked by this footnote.

TABLE I  
OVERVIEW OF DAG-BASED CONSENSUS PROTOCOLS

Protocol	Participation	Additional Broadcast Primitive	Additional Mempool	Leader-based	Dynamic Availability	Ordering	Finality <sup>2</sup>	Network Model <sup>3</sup>	Innovation/Relation
Nakamoto [5]	PoW	No	Yes	Yes	Yes	Total	Probabilistic	Synchronous	Genesis of blockchain and cryptocurrencies
GHOST [30]	PoW	No	Yes	Yes	Yes	Total	Probabilistic	Synchronous <sup>4</sup>	Improved upon Nakamoto's chain structure for higher throughput using heaviest chain selection
Inclusive [31]	PoW	No	Yes	Yes	Yes	Total	Probabilistic	Synchronous <sup>4</sup>	Evolved from GHOST, presents a game-theoretic model for reward among miners.
Byteball [32]	Open	No	Yes	No	Yes	Total	Deterministic	Synchronous <sup>4</sup>	Witness nodes for syncing and ordering transactions using main chain index (MCI)
SPECTRE [16]	PoW	No	Yes	No	Yes	Partial	Probabilistic	Partial Synchronous	Further evolved from GHOST, prioritizing fast confirmations using pairwise voting
PHANTOM [33]	PoW	No	Yes	No	Yes	Total	Probabilistic	Synchronous	Introduced k-cluster for better structure over SPECTRE's DAG
GHOSTDAG [33]	PoW	No	Yes	No	Yes	Total	Probabilistic	Synchronous	Focused on reducing latency and refined k-cluster approach
DAG KNIGHT [34]	PoW	No	Yes	No	Yes	Total	Probabilistic	Partial Synchronous <sup>3</sup>	Parameterless expansion from GHOSTDAG without latency bounding
Prism [35]	PoW	No	Yes	Yes	Yes	Total	Probabilistic	Synchronous	Separated consensus functions into multiple parallel chains
GHASt [36]	PoW	No	Yes	Yes	Yes	Total	Probabilistic	Synchronous	Tree-graph approach and adaptive mechanisms for performance and security
Hashgraph [6]	Permissioned	No	No	No	No	Total	Deterministic	Asynchronous	First optimistic DAG with fair ordering derived from block's timestamps
Jointgraph [37]	Permissioned	No	No	Yes	No	Partial	Deterministic	Synchronous <sup>4</sup>	Improves throughput and latency of Hashgraph using a supervisor node with additional events
Aleph [38]	Permissioned	Yes	Yes	No	No	Total	Deterministic	Asynchronous	Round-based DAG with retrospective randomized leader selection
DAG-Rider [39]	Permissioned	Yes	Yes	No	No	Total	Deterministic	Asynchronous	Reduced latency compared to Aleph by a new commit rule
Avalanche [18]	Permissioned	No	No	No	No	Partial	Probabilistic	Synchronous	Introduced leaderless consensus using repeated sub-sampled querying
OHIE [40]	PoW	No	Yes	No	Yes	Total	Probabilistic	Synchronous	Many parallel instances of the Nakamoto consensus
Parallel Chains [41]	PoW / PoS	No	Yes	No	Yes	Total	Probabilistic	Partial synchronous	theoretical framework for optimistic throughput
Expected consensus [42]	PoS	No	Yes	No	Yes	Total	Probabilistic	Synchronous	Increased throughput and adapted consensus for storage resource
Nano [43]	Open	No	Yes	No	Yes	Partial	Deterministic	Synchronous <sup>4</sup>	Block-lattice structure allowing asynchronous updates
Chainweb [44]	PoW	No	Yes	No	No	Partial	Probabilistic	Synchronous <sup>4</sup>	Employs simple payment verification (SPV) for token transfer
Meshcash [45]	PoW	No	Yes	No	Yes	Total	Deterministic	Asynchronous	Ebb-and-flow type protocol
Tusk [46]	Permissioned	Yes	Yes	No	No	Total	Deterministic	Asynchronous	Mempool management by Narwhal, resulting in high throughput; Tusk orders digests of blocks
Bullshark [47]	Permissioned	Yes	Yes	Both	No	Total	Deterministic	Partially Synchronous / Asynchronous	Improved latency by introducing new commit rule; garbage collection and timely fairness after GST
Cordial Miner [48]	Permissioned	No	No	Both	No	Total	Deterministic	Asynchronous	Reduced latency by using best effort broadcast and committing leader blocks
Tangle 2.0 [17]	PoS	No	No	No	No	Partial	Deterministic	Synchronous <sup>4</sup>	Conflict resolution using the DAG and reality-based ledger state
BBCA-Chain [49]	Permissioned	Yes	Yes	Yes	No	Total	Deterministic	Partially Synchronous	Latency reduction by several network trips by using a new BBCA primitive for leader blocks
Sailfish [52]	Permissioned	No	No	Yes	No	Total	Deterministic	Partially Synchronous	Reduced latency compared to Shoal by having multiple leaders every round and novel commit rule
Mysticeti [23]	Permissioned	No	No	Yes	No	Total	Deterministic	Partially Synchronous	Pipelined leader schedule, fast finality for owned object transactions, epoch-closing mechanism
Shoal [50]	Permissioned	No	No	Yes	No	Total	Deterministic	Partially Synchronous	Improved the latency compared to Bullshark by pipelining leader blocks
Shoal++ [51]	Permissioned	No	No	Yes	No	Total	Deterministic	Partially Synchronous	Optimised latency by operating multiple DAGs in parallel.
Sailfish [52]	Permissioned	No	No	Yes	No	Total	Deterministic	Partially Synchronous	Reduced latency compared to Shoal by having multiple leaders every round and novel commit rule
Mahi-Mahi [53]	Permissioned	No	No	Yes	No	Total	Deterministic	Asynchronous	Reduced latency compared to Cordial miners through boost rounds
Slipstream [54]	Permissioned	No	No	No	Yes	Total	Deterministic	Partially Synchronous	Ebb-and-flow DAG-based protocol and fast path confirmation on DAG

from at least a supermajority of the validators, making it irrevocably part of the ledger.

Additionally, some DLTs with DAG structures may initially only offer probabilistic finality. However, they can obtain deterministic finality by designating special roles; for instance, witness nodes in Byteball [32] or nodes creating snapshot chains in Vite [55] and IOTA 2.0 [56].

### K. Overview

Table I presents a diverse range of DAG-based consensus protocols, their underlying models, technical specifications/features, and key innovation ideas. The column “Innovation/Relation” highlights the novel contributions of each protocol or delineates its evolution from preexisting technologies.

## III. DAG-BASED CONSENSUS PROTOCOLS

In distributed systems, consensus is crucial for agreeing on a single version of the system's state. In DAG-based DLTs, nodes collaboratively build a blockDAG and agree on a subset of interconnected blocks called a prefix<sup>4</sup> to ensure a shared understanding of the ledger's history. Once nodes agree on this prefix, they can establish a total ordering for the included transactions. In certain protocols, they might only ascertain a partial ordering, which can suffice for transaction execution depending on the chosen ledger model, see Section II-B.

As we develop our systematization, it's essential to recognize the fundamental limitations imposed by the CAP theorem. This theorem states that any distributed system can achieve a

<sup>4</sup>A prefix in this context is a set of blocks in a DAG which is closed by causal order relations imposed by the DAG.

maximum of two out of three key properties - consistency, availability, and partition tolerance - at the same time. Our SoK is based on this principle, which guides our analysis and comprehension of the inherent trade-offs and design choices. It's important to note that these properties are defined originally in [57], but we emphasize the need to interpret them in the current context, e.g., [58]. The exact definition varies, and we give the definition provided by Brewer in [58].

The CAP theorem states that any networked shared-data system can have at most two of three desirable properties:

- *Consistency (Safety)*: consistency (C) equivalent to having a single up-to-date copy of the data;
- *Availability (Liveness)*: high availability (A) of that data (for updates); and
- *Partition Tolerance*: tolerance to network partitions (P).

In the above, we informally correlate “consistency” with “safety” and “availability” with “liveness,” despite recognising that these terms traditionally describe distinct concepts within distributed systems literature. This correlation hopefully highlights the most essential trade-offs in DAG-based DLT systems.

Given these constraints, DAG-based consensus protocols face trade-offs and are classified into two primary classes.

### 1) Availability-focused Protocols:

Focusing primarily on maintaining the continuous progress of the ledger, most of these protocols use variations of a weighted<sup>5</sup> lottery, where nodes are randomly selected to propose new blocks to the DAG and vote on the existing ones. Such protocols are typically suited for permissionless

<sup>5</sup>For instance, the weight could be computing power, stake, or storage.

TABLE II  
CATEGORIZATION OF CONSENSUS PROTOCOLS IN DAG-BASED DLTs

DAG-based consensus protocols			
Availability-Focused		Consistency-Focused	
Structured DAG	Unstructured DAG	Optimistic DAG	Certified DAG
Expected consensus [42]	SPECTRE [16]	Hashgraph [6], [59]	Aleph [38]
Inclusive [31]	PHANTOM [60]	Jointgraph [37]	DAG-Rider [39]
GHOST/Conflux [36], [61]	GHOSTDAG [33]	Cordial Miners [48]	Tusk [46]
OHIE [40]	DAG KNIGHT [34]	Tangle 2.0 [17]	Bullshark [47]
Chainweb [44]	Slipstream [54]	Mysticeti [23]	BBCA-Ledger [62]
Parallel chains [41]	Byteball [32]	Slipstream [54]	BBCA-Chain [49]
Prism [35]	Meshcash [45]	Mahi-Mahi [53]	Shoal, Shoal++ [50], [63]
	Graphchain [64]		Sailfish [65]
	Avalanche [66]		

or dynamically available settings [24]. In this setting, a potentially large number of entities participate in the consensus process, and the list of active participants is unknown and changing over time. Canonical examples of availability-focused consensus protocols among linear blockchains are Nakamoto consensus [5] and Ouroboros [67]. When applied in a DAG-based architecture, this approach manifests in different forms, including structured and unstructured DAGs. It is worth noting that, if guaranteed, finalization in availability-focused protocols is often probabilistic.

