Technical Whitepaper

Intelchain Team

Version 1.0

1.Introduction

Since the release of the Bitcoin whitepaper in 2008, the idea of blockchain has gained global recognition. Despite the increasing awareness of decentralized money and applications, design constraints have hindered Bitcoin's primary goals. Originally intended as a peer-to-peer payment system to facilitate transactions without intermediaries like banks, Bitcoin's popularity exposed its limitations, notably a throughput of approximately 7 transactions per second (TPS) and high transaction costs.

In 2014, Buterin et al. introduced Ethereum, a new blockchain framework that allowed developers to build diverse blockchain applications using "smart contracts." However, Ethereum also faced scalability issues, achieving only around 15 TPS, which was inadequate for high-throughput applications such as gaming or decentralized exchanges.

Given these performance constraints, numerous blockchain projects have sought to enhance transaction throughput. Some proposed replacing the Proof-of-Work (PoW) consensus with Proof-of-Stake (PoS), while others, like EOS, implemented Delegated Proof of Stake (DPoS), where block proposers are elected through voting rather than an on-chain algorithmic process. Projects such as IOTA adopted a Directed Acyclic Graph (DAG) structure to circumvent the limitations of sequential transaction processing.

Nevertheless, these solutions often compromise essential aspects such as security and decentralization. A promising approach that maintains both is sharding, which involves dividing the network into multiple groups (shards) of validators that process transactions concurrently. This method increases total transaction throughput proportionally with the number of shards. Zilliqa was the first public blockchain to tackle the scalability issue using sharding. However, it has two main drawbacks: it does not implement state sharding, which limits participation from devices with restricted resources, and its sharding mechanism is vulnerable to single-shard takeover attacks due to its PoW-based randomness generation.

We present Intelchain, an advanced sharding-based blockchain that is fully scalable, provably secure, and energy efficient. Intelchain addresses the shortcomings of existing blockchains by integrating cutting-edge research and engineering practices into a finely-tuned system. Intelchain's key innovations include:

- Fully Scalable
  Intelchain achieves full scalability by sharding not only network communication and transaction validation, as seen in Zilliqa, but also the blockchain state. This comprehensive sharding approach ensures that Intelchain can efficiently manage and distribute data across the network, making it a fully scalable blockchain.

* Secure Sharding
Intelchain's sharding process is provably secure due to its Distributed Randomness Generation (DRG) process, which is designed to be unpredictable, unbiasable, verifiable, and scalable. Furthermore, Intelchain periodically reshards the network in a non-disruptive manner, safeguarding against slowly adaptive Byzantine adversaries.
* Efficient and Fast Consensus
Intelchain employs a Proof-of-Stake (PoS) mechanism to select validators, making it energy efficient compared to other sharding-based blockchains that rely on Proof-of-Work (PoW). Consensus is achieved through a linearly scalable Byzantine Fault Tolerance (BFT) algorithm, which is 100 times faster than the traditional Practical Byzantine Fault Tolerance (PBFT).
* Adaptive-Thresholded PoS
The threshold of stakes required for a node to join the network is dynamically adjusted based on the total staking volume. This prevents malicious stakers from concentrating their power in a single shard while ensuring the threshold is low enough to allow participation from small stakers, enabling them to earn rewards.
* Scalable Networking Infrastructure
Utilizing RaptorQ fountain code, Intelchain can swiftly propagate blocks within and across shards using the Adaptive Information Dispersal Algorithm. Additionally, Intelchain adopts Kademlia routing to facilitate cross-shard transactions, scaling logarithmically with the number of shards.
Consistent Cross-Shard Transactions
Intelchain supports cross-shard transactions by enabling direct communication between shards. An atomic locking mechanism ensures the consistency and reliability of these transactions, maintaining system integrity.

By innovating on both the protocol and network layers, Intelchain provides a scalable and secure blockchain system capable of supporting the emerging decentralized economy. Intelchain enables applications previously impractical on blockchain, including high-volume decentralized exchanges, interactive fair games, Visa-scale payment systems, and Internet-of-Things transactions. Intelchain aims to scale trust for billions of people and foster a radically fair economy.

2. Consensus Mechanism

The consensus protocol is fundamental to any blockchain, determining how validators securely and efficiently agree on the next block. The first blockchain consensus protocol, powering Bitcoin, is Proof-of-Work (PoW). In PoW, miners compete to solve cryptographic puzzles, and the first to solve the puzzle gets to propose the next block and earn token rewards. PoW relies on the assumption that the majority of hashing power is controlled by honest nodes. The longest chain is considered the canonical one, which is why PoW is also known as chain-based consensus.

Practical Byzantine Fault Tolerance (PBFT) is another consensus protocol, studied for over two decades in academia. In PBFT, a "leader" node is elected, with the other nodes acting as "validators." The PBFT process involves two main phases: the prepare phase and the commit phase.

involves two major phases: the prepare phase and the commit phase. In the prepare phase, the leader broadcasts its proposal to all of the validators, who in turn broadcast their votes on the proposal to everyone else. The reason for the rebroadcasting to all validators is that the votes of each validator need to be counted by all other validators. The prepare phase finishes when more than $2f + 1$ consistent votes are seen, where f is the number of malicious validators, and the total number of validators plus the leader is $3f + 1$. The commit phase involves a similar vote counting process, and consensus is reached when $2f + 1$ consistent votes are seen. Due to the rebroadcasting of votes among validators, PBFT has $O(N^2)$ communication complexity, which is not scalable for a blockchain system with hundreds or thousands of nodes.

