

DAG-based Efficient Parallel Scheduler for Blockchains: Hyperledger Sawtooth as a Case Study

Manaswini Piduguralla

Saheli Chakraborty
Sathya Peri

Parwat Singh Anjana

Euro-Par 2023

01 September 2023



भारतीय प्रौद्योगिकी संस्थान हैदराबाद
Indian Institute of Technology Hyderabad


- 1 Research Objective and Introduction
- 2 Sawtooth Introduction
- 3 Motivation
- 4 Framework design and implementation
- 5 Results
- 6 Conclusion and Future Work
- 7 Related Work

- 1 Research Objective and Introduction
- 2 Sawtooth Introduction
- 3 Motivation
- 4 Framework design and implementation
- 5 Results
- 6 Conclusion and Future Work
- 7 Related Work

- The objective is to develop an Efficient Distributed and Secure framework for the Execution of Smart Contracts in Blockchains.
 - Develop framework for concurrent execution of transactions.

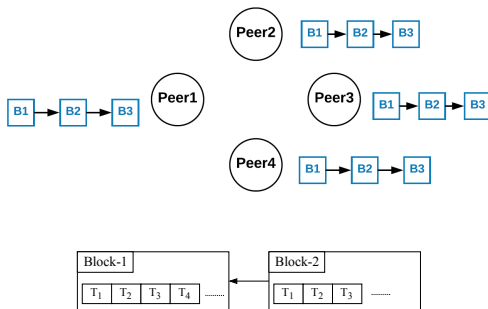


- Blockchain is a decentralized distributed immutable ledger shared among untrusted parties¹.



¹Nakamoto:Bitcoin:2009.

- Blockchain is a decentralized distributed immutable ledger shared among untrusted parties¹.



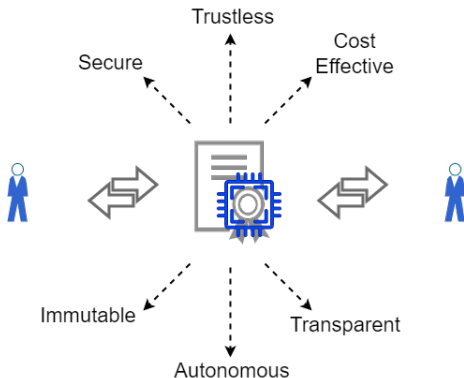
- Each block contains a set of transactions.

¹Nakamoto:Bitcoin:2009.



- Smart contracts (SCs) are self executing contracts with agreement between two or more parties that are written in the form of computer code.
- Smart contracts are like a 'class' that encapsulates data and methods.

- Smart contracts (SCs) are self executing contracts with agreement between two or more parties that are written in the form of computer code.
- Smart contracts are like a 'class' that encapsulates data and methods.





- Blockchain nodes form a peer-to-peer system.
- Clients (external to the system) wishing to execute transactions, contact a peer of the system.

- Blockchain nodes form a peer-to-peer system.
- Clients (external to the system) wishing to execute transactions, contact a peer of the system.

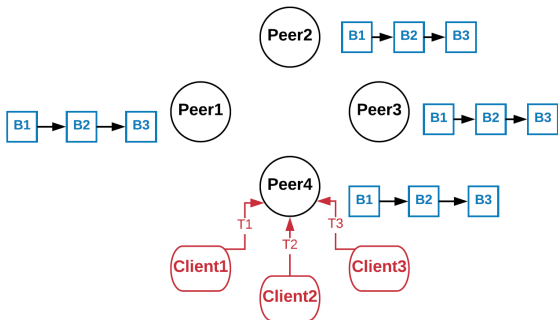


Figure: Clients send Transaction T1, T2 and T3 to block producer (Peer4)²

²Parwat+:Netys:2020.