- a) *Structured DAG*: These protocols enforce strict block addition rules, keeping an organized DAG topology.
- b) *Unstructured DAG*: These protocols represent a more flexible approach to DAG construction, where nodes have greater liberty in attaching new blocks without stringent adherence to predefined rules. However, the trade-off often lies in increased complexity for achieving consensus and validating transactions.

## 2) Consistency-focused Protocols:

The protocols of this type prioritize ledger consistency by drawing from classical Byzantine Fault Tolerant (BFT) protocols [28], [68], [69]. Consistency is mostly achieved in permissioned and quasi-permissionless environments, e.g., see [24]. These protocols can be further categorized based on how participants disseminate their blocks.

- a) *Optimistic DAG*: In these protocols, block dissemination is done through the best-effort broadcast, i.e. sending blocks to other nodes without providing any guarantees about the uniqueness of the blocks. Thereby, such protocols require mechanisms to handle equivocation.
- b) *Certified DAG*: These protocols employ reliable and consistent broadcast primitives for block dissemination. This means that every block in a locally maintained DAG can be seen as a certificate signed by a quorum of nodes. Such protocols ensure participants maintain a consistent DAG view.

We depict Table II to show which consensus protocols fall in which category. The subsequent subsections will explore these categories, analyzing key protocols and their innovations.

## A. Availability-focused Protocols with Structured DAG

1) *Parallel instances of Nakamoto consensus*: The most straightforward approach to improve Bitcoin’s scalability is to execute  $m$  concurrent instances of the Nakamoto consensus protocol as proposed in *Parallel Chains* [41] and *OHIE* [40]. Mining simultaneously and uniformly for all  $m$  chains is achieved by a similar technique in both protocols. A newly mined block must include the digests of the latest blocks from all chains, and a miner remains unaware of which chain a new block will extend until it solves the PoW puzzle. After successful mining, the block’s hash value decides which chain it belongs to.

Establishing total ordering on parallel chain architectures can be achieved in different ways. In *OHIE*, a ranking system for blocks is used. Each block is equipped with a *current rank* and a *next rank*; the current rank mirrors the next rank of its referenced predecessor, while the next rank represents the highest current rank observed across all chains. *Parallel Chains* employs a different strategy, where one specific chain is designated for synchronization. Each block in the parallel architecture references a block from this synchronization chain. The chronological order of the entire network is then inferred based on the sequence of blocks in this special chain.

*Parallel Chains* and *OHIE* guarantee the same safety properties as the Nakamoto Consensus, [5], and improve either the throughput or latency, as their scalability does not offer both desired features simultaneously. Multiple other protocols can also fall into this category, e.g., *Chainweb* [44], but they degrade in security or latency compared to Nakamoto consensus. *Nano* [43], *Vite* [55] protocols follow a similar concept where participants maintain their parallel chains.

2) *Decoupling block’s functionalities*: A conceptual approach suggested in *Prism* [35] to augment scalability in parallel chain architectures involves deconstructing the multifaceted functionalities of blocks, as seen in Nakamoto’s consensus. Traditionally, blocks in Bitcoin simultaneously fulfil three key roles: 1) leader election, 2) transaction proposing, and 3) voting on the causal history through parent links. *Prism* proposes a clear segregation of these functions into distinct types of blocks, each dedicated to a specific purpose, such as transaction processing, block proposal, and voting.



This segregation results in the formation of different types of chains within the architecture. Primarily, it creates one *proposer* chain, where each block references a preceding block within the proposer chain and several *transaction* blocks. In parallel,  $m$  *vote* chains emerge, with each chain functioning to cast votes determining the total order of blocks within the proposer chain. Such a structural reorganization clarifies the roles within the blockchain and allows high throughput, low (constant) latency, and the same safety guarantees as in the Nakamoto consensus.

Note that there are predecessors of Prism where functionalities of blocks are decoupled, e.g., FruitChain [70], Bitcoin-NG [71], but have not improved some other aspects, e.g., latency.

3) *Tree-structure in a DAG*: In *GHOST* (Greedy Heaviest Adaptive SubTree) [36], [61], each block except genesis has precisely one outgoing parent edge and can have multiple outgoing reference edges. The parent edges construct a tree within the broader DAG, and the reference edges establish temporal precedence, indicating that the referenced block predates the block containing the reference. Once a winning chain within the tree is selected, a deterministic order over all referenced blocks can be established. The selection of this winning chain is influenced by the weights assigned to each block, akin to the *GHOST* protocol [30]. Here, the fork choice rule systematically favours the heaviest subtree at every fork, typically equating block weight with the extent of computational work done.

*GHOST* diverges from traditional protocols in its adaptive response to potential liveness attacks. During such scenarios, it modifies the weight distribution among blocks. While normally all blocks carry equal weight, under attack detection, the protocol randomly selects a few blocks to assign non-zero weights and keeps the expected block weight at the same level.

In *Inclusive* [31], the main idea was to include transactions from all the blocks in the ledger. An inclusive rule is defined to select a main chain from the DAG, and further non-conflicting blocks from the outside DAG are appended to the final block structure. The Inclusive protocol also rewards miners whose blocks are not in the main chain but are contained within the DAG.

4) *Round-based chain of tipsets*: Another attempt to improve the latency of Nakamoto’s consensus was suggested in *Expected Consensus* [42]. This heaviest-chain style protocol operates in rounds. Each round involves a cryptographic sortition on the weighted list of participants. It is parametrized such that, on average, a given number of  $m$  participants may be eligible to propose a block every round. Every block must reference a *tipset*, a set of blocks sharing the same round and a parent tipset. Every block adds weight (determined by the sortition) to the chain of tipsets, and the consensus is achieved by following the heaviest chain of tipsets.

## B. Availability-focused Protocols with Unstructured DAG

1) *Unstructured Block-DAG protocols*: The *GHOST* protocol [30] diverges from Nakamoto’s longest chain rule in

selecting the Greedy Heaviest-Observed Sub-Tree. This enables shorter intervals between block creations without compromising security, a concept explored in [72] in further detail. *GHOST* capitalizes on the proof of work in “off-chain” blocks, traversing the tree-like structure that emerges from chain forks. This alternate selection method for the main chain is specifically tailored to mitigate issues associated with network latency.

The *SPECTRE* [16] protocol focuses on building a blockDAG structure with separate mechanisms for mining and consensus, leading to a notion of weak liveness for higher scalability. *SPECTRE* uses a recursive weighted voting technique where each new block in the DAG submits votes (preference ordering) over every pair of blocks based on which block they believe occurred first. The final ordering in *SPECTRE* is based on the majority vote on pairwise ordering across all the blocks. However, this ordering is not extendable to a total order due to Condorcet paradox [73]. While *SPECTRE* only achieves partial ordering, it works under a partial synchronous network model. This partial ordering makes *SPECTRE* suitable for payment systems but not for systems that rely on account-based ledgers.

*PHANTOM* [60] is an extension of the *SPECTRE* protocol. It aims to achieve a total ordering but under a synchronous network model. *PHANTOM* works by separating block creation or mining from the consensus mechanisms. It builds on the intuition that blocks created by honest miners are well-connected, while adversarial blocks are not well-connected and thus should be excluded. The notion of  $k$ -cluster expresses this well-connectedness. Once this  $k$ -cluster is determined, a total block order is established via a topological ordering. The synchronicity assumptions become apparent in the choice of the parameter  $k$ , which depends on the latency. Moreover, the identification of such  $k$ -cluster is NP-hard, *GHOSTDAG*, [33], is proposed as a practical alternative; it also addressed possible attack vectors, pointed out by [74].

*DAG KNIGHT* [34] is an evolution of *GHOSTDAG*. Unlike *GHOSTDAG*, *DAG KNIGHT* assumes no upper bound on network latency. It leverages a dual min-max optimization approach to dynamically find the largest  $k$ -cluster for each  $k$ , selecting the minimal  $k$  covering at least 50% of the DAG. This design accommodates latency variations for safety and provides a responsive protocol to actual network latency rather than a hard-coded latency bound as in *GHOSTDAG*.

*Meshcash* [45] is a framework containing two layered protocols to reach a consensus on the blocks added to the DAG using PoW. The blocks are arranged in layers; each block belongs to a layer and refers to blocks in the previous layer. The consensus protocol is an ebb-and-flow style consensus protocol [75] with a slower protocol that favors safety and a faster one that favors liveness. The slower protocol works for the blocks in the far past (in old layers), where a weighted voting is performed to decide on the consistency of the blocks.

Tangle 2.0 [17] employs a unique approach to conflict resolution, preventing the locking of funds by actively voting

on conflicts within the DAG. Leveraging the reality-based ledger [21], transactions can be treated optimistically, allowing nodes to build on unresolved or unconfirmed transactions. This approach enables concurrent versions to coexist within the ledger without requiring strict total ordering.

