PriCLeSS

Privacy-Conscious Legally-Sound blockchain Storage Proof-of-Concept

WIDE, Inria Rennes, IRISA, UR1
CIDRE, Inria Rennes, IRISA, UR1
GDD, LS2N, University of Nantes
IODE, University of Rennes 1

Current Concept

SplitChain: Sharding in presence of an adaptive adversary

Introduction

Blockchains suffer from various technical issues, such as their inability to scale due to problematic communication costs and latency and a fairly low transaction throughput. We present SplitChain, a protocol intended to support the creation of scalable account-based blockchains without undermining decentralization and security. This is achieved by using sharding , i.e. by splitting the blockchain into several lighter chains managed by their own disjoint sets of validators called shards. These shards balance the load by processing disjoint sets of transactions in parallel.

Contributions

- A distributed validator attribution mechanism.
- Chain merging and chain splitting protocols to automatically adapt the number of chains to the current payload of the system.
- A routing protocol allowing chains to efficiently redirect transactions and deliver messages.

Chain Structure and Block Pruning

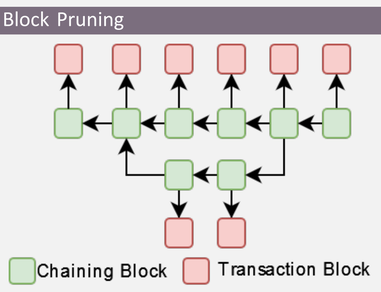

Most transactions inside a blockchain are independent from each other and do not need to be totally ordered inside a single chain of blocks. Instead, the storage of transactions and user accounts and the computation of new blocks are subdivided into multiple independent chains. Chains contain two types of blocks:

- Transaction blocks contain the actual transactions stored only inside the chain.
- Chaining blocks link the blocks of a chain together. They are only a few hundred bytes in size and are broadcasted through the entire system to provide synchronization data.

Periodically (e.g. 1000 blocks), a chaining block called cue block serves as a checkpoint. This allows to gradually prune the previous chaining blocks and transaction blocks.

Assigning Validators to Chains

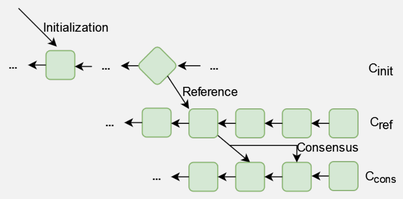

Partitioning the system into smaller chains reduces the amount of colluding validators required to corrupt them. To prevent the adaptive adversary from devising complex strategies to concentrate its manipulated validators in a targeted chain, the attribution of validators to the consensus of a chain is a three-stage process that makes validator attribution unpredictable and ephemeral. Specifically, each chain of SplitChain has three roles in relation to validators: Initialization, reference and consensus.

- The initialization chain is the chain that processed the transaction creating the validator's account. The validator becomes active when the next cue block of the chain is created and a chain is randomly selected as its referencing chain.
- The referencing chain provides validators with a stable foothold in the network, and attributes them to their consensus chain. The referenced validators are distributed fairly among the different chains of the system using a pseudo-random attribution algorithm executed during each consensus execution. A Merkle root of the attribution of the validators is embedded to the newly created chaining block to serve as proof of their assignment.
- A validator can only participate to at most $T_A - 1$ consecutive consensus executions of its consensus chain. That is because the adversary is adaptive over a period T_A, i.e. it can target any validator in the system, which will be corrupted after a delay of T_A consensus executions. To ensure that validator credentials expire whether or not they participate in the consensus, only the credentials referenced by the latest chaining blocks of a chain are accepted for the consensus. Once its credentials have expired, a validator must issue a new request to its referencing chain with new credentials to obtain a new consensus chain.

.

Hypercube Network Topology and Routing Protocol

All the chains of SplitChain are organized in a loosely hypercubic topology. A hypercube of dimension n (or n-cube) has 2^n vertices, each of which is connected to n neighbors and at a maximum distance of n edges of their furthest vertex of the n-cube. Each chain is a vertex of the hypercube.

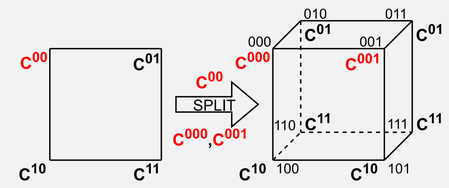

Initially, SplitChain consists merely of an unique chain. When the number of transactions per block of the chain exceeds a certain threshold, the chain is replaced by two new chains called siblings. We call this operation a split. Similarly, if the number of transactions per block of a chain is too low, this chain will merge with its sibling.