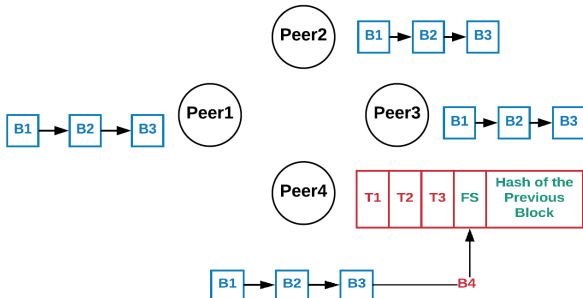


Figure: Block producer forms a block B4 and computes final state (FS) sequentially




Figure: Block producer broadcasts the block B4

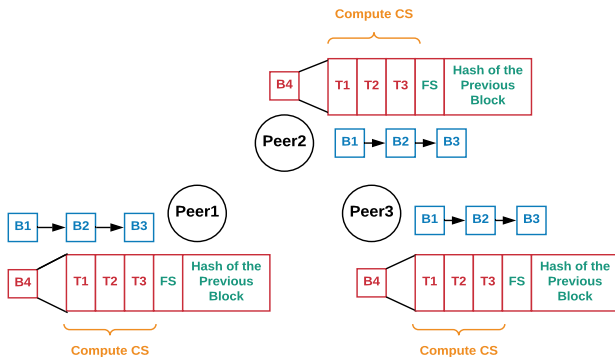


Figure: Validators (Peer 1, 2, and 3) compute current state (CS) sequentially

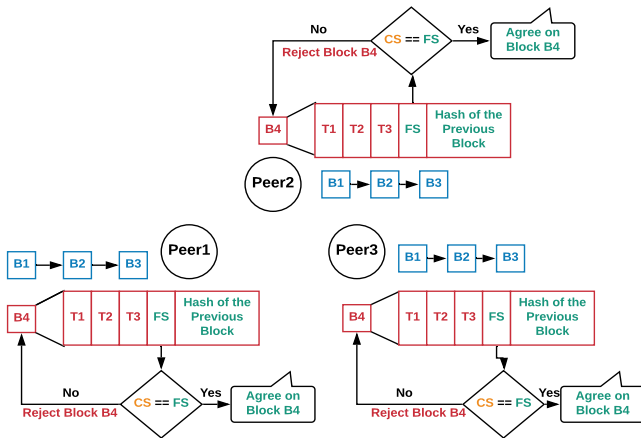


Figure: Validators verify the FS and reach the consensus protocol

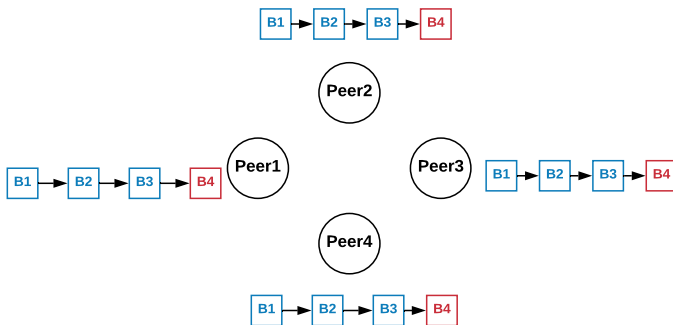


Figure: Block B4 successfully added to the blockchain

- 1 Research Objective and Introduction
- 2 Sawtooth Introduction**
- 3 Motivation
- 4 Framework design and implementation
- 5 Results
- 6 Conclusion and Future Work
- 7 Related Work

- **Hyperledger Sawtooth³** is a modular platform for building, deploying, and running distributed ledgers.
- The modular property of sawtooth, enables enterprises and consortiums to make decisions about their blockchain applications for themselves.

³[sawtooth:url](#).

- **Hyperledger Sawtooth³** is a modular platform for building, deploying, and running distributed ledgers.
- The modular property of sawtooth, enables enterprises and consortiums to make decisions about their blockchain applications for themselves.