The *Slipstream* protocol, [54], is also ebb-and-flow style consensus protocol [75]. As such, it offers two types of block orderings: an optimistic ordering, which is live and secure in a sleepy model under up to 50% Byzantine nodes, and a final ordering, which is a prefix of the optimistic ordering and ensures safety and liveness in an eventual lock-step synchronous model under up to 33% Byzantine nodes. Slipstream utilizes wall clocks instead of Lamport clocks, allowing for dynamic availability in the network. Each node in Slipstream generates commitments based on its view of the local DAG and the online nodes within fixed-duration time intervals (rounds). Each commitment represents a hash of the accepted data (e.g., blocks and transactions) within that round, linking to the commitment of the previous round to form a commitment chain. This chain, representing the prefix of a DAG, allows the network to progress and maintain the available ledger state even during partitions. When a partition resolves and a supermajority of validators is online, nodes merge the commitment chains, resuming finalization. Slipstream also incorporates a UTXO payment system that enables fast transaction confirmation independently of block ordering. During synchrony, transactions can be confirmed in three rounds, with unconfirmed double spends resolved using the DAG structure in a novel approach.

2) *Unstructured Transaction-DAG protocols*: This class includes protocols that append transactions directly in their DAGs without the block as a wrapper, e.g., Graphchain [64].

The *Byteball* [32] protocol allows users to add transaction units to the DAG by signing them. These units are linked in the DAG, containing hashes of prior units to confirm their validity and create a partial order. Users choose trustworthy nodes, called witnesses, based on their reputation. These witness nodes periodically generate transaction units that help compute a main chain in the DAG. Consensus is then found by the total ordering of the transaction using the main chain.

*Avalanche* [66] protocol allows nodes to repeatedly query a random group of nodes about the validity of a transaction. Further, to be able to vote on several transactions at the same time, the nodes employ a DAG structure. When a node promptly accumulates sufficient positive responses for its query, it solidifies its decision. Critiques have emerged regarding the security (especially liveness) of the protocol [76], [77], and the Avalanche project stopped working on the DAG-version of the consensus and maintains a chain version of it.

### C. Consistency-focused Protocols with Optimistic DAG

The protocols in this category rely on validators, a special subset of nodes with a non-zero weight in the voting process. Blocks are propagated using best-effort broadcast or gossiping protocol. During this propagation, the nodes need to ensure that the recipient of the new block also knows the causal history of

the block. In optimistic cases, this can result in fewer network trips to achieve block finality. Typically, these protocols can tolerate adversarial validators holding less than 1/3 of the total weight. It will be convenient to refer to validators holding more than 2/3 of the total weight as a *supermajority*, and the main assumption in the following protocols is there exists an honest supermajority of validators.

One main problem of reliable or consistent broadcasting blocks among the validators is the problem of equivocations. This describes the situation where a validator presents different blocks to different parts of the network. In the following protocols, the DAG structure is used to avoid equivocations using a mechanism of double approvals. Specifically, a block  $X$  is *approved* by a block  $Y$  if  $X$  is reachable from  $Y$  in the DAG and the block creator of  $X$  does not equivocate in the causal history of block  $Y$ . A block  $X$  is *doubly approved* by a block  $Y$  if a supermajority of blocks exists in the causal history of block  $Y$ . Each validator can be shown to have at most one block in a given round, which is doubly approved by a supermajority of nodes. The idea was introduced in *Hashgraph* [6], [59], and its variation was used in other protocols.

Hashgraph suggests constructing a DAG by gossiping over gossiping events<sup>6</sup>. In this protocol, each node maintains its chain of events. When receiving an event block from another node, it randomly selects at least one peer to disseminate information about its current knowledge about blocks, e.g., a new block within the maintained chain is created which in addition references the received block. A node advances to the next round once its block doubly approves a supermajority of blocks from the previous round.

By employing this idea of double approval in the resulting DAG structure, the Hashgraph protocol decides when the network reliably receives each event. The *received* timestamp is then assigned, and the total ordering is achieved by sorting events over the received timestamps.

*Jointgraph* [37] reduces the number of voting rounds in Hashgraph by introducing a *supervisor node*. An event in Jointgraph is final if the event receives more than 2/3 votes from the nodes in which one of the votes is from the supervisor node. The supervisor node also creates snapshots and storage events to finalize the events from the ordinary nodes and provide a total order. When the supervisor node is offline, Jointgraph switches to the Hashgraph process.

*Cordial Miners* [48] is a family of simple and efficient DAG-based consensus protocols with instances for the asynchronous and eventual synchronous models. Similar to Hashgraph, the core idea of Cordial Miners is to utilize the constructed optimistic DAG, called the *blocklace*, for different tasks such as block distribution, equivocation-exclusion, and block ordering. While Hashgraph uses timestamps for ordering, Cordial Miners commit leader blocks, and the slices between committed leader blocks result in ordering. The blocklace is fragmented into waves using the rounds assigned to blocks. Each wave could have at most one committed leader block

<sup>6</sup>In Hashgraph, the term an *event* is used instead of a block

chosen using either a round-robin for a partially synchronous network model or a retrospective random selection for an asynchronous network model.

It has been shown [13] that a system that enables participants to make simple payments needs to solve a simpler task. In cases where payments are independent of one another (e.g. single-owned token assets or UTXO transactions), there is no necessity to order them, and partial ordering is sufficient. Hence, consensus is not required, and the payment system can be realized deterministically in an asynchronous network setting. *Flash* [15] is built by encoding payment transactions into the blocks of the blocklace (see Cordial Miners). The resulting ordering is only partial and hence weaker than the total ordering achieved by Cordial Miners.

*Mysticeti* [23] improves Cordial Miners for a partially synchronous network. Compared to its predecessor, this consensus protocol allows pipelined leader blocks to be used, improving the finality time for average transactions and faster confirmation in the presence of crashed nodes. It integrates fast payment for owned object transactions using the same underlying DAG and enables checkpoints and epoch-closing mechanisms embedded into the DAG.

*Mahi-Mahi* [53] is a practical improvement of Cordial Miners for an asynchronous network. Like *Mysticeti*, this protocol allows multiple leaders to be committed per round. In addition, *Mahi-Mahi* developed the commit rule of Cordial Miners, which resulted in an improved average latency to commit a leader.

#### D. Consistency-focused Protocols with Certified DAG

These consensus protocols are based on classical BFT algorithms. These protocols are round-based, i.e., a node can increase the round number only when the DAG contains a quorum of blocks with the current round number. These protocols optimize their performance by assigning leaders for certain rounds and using special commit rules for leader blocks. Committed leader blocks form a backbone sequence that allows for the partitioning of the DAG into slices and the deterministic sequencing of blocks in the slices.

*Aleph* [38] protocol improves the Hashgraph complexity by building a round-based structured DAG and employing an efficient binary agreement protocol [78]. Like the Hashgraph protocol, *Aleph* separates the network layer (communication DAG) from the protocol logic (virtual voting and ordering). Every block created in *Aleph* is disseminated using a Byzantine Reliable Broadcast (BRB) primitive. Thereby, nodes build a certified DAG. The protocol operates in an asynchronous network and utilizes a trustless ABFT Randomness Beacon to circumvent the FLP-impossibility result, see [29], to reach a total ordering.

*DAG-Rider* [39] is an asynchronous Byzantine Atomic Broadcast (BAB) protocol. *DAG-Rider* creates the round-based structured DAG similar to *Aleph*. *DAG-Rider* is constructed in two layers: communication and ordering layer. In the communication layer, nodes reliably broadcast their proposals (messages) and form a DAG of the messages they deliver.

In each round of *DAG-Rider*, each node broadcasts at most one message, which should contain references (strong edges) to messages (i.e., at least  $2f + 1$ ) of previous rounds. The message can also have references (weak edges) to messages of rounds before the previous round. The weak edges ensure the validity property of BAB. Furthermore, in the ordering layer of *DAG-Rider*, each node observes its local DAG and locally orders all the messages in DAG by employing randomization. This randomization is achieved through a global perfect coin, implemented using threshold signatures that circumvent the FLP-impossibility result.

*Narwhal & Tusk* [46]: *Narwhal* is a DAG-based, structured mempool incorporating concepts from reliable broadcast [79] and reliable storage [80], with the addition of a byzantine fault-tolerant threshold clock [81] for round advancement. The *Tusk* protocol is built atop *Narwhal* to achieve asynchronous consensus with minimal latency. This approach empowers each node to reach a consensus on agreed-upon values by examining its local DAG without additional messages.

*Tusk* refines *DAG-Rider*, transforming it from theory into an implementable system. This is achieved through three pivotal steps: Firstly, *Tusk* adopts a Quorum-based reliable broadcast instead of the conventional reliable broadcast utilized in *DAG-Rider*. Secondly, refining the commit rules enhances the latency in common cases. Lastly, *Tusk* removes the weak links of *DAG-Rider*, enabling efficient garbage collection.

*Bullshark* [47] is a zero-overhead BFT protocol on top of the *Narwhal's* DAG optimized for the synchronous case and achieves substantially reduced latency compared to both *Tusk* and *DAG-Rider*. In the partially synchronous version, *Bullshark* is the most performant and robust compared to existing protocols. *Bullshark* upholds all the desired properties of *DAG-Rider* while reducing the common-case latency during the period of synchrony. The main feature of the *Bullshark* protocol is to exploit the synchronous periods and remove the complex processes of view-change and view-synchronization. *Bullshark* attains amortized complexity and addresses the imperative facet of fairness. *Shoal* [50] reduces the latency of non-leader blocks by interleaving two instances of *Bullshark*.

*BBCA-Chain* [49] introduces a new broadcast primitive, called Byzantine Broadcast with Complete-Adopt (BBCA) to reduce this latency and simplify the consensus logic. BBCA is applied only for leader blocks, while Best-Effort Broadcast (BEB) is applied for all other blocks. The nodes build the DAG as a mixture of certified and optimistic approaches. In addition, *BBCA-Chain* makes all rounds symmetric by assigning leaders every round. Note that there has been a similar effort; specifically, *Sailfish* [52] and *Shoal++* [51] assign a leader node for every round and allow committing even before BRB instances deliver voting blocks.