As an improvement on PBFT , Intelchain's consensus protocol is linearly scalable in terms of communication complexity, and thus we call it Fast Byzantine Fault Tolerance (FBFT). In FBFT, instead of asking all validators to broadcast their votes, the leader runs a multi-signature signing process to collect the validators' votes in a $O(1)$ -sized multi-signature and then broadcast it. So instead of receiving $O(N)$ signatures, each validator receives only one multi-signature, thus reducing the communication complexity from $O(N^2)$ to $O(N)$.

The idea of using $O(1)$ -sized multi-signature is inspired by ByzCoin's BFT which uses the Schnorr signature scheme for constant-sized multi-signature aggregation and forms a multicast tree among validators to facilitate the message delivery. However, a Schnorr multi-signature requires a secret commitment round, which leads to a total of two round-trips for a single multi-signature. Intelchain improves upon that by using BLS (Boneh–Lynn–Shacham) multi-signature, which only requires one round-trip. Therefore, FBFT is at least 50% faster than ByzCoin's BFT. Besides, Intelchain adopts RaptorQ fountain code to speed up the block broadcasting process (discussed in. The fountain code broadcasting technique also avoids a security issue in ByzCoin's original tree-based multicasting design.

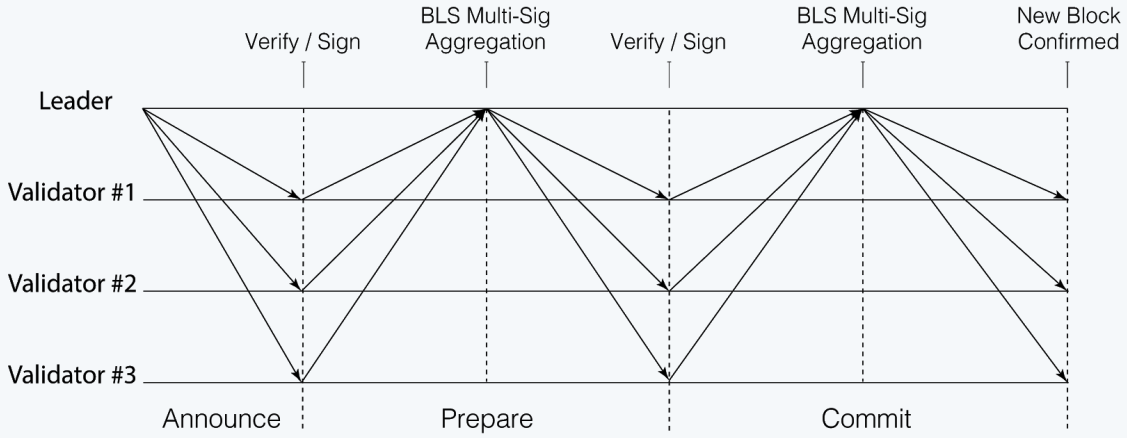

communication during a single round of consensus.

Specifically, Intelchain's FBFT consensus involves the following steps:

1. The leader constructs the new block and broadcasts the block header to all validators. Meanwhile, the leader broadcasts the content of the block with erasure coding (details discussed in. This is called the "announce" phase.
2. The validators check the validity of the block header, sign the block header with a BLS signature, and send the signature back to the leader.
3. The leader waits for at least $2f + 1$ valid signatures from validators (including the leader itself) and aggregates them into a BLS multi-signature. Then the leader broadcasts the aggregated multi-signature along with a bitmap indicating which validators have signed. Together with Step 2, this concludes the "prepare" phase of PBFT.
4. The validators check that the multi-signature has at least $2f + 1$ signers, verify the transactions in the block content broadcasted from the leader in Step 1, sign the received message from Step 3, and send it back to the leader.
5. The leader waits for at least $2f + 1$ valid signatures (can be different signers from Step 3) from Step 4, aggregates them together into a BLS multi-signature, and creates a bitmap logging all the signers. Finally, the leader commits the new block with all the multi-signatures and bitmaps attached, and broadcasts the new block for all validators to commit. Together with Step 4, this concludes the "commit" phase of PBFT.

The validators of Intelchain's consensus are elected based on Proof-of-Stake. Therefore, the actual protocol differs slightly from the one described above in a sense that a validator with more voting shares has more votes than others, rather than one-signature-one-vote. So instead of waiting for at least $2f + 1$ signatures from validators, the leader waits for signatures from the validators who collectively possess at least $2f + 1$ voting shares. The details of the proof-of-stake election mechanism will be discussed.

3. Sharding

Blockchain sharding as a scalability solution has gained lots of attention since late 2017. Various sharding solutions have been proposed both in industry and academia.

In industry, Zilliqa was the first sharding-based public blockchain that claimed a throughput of 2,800 TPS. Zilliqa uses PoW as identity registration process (i.e. Sybil attack prevention). Zilliqa's network contains a single directory-service committee and multiple shard committee (i.e. *network sharding)*, each containing hundreds of nodes. Transactions are assigned to different shards and processed separately (i.e. *transaction sharding)*. The resulting blocks from all shards are collected and merged at the directory-service committee. Zilliqa is not a *state sharding* solution because each node has to hold the entire blockchain state to be able to process transactions.