Sawtooth Key Features:

- Modular
- Permissioned as well as permissionless infrastructure
- Parallel transaction execution
- Pluggable consensus algorithms
- Multi language support
- Dynamic consensus

Sawtooth

³[sawtooth:url.](https://sawtooth.hyperledger.org/)

- 1 Research Objective and Introduction
- 2 Sawtooth Introduction
- 3 Motivation**
- 4 Framework design and implementation
- 5 Results
- 6 Conclusion and Future Work
- 7 Related Work



Smart contract in blockchain are executed in two different contexts :

Smart contract in blockchain are executed in two different contexts :

- While block producers are composing the blocks.

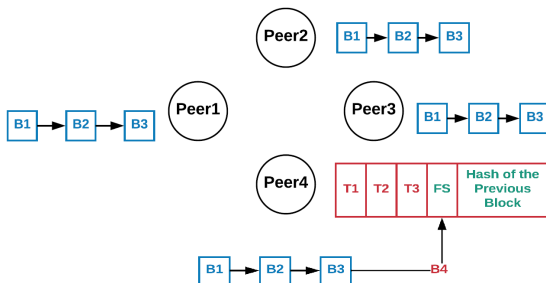


Figure: Block producer create the block and broadcast it to others



Smart contract in blockchain are executed in two different contexts :

- While block producers are composing the blocks.
- While nodes/peers are validating the blocks.

Smart contract in blockchain are executed in two different contexts :

- While block producers are composing the blocks.
- While nodes/peers are validating the blocks.

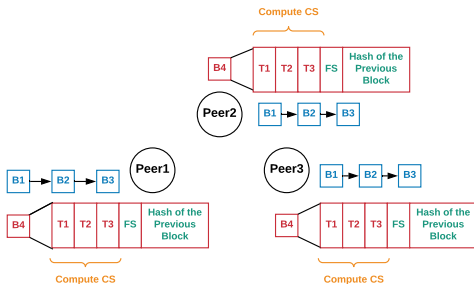


Figure: Validators (Peer 1, 2, and 3) compute current state (CS) sequentially



In both stages, the transactions in the block are executed serially in most blockchains⁴:

- Not utilizing the multi-core processors efficiently.
- Results in lower throughput.

⁴Dickerson+:ACSC:PODC:2017.



In both stages, the transactions in the block are executed serially in most blockchains⁴:

- Not utilizing the multi-core processors efficiently.
- Results in lower throughput.

Solution: *Concurrent execution of SCTs*

⁴Dickerson+:ACSC:PODC:2017.

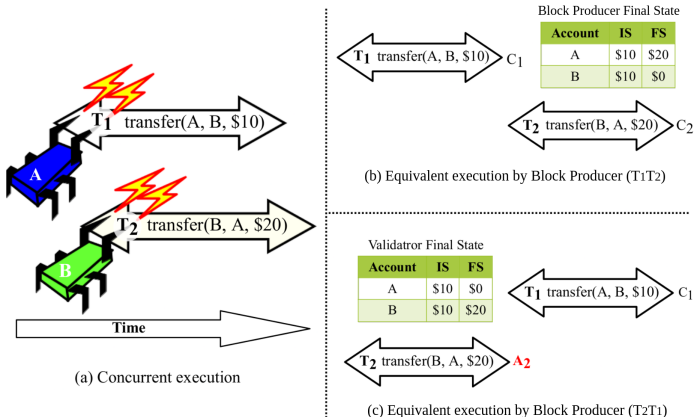


- Validator may incorrectly reject a valid block proposed by the block producer. We call such error as **False Block Rejection (FBR)** error^{5, 6}.

⁵Dickerson+:ACSC:PODC:2017.

⁶Parwat+:Netys:2020.

- Validator may incorrectly reject a valid block proposed by the block producer. We call such error as **False Block Rejection (FBR)** error.



Solution: The concurrent execution is always conflict equivalent to a serial schedule.



- We have proposed a concurrent transaction execution framework for blockchains.
- The proposed approach has been thoroughly tested in Hyperledger Sawtooth 1.2.6.
- It is flexible enough for implementation in any blockchain that follows the order-execute paradigm.