## IV. SECURITY AND ATTACK VECTORS

Consensus protocols must satisfy two essential properties: *safety* and *liveness*. *Safety* ensures that all honest nodes agree on a consistent state, free of conflicting information, while

*liveness* guarantees the continuous addition of new transactions to the ledger.

Security assessments of consensus protocols are conducted under specific theoretical assumptions, primarily defined by the network model (see Section II-D) and the adversary model, which is typically described by the proportion of faulty or malicious nodes or weights in the system.

In academic analysis, security properties are proven to hold under these assumptions. Often, specific attack vectors reveal protocol vulnerabilities and help establish tight security bounds, highlighting optimal adversarial strategies. Given the varying theoretical approaches to analyzing protocol security, we categorize only the main attack vectors relevant to DAG-based consensus in the following section.

#### A. Attack Vectors

We focus here on attack vectors specific to the DAG structure in consensus protocols, excluding general attacks like Sybil, DDoS and Spamming, and double-spending attacks, common across blockchain systems.

DAG-based protocols are susceptible to attacks that leverage the following adversarial actions:

- Interfering with network communication (as constrained by the network model) and withholding blocks selectively.
- Sending blocks to only a subset of honest nodes.
- Producing blocks beyond the protocol’s allowed rate (relevant in PoS, committee-based, and permissioned systems, but generally not in PoW protocols).

We categorize attacks by protocol focus—*availability-focused* or *consistency-focused*—and by DAG structure (structured vs. unstructured) and certification type (optimistic vs. certified). Availability-focused protocols are generally more vulnerable to safety attacks, while consistency-focused protocols often encounter liveness issues.

Attacks are valuable in theoretical analysis, as they help determine security bounds, and understanding their impact ensures better security resilience. Below, we summarize several key attack types relevant to DAG-based structures:

*Balance Attack*: A balance attack aims to keep a distributed ledger system in an undecided or “balanced” state, where the network is split between at least two competing subDAGs. This can lead to different outcomes depending on the protocol type: for consistency-focused protocols, the attack causes *liveness* issues by stalling confirmation; for availability-focused protocols, it creates *safety* issues by risking conflicting views.

The balance attack on Conflux [61] demonstrates this vulnerability under high throughput conditions [40]. By maintaining balanced subDAGs, adversaries can manipulate the protocol’s confirmation rules to induce either liveness or safety issues based on the timing and visibility of the attack.

GHOST, although secure at low mining rates [72], is susceptible to balance attacks when mining rates increase [82], [83], effectively limiting its throughput in a manner similar to Bitcoin. Protocols that rely on GHOST, such as Inclusive [31] and Conflux [61], inherit this limitation.

Recent advancements address this challenge in different ways. GHOSTDAG [33] introduces a parameter to handle high throughput, enhancing resilience against balance attacks, while DAG-Knight [34] offers an adaptive solution by dynamically adjusting protocol parameters based on the system’s throughput.

*Parasite-Chain Attack*: In this attack, the adversary constructs a hidden subDAG in parallel to the honest DAG, intending to replace it at some point. This is particularly effective in protocols like IOTA’s Tangle and SPECTRE, which rely on recent tip selection strategies [7], [16].

*Equivocation Attack*: Equivocation involves presenting different information to different nodes, leading to conflicting states. This attack is challenging for deterministic systems that aim for linearizability, as demonstrated in foundational work on Byzantine fault tolerance [69], [84]. Every PoS or permission system is concerned, and there is a natural connection to the balancing attack.

*Liveness Attack*: Liveness attacks target the confirmation or finalization of transactions by selectively slowing down consensus. Leader-based DAG protocols face liveness risks from attacks that delay leader block broadcasts, e.g., [53], or introduce equivocation of the leader blocks, both of which disrupt consensus progression.

*Censorship Attack*: Censorship attacks aim to exclude specific blocks or transactions by withholding blocks from particular nodes, hindering consensus progress. In Avalanche, for example, selective censorship can cause liveness issues, see [85], although recent work has proposed mitigations [18], [86].

*Fork-Bomb Attack*: Initially described in [38], this spam attack forces honest nodes to process excessive data, risking system overload.

*Data-Availability Attack*: In this attack, adversaries withhold blocks from certain parts of the network, risking protocol stalling [38].

Table III provides an overview of these attack vectors across different DAG-based protocols, categorizing protocols analyzed under each type of attack. This categorization does not imply vulnerability but indicates that the protocol has been analyzed for resilience against these attacks. Additionally, the presence of a protocol in a specific category suggests that the respective attack vector may be relevant or of interest, given the protocol’s design.

We also note the absence of attacks for certified DAGs. This stems from the additional use of certificates and reliable broadcast primitives. These external certificates present a possible attack vector, as the block signature flood attack for Tusk studied in [87].

## V. DESIRABLE PROPERTIES

### A. Ordering

Ordering in distributed ledgers refers to the organization of vertices (blocks/transactions). In blockchain, blocks are typically organized linearly, while in DAG-based systems,

TABLE III  
STUDIED ATTACK VECTORS IN DAG-BASED DLTs

	Availability-Focused		Consistency-Focused	
	Structured DAG	Unstructured DAG	Optimistic DAG	Certified DAG
Balance Attack	[31], [36], [40], [42], [61]	[7], [17], [32], [88], [89]		
Parasite-Chain Attack		[7], [16]		
Equivocation Attack	[42]		[6], [23], [37], [48], [90]	
Liveness Attack		[85], [91], [92]	[53]	
Censorship Attack		[85], [91], [92]		
Fork-Bomb Attack	[38]	[38]	[38], [90]	
Data-Availability Attack	[38]	[38]	[38]	

ordering depends on factors like graph structure and the underlying consensus protocol. Ordering in DAG-based DLTs can be classified into two types: Partial order and Total order.

*Partial order* arranges DAG blocks through topological sorting, common in protocols like IOTA [7], Graphchain [64], SPECTRE [16], and Nano [43], where not all blocks are reachable from each other. This approach enables efficient handling of simple payment transactions but limits the use of smart contracts, as shared global states are not inherently supported.

*Total order* organizes blocks in a linear sequence based on parameters such as weight, votes, or timestamp. Some protocols, like GHOST [30], PHANTOM [33], and DAG KNIGHT [34], determine total order as blocks are added to the DAG by constructing a backbone chain. Others, including the consistency focused protocols, calculate order after rounds of leader selection, which leads to a total order. In Slipstream [54], total ordering is achieved during synchronous network conditions without a dedicated leader.

*Fast Path*: To optimize transaction processing, some DAG-based systems combine partial and total ordering. Simple payment transactions, which do not require global consensus, are confirmed through partial ordering for rapid verification. In contrast, transactions needing a consistent global state, such as those in shared smart contracts, use total ordering.

Sui-Lutris [93], Mysticeti-FPC [23], and Slipstream [54], implement this approach with variations. Sui-Lutris includes a subprotocol allowing nodes to confirm owned-object transactions *before* they are included in consensus blocks by constructing explicit transaction certificates. In both Mysticeti-FPC and Slipstream, confirmation of owned-objects and UTXO transactions happens *after* including them in consensus blocks by interpreting the local DAGs.

## B. Fairness

Fairness is a fundamental consideration in consensus protocols within DLT systems, addressing technical, economic, and social dimensions. Unfair conditions can foster dissatisfaction among participants, potentially slowing technology adoption. For DAG-based protocols, fairness often involves accepting contributions from slower but honest nodes, defined as those with lower computing power in PoW-based protocols, lower stakes in PoS-based protocols, or substantial communication delays.

In traditional blockchain protocols, centralized, leader-based designs often challenge fairness where a single leader per round has disproportionate influence over block ordering and transaction inclusion. By contrast, DAG-based consensus protocols distribute the role of advancing the ledger across multiple nodes, reducing individual influence over state progression and enhancing fairness.

Raikwar et al. [94] examine fairness in DAG-based DLTs, focusing on participant inclusion, consensus ordering, and component roles within the protocol. They define various fairness criteria and propose methods to support fair participation across different aspects of the system. We discuss two important topics of fairness in the following two sections.

## C. MEV Protection

Ordering in a consensus protocol can be exploited through Miner Extractable Value (MEV) attacks, where an adversary attempts to include, exclude, or reorder clients' transactions to maximize profit. MEV, initially introduced as a measure in Flash Boys 2.0 [95], was later renamed Maximal Extractable Value to reflect its broader applicability. These attacks allow validators to extract value at the cost of regular users, impacting the fairness and efficiency of the protocol.

In major DLTs like Ethereum, the urgency for MEV protection is well-known. There have been a few studies to countermeasure these MEV attacks. Yang et al. [96] present a SoK on MEV attacks and its countermeasures. A recent paper [97] presents important insights about MEV extraction on Ethereum through MEV bot bidding strategies where a bid is the highest bribe per computation.

MEV attacks can be executed on leader-based DAG protocols, yet they can also occur in unstructured DAG protocols where adversaries can act as block proposers. Fino [98] protocol integrates MEV protection into a BFT-based DAG protocol by employing  $k$ -out-of- $n$  secret-sharing technique. Nasrulin et al. [99] present an accountable base layer protocol, L $\emptyset$ , which detects and mitigates transaction manipulations by creating a secure mempool of verifiable transactions.

## D. Garbage Collection

In DAG-based consensus protocols, ensuring validity and fairness demands nodes to have unlimited (fast) memory, posing challenges for deploying these protocols. Due to the necessity for boundless memory, nodes can't garbage collect old data, risking the loss of honest but slower nodes' blocks.

Garbage collection clashes with fairness, especially within an asynchronous network framework where blocks may experience indefinite delays. This creates a crucial trade-off between ensuring fairness and optimizing performance.