In academia, publications like Omniledger and RapidChain have proposed solutions that feature *state sharding* where each shard holds a subset of the blockchain state. Omniledger employs a multi-party computation scheme called RandHound to generate a secure random number, which is used to randomly assign nodes into shards. Omniledger assumes a *slowly*

adaptive corruption model where attackers can corrupt a growing portion of the nodes in a shard over time. Under such security model, a single shard can be corrupted eventually. Omniledger prevents the corruption of shards by reshuffling all nodes in the shards at a fixed time interval called *epoch*. RapidChain builds on top of Omniledger and proposes the use of the *Bounded Cuckoo Rule* to reshuffle nodes without interruptions.

Intelchain draws inspiration from these three previous solutions and designs a PoS-based full sharding scheme that's linearly scalable and provably secure. Intelchain contains a beacon chain and multiple shard chains. The beacon chain serves as the randomness beacon and identity register, while the shard chains store separate blockchain states and process transactions concurrently. Intelchain proposes an efficient algorithm for randomness generation by combining Verifiable Random Function (VRF) and Verifiable Delay Function (VDF). Intelchain also incorporates PoS in the sharding process which shifts the security consideration of a shard from the minimum number of nodes to the minimum number of voting shares.

3.1 Distributed Randomness Generation

Background

Various approaches have been proposed to assign nodes into shards such as randomness-based sharding location-based sharding and centrally-controlled sharding . Out of all the approaches, randomness-based sharding has been recognized as the most secure solution. In randomness-based sharding, a mutually agreed random number is used to determine the sharding assignment for each node. The random number must have the following properties:

1. Unpredictable: No one should be able to predict the random number before it is generated.
2. Unbiaseable: The process of generating the random number should not be biasable by any participant.
3. Verifiable: The validity of the generated random number should be verifiable by any observer.
4. Scalable: The algorithm of randomness generation should scale to a large number of participants.

Omniledger uses the RandHound protocol, which is a leader-driven distributed randomness generation (DRG) process that involves PVSS (Publicly Verifiable Secret Sharing) and Byzantine Agreement. RandHound is an $O(n * c^2)$ protocol that divides participant nodes into multiple groups of size c. It achieves the first three properties above but is impractically slow to qualify as scalable.

RapidChain takes a simpler approach by letting each participant perform VSS (Verifiable Secret Sharing) and using the combined secret shares as the resulting randomness. Unfortunately, this protocol is not secure because the malicious nodes can send inconsistent shares to different nodes. Besides, RapidChain does not describe how the nodes reach consensus on the multiple possible versions of reconstructed randomness.

In addition, Algorand relies on the VRF-based (Verifiable Random Function) cryptographic sortition to select the group of consensus validators. The Ethereum 2.0 design proposes the use of VDF (Verifiable Delay Function) to delay the revelation of the actual random number so as to prevent last-revealer attack. The VDF is a newly invented cryptographic primitive; it takes an adjustable minimum amount of time to compute and the result can be verified immediately.

Scalable Randomness Generation with VRF and VDF

Intelchain's approach combines the strengths of the solutions above. First, Intelchain's DRG protocol complexity is $O(n)$, which in practice is at least an order of magnitude faster than RandHound. Second, unlike RapidChain's simple VSS-based approach, ours is unbiasable and verifiable. Third, compared to Ethereum 2.0's solution, our approach uses BFT consensus to provide finality to the random number. Specifically, the protocol includes the following steps:

1. A leader sends an *init*message with the hash of the last block $H(B_{n-1})$ to all the validators.
2. For each validator i, after receiving the *init* message, a VRF is computed to create a random number r_i and a proof p_i: $(r_i, p_i) = VRF(sk_i, H(B_{n-1}), v)$, where sk_i is the secret key of validator i and v is the current view number of consensus. Then, each validator sends back (r_i, p_i) to the leader.
3. The leader waits until it receives at least $f+1$ valid random numbers and combines them with an XOR operation to get the preimage of the final randomness $pRnd$.
4. The leader runs BFT among all the validators to reach consensus on the $pRnd$ and commit it in block B_n.
5. After $pRnd$ is committed, the leader starts computing the actual randomness $Rnd = VDF(pRnd, T)$, where T is the VDF difficulty and is set algorithmically such that the randomness can only be computed after k blocks.
6. Once Rnd is computed, the leader initiates a BFT among all validators to agree on the validity of Rnd and finally commit the randomness into the blockchain.