- 1 Research Objective and Introduction
- 2 Sawtooth Introduction
- 3 Motivation
- 4 Framework design and implementation**
- 5 Results
- 6 Conclusion and Future Work
- 7 Related Work




Figure: Proposed framework in the blockchain



- The parallel scheduler performs the operations of DAG creation and conflict-free transaction execution.



- The parallel scheduler performs the operations of DAG creation and conflict-free transaction execution.
- DAG (Direct Acyclic Graph) to represent the dependencies among the transactions:
 - The nodes of the graph will be the transactions in the block.

- The parallel scheduler performs the operations of DAG creation and conflict-free transaction execution.
- DAG (Direct Acyclic Graph) to represent the dependencies among the transactions:
 - The nodes of the graph will be the transactions in the block.
 - An edge represents the dependency between two transactions.




Figure: DAG representation of dependencies in the block.



- Dependency:
When two transactions are accessing the same data while at least one of which is modifying it.
 - T1 : View(A)
 - T5: Send(B, A, 10)




Figure: DAG representation of dependencies in the block.

- Scheduling:

- Transactions with zero indegree are not dependent on current set of executing transactions.
- These transactions can be scheduled for execution.




Figure: DAG representation of dependencies in the block.

- Scheduling:
 - The out edges are removed of the transactions that have completed execution.




Figure: DAG representation of dependencies in the block.




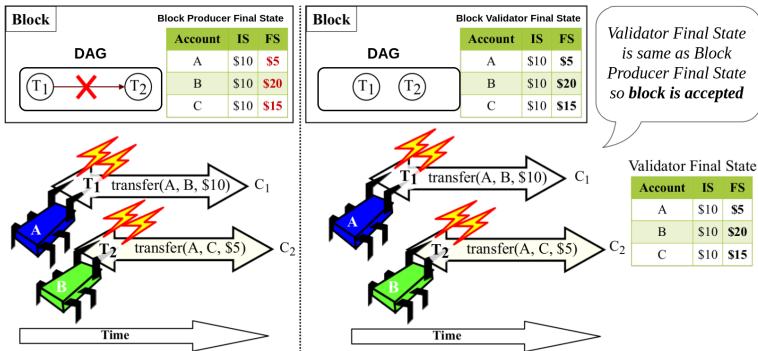
Figure: Proposed framework in the blockchain



- The *Malicious block producer* can send an incorrect Block Graph to harm the blockchain, missing some edges, to cause *double spending*. We call such error as **Edge Missing BG (EMB)** error⁷.

⁷Parwat+:Netys:2020.

- The *Malicious block producer* can send an incorrect Block Graph to harm the blockchain, missing some edges, to cause *double spending*. We call such error as **Edge Missing BG (EMB)** error.

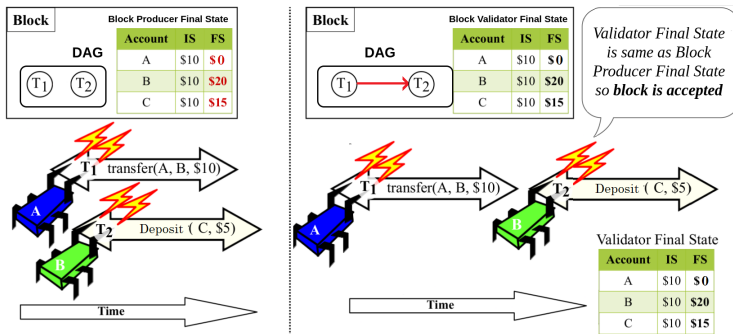


Solution: We propose a *Secure Multi-threaded Validator (SMV)* to detect EMB error and rejects the corresponding blocks.



- The *Malicious block producer* can send an incorrect Block Graph with extra edges, to intentionally slow the validation process. We call such error as **Extra Edge BG (EEB)** error.

- The *Malicious block producer* can send an incorrect Block Graph with extra edges, to intentionally slow the validation process. We call such error as **Extra Edge BG (EEB)** error.



