THE UNIVERSITY OF SYDNEY

Blockchain Evaluation

Andrei Lebedev Vincent Gramoli

School of Computer Science, the University of Sydney

Context

Blockchain

- Distributed Ledger Technology
- Append-only log of blocks comprised of transactions
- Provides immutability, data integrity, and fault-tolerance, in a decentralized environment

Transaction performance metrics

- Throughput
- Latency

Motivation and Research Gap

- Different blockchain proposals exist and appear with impressive performance and fault-tolerance claims
- Need to compare the protocols on the same ground and test their different aspects
- Existing work focuses on individual protocols, uses different hardware in the experiments making results hard to compare, or does not reflect realistic network behaviors and faults

AWS Findings

- Evaluated blockchains not yet suitable for demanding centralized service workloads
- Eventual consistency in blockchains aids scalability in networks with many nodes
- Protocols may struggle with challenging configurations, especially high RTT and large hardware resources
- Deterministic BFT consensus protocols in blockchains impacted by consistently high workloads
- Leader-based BFT consensus design known for scalability limitations
- Protocols with probabilistic or eventual consistency guarantees maintain throughput under high workloads
- Tradeoff between robustness and availability observed in experimental results

Claimed and Observed Performance

Differences between the claimed performance and the actual performance of different blockchains. The results we observed for each blockchain are the best performances we obtained among all used configurations. These were obtained in the `testnet` and `datacenter` configurations

Blockchain	Claimed results		Observed results	
	throughput	latency	throughput	latency
Algorand	1K–46K TPS	2.5–4.5 s	885 TPS	8.5 s
Avalanche	4.5K TPS	2 s	323 TPS	49 s
Solana	200K TPS	<1 s	8845 TPS	12 s

Diablo

The architecture of Diablo comprises configuration files for the Primary to send the right workload for the right blockchain to a set of Secondaries that then send requests to blockchain nodes and collect performance results from these blockchain nodes

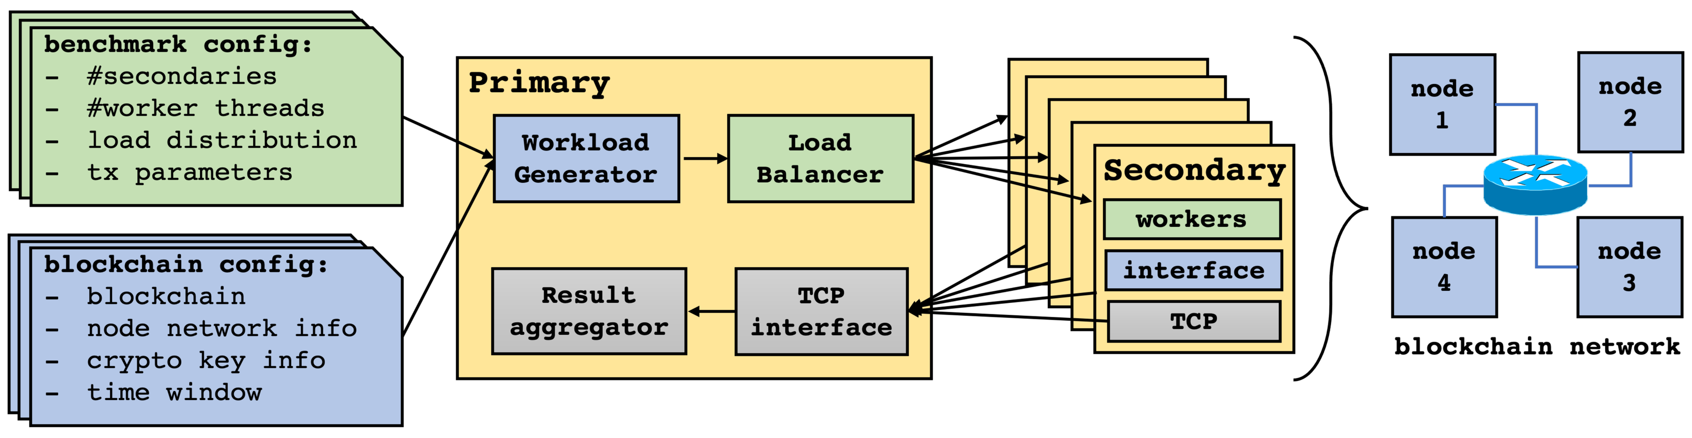

Client Scalability

Average throughput and average latency of each blockchain when stressed with a constant workload of 1,000 TPS on different configurations, from the least challenging (`datacenter`) to the most challenging (`community`)