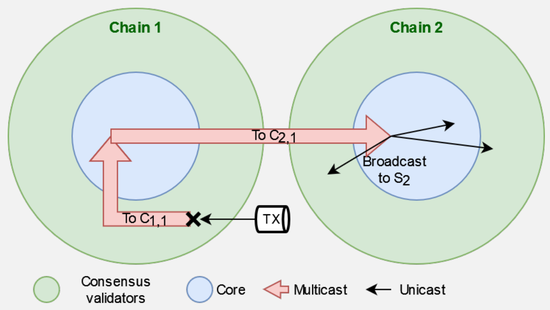

A small subset of a chain's validators call core stores and maintains the routing tables used for intra-chain communication and message forwarding between the core validators of neighboring chains in the hypercube. Whenever a core validator receives a message, it either broadcasts it to the validators of its chain or forwards it to the core of one of its neighboring chains depending of the message's destination.

Conclusion

SplitChain supports the creation of scalable account-based blockchains without undermining decentralization and security. It distinguishes itself from other sharded blockchains by minimizing the synchronization constraints among shards while maintaining security guarantees. Specifically, SplitChain is the first sharded blockchain that do not require a dedicated shard or a global blockchain to attribute validators to their consensus chain. This avoids the need for a global reconfiguration of the shards each time a new batch of validators is added to the system. A dedicated routing protocol enables transactions to be redirected between shards with a low number of hops and messages. Finally, SplitChain dynamically adapts the number of shards to the system load to avoid over-dimensioning issues encountered in static sharding-based solutions.

Development Projects

Risk Management: Plan Modification

Difficulty in recruitment did not allow us to start development by Summer 2023 as announced.
We are now planning to use the 18 month funding as follows.

- One engineer on original project for 12 months: January 2024 - December 2024.
- A second engineer on Blockchain-free SSI also for 12 months: January 2024 - December 2024, funded for 6 months by the follow up and for 6 months on WIDE's own funds (resulting from SOTERIA H2020 overhead).

Privacy Preserving Geolocalized Storage

Tradeoff between Legal Compliance and Resilience

How address how to best store data in the context of storage-oriented applications, while respecting the legal requirements set by the GDPR.

- Initially planned by September 2023.
- Current target December 2023.

Human Resources

- One engineer for 12 months starting December 2023 - January 2024.

RoadMap

- M0: Identify how to best store data in the context of storage-oriented applications, while respecting the legal requirements set by the GDPR.
  - Initially planned by September 2023.
  - Current target December 2023.
- M1-M3: Analysis of the state of the art:
  - Sharding and Proof-of-Eligibility consensus implementations.
  - Current specification of SplitChain.
  - Storage-oriented features (M0).
- M4: High-level design of the implementation.
- M6: Single shard implementation, with features for multiple shards and privacy (indirection, encryption, etc.).
- M8: Hypercube routing for inter-shard communication.
- M12: Prototype implementation of the complete system.

Original plan included additional months for experimentation. We will postpone this to a future PhD thesis/master internship.

Blockchain-Free Self-Sovereign Identity

From Theory to Practice

Leverage our work on Decentralized Identity, and Broadcast-Based Distributed Objects.

- New plan combining PriCLeSS FollowUP with H2020 overhead funds.

Our work [1] shows that the implementation of SSI solutions does not require global consensus. Similar to a sharded cryptocurrency system, an SSI platform can operate without a global synchronized ledger.

Human Resources

One engineer for 12 months starting December 2023 - January 2024 with additional funding on existing projects.

RoadMap

- M0: Design of a dynamically adaptable algorithm that can perform consensus a the required scale.
  - We have an algorithm for permissioned systems.
  - Extension to permissionless systems will be subject of research in related projects SOTERIA H2020 and Byblos ANR.
- M1-M3: Analysis of the state of the art:
  - Self-Sovereign-Identity Systems.
  - Broadcast-based blockchain alternatives (see other poster).
  - Verifiable Credentials (from our work on SOTERIA H2020).
- M4: High-level design of blockchain-free SSI.
- M6: First Prototype implementing Broadcast/Consensus.
- M9: Protytpe of Credential Management.
- M12: Prototype implementation of the complete system.

[1] Davide Frey, Mathieu Gestin, and Michel Raynal. "The Synchronization Power (Consensus Number) of Access-Control Objects: The Case of AllowList and DenyList". In: DISC 2023. Oct. 2023.

Workshop Organization

Blockchain & Privacy: International and Comparative Aspects

16 November 2023, University of the French Antilles - organized by Brunessen Bertrand et Sandrine Turgis
Goals:

- To have a very international approach to Blockchains & Privacy issues by crossing the eyes on the different legal regulations in the world.
- International and comparative aspects: not only the European point of view.
- Guadeloupe is well from this point of view because geographically we are close to geographical eras where a lot is happening.
- The workshop will take place the 16 November with American, Canadian, Brazilian and African colleagues.
- We are planning to publish the work in English with an Anglo-Saxon publisher to ensure maximum dissemination of this work.