Solution: We propose a *Secure Multi-threaded Validator (SMV)* to detect EEB error and rejects the corresponding blocks.




Figure: Linked list address data for secure validator

- **Edge Validation:** There should be an edge connecting each transaction in the read list to every transaction in the write list.
- An edge should also connect any two transactions within the write list.


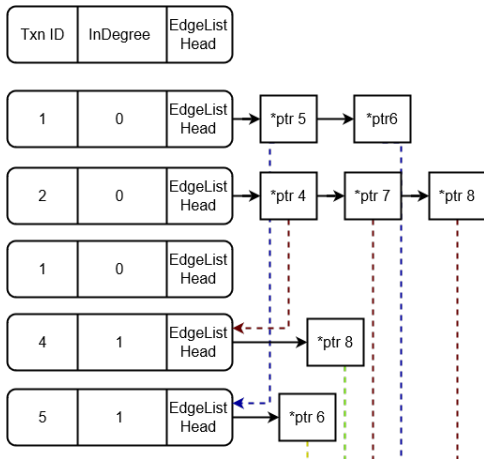


Figure: Linked list address data for secure validator

- **Edge Validation:** There should be an edge connecting each transaction in the read list to every transaction in the write list.
- An edge should also connect any two transactions within the write list.
- **In-degree validation:** We track the in-degree of each transaction during edge validation and cross check it with DAG in-degree.

- We have developed two implementations for the proposed DAG scheduler:
 - Linked List
 - Adjacency Matrix





We implemented four transaction families to test the performance of our approach:



We implemented four transaction families to test the performance of our approach:

- IntKey Transaction Family
- SmallBank Transaction Family
- Voting Transaction Family
- Insurance Transaction Family
- Mixed block containing all three



- 1 Research Objective and Introduction
- 2 Sawtooth Introduction
- 3 Motivation
- 4 Framework design and implementation
- 5 Results**
- 6 Conclusion and Future Work
- 7 Related Work

- Voting Transaction Family
 - Y-axis: Time in seconds
 - X-axis: Number of blocks
 - Number of threads: 56
 - Degree of dependency: 50% CP3
 - Number of txns: 1000




Figure: Experiment one: Voting

Curves: Serial Scheduler, Tree Parallel Scheduler, LLDAG Scheduler, Adj DAG Scheduler.

- SimpleWallet Transaction Family
 - Y-axis: Time in seconds
 - X-axis: Number of transactions per block
 - Number of threads: 56
 - Degree of dependency: 50% CP2
 - Number of blocks: 20




Figure: Experiment two: SimpleWallet

Curves: Serial Scheduler, Tree Parallel Scheduler, LLDAG Scheduler, Adj DAG Scheduler.

- SimpleWallet Transaction Family
 - Y-axis: Time in seconds
 - X-axis: Number of transactions per block
 - Number of threads: 56
 - Number of blocks: 20
 - Number of txns: 1000




Figure: Experiment three: Mixed Block

Curves: Serial Scheduler, Tree Parallel Scheduler, LLDAG Scheduler, Adj DAG Scheduler.

- Dependency data structure creation or validation time for SimpleWallet Transaction Family
 - Y-axis: Time in seconds
 - X-axis: Number of transactions per block
 - Number of threads: 56
 - Degree of dependency: 50% CP2
 - Number of blocks: 20




Figure: Comparison of data structure creation time.

Curves: Tree Parallel Scheduler, LLDAG Scheduler, Adj DAG Scheduler, Secure Validator



- 1 Research Objective and Introduction
- 2 Sawtooth Introduction
- 3 Motivation
- 4 Framework design and implementation
- 5 Results
- 6 Conclusion and Future Work**
- 7 Related Work

- We have introduced a concurrent transaction execution framework for blockchain systems.
- Rigorous testing of our approach has taken place within Hyperledger Sawtooth 1.2.6.

