Blockchains: Past, Present, and Future

Arvind Narayanan
Princeton University
Princeton, NJ, USA
arvindn@cs.princeton.edu

ABSTRACT

Blockchain technology is assembled from pieces that have long pedigrees in the academic literature, such as linked timestamping, consensus, and proof of work. In this tutorial, I'll begin by summarizing these components and how they fit together in Bitcoin's blockchain design.

Then I'll present abstract models of blockchains; such abstractions help us understand and reason about the similarities and differences between the numerous proposed blockchain designs in a succinct way. Here is one such abstraction. Blockchains can be understood in terms of (1) a log of messages: for example, a ledger of financial transactions; (2) the state that summarizes the result of processing the log: for example, a set of account balances; (3) a set of validity rules for messages/state updates: for example, transactions must spend no more than the available balances, must have verifiable signatures, etc; (4) consistency rules that determine whether two views of the log by different participants on the network are consistent with each other.

In the second half of the tutorial I'll describe several research directions, focusing on those likely to be of interest to the PODS community. Here are a few examples.

- *Efficient verification of state.* A participant might want to verify a statement about a small part of the global state, such as the inclusion of a particular transaction in the blockchain. While the basics have been worked

out, and involve hash pointers, Merkle trees, and other “authenticated data structures”, many interesting questions remain.

- *Reconciling different views of consensus.* In the game theory view of blockchains, all players are rational and follow their incentives; there are no honest, faulty, or malicious players. How do the predictions of this view compared to those of the traditional consensus literature? Can we come up with hybrid models that reconcile these assumptions?
- *Scaling and sharding.* In traditional designs, the blockchain is fully replicated by every node, leading to massive inefficiency and severely limiting transaction throughput. What are the fundamental limits to scaling, and how can we improve scalability without weakening security? In particular, is it possible to shard the blockchain, that is, partition it among subsets of nodes, given the Byzantine setting?

ACM Reference Format:

Arvind Narayanan. 2018. Blockchains: Past, Present, and Future. In Proceedings of 2018 International Conference on Management of Data (PODS'18), ACM, New York, NY, USA, 1 page.
<https://doi.org/10.1145/3196959.3197545>

Author Keywords

blockchain, cryptocurrency, consensus, state replication.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

PODS'18, June 10-15, 2018, Houston, TX, USA.

© 2018 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-4706-8/18/06...\$15.00.

DOI: <https://doi.org/10.1145/3196959.3197545>