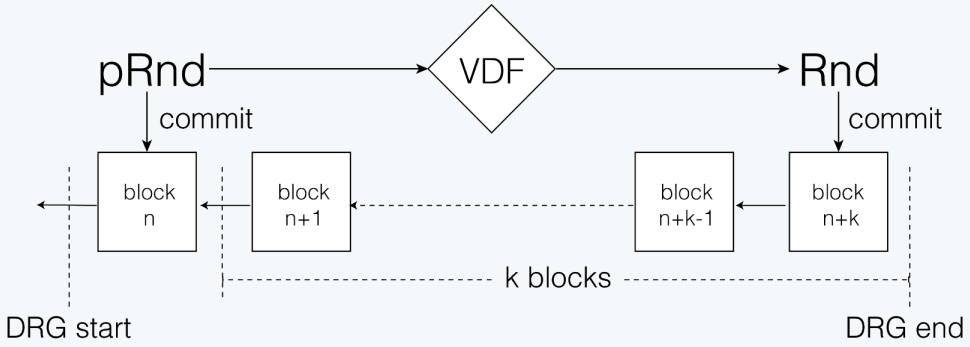

The VDF (Verifiable Delay Function) delays the revelation of the final randomness.

The VDF is used to provably delay the revelation of Rnd and prevent a malicious leader from biasing the randomness by cherry-picking a subset of the VRF random numbers. Because of the

VDF, the leader won't be able to know the actual final randomness before $pRnd$ is committed to the blockchain. By the time Rnd is computed with the VDF, $pRnd$ is already committed in a previous block so the leader cannot manipulate it anymore. Therefore, the best a malicious leader can do is to either blindly commit the randomness $pRnd$, or stall the protocol by not committing $pRnd$. The former is the same as the honest behavior. The latter won't cause much damage as the same timeout mechanism in PBFT will be used to switch the leader and restart the protocol.

We assume, in the long run, the existence of ASICs to compute VDFs, where a few altruistic nodes running an ASIC (Application-Specific Integrated Circuit) will publish the result, and no one could game the system. It is possible, before VDF ASICs are in production, that an attacker with a faster computing device could calculate the result before other honest nodes. Until this happens, the attacker can only know the randomness slightly before the honest nodes. While in principle the attacker could take advantage of this (e.g. withdrawing its fund if the bet on a smart contract was unfavorable to him), this problem can be mitigated on the smart contract layer with a proper delay, such that there should be a waiting period for the randomness to be committed to the protocol before a fund withdrawal is made possible.

3.2 Epochs

In Intelchain, the consensus and sharding process is orchestrated by the concept of epochs. An epoch is a predetermined time interval (e.g. 24 hours) during which the sharding structure is fixed and each shard continuously runs consensus with the same set of validators. At the beginning of each epoch, a random number will be generated using the DRG protocol and the sharding structure will be determined based on that randomness. Validators who want to validate transactions in epoch e need to stake their tokens during epoch $e-1$. The cutoff time for staking is before the randomness preimage $pRnd$ is committed into the blockchain.

3.3 Staking-based Sharding

Validator Registration

Sybil attack prevention is a key security consideration in public blockchains. Bitcoin and Ethereum require the miners to compute a cryptographic puzzle (PoW) before they can propose a block. Similarly, sharding-based blockchains like Zilliqa or Quarkchain also use PoW to prevent Sybil attacks. Intelchain adopts a different approach with proof-of-stake (PoS) as the validator registration or Sybil attack prevention mechanism. In order to become a Intelchain validator, prospective participants (or stakers) have to stake a certain amount of tokens to be eligible. The number of tokens staked will determine the number of voting shares assigned to the validator. Each voting share corresponds to one vote in the BFT consensus.

Sharding by Voting Shares

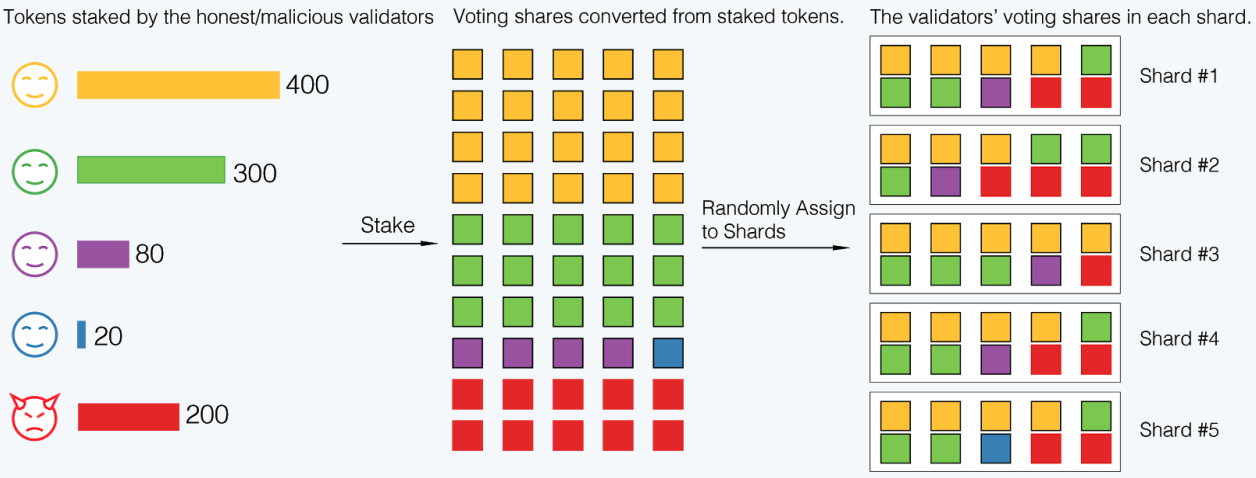

The stakers obtain voting shares proportional to their staked tokens. Voting shares are then randomly assigned to shards. Stakers become validators for the shard(s) where their voting shares are assigned.