- We have introduced a concurrent transaction execution framework for blockchain systems.
- Rigorous testing of our approach has taken place within Hyperledger Sawtooth 1.2.6.
- Our adaptable framework is compatible with any blockchain adhering to the order-execute paradigm.
- The artifact of our framework is evaluated and is available in Figshare repository⁸.

⁸Piduguralla:artifact:2023.



- Optimize the framework by improving DAG creation and secure validation is our immediate next step.
- We will be exploring enhancing the fault tolerance and scalability for each blockchain node.

Thank you



Thank you to ACM-W and Google for sponsoring my attendance and paper presentation at this conference


Extras




Transaction Families \approx Smart Contracts



Creation



Data Model



Business Logic

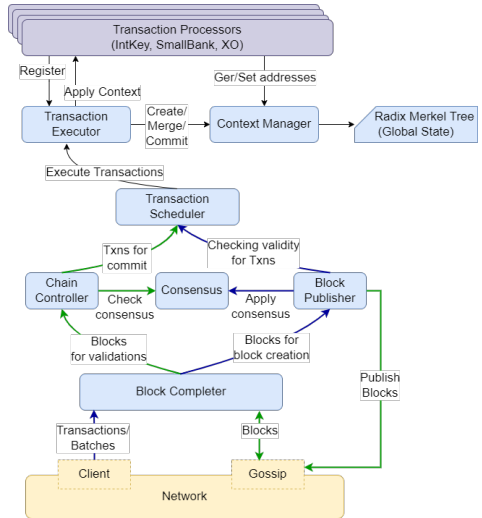
Intro



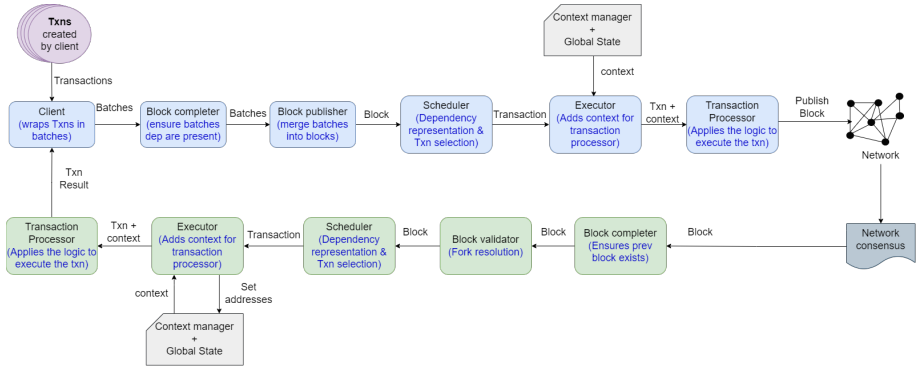
- The **journal** is the group of validator subcomponents that work together to handle batches and proposed blocks.
- The **Block Completer** initially receives the blocks and batches. It guarantees that all dependencies for the blocks have been satisfied.
- Completed batches go to the **Block Publisher** for batch validation and inclusion in a block.
- Completed blocks go to the **Block Validator** for validation and fork resolution.

Sawtooth architecture

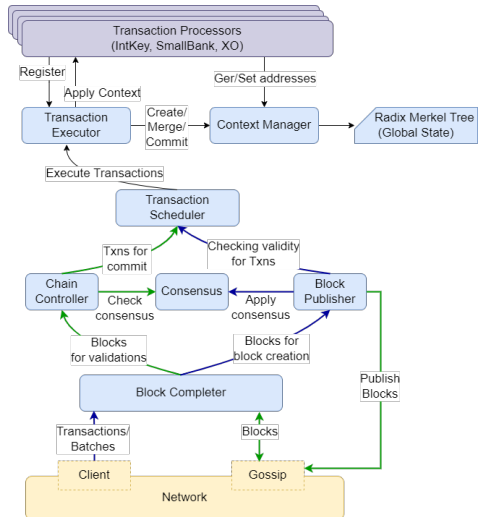
- Journal
- Transaction Scheduler
- Transaction Executor
- Transaction Processor
- Global State