Garbage collection methods vary across DAG-based consensus protocols, aiming to balance the need for memory optimization. The blocks of the DAG can be garbage collected based on their depth (or round), timestamp (or age), or finality. It can also be performed in DAG-based consensus protocols Vite [55] and Jointgraph [37], which maintain snapshot blocks (or chain). Consequently, older blocks can be garbage collected since their history is stored within the snapshot blocks. Jointgraph uses its snapshot and storage events to release the previous memory.

Narwhal [46] cleans up its DAG by discarding information up to a specific round from the genesis. However, this compromises fairness at the block level, risking disposal of slower nodes' data before proper ordering. In contrast, Bullshark [47] employs a garbage collection mechanism while mitigating fairness issues. Bullshark and IOTA 2.0 [56] garbage collect blocks that are not committed for a long period of time, ensuring fairness post-GST during synchronous periods.

## VI. DISCUSSION

### A. Broadcast Techniques

Broadcast techniques in DAG-based consensus protocols ensure the propagation of blocks. The protocols using a certified DAG, see Table II, rely on underlying reliable and consistent broadcast primitives to disseminate blocks, whereas the protocols using an optimistic DAG (this includes, in addition, most availability-focused protocols, see Table II) use best-effort broadcast.

- *Byzantine Consistent Broadcast (BCB)* guarantees that a sender cannot send conflicting blocks (equivocate). It employs a multi-step communication pattern between the sender and the receiver. In DAG-based protocols, e.g. Tusk [46], BCB is treated as an abstract entity. For each block, the block creator calls its own BCB broadcast instance. In [49], [62], an abortable variant of BCB, called BBCA, is introduced that is then used to broadcast only the leader blocks.
- *Byzantine Reliable Broadcast (BRB)* ensures the reliable dissemination of messages, even when a certain number of nodes in the system may be malicious or faulty. BRB requires at least  $f + 1$  honest nodes to unanimously agree on a value  $v$  before committing across all nodes. While BRB adds an extra step compared to BCB, DAG-based protocols, e.g., Bullshark [47], use BRB to ensure the liveness of the protocol.
- *Best-Effort Broadcast (BEB)* facilitates block dissemination optimistically. In this approach, when a node intends to share a block, it simply transmits it with some (potentially unknown) part of its causal history to all other nodes, requiring only a single round of communication instead of at least two rounds when utilizing CB. This approach is

demonstrated in Hashgraph [6] and Cordial Miners [48]. To minimize the communication overhead, many actual implementations of DAG-based protocols do not send the causal history of the block; instead, they use a pull strategy to minimize the communication overhead in the common case.

### B. Latency and Throughput

Network latency and round trip time (RTT) are essential for consensus. Latency is the time for a packet to travel from point A to point B, while RTT encompasses the time for a packet to go from A to B and back, including encoding, queuing, processing, decoding, and propagation delays. These factors, typically consistent for a given pair of endpoints, can be affected by network congestion, adding variability to RTT. It's important to note that latency is not necessarily half of RTT due to potential differences in delays between endpoints.

Consistency-focused consensus protocols based on BFT incur high latency in asynchronous networks. Recent works aim to minimize the latency of these protocols. Spiegelman et al. [50] introduced Shoal, a protocol-agnostic framework that improves the latency of the BFT-based consensus protocols. Shoal [50] enhances the latency of Narwhal-based consensus protocols by employing a leader reputation mechanism to prevent failures and introducing pipelining to ensure a well-ordered DAG construction without timeout requirements.

Liu et al. [100] present a novel DAG-based commit rule that effectively reduces transaction latency in the UTXO model. Their novel commit rule accelerates transaction confirmation and reduces the confirmation latency. Another recent work [101] presents a flexible advancement in asynchronous BFT consensus protocols, bridging the ordering and agreement components and reduces the latency of the consensus.

The primary advantage of DAG-based protocols over blockchains is their enhanced throughput. Various DAG-based consensus protocols demonstrate high throughput. Amores-Sesar and Cachin [102] present a construction that takes a DAG-based consensus protocol  $\Pi$  as input and outputs a new DAG-based protocol  $\Pi'$  which has superior throughput and latency.

### C. Execution Bottlenecks in High-Throughput DAG-Based DLTs

In DAG-based DLTs, the increased throughput shifts the primary bottleneck from data writing and consensus to transaction execution, necessitating an efficient, parallelized execution layer. Prism addresses this by decoupling consensus from transaction validation, achieving high throughput independently of consensus finality [103]. This design emphasizes the need for scalable execution frameworks. Block-STM [104] scales smart contract processing through speculative execution, dynamically resolving conflicts to maintain high throughput. Pilotfish [105], a distributed execution engine, leverages a novel crash-recovery protocol and versioned-queues scheduling to ensure state consistency while maximizing parallelism, even with complex read-write dependencies.

#### D. Mathematical Models and Simulations

A mathematical model can analyze DAG-based consensus protocol metrics, but constructing a generic model to simulate and evaluate these protocols is a challenging task. Some attempts have been made to build such models and simulations.

S. Popov’s initial mathematical model for a DAG-based protocol [7] assumed new message arrivals following a unique Poisson process, leading to a conjecture about the stationary rate of unapproved transactions. Simulations supporting this conjecture with homogeneous delays spurred interest in developing new mathematical models for DAG-based protocols, as seen in studies such as [106], [107]. Works like [108] explored non-Poisson message arrival scenarios, often assuming a central node managing ledger records with other nodes accessing the ledger state through it. Penzkofer et al. [109] extended Popov’s model, introducing heterogeneous delays instead of homogeneous ones. Zander et al. introduced DAGsim [110], a multi-agent simulator based on a heterogeneous delay model for DAG-based protocols. Recent research [111], [112] employed a discrete-time Markov chain to model the evolution of unapproved transactions under heterogeneous delay. Lin et al. developed TangleSim [113] to implement leaderless Tangle 2.0 in various network scenarios and byzantine environments. Bullshark [47].

For the simulation of deterministic protocols, Schett and Danezis [114] formalized a blockDAG protocol implementing a reliable point-to-point channel, embedding deterministic BFT protocols while maintaining safety and liveness properties. They utilized deterministic state machines for node communication, affirming properties claimed by Hashgraph [6], Blockmania [115], and Flare [116] within their framework. Attiya et al. [117] extended this work, faithfully simulating non-deterministic (Randomized) BFT protocols on blockDAG, capturing probabilistic guarantees and enabling analysis of protocols like Aleph [38], DAG-Rider [39], and Bullshark [47].

We want to mention that a series of recent performance evaluations are based on similar testbeds. For instance, the performance of optimistic DAGs is evaluated in [23] (Mysticeti), [53] (Mahi-Mahi), and in [25] (Obelia). The comparison of certified DAG-based consensus protocols such as Bullshark, Shoal, and Sailfish is provided in [52].

#### E. Lightweight DAG-based Protocols

DAG-based protocols provide advantages in performance and scalability over traditional blockchains but encounter challenges related to storage scalability due to data redundancy. This issue increases costs for developers of decentralized applications (dApps), who must efficiently track and verify state changes. This is a general challenge for high-throughput DLTs, and here we provide a brief overview of solutions specifically proposed for DAG-based protocols.

To support diverse applications like Vehicular Social Networks (VSN) and the Internet-of-Things (IoT), lightweight protocols have been developed to address resource constraints. Yang et al. [118] proposed LDV, a DAG-based protocol

optimized for VSNs, which selectively retains necessary information and prunes historical data, making it suitable for resource-limited vehicles. Similarly, Cherupally et al. [119] introduced LSDI, a scalable DAG-based ledger designed for verifying IoT data integrity within cloud-based architectures. LSDI enhances lightweight functionality by pruning outdated segments.

Recently, Dai et al. [120] proposed GeckoDAG, a lightweight DAG protocol that addresses data redundancy issues. GeckoDAG reduces account and reference redundancy by consolidating transactions and introducing *reference override* for efficient reference management. These optimizations reduce storage requirements without compromising security.

### VII. FUTURE RESEARCH

The following presents an overview of emerging discussion points and future research questions.

*Performance Evaluation:* No standardized benchmarks exist for evaluating DAG-based consensus protocol performance. Research papers often focus on specific metrics, emphasizing protocol benefits. Developing standardized performance metrics is crucial for objectively evaluating whether DAG-based approaches mitigate performance bottlenecks in chain-based protocols. Actual protocol implementations, when available, frequently diverge from their theoretical designs due to practical considerations. This discrepancy underscores the importance of real-world testing in validating the effectiveness of DAG-based systems.

*Security Analysis:* Formal security proofs for DAG-based protocols appear sporadically and vary in scope and depth. Developing a comprehensive framework is necessary to understand and compare these protocols’ security features. Furthermore, the confluence of ideas and concepts within these protocols highlights the need for a generalized abstraction. This abstraction would facilitate a universal treatment of security measures and attack vectors, allowing for a more structured analysis of protocol security.

*Incentive Attacks and Game-Theoretic Challenges:* In DAG-based distributed ledgers, traditional blockchain incentive concerns—such as protocol deviations for selfish gain—remain underexplored. Analyzing these challenges in the context of availability- and consistency-focused protocols could reveal generalized principles for ensuring fair participation. For instance, availability-focused systems might require incentive mechanisms that prevent balancing attacks, while consistency-focused systems may benefit from mechanisms that discourage strategic manipulations of block ordering and equivocations.

*Transaction duplicates:* Transaction duplicates pose a challenge for DAG-based protocols to enhance system efficiency. The issue arises when multiple validators pull transactions from a shared mempool for parallel block proposals, risking transaction duplication and negating parallel execution benefits. A solution like assigning transactions to specific validators could mitigate repetition but potentially enable censorship and centralized control, undermining decentralization principles. Some DAG-based projects address it through ad hoc methods

as the random association between transactions and nodes. However, there is a need for research focused on strategies to prevent transaction duplication without sacrificing system fairness or decentralization.