A voting share is a virtual ticket that allows a validator to cast one vote in the consensus. Validators can acquire voting shares by staking tokens. The amount of tokens required for a voting share is algorithmically adjusted. At the beginning of each epoch, new validators' voting shares will be randomly assigned to shards. The new validators join the shard(s) where their voting shares get assigned. The consensus in a shard is reached by validators who collectively possess at least $2f + 1$ voting shares to sign the block.

To guarantee the security of a single shard, the amount of voting shares by malicious validators needs to be kept below $\frac{1}{3}$ of all the voting shares in that shard. This is required due to the nature of BFT consensus. Intelchain's adaptive thresholded PoS guarantees the above security requirement by adaptively adjusting the price of a voting share and assigning individual voting shares to shards rather than individual validators.

Our security assumption is that across all the staked tokens, up to $\frac{1}{4}$ of them belong to malicious validators. If we shard by validators (i.e. assign one validator to one shard), in the worst case where a single malicious validator holds $\frac{1}{4}$ of all the staked tokens (or the voting shares), it will easily possess more than $\frac{1}{3}$ voting shares in that shard. The reason is that the stakes at each shard is m times less than the stakes of the whole network, where m is the number of shards. We call this attack scenario a *large-stake attack* (a special type of single-shard takeover attack).

To prevent *large-stake attack*, instead of sharding by validators, we shard by voting shares (i.e. assign one voting share to one shard). Specifically, after the Rnd is revealed at the start of the current epoch, a random permutation (seeded with Rnd) on all the voting shares will be done and the permuted list of voting shares will be divided evenly into m buckets, where m is the number of shards. The voting shares falling in the ith bucket are assigned to shard i, so are the

corresponding validators. In practice, a single validator may be assigned to multiple shards if he possesses voting shares assigned to those shards. The shard leader is determined as the validator who possess the first voting share in the bucket.

It's worth noting that validators with larger stakes will have more chance of being selected as the leader. We argue that it's actually a desirable scenario because large stakers have more incentive to follow the protocol due to the fear of their stake being slashed, In addition, they are also more likely to possess more powerful machines with fast and stable network.

Adaptive-Thresholded PoS

The price of a voting share is set algorithmically so that it's small enough that malicious stakers can not concentrate their voting power in a single shard. Specifically, we set the price of a voting share to be P_{vote} tokens:

$$P_{vote} = \frac{TS_{e-1}}{NumShard * \lambda}$$

Here λ is a security parameter, $NumShard$ is the number of shards and TS_{e-1} is the total amount of tokens staked during epoch $e-1$.

Now we prove that when $\lambda > 600$, the chance of a single shard having more than $\frac{1}{3}$ malicious voting shares (i.e. probability of failure) is negligible.

Given the definition of P_{vote}, the total number of voting shares will be $N = \frac{TS_{e-1}}{P_{vote}} = NumShard * \lambda$. Given a trustable randomness source (discussed in §3.1) and the sharding process based on the randomness, the probability distribution of the number of malicious voting shares in each shard can be modeled as a hypergeometric distribution (i.e. random sampling without replacement):

$$P(X = k) = \frac{\binom{K}{k}\binom{N-K}{n-k}}{\binom{N}{n}}$$

Here N is the total number of voting shares, $K = \frac{N}{4}$ is the maximum number of malicious voting shares, $n = \frac{N}{NumShard}$ is the number of voting shares in each shard, and k is the number of malicious voting shares in a shard. The actual failure rate of a shard $P(X \leq k)$ follows cumulative hypergeometric distribution $CDF_{hg}(N, K, n, k)$ which, when N is large, degrades to binomial distribution (i.e. random sampling with replacement):

$$P(X \leq k) = \sum_{i=0}^{k} \binom{n}{i} p^i (1-p)^{n-i}$$

We can show that when n is large enough, the probability that a shard contains more than $\frac{1}{3}$ tokens held by malicious entities is negligible. In fact, when $n = 600$, the probability that a shard contains less than $_3$ malicious voting shares is $P(X \leq 200) = 0.999997$, which translates to a shard failure (i.e. consensus cannot be reached) rate of "once in around 1000 years" (given an epoch interval of 24 hours). Therefore, we will set $\lambda = 600$ to guarantee the high security of our shards. (Intuitively, λ governs minimum number of voting shares a single shard should contain. This is functionally similar to the minimum number of nodes in a shard as described in other PoW-based sharding solutions)

This approach is resistant to the fluctuation of the number of validators. We are not setting a lower limit on number of validators in each shard as in other solutions like Zilliqa. Instead, we adopt an adaptive PoS-based model to ensure that the malicious people can never occupy more than $\frac{1}{3}$ of the voting shares in a single shard, thus making it secure.