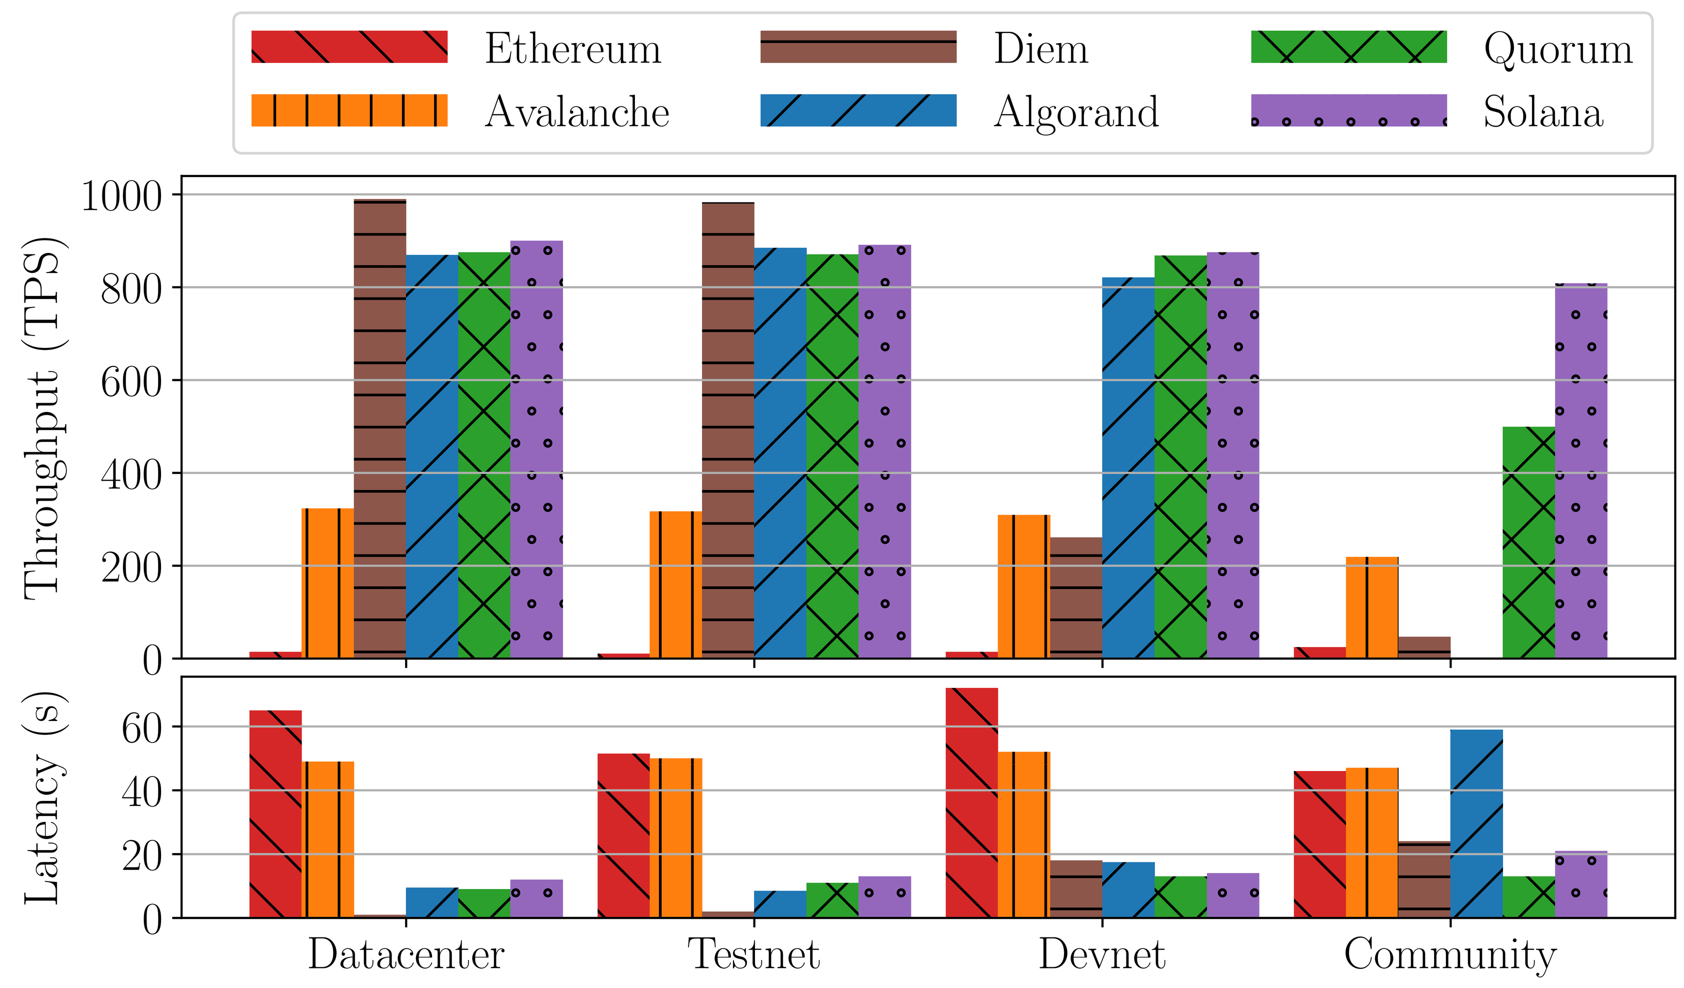

Emulated Delay

Throughput and latency, 1,000 TPS workload, varied number of isles and switches