Life Cycle of a Sawtooth Transaction



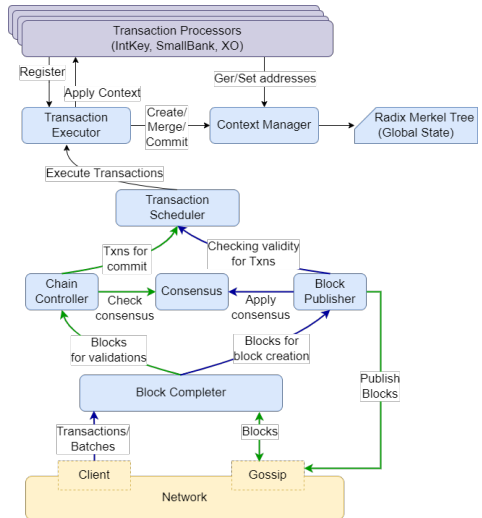
Intro



- Block Completer:**
 Receives the blocks and batches from the network




- **Block Publisher:** Creation of blocks after validation

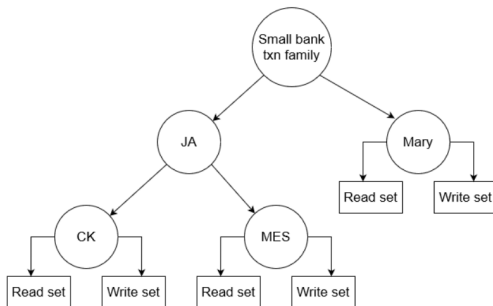




- Block Validator:** Performs validation and fork resolution for transactions in the block




- Sawtooth tree structure to keep track of dependencies using merkle tree.
- Each data location has read and write lists that contain the ID of the transactions that are accessing them.
- David, Daniel, Dylan.

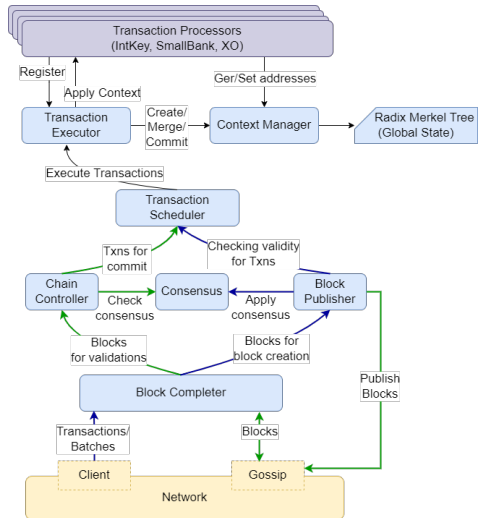


Transaction Executor

- The executor obtains the next transaction from the scheduler.
- Acquires context reference from context manager and state.
- Combines transaction and context for transaction processor.
- The Executor updates the scheduler with the transactions's result with the updated context using locks.



- Sawtooth uses an addressable Radix Merkle tree to store data for transaction families.
- Merkle tree : stores successive node hashes from leaf-to-root.
- Radix tree: addresses uniquely identify the paths to leaf nodes.
- An address is a hex-encoded string of 35 bytes, each node has 2^8 possible children.





- Sawtooth architecture uses tree with addresses as nodes to represent the dependencies.

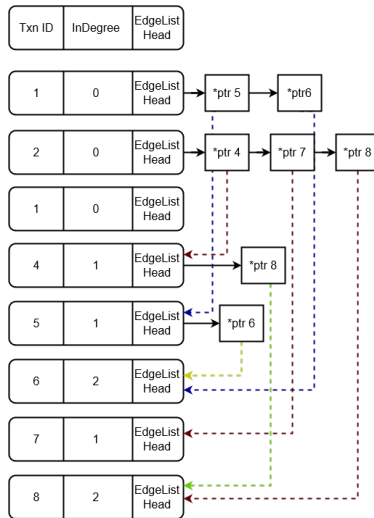
(Total number of addresses \gg Transactions in a block)

- The construction of the tree is done serially.
- Status of all the dependent transactions is checked before scheduling a transaction.