*Fairness:* Research on fairness within DAG-based protocols should assess the implications of more open writing access on transaction equity, particularly in the Maximal Extractable Value (MEV) context. It is essential to consider how the total ordering of transactions, despite more open writing access, can influence MEV opportunities. This inquiry should extend to the role of entry points, the underlying peer-to-peer (P2P) and Remote Procedure Call (RPC) nodes, which are critical in MEV dynamics. Investigating these aspects will provide insight into whether decentralization in writing access successfully mitigates fairness issues. Moreover, it is essential to understand the relationship between access control, tokenomics, and consensus mechanisms to understand fairness and evaluate the level of decentralization.

*Privacy:* DAG-based consensus protocols should prioritize constructing privacy-preserving mechanisms that leverage higher bandwidth in DAG structures. Implementing privacy in DAG-based protocols introduces challenges. The tip selection process may become costlier as validating transactions in the tips requires computational effort due to cryptographic privacy measures on transaction data encoded within the tips. This, alongside concerns about transaction ordering and transaction auditability in the DAG, underscores the significance of selecting the appropriate privacy technique. Therefore, efforts should be directed at developing protocols that protect identities and transaction data without sacrificing performance, exploring ways to balance confidentiality with the enhanced capacity of the DAG network.

*Mempool Abstraction:* Mempool abstraction allows decoupling transaction dissemination from the consensus protocol, streamlining the ordering process for better efficiency. The selection of transactions can be optimized to prevent redundancy upon their block inclusion. Additionally, leveraging load-balancing protocols for distribution across nodes can contribute to equitable transaction handling. Consequently, further investigation is required to explore how mempool abstraction affects fairness within DAG-based consensus protocols and its implications for the cost of transaction ordering.

## VIII. CONCLUSION

In this SoK, we analyzed DAG-based consensus protocols, classifying them by availability and consistency requirements. We outlined key models, components, security aspects, potential attack vectors, and relevant countermeasures. The study also highlighted recent developments in execution scalability and fairness, reflecting the unique demands of high-throughput DLTs.

## IX. CALL FOR CONTRIBUTIONS AND FUTURE UPDATES

While we have aimed to provide a comprehensive overview of relevant works in DAG-based consensus protocols, some research may have been unintentionally omitted. We encourage

researchers to contact us if any important work has been missed. We also intend to keep this study up-to-date as long as there is interest. We invite you to share new developments with us or cite this work to facilitate discovery.

## REFERENCES

- [1] A. Kiayias and G. Panagiotakos, "Speed-security tradeoffs in blockchain protocols," *Cryptology ePrint Archive*, 2015.
- [2] S. Wu, P. Wei, R. Zhang, and B. Jiang, "Security-Performance Tradeoff in DAG-based Proof-of-Work Blockchain Protocols," *Cryptology ePrint Archive*, 2023.
- [3] NeonVest, "The scalability trilemma in blockchain," 2018, accessed: 23.08.2023. [Online]. Available: <https://aakash-111.medium.com/the-scalability-trilemma-in-blockchain-75fb57f646d1>
- [4] J. Kalajdjieski, M. Raikwar, N. Arsov, G. Velinov, and D. Gligoroski, "Databases fit for blockchain technology: A complete overview," *Blockchain: Research and Applications*, p. 100116, 2022.
- [5] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Decentralized business review*, p. 21260, 2008.
- [6] L. Baird and A. Luykx, "The hashgraph protocol: Efficient asynchronous BFT for high-throughput distributed ledgers," in *2020 International Conference on Omni-layer Intelligent Systems (COINS)*. IEEE, 2020, pp. 1–7.
- [7] S. Popov, "The tangle," *White paper*, vol. 1, no. 3, p. 30, 2018.
- [8] Q. Wang, J. Yu, S. Chen, and Y. Xiang, "Sok: Dag-based blockchain systems," *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–38, 2023.
- [9] X. Lu, C. Jiang, and P. Wang, "A survey on consensus algorithms of blockchain based on dag," in *Proceedings of the 2024 6th Blockchain and Internet of Things Conference*, ser. BIOTC '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 50–58. [Online]. Available: <https://doi.org/10.1145/3688225.3688232>
- [10] B. Bellaj, A. Ouaddah, E. Bertin, N. Crespi, and A. Mezrioui, "SOK: a comprehensive survey on distributed ledger technologies," in *2022 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, 2022, pp. 1–16.
- [11] N. Kannengießer, S. Lins, T. Dehling, and A. Sunyaev, "Mind the gap: Trade-offs between distributed ledger technology characteristics," *arXiv preprint arXiv:1906.00861*, 2019.
- [12] H. Y. Wu, X. Yang, C. Yue, H.-Y. Paik, and S. S. Kanhere, "Chain or dag? underlying data structures, architectures, topologies and consensus in distributed ledger technology: A review, taxonomy and research issues," *Journal of Systems Architecture*, vol. 131, p. 102720, 2022.
- [13] R. Guerraoui, P. Kuznetsov, M. Monti, M. Pavlović, and D.-A. Seredinschi, "The consensus number of a cryptocurrency," in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, 2019, pp. 307–316.
- [14] M. Baudet, G. Danezis, and A. Sonnino, "Fastpay: High-performance byzantine fault tolerant settlement," in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 2020, pp. 163–177.
- [15] A. Lewis-Pye, O. Naor, and E. Shapiro, "Flash: An asynchronous payment system with good-case linear communication complexity," *arXiv preprint arXiv:2305.03567*, 2023.
- [16] Y. Sompolinsky, Y. Lewenberg, and A. Zohar, "Spectre: A fast and scalable cryptocurrency protocol," *Cryptology ePrint Archive*, 2016.
- [17] S. Müller, A. Penzkofer, N. Polyanskii, J. Theis, W. Sanders, and H. Moog, "Tangle 2.0 leaderless nakamoto consensus on the heaviest dag," *IEEE Access*, vol. 10, pp. 105 807–105 842, 2022.
- [18] T. Rocket, M. Yin, K. Sekniqi, R. van Renesse, and E. G. Sirer, "Scalable and probabilistic leaderless bft consensus through metastability," *arXiv preprint arXiv:1906.08936*, 2019.
- [19] M. M. T. Chakravarty, J. Chapman, K. MacKenzie, O. Melkonian, M. Peyton Jones, and P. Wadler, "The Extended UTXO Model," in *Financial Cryptography and Data Security*, M. Bernhard, A. Bracciali, L. J. Camp, S. Matsuo, A. Maurushat, P. B. Rønne, and M. Sala, Eds. Cham: Springer International Publishing, 2020, pp. 525–539.
- [20] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, no. 2014, pp. 1–32, 2014.
- [21] S. Müller, A. Penzkofer, N. Polyanskii, J. Theis, W. Sanders, and H. Moog, "Reality-based utxo ledger," *Distrib. Ledger Technol.*, vol. 2, no. 3, sep 2023. [Online]. Available: <https://doi.org/10.1145/3616022>



- [22] The Mysten Labs Team, “The sui smart contracts platform,” 2022, accessed: 22.08.2023. [Online]. Available: <https://github.com/MystenLabs/sui/blob/main/doc/paper/sui.pdf>
- [23] K. Babel, A. Chursin, G. Danezis, L. Kokoris-Kogias, and A. Sonnino, “Mysticeti: Low-latency dag consensus with fast commit path,” 2023.
- [24] A. Lewis-Pye and T. Roughgarden, “Permissionless consensus,” *arXiv preprint arXiv:2304.14701*, 2023.
- [25] G. Danezis, L. Kokoris-Kogias, A. Sonnino, and M. Tian, “Obelia: Scaling dag-based blockchains to hundreds of validators,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.08701>
- [26] J. Neu, E. N. Tas, and D. Tse, “The availability-accountability dilemma and its resolution via accountability gadgets,” in *International Conference on Financial Cryptography and Data Security*. Springer, 2022, pp. 541–559.
- [27] R. Pass and E. Shi, “The sleepy model of consensus,” in *Advances in Cryptology—ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II 23*. Springer, 2017, pp. 380–409.
- [28] G. Bracha, “Asynchronous byzantine agreement protocols,” *Information and Computation*, vol. 75, no. 2, pp. 130–143, 1987.
- [29] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process,” *J. ACM*, vol. 32, no. 2, p. 374–382, apr 1985. [Online]. Available: <https://doi.org/10.1145/3149.214121>
- [30] Y. Sompolinsky and A. Zohar, “Secure high-rate transaction processing in bitcoin,” in *Financial Cryptography and Data Security: 19th International Conference, FC 2015, San Juan, Puerto Rico, January 26-30, 2015, Revised Selected Papers 19*. Springer, 2015, pp. 507–527.
- [31] Y. Lewenberg, Y. Sompolinsky, and A. Zohar, “Inclusive block chain protocols,” in *Financial Cryptography and Data Security: 19th International Conference, FC 2015, San Juan, Puerto Rico, January 26-30, 2015, Revised Selected Papers 19*. Springer, 2015, pp. 528–547.
- [32] A. Churyumov, “Byteball: A decentralized system for storage and transfer of value,” *URL https://byteball.org/Byteball.pdf*, p. 11, 2016.
- [33] Y. Sompolinsky, S. Wyborski, and A. Zohar, “Phantom ghostdag: a scalable generalization of nakamoto consensus: September 2, 2021,” in *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, 2021, pp. 57–70.
- [34] Y. Sompolinsky and M. Sutton, “The DAG KNIGHT Protocol: A Parameterless Generalization of Nakamoto Consensus,” *Cryptology ePrint Archive*, 2022.
- [35] V. Bagaria, S. Kannan, D. Tse, G. Fanti, and P. Viswanath, “Prism: Deconstructing the blockchain to approach physical limits,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’19. New York, NY, USA: Association for Computing Machinery, 2019, p. 585–602.
- [36] C. Li, F. Long, and G. Yang, “Ghast: Breaking confirmation delay barrier in nakamoto consensus via adaptive weighted blocks,” 2020.
- [37] F. Xiang, W. Huaimin, S. Peichang, O. Xue, and Z. Xunhui, “Joint-graph: A dag-based efficient consensus algorithm for consortium blockchains,” *Software: Practice and Experience*, vol. 51, no. 10, pp. 1987–1999, 2021.
- [38] A. Gagol, D. Leśniak, D. Straszak, and M. Świątek, “Aleph: Efficient atomic broadcast in asynchronous networks with byzantine nodes,” in *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, 2019, pp. 214–228.
- [39] I. Keidar, E. Kokoris-Kogias, O. Naor, and A. Spiegelman, “All you need is dag,” in *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, 2021, pp. 165–175.
- [40] H. Yu, I. Nikolić, R. Hou, and P. Saxena, “Ohie: Blockchain scaling made simple,” in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 90–105.
- [41] M. Fitz, P. Ga, A. Kiayias, and A. Russell, “Parallel chains: Improving throughput and latency of blockchain protocols via parallel composition,” *Cryptology ePrint Archive*, 2018.
- [42] X. Wang, S. Azouvi, and M. Vukolić, “Security analysis of filecoin’s expected consensus in the byzantine vs honest model,” *arXiv preprint arXiv:2308.06955*, 2023.
- [43] C. LeMahieu, “Nano: A feeless distributed cryptocurrency network,” *Nano [Online resource]*. URL: <https://nano.org/en/whitepaper> (date of access: 24.03. 2018), vol. 16, p. 17, 2018.
- [44] W. Martino, M. Quaintance, and S. Popejoy, “Chainweb: A proof-of-work parallel-chain architecture for massive throughput,” *Chainweb whitepaper*, vol. 19, 2018.
- [45] I. Bentov, P. Hubáček, T. Moran, and A. Nadler, “Tortoise and hares consensus: the meshcash framework for incentive-compatible, scalable cryptocurrencies,” in *Cyber Security Cryptography and Machine Learning: 5th International Symposium, CSCML 2021, Be’er Sheva, Israel, July 8–9, 2021, Proceedings 5*. Springer, 2021, pp. 114–127.
- [46] 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*, 2022, pp. 34–50.
- [47] A. Spiegelman, N. Girdharan, A. Sonnino, and L. Kokoris-Kogias, “Bullshark: DAG BFT protocols made practical,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 2705–2718.
- [48] I. Keidar, O. Naor, and E. Shapiro, “Cordial miners: A family of simple, efficient and self-contained consensus protocols for every eventuality,” *arXiv preprint arXiv:2205.09174*, 2022.
- [49] D. Malkhi, C. Stathakopoulou, and M. Yin, “BBCA-CHAIN: One-Message, Low Latency BFT Consensus on a DAG,” *arXiv preprint arXiv:2310.06335*, 2023.
- [50] A. Spiegelman, B. Aurn, R. Gelashvili, and Z. Li, “Shoal: Improving DAG-BFT Latency And Robustness,” *arXiv preprint arXiv:2306.03058*, 2023.
- [51] B. Arun, Z. Li, F. Suri-Payer, S. Das, and A. Spiegelman, “Shoal++: High throughput dag bft can be fast!” *arXiv preprint arXiv:2405.20488*, 2024.
- [52] N. Shrestha, R. Shrothrium, A. Kate, and K. Nayak, “Sailfish: Towards improving latency of dag-based bft,” *Cryptology ePrint Archive*, 2024.
- [53] P. Jovanovic, L. K. Kogias, B. Kumara, A. Sonnino, P. Tennage, and I. Zabolotchi, “Mahi-mahi: Low-latency asynchronous bft dag-based consensus,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.08670>
- [54] N. Polyanski, S. Muller, and M. Raikwar, “Slipstream: Ebb-and-flow consensus on a dag with fast confirmation for utxo transactions,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.14876>
- [55] C. Liu, D. Wang, and M. Wu, “Vite: A high performance asynchronous decentralized application platform,” *White Paper*, 2018.
- [56] IOTA Foundation, “Consensus on a DAG,” 2023, accessed: 19.11.2023. [Online]. Available: <https://wiki.iota.org/learn/protocols/iota2.0/core-concepts/consensus/introduction/>
- [57] E. A. Brewer, “Towards robust distributed systems,” in *PODC*, vol. 7, no. 10.1145. Portland, OR, 2000, pp. 343 477–343 502.
- [58] E. Brewer, “Cap twelve years later: How the “rules” have changed,” *Computer*, vol. 45, no. 2, pp. 23–29, 2012.
- [59] L. Baird, M. Harmon, and P. Madsen, “Hedera: A governing council & public hashgraph network,” *The trust layer of the internet, whitepaper*, vol. 1, pp. 1–97, 2018.
- [60] Y. Sompolinsky and A. Zohar, “Phantom,” *IACR Cryptology ePrint Archive, Report 2018/104*, 2018.
- [61] C. Li, P. Li, D. Zhou, Z. Yang, M. Wu, G. Yang, W. Xu, F. Long, and A. C.-C. Yao, “A decentralized blockchain with high throughput and fast confirmation,” in *2020 {USENIX} Annual Technical Conference ({USENIX}{ATC} 20)*, 2020, pp. 515–528.
- [62] C. Stathakopoulou, M. Wei, M. Yin, H. Zhang, and D. Malkhi, “BBCA-LEDGER: High Throughput Consensus meets Low Latency,” *arXiv preprint arXiv:2306.14757*, 2023.
- [63] B. Arun, Z. Li, F. Suri-Payer, S. Das, and A. Spiegelman, “Shoal++: High throughput dag bft can be fast!” 2024. [Online]. Available: <https://arxiv.org/abs/2405.20488>
- [64] X. Boyen, C. Carr, and T. Haines, “Graphchain: A blockchain-free scalable decentralised ledger,” in *Proceedings of the 2nd ACM Workshop on Blockchains, Cryptocurrencies, and Contracts*, 2018, pp. 21–33.
- [65] N. Shrestha, A. Kate, and K. Nayak, “Sailfish: Towards improving latency of dag-based bft,” *IACR Cryptol. ePrint Arch.*, vol. 2024, p. 472, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:268678908>
- [66] S. Buttolph, A. Moin, K. Sekniqi, and E. G’un Sirer, “Avalanche native token (\$avax) dynamics,” 2020, accessed: 14.11.2023. [Online]. Available: <https://assets.website-files.com/5d80307810123f5ffb34d6e/6008d7bc56430d6b87291>
- [67] B. David, P. Gaži, A. Kiayias, and A. Russell, “Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain,” in