3.4 Resharding

We've described a secure sharding scheme that prevents malicious validators from overtaking a single shard. Nonetheless, if the sharding structure stays fixed, malicious attackers can still overtake a shard by corrupting the validators in that shard. There are three models of attackers:

1. Static Round-Adaptive: where attackers can only corrupt a subset of nodes at a predetermined stage. Elastico assumes attackers can only corrupt nodes at the beginning of each epoch.
2. Slowly Adaptive: where attackers can corrupt a subset of nodes over time during the epoch.
3. Fully Adaptive: where attackers can corrupt a subset of nodes instantaneously and at any time

Intelchain assumes the slowly adaptive corruption model under which the attacker can corrupt a constant number of nodes and it takes a certain amount of time. Omniledger assumes the same corruption model and it prevents the attack by replacing validators in all shards every epoch. This approach has two major problems. The first is the high cost of bootstrapping at every epoch. The second is the security concern when all nodes are being replaced during the consensus.

Intelchain mitigates these problems by adopting the Cuckoo-rule based resharding mechanism. After the end of an epoch, the validators who withdrew their stake will be evicted from the network, while those who keep their stakes stay. The new validators who staked during this epoch get new voting shares. These voting shares will be randomly assigned to the shards who have more than the median of the total voting shares. Next, a constant number of the voting shares from all shards will be randomly re-distributed to the other half of the shards who have less than the median of total voting shares. It's proven in that this resharding scheme can keep the voting shares in all shards balanced while fulfilling the security requirement.

3.5 Fast State Synchronization

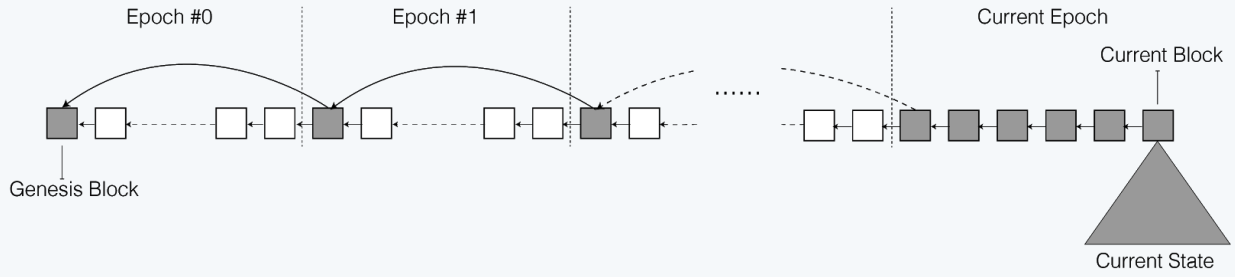

The first block of an epoch contains a hash link to the first block of last epoch. This allows fast state synchronization of new nodes where they can rely only on the blocks in grey to quickly verify the current state.

When validators join a new shard, they will need to quickly synchronize to the current state of the shard in order to validate new transactions. The traditional procedure of downloading the blockchain history and reconstructing the current state is too slow for resharding to be possible (it takes days to fully synchronize the Ethereum blockchain history). Fortunately, the current state is orders of magnitude smaller than the whole blockchain history. Downloading the current state within the time window of an epoch is feasible compared to downloading the whole history.

In Intelchain, new validators joining a shard first download the current state trie of that shard so they can start validating transactions quickly. To ensure the current state downloaded is valid, the new node needs to do proper verification. Instead of downloading the whole blockchain history and replaying all the transactions to validate the current state, the new node downloads historical block headers and validates the headers by checking their signatures. As long as there is a cryptographic trace (e.g. hash pointers and signatures) from the current state back to the genesis block, the state is valid. Nonetheless, signature verification is not computationally free and it takes a significant amount of time to verify all the signatures starting from the genesis block. To mitigate this problem, the first block of each epoch will include an additional hash pointer to the first block of the last epoch. This way, the new node can jump across the blocks within an epoch when tracing hash pointers to genesis block. This will significantly speed up the verification of the current blockchain state.

To further optimize the state synchronization process, we will make the blockchain state itself as small as possible. One observation from the Ethereum blockchain state is that a lot of accounts are empty and wasting the precious space of blockchain state. In Ethereum, the empty accounts with a specific nonce cannot be deleted because of potential replay attacks where old transactions are re-submitted on the deleted account. Intelchain will adopt a different model of avoiding replay attacks by letting the transactions specify the hash of the current block: a transaction is only valid before a certain number (e.g. 100) of blocks following the block of the specified hash. This way, the old accounts can be safely deleted and the blockchain state can be kept slim.

4. Shard Chain and Beacon Chain

4.1 Shard Chain

A shard chain is a blockchain that processes and validates its own transactions and stores its own state. A shard only processes transactions that is relevant to itself. Although a shard chain is relatively independent, it will communicate with other shard chains through cross-shard communication.

Cross-shard Communication

Cross-shard communication is a key component of any sharding-based blockchain. Cross-shard capability breaks the barrier between shards and extends the utility of a single shard beyond itself. Overall, there are three categories of cross-shard communication:

1. Main-chain-driven: Projects like Zilliqa rely on the main chain to achieve transactions across shards.
2. Client-driven: Omniledger proposed a client-driven cross-shard transaction mechanism where the messages between shards are collected and sent to shards by clients. This adds an extra burden to the client that is not desirable for an adhoc light client.
3. Shard-driven: RapidChain proposed that the messages between shards are directly sent by the nodes in the shard without external help.