(This generates a delay in scheduling the transaction)

- The available transaction are searched from the start of the array each time a transaction is scheduled.
- Locks are used to update transaction status.


- We have developed two implementations for the proposed DAG scheduler:
 - Linked List
 - Adjacency Matrix



- For DAG scheduling the function “Next_transaction” is modified.
- Transactions are selected for scheduling using the DAG created.
- In selecting a transaction there are multiple C++ modules initiated by python threads.



- For DAG creation we have modified the “add_batch” function in Sawtooth scheduler.
- The creation of DAG is done through multiple threads in C++ module.



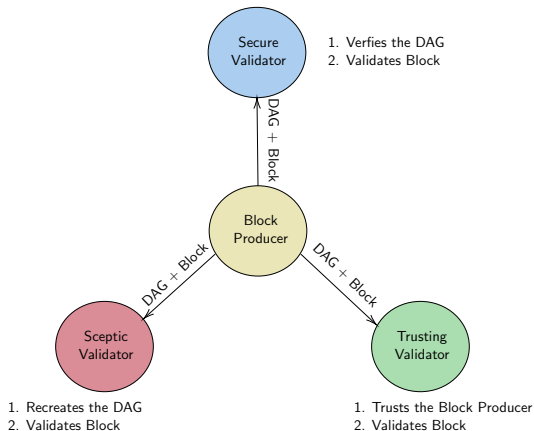


Figure: Possible validators in the blockchain



Consensus protocols and transaction throughput are the two significant bottlenecks of blockchain performance.

Consensus:

- Proof of Stake (*PoS*)⁹
- Proof of Elapsed Time (*PoET*)¹⁰
- Practical Byzantine Fault Tolerance (*PBFT*)¹¹

⁹Vasim:2014:Blackcoin.

¹⁰sawtooth:url.

¹¹Castro:2002:ACM.



Consensus protocols and transaction throughput are the two significant bottlenecks of blockchain performance.

Sharding:

- Network Sharding¹²
- State Sharding¹³
- Transaction Sharding¹⁴

¹²DiPETrans:CPE:2022.

¹³Zheng+:IEEE:TII:2022.

¹⁴Dang+:SIGMOD:2019.



- 1 Research Objective and Introduction
- 2 Sawtooth Introduction
- 3 Motivation
- 4 Framework design and implementation
- 5 Results
- 6 Conclusion and Future Work
- 7 Related Work**



Consensus protocols and transaction throughput are the two significant bottlenecks of blockchain performance.

¹⁵ **Dickerson+:**ACSC:PODC:2017.

¹⁶ **Parwat+:**Netys:2020.

¹⁷ **blockSTM:**2022.



Consensus protocols and transaction throughput are the two significant bottlenecks of blockchain performance. **Transaction throughput:**

STMs

- Block Graph (BG)¹⁵
- BG + Smart Validator¹⁶
- BlockSTM¹⁷

¹⁵Dickerson+:ACSC:PODC:2017.

¹⁶Parwat+:Netys:2020.

¹⁷blockSTM:2022.



- We have developed two implementations for the proposed DAG scheduler:
 - Linked List
 - Adjacency Matrix

Adjacency Matrix

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

In-Degree Matrix

$$\begin{bmatrix} 0 & 0 & 0 & 1 & 2 & 1 & 1 & 2 \end{bmatrix}$$



In order to assess the framework's performance across different scenarios, we have devised three conflict parameters (CP):

- CP1 measures the proportion of transactions in the DAG that have at least one dependency
- CP2 represents the ratio of dependencies to the total number of transactions in the DAG
- CP3 quantifies the number of disjoint components, which are subgraphs without interconnections in the DAG.



Parameter that we can change in Sawtooth architecture:

- Number of blocks
- Number of transactions in a block
- Metric for degree of dependency
- Number of threads