- Advances in Cryptology—EUROCRYPT 2018: 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29–May 3, 2018 Proceedings, Part II* 37. Springer, 2018, pp. 66–98.
- [68] M. Castro, B. Liskov *et al.*, “Practical byzantine fault tolerance,” in *OsDI*, vol. 99, no. 1999, 1999, pp. 173–186.
- [69] L. Lamport, R. Shostak, and M. Pease, “The byzantine generals problem,” *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 4, no. 3, pp. 382–401, 1982.
- [70] R. Pass and E. Shi, “Fruitchains: A fair blockchain,” in *Proceedings of the ACM symposium on principles of distributed computing*, 2017, pp. 315–324.
- [71] I. Eyal, A. E. Gencer, E. G. Sirer, and R. Van Renesse, “{Bitcoin-NG}: A scalable blockchain protocol,” in *13th USENIX symposium on networked systems design and implementation (NSDI 16)*, 2016, pp. 45–59.
- [72] A. Kiayias and G. Panagiotakos, “On trees, chains and fast transactions in the blockchain,” in *Progress in Cryptology—LATINCRYPT 2017: 5th International Conference on Cryptology and Information Security in Latin America, Havana, Cuba, September 20–22, 2017, Revised Selected Papers 5*. Springer, 2019, pp. 327–351.
- [73] W. V. Gehrlein, “Condorcet’s paradox,” *Theory and Decision*, vol. 15, no. 2, pp. 161–197, 1983.
- [74] C. Li, P. Li, D. Zhou, W. Xu, F. Long, and A. Yao, “Scaling nakamoto consensus to thousands of transactions per second,” *arXiv preprint arXiv:1805.03870*, 2018.
- [75] J. Neu, E. N. Tas, and D. Tse, “Ebb-and-flow protocols: A resolution of the availability-finality dilemma,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 446–465.
- [76] S. Popov and S. Müller, “Voting-based probabilistic consensus and their applications in distributed ledgers,” *Annals of Telecommunications*, vol. 77, no. 1, pp. 77–99, 2022. [Online]. Available: <https://doi.org/10.1007/s12243-021-00875-7>
- [77] I. Amores-Sesar, C. Cachin, and E. Tedeschi, “When is spring coming? a security analysis of avalanche consensus,” 2022.
- [78] C. Cachin, K. Kursawe, and V. Shoup, “Random oracles in constant-time: practical asynchronous byzantine agreement using cryptography,” in *Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing*, 2000, pp. 123–132.
- [79] G. Bracha and S. Toueg, “Asynchronous consensus and broadcast protocols,” *Journal of the ACM (JACM)*, vol. 32, no. 4, pp. 824–840, 1985.
- [80] M. Abd-El-Malek, G. R. Ganger, G. R. Goodson, M. K. Reiter, and J. J. Wylie, “Fault-scalable byzantine fault-tolerant services,” *ACM SIGOPS Operating Systems Review*, vol. 39, no. 5, pp. 59–74, 2005.
- [81] B. Ford, “Threshold logical clocks for asynchronous distributed coordination and consensus,” *arXiv preprint arXiv:1907.07010*, 2019.
- [82] L. Kiffer, R. Rajaraman, and a. shelat, “A better method to analyze blockchain consistency,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 729–744. [Online]. Available: <https://doi.org/10.1145/3243734.3243814>
- [83] C. Natoli and V. Gramoli, “The balance attack against proof-of-work blockchains: The r3 testbed as an example,” 2016. [Online]. Available: <https://arxiv.org/abs/1612.09426>
- [84] B.-G. Chun, P. Maniatis, S. Shenker, and J. Kubiatowicz, “Attested append-only memory: Making adversaries stick to their word,” *ACM SIGOPS Operating Systems Review*, vol. 41, no. 6, pp. 189–204, 2007.
- [85] I. Amores-Sesar, C. Cachin, and P. Schneider, “An analysis of avalanche consensus,” in *Structural Information and Communication Complexity: 31st International Colloquium, SIROCCO 2024, Vietri Sul Mare, Italy, May 27–29, 2024, Proceedings*. Berlin, Heidelberg: Springer-Verlag, 2024, p. 27–44. [Online]. Available: [https://doi.org/10.1007/978-3-031-60603-8\\_2](https://doi.org/10.1007/978-3-031-60603-8_2)
- [86] A. Buchwald, S. Buttolph, A. Lewis-Pye, P. O’Grady, and K. Sekniqi, “Frosty: Bringing strong liveness guarantees to the snow family of consensus protocols,” 2024. [Online]. Available: <https://arxiv.org/abs/2404.14250>
- [87] G. Giuliani, A. Sonnino, M. Frei, F. Streun, L. Kokoris-Kogias, and A. Perrig, “An empirical study of consensus protocols’ dos resilience,” in *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security*, ser. ASIA CCS ’24. New York, NY, USA: Association for Computing Machinery, 2024, p. 1345–1360. [Online]. Available: <https://doi.org/10.1145/3634737.3656997>
- [88] S. Popov, H. Moog, D. Camargo, A. Capossele, V. Dimitrov, A. Gal, A. Greve, B. Kusmierz, S. Mueller, A. Penzkofer *et al.*, “The coordinate,” *Accessed Jan*, pp. 1–30, 2020.
- [89] L. Kovalchuk, M. Rodinko, and R. Oliynykov, “Upper bound probability of double spend attack on spectre,” in *Proceedings of the 3rd Workshop on Cryptocurrencies and Blockchains for Distributed Systems*, 2020, pp. 18–22.
- [90] A. Chursin, “Adelie: Detection and prevention of byzantine behaviour in dag-based consensus protocols,” 2024. [Online]. Available: <https://arxiv.org/abs/2408.02000>
- [91] P. Ferraro, A. Penzkofer, C. King, and R. Shorten, “Feedback control for distributed ledgers: An attack mitigation policy for dag-based dlts,” *IEEE Transactions on Automatic Control*, vol. 69, no. 8, pp. 5492–5499, 2024.
- [92] D. Camargo, A. Penzkofer, S. Müller, and W. Sanders, “Mitigation of liveness attacks in dag-based ledgers,” in *2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2023, pp. 1–9.
- [93] S. Blackshear, A. Chursin, G. Danezis, A. Kichidis, L. Kokoris-Kogias, X. Li, M. Logan, A. Menon, T. Nowacki, A. Sonnino *et al.*, “Sui lutris: A blockchain combining broadcast and consensus,” *arXiv preprint arXiv:2310.18042*, 2023.
- [94] M. Raikwar, N. Polyanskii, and S. Müller, “Fairness Notions in DAG-Based DLTs,” in *2023 5th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*, 2023, pp. 1–8.
- [95] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, “Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability,” in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 910–927.
- [96] S. Yang, F. Zhang, K. Huang, X. Chen, Y. Yang, and F. Zhu, “SoK: MEV countermeasures: Theory and Practice,” *arXiv preprint arXiv:2212.05111*, 2022.
- [97] C. Raun, B. Estermann, L. Zhou, K. Qin, R. Wattenhofer, A. Gervais, and Y. Wang, “Leveraging machine learning for bidding strategies in miner extractable value (mev) auctions,” *Cryptology ePrint Archive, Paper 2023/1281*, 2023. [Online]. Available: <https://eprint.iacr.org/2023/1281>
- [98] D. Malkhi and P. Szalachowski, “Maximal extractable value (mev) protection on a dag,” *arXiv preprint arXiv:2208.00940*, 2022.
- [99] B. Nasrulin, G. Ishmaev, J. Decouchant, and J. Pouwelse, “Lo: An accountable mempool for mev resistance,” in *Proceedings of the 24th International Middleware Conference*, ser. Middleware ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 98–110. [Online]. Available: <https://doi.org/10.1145/3590140.3629108>
- [100] K. Liu, M. Jourenko, and M. Larangeira, “Reducing Latency of DAG-based Consensus in the Asynchronous Setting via the UTXO Model,” *arXiv preprint arXiv:2307.15269*, 2023.
- [101] S. Liu, W. Xu, C. Shan, X. Yan, T. Xu, B. Wang, L. Fan, F. Deng, Y. Yan, and H. Zhang, “Flexible advancement in asynchronous bft consensus,” in *Proceedings of the 29th Symposium on Operating Systems Principles*, 2023, pp. 264–280.
- [102] I. Amores-Sesar and C. Cachin, “We will dag you,” *arXiv preprint arXiv:2311.03092*, 2023.
- [103] G. Wang, S. Wang, V. Bagaria, D. Tse, and P. Viswanath, “Prism removes consensus bottleneck for smart contracts,” in *2020 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 2020, pp. 68–77.
- [104] R. Gelashvili, A. Spiegelman, Z. Xiang, G. Danezis, Z. Li, D. Malkhi, Y. Xia, and R. Zhou, “Block-stm: Scaling blockchain execution by turning ordering curse to a performance blessing,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.06871>
- [105] Q. Knip, L. Kokoris-Kogias, A. Sonnino, I. Zabolochi, and N. Zhang, “Pilotfish: Distributed transaction execution for lazy blockchains,” 2024. [Online]. Available: <https://arxiv.org/abs/2401.16292>
- [106] B. Kusmierz, W. Sanders, A. Penzkofer, A. Capossele, and A. Gal, “Properties of the tangle for uniform random and random walk tip selection,” in *2019 IEEE International Conference on Blockchain (Blockchain)*. IEEE, 2019, pp. 228–236.
- [107] S. Park, S. Oh, and H. Kim, “Performance analysis of dag-based

- cryptocurrency,” in *2019 IEEE International Conference on Communications workshops (ICC workshops)*. IEEE, 2019, pp. 1–6.
- [108] Y. Li, B. Cao, M. Peng, L. Zhang, L. Zhang, D. Feng, and J. Yu, “Direct acyclic graph-based ledger for internet of things: Performance and security analysis,” *IEEE/ACM Transactions on Networking*, vol. 28, no. 4, pp. 1643–1656, 2020.
- [109] A. Penzkofer, O. Saa, and D. Dziubałowska, “Impact of delay classes on the data structure in iota,” in *International Workshop on Data Privacy Management*. Springer, 2021, pp. 289–300.
- [110] M. Zander, T. Waite, and D. Harz, “Dagsim: Simulation of dag-based distributed ledger protocols,” *ACM SIGMETRICS Performance Evaluation Review*, vol. 46, no. 3, pp. 118–121, 2019.
- [111] N. Kumar, A. Reiffers-Masson, I. Amigo, and S. Ruano Rincón, “The effect of network delays on Distributed Ledgers based on Direct Acyclic Graphs: A mathematical model,” *Available at SSRN 4253421*, 2022.
- [112] S. Müller, I. Amigo, A. Reiffers-Masson, and S. Ruano-Rincón, “Stability of local tip pool sizes,” 2023. [Online]. Available: <https://arxiv.org/abs/2302.01625>
- [113] B.-Y. Lin, D. Dziubałowska, P. Macek, A. Penzkofer, and S. Müller, “Tanglesim: An agent-based, modular simulator for dag-based distributed ledger technologies,” in *2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, 2023, pp. 1–5.
- [114] M. A. Schett and G. Danezis, “Embedding a deterministic bft protocol in a block dag,” in *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, 2021, pp. 177–186.
- [115] G. Danezis and D. Hrycyszyn, “Blockmania: from block dags to consensus,” *arXiv preprint arXiv:1809.01620*, 2018.
- [116] N. S. Rowan and N. Usher, “The flare consensus protocol: Fair fast federated byzantine agreement consensus,” Tech. rep, Tech. Rep, Tech. Rep., 2019.
- [117] H. Attiya, C. Enea, and S. Nassar, “Faithful Simulation of Randomized BFT Protocols on Block DAGs,” *Cryptology ePrint Archive*, 2023.
- [118] W. Yang, X. Dai, J. Xiao, and H. Jin, “Ldv: A lightweight dag-based blockchain for vehicular social networks,” *IEEE Transactions on Vehicular Technology*, vol. 69, no. 6, pp. 5749–5759, 2020.
- [119] S. R. Cherupally, S. Boga, P. Podili, and K. Kataoka, “Lightweight and scalable dag based distributed ledger for verifying iot data integrity,” in *2021 International Conference on Information Networking (ICOIN)*. IEEE, 2021, pp. 267–272.
- [120] X. Dai, Y. Zhou, J. Xiao, F. Cheng, X. Xie, H. Jin, and B. Li, “Geckodag: Towards a lightweight dag-based blockchain via reducing data redundancy,” in *2023 IEEE 43rd International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2023, pp. 451–462.