Intelchain adopts the shard-driven approach for its simplicity and the absence of burden on clients. We believe the benefits of shard-driven communication outweighs its drawbacks. The cost on the overall network for shard-driven communication can be considerable because every cross-shard message is a network-level broadcast, which incurs a $O(N)$ network cost. To solve this problem, Intelchain uses the Kademlia routing protocol to reduce the communication complexity to $O(log(N))$. In addition, the data being communicated will be encoded with erasure code to ensure the robustness of cross-shard communication.

4.2 Beacon Chain

The Intelchain beacon chain is a special blockchain that serves additional purposes compared to the shard chains. In effect, the beacon chain is also a shard. Besides processing transactions, like other shard chains do, the beacon chain is in charge of two additional key functionalities: generating the random number and accepting stakes, which means that the beacon chain is the chain where stakers deposit their tokens to become validators.

The validators for the beacon chain are determined similarly as the other shard chains are. During the sharding assignment, the voting shares are randomly divided into $NumShard + b$ buckets, where the extra b buckets are for the beacon chain.

Hash Link from Shard Chain

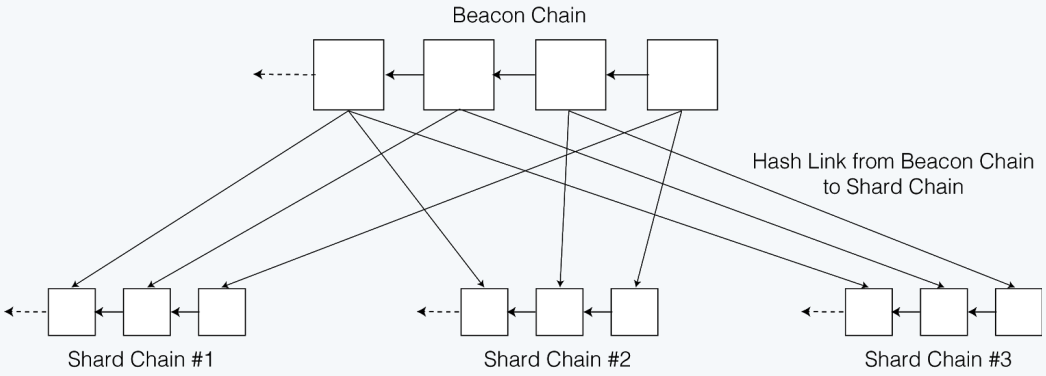

Hash link from beacon chain block to shard chain block.

The beacon chain helps strengthen the security and consistency of the shard chains' states by including the block header from each shard chain. Specifically, after a new block is committed to a shard chain, its block header will be sent (via Kademlia-based inter-shard communication) to the beacon chain. The beacon chain checks the validity of the block header by:

1. The hash of its previous block, which must have already been committed in the beacon chain;
2. The signers of the block's multi-signature, which must be the correct validators for that shard.

The committed block headers at the beacon chain will then be broadcasted to the whole network. Each shard will keep a chain of valid block headers for all other shards, which will be used to check the validity of transactions from other shards (i.e. simple payment verification). Adding the shard chains' block headers into the beacon chain serves two main purposes:

1. Increases the difficulty of attacking a single shard.
   Attackers have to corrupt both the shard chain and beacon chain in order to convince others that an alternative block in the shard chain is valid.
2. Reduce the network cost of broadcasting the block headers among shards.
   There will be a $O(N^2)$ network communication if we let each shard broadcast its headers separately. With the beacon chain as a central relay, the complexity is reduced to $O(N)$.

5. Blockchain State Sharding

Unlike other state-sharding blockchains that adopted UTXO (Unspent Transaction Output) data model, Intelchain's state sharding is applied on account-based data model. Each shard chain contains its own account state, and all the tokens in existence are spread among all the shard.

We treat the user account and the smart contract account differently in sharding. An user account can have multiple balances at different shards (e.g. 100 tokens at Shard A and 50 tokens at Shard B). A user account can move its balance between shards by issuing a cross-shard transaction. A smart contract account is limited to the specific shard where the contract was created. However, for a decentralized application that requires more throughput than a single shard can handle, the Dapp (Decentralized Application) developer can instantiate multiple instances of the same smart contract in different shards and let each instance handle a subset of the incoming traffic. Note that the different instances of the same smart contract do not share the same state, but they can talk to each other via cross-shard communication.

6. Networking

Previous research has pointed out that network capacity is one of the major bottlenecks for blockchain systems. In order to increase performance, Intelchain focuses on improving the efficiency of network utilization. Intelchain also proposes a number of improvements to deal with real-world networking scenarios.

6.1 Kademlia-based Routing

Inspired by RapidChain , we will adopt Kademlia as the routing mechanism for cross-shard messages. Each node in the Intelchain's network maintains a routing table that contains nodes from different shards. The distance between shards is defined as the XOR distance of the shard IDs. When a message from shard A needs to be sent to shard B, the nodes in shard A will look at the routing table and send the message to the nodes with the closest shard ID. With Kademlia-based routing, a message only travels across $O(logN)$ nodes before it reaches the destination shard. Compared to normal gossip broadcasting, which requires a $O(N)$ network complexity, the Kademlia routing mechanism can significantly reduce the overall network load in a sharded blockchain.

6.2 Efficient Broadcasting with Erasure Code

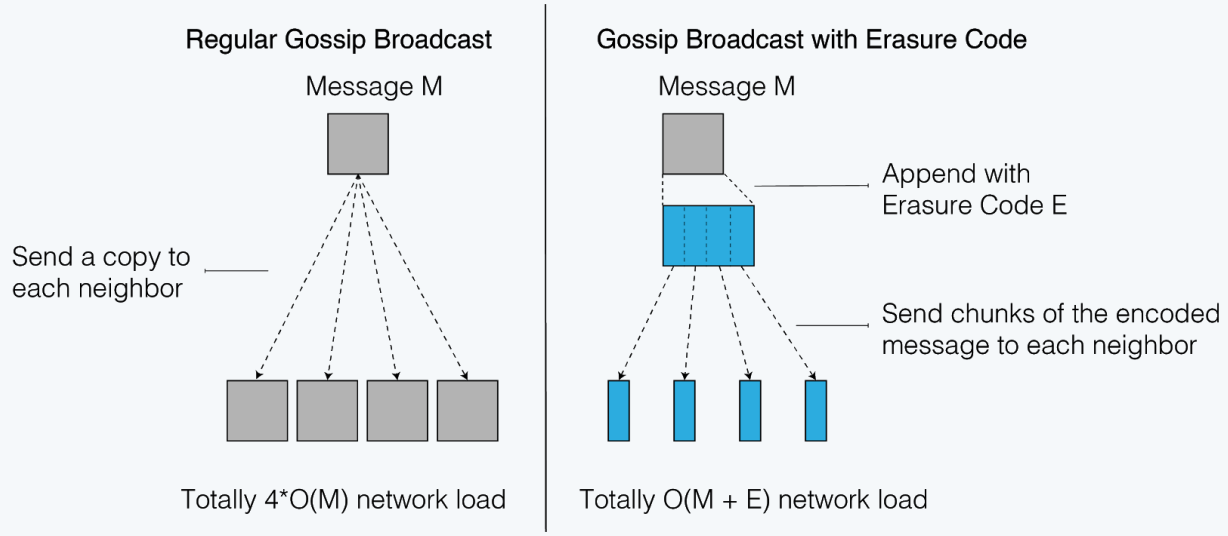

Comparison between normal gossip broadcast with gossip broadcast with erasure code.

Broadcast is a frequent network action in any blockchain system built on P2P (Peer-to-Peer) network overlay. Specifically in our consensus protocol, there are three scenarios where broadcasting is needed:

1. A newly proposed block needs to be broadcasted by the leader to all validators.
2. A newly generated master chain block needs to be broadcasted to the whole network.
3. The cross shard communication requires the broadcast of a message between shards.

In a normal P2P broadcasting, the original sender needs to send a copy of the message to each of its neighbors. This will incur $O(d * M)$ network load on the sender, where d is the average number of neighbors of the sender and M is the message size. Instead, in Intelchain a sender first encodes the message with erasure code and then send chunks of the encoded message to each neighbor. This reduces the load on the sender to $O(M + e)$ where e is the size of erasure code and it is usually smaller than the size of the original message M. Therefore, Intelchain's network broadcasting mechanism significantly lowers the network load of the broadcast sender. In addition, Intelchain proposes to improve IDA's robustness by replacing the original Reed-Solomon erasure code with RaptorQ fountain code so that the broadcaster can always send more erasure codes to further ensure the data is eventually received.

6.3 FEC-based Unicast

Traditional reliable transports such as TCP relies upon retransmission and ACK-based signaling in order to deal with lost packets. This is known to introduce latency spikes proportional to the round-trip time between the sender and receiver. Also, window-based congestion control—such as Reno, NewReno, and CUBIC in use by most TCP implementations—are all

additive increase/multiplicative decrease (AIMD) algorithms, whose bandwidth is known to be severely impacted by transient packet losses.

Intelchain uses the RaptorQ fountain code to combat these two problems. Each message is encoded into symbols, and symbols are sent over the wire until the receiver acknowledges successful decoding of the message using the symbols that it received. Unlike using fixed-rate codes such as Reed–Solomon where the transmission fails once the symbols have been exhausted, fountain code enables infinite, just-in-time generation and use of encoding symbols.

6.4 Support for Home Nodes

P2P nodes on a typical residential network pose a major, distinctively unique problem: They cannot be reached from the outside unless mediated by their residential internet router, which employs a technique called network address translation (NAT). Support for inbound traffic by these routers vary, and different approaches have been developed to work around different types of routers. In particular, routers implementing symmetric NAT cannot easily be worked around unless explicitly configured to support other hole-punching mechanisms such as Internet Gateway Device Protocol (IGDP).

Intelchain's P2P layer tries to detect the NAT mechanism behind which a node operates and employs the right workaround mechanism, such as STUN, TURN, IGDP, etc. In particular, Intelchain implements the overall detection and mitigation protocol named ICE (Interactive Connectivity Establishment).

6.5 Support for Locator Mobility

Nodes may change their IP addresses, with some type of nodes more so than others. One such example is a laptop, which may frequently hop between different Wi-Fi networks, with its IP address changing each time. When an IP address of a node changes, all existing transport connections that use the IP address as a local or remote endpoint are interrupted, and applications directly using such transport connections need to re-establish connections using the new IP address in order to continