



Towards Formal Verification of DAG-Based Blockchain Consensus Protocols

Nathalie Bertrand  



Univ Rennes, Inria, CNRS, IRISA

Pranav Ghorpade  



The University of Sydney

Sasha Rubin  

The University of Sydney

Bernhard Scholz  

Fantom Research

Pavle Subotić  

Fantom Research

Abstract

There is a trend in blockchains to switch to DAG-based consensus protocols to decrease their energy footprint and improve security. A DAG-based consensus protocol orders transactions for delivering blocks, and relies on built-in fault tolerance communications via Byzantine Atomic Broadcasts. The ubiquity and strategic importance of blockchains call for formal proof of their correctness. We formalize the DAG-based consensus protocol called DAG-Rider in TLA+ and prove its safety properties with the TLA+ proof system. The formalization requires a refinement approach for modelling the consensus. In an abstracted model, we first show the safety of DAG-based consensus on leaders and then further refine the specification to encompass all messages for all processes. The specification consists of 683 lines and the proof system verifies 1922 obligations in about 5 minutes.

2012 ACM Subject Classification Theory of computation \rightarrow Logic and verification; Theory of computation \rightarrow Proof theory; Theory of computation \rightarrow Distributed algorithms

Keywords and phrases Formal verification, Consensus, Blockchain, Theorem Proving, TLA+

Consensus and DAG-based protocols. Consensus is a fundamental problem in distributed computing. It aims to coordinate processes so that they agree on some value(s). Consensus algorithms have recently become an important topic in “proof-of-stake” blockchains that collaboratively build an order for submitted transactions. Of particular interest are consensus algorithms that assume little about the environment, namely, asynchronous communications with malicious processes, namely Byzantine Fault Tolerant (BFT) [13].

Early blockchain consensus protocols assume degrees of synchrony in the environment to ensure safety and liveness [14, 2, 9, 6]. Recently, a family of probabilistic asynchronous consensus protocols have been introduced that are based on Directed Acyclic Graphs (DAG-based protocols) [8, ?, 5]. These protocols report high performance while guaranteeing BFT, utilize processes fairly, and exhibit low communication complexity. Several leading blockchains thus have adopted DAG-based protocols as their main consensus mechanism [1, 4, 7].

DAG-Rider [8] is such a DAG-based protocol and has two main components: (1) a communication layer and (2) an offline ordering layer. The communication layer asynchronously exchanges messages between processes in rounds using *reliable broadcast*. Messages contain transaction proposals and metadata forming a DAG for each node. For a process, the DAG provides a local view of the order of blocks with respect to happened-before relation [10]. Due to the asynchronous nature of the network, processes do not necessarily have the same local DAGs at any point in time. However they are guaranteed to have same DAGs eventually. The ordering layer selects anchor points, guaranteeing consistent selection across all the processes. This allows the DAGs to be locally totally ordered while guaranteeing that all the

processes agree on the same total order of messages.

Formal verification of DAG-based consensus. Blockchains provide mission-critical financial services and hence require rigour to show correctness. The verification challenges arise from the large number of possible interleaving in an asynchronous environment, the behaviours of Byzantine processes, and perhaps even more importantly the fact that correctness should hold for any number of participating processes.

We report here on our TLA+ [11] specification and proof –both publicly available at [?]¹– for a DAG-based consensus protocol using the TLA+ Proof System (TLA-PS) [3].

Procedural code is commonly modeled in TLA+ by a discrete transition system whose traces correspond to possible executions of the code. The naïve translation from the pseudo-code (by setting every variable from the protocol, including a variable for each process’s current line number, to be a variable in the specification) into a TLA+ specification is not viable. While direct, this model is very fine-grained and renders the proofs extremely tedious.

To obtain a more succinct and tractable model, we employ several abstraction techniques: they remove unnecessary details and produce a specification that is amenable to proofs. First, we employ a *procedural abstraction* that ignores all states that are internal to a procedure and only represents the input/output behaviour of each procedure in the DAG-Rider protocol. For instance, in the *wave_ready* procedure of [8], the relevant variables are *decidedWave*, *deliveredVertices*, *leadersStack*, but not the loop variable w' or the auxiliary variable v' . Second, because we focus on safety properties, we remove component features that are only required for liveness and have no impact on the safety proof. For instance, random coin tosses can be replaced with deterministic ones. Third, we use memoization to efficiently compute the values taken by recursive functions, by introducing a fresh state variable that stores the needed information to evaluate recursive functions in a single step. Finally, we separate the concerns and break the safety property into two, namely (1) consistent communication and (2) consistent leader election. For (1) we model the DAG-construction and show that the causal histories agree for a same vertex in the DAG of two different processes¹. For (2), we model the consensus protocol and prove that the same leaders are elected and in the same order. To obtain a complete yet simple model of the consensus protocol, we observe that it only needs reachability information associated with wave leader vertices to commit leaders and, therefore, abstract the content of DAG into the so-called *leaderReachability* record. We combine consensus protocol specifications in DAG construction specifications to obtain one of the DAG-Rider protocols. This abstraction is not only interesting for DAG-Rider but could be helpful to generalize to other DAG-based protocols.