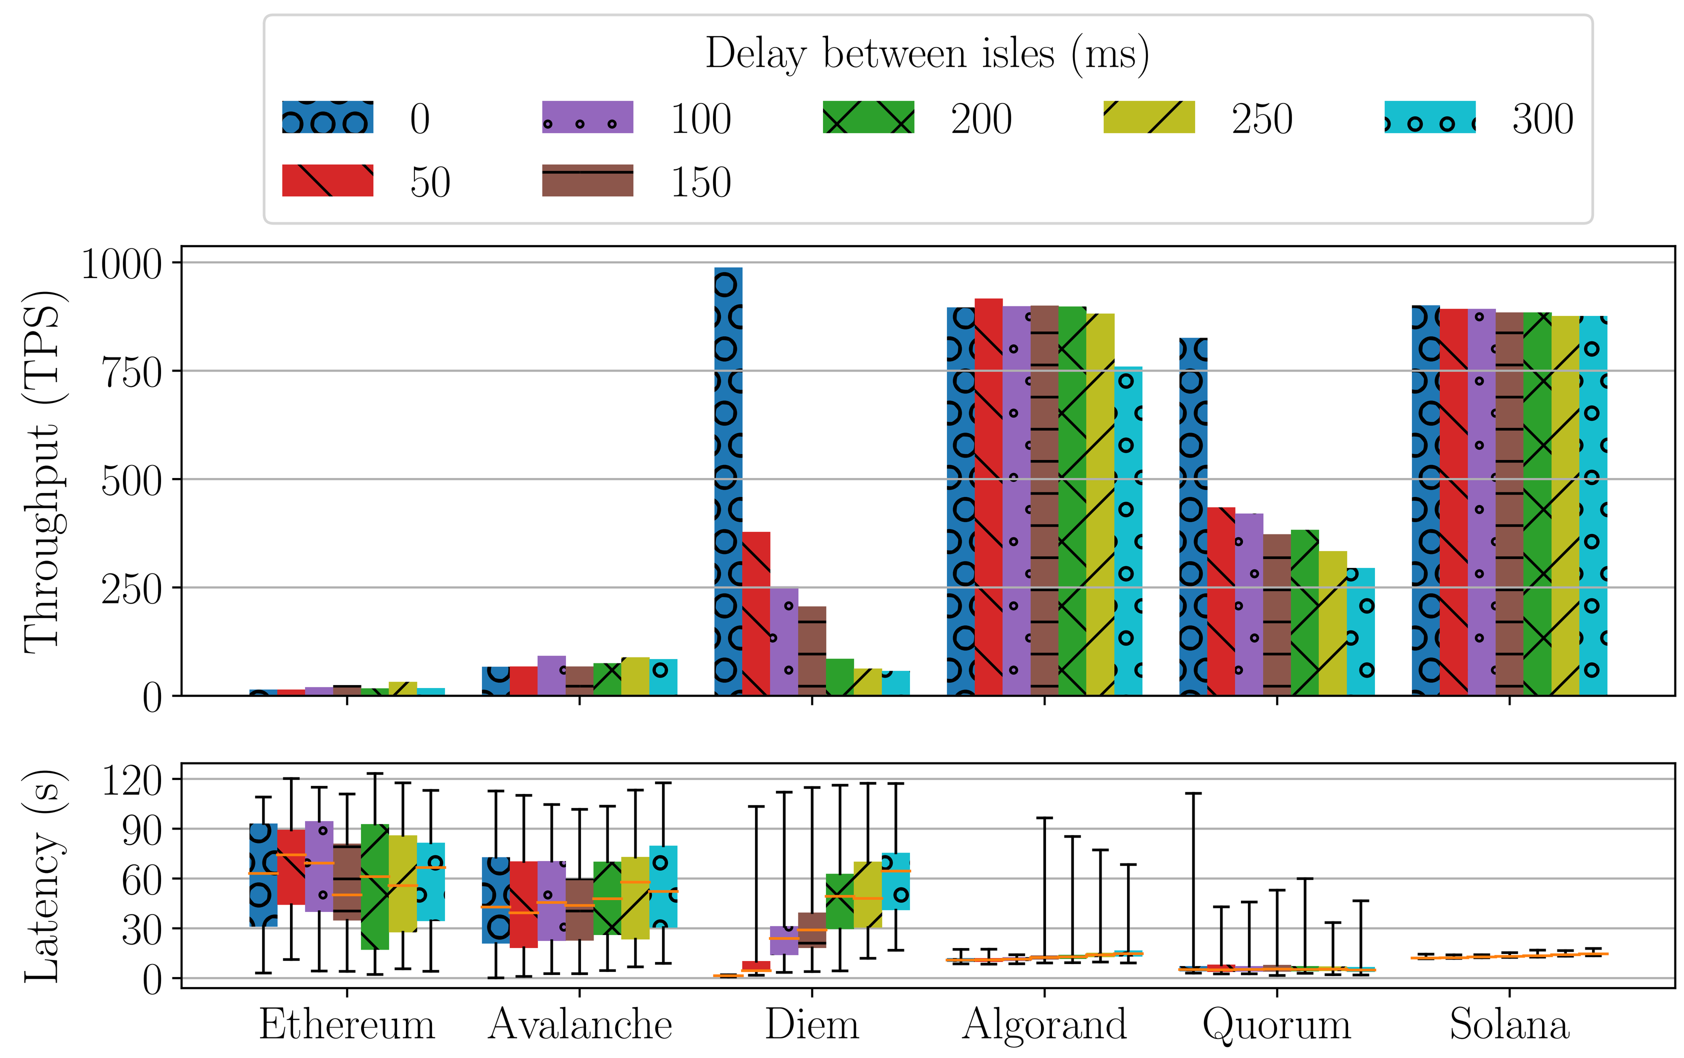

AWS Environment

The experimental settings range from a `datacenter` scenario with extensive resources to a `testnet` of collocated machines, to a geo-distributed `devnet`, to a large-scale `community` of machines, to a large-scale `consortium` of modern machines