Given our faithful specification of DAG-Rider in TLA+, we prove its expected safety properties by identifying invariants and proving them within TLA-PS. When using TLA-PS and similar proof systems, the most challenging task is to come up with relevant inductive invariants (that hold initially and are preserved when taking transitions), see for instance [12]. For DAG-Rider, to prove the consistency of communication during the DAG construction we identified 6 new invariants, and to prove the consistency of leader election we identified 10 new invariants. We prove each one of the invariants hierarchically by induction.

Table 1 provides some metrics on our experiments, showing quite reasonable performances in terms of verification time. Most importantly, due to the modularity of our specification, we argue the effort to adapt proofs is minimal when making small changes to the specification.

Conclusion. Our work on DAG-Rider is an important and promising step towards a general library for specifying and verifying DAG-based consensus protocols. Beyond the

¹ The non equivocation of blocks is guaranteed by reliable broadcast abstraction

■ **Table 1** Summary of experiments. An obligation is a condition that TLA-PS checks. The time to check is on a 2.10 GHz CPU with 8 GB of memory, running Windows 11 and TLA-PS v1.4.5.

Metric	DAG-Constr. Spec.	Consensus Spec.	DAG-Rider Spec.
Size of spec. (# loc)	460	250	710
Size of proof (# loc)	521	782	1303
Max level of proof tree nodes	10	9	10
Max degree of proof tree nodes	7	7	7
# obligations in TLA-PS	722	1205	1927
Time to check by TLA-PS (s)	224	87	311

specification of DAG-Rider, our specification reveals interesting insights into developing a modular and efficient TLA+ specification that is amenable to proofs in TLA-PS.

References

- 1 Aptos Foundation. Understanding Aptos: A comprehensive overview, 2024. URL: <https://messari.io/report/understanding-aptos-a-comprehensive-overview>.
- 2 Vitalik Buterin. Ethereum white paper: A next generation smart contract & decentralized application platform, 2013. URL: <https://github.com/ethereum/wiki/wiki/White-Paper>.
- 3 Denis Cousineau, Damien Doligez, Leslie Lamport, Stephan Merz, Daniel Ricketts, and Hernán Vanzetto. TLA + proofs. In *Proceedings of FM 2012*, volume 7436 of *Lecture Notes in Computer Science*, pages 147–154. Springer, 2012. doi:10.1007/978-3-642-32759-9_14.
- 4 Fantom Foundation. Lachesis aBFT, 2024. URL: <https://docs.fantom.foundation/technology/lachesis-abft>.
- 5 Adam Gagol, Damian Lesniak, Damian Straszak, and Michal Swietek. Aleph: Efficient atomic broadcast in asynchronous networks with Byzantine nodes. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies, AFT 2019, Zurich, Switzerland, October 21-23, 2019*, pages 214–228. ACM, 2019. doi:10.1145/3318041.3355467.
- 6 Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling Byzantine agreements for cryptocurrencies. In *Proceedings of SOSP 2017*, pages 51–68. ACM, 2017. doi:10.1145/3132747.3132757.
- 7 Hedra. Streamlining consensus, 2024. URL: <https://hedera.com/blog/streamlining-consensus-throughput-and-lower-latency-with-about-half-the-events>.
- 8 Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is DAG. In *Proceedings of PODC 2021*, pages 165–175. ACM, 2021. doi:10.1145/3465084.3467905.
- 9 Jae Kwon. Tendermint: Consensus without mining. <https://tendermint.com/docs/tendermint.pdf>, 2014.
- 10 Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, jul 1978. doi:10.1145/359545.359563.
- 11 Leslie Lamport. *Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2002. URL: <http://research.microsoft.com/users/lamport/tla/book.html>.
- 12 Leslie Lamport. Teaching concurrency, 2009. URL: <https://lamport.azurewebsites.net/pubs/teaching-concurrency.pdf>.
- 13 Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, jul 1982. doi:10.1145/357172.357176.
- 14 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.