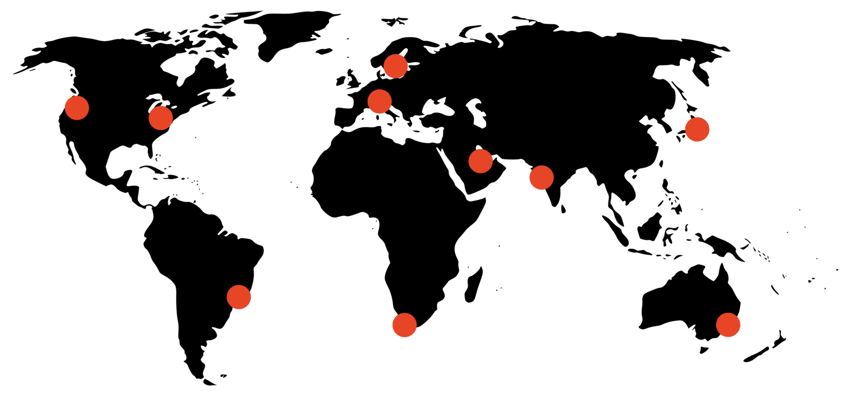

Bare Metal Setup

The bare metal cluster consists of 7 groups of machines called isles (named A, B, C, D, E, F, S) of 6 machines each, plus an isle of 3 machines (named R), giving 8 isles and 45 machines in total. Every two isles (A and S, B and R, C and D, E and F) are connected to a switch, and there are overall 4 switches, and all of them are connected to each other.

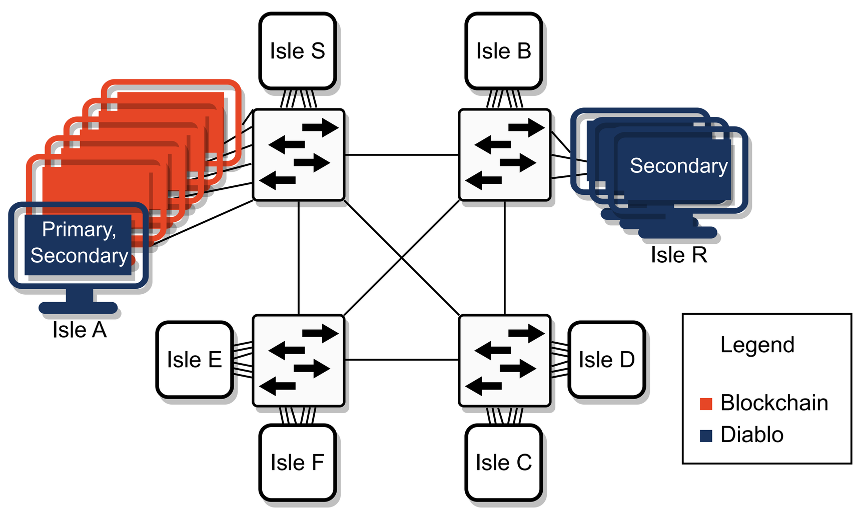

Bare Metal Setup Findings

- Switches within the LAN showed minimal influence on blockchain performance in our study
- The average transaction latency of blockchains did not reliably reflect their tail latency
- Artificial network delays can present the same performance patterns as the geo-distributed network

References

- Vincent Gramoli, Rachid Guerraoui, Andrei Lebedev, Chris Natoli, and Gauthier Voron. Diablo: A benchmark suite for blockchains. In *Proceedings of the Eighteenth European Conference on Computer Systems*, EuroSys '23, 2023.
- Andrei Lebedev and Vincent Gramoli. On the relevance of blockchain evaluations on bare metal. In *7th Symposium on Distributed Ledger Technology (SDLT 2023)*, page to appear. Springer, 2023.

andrei.lebedev@sydney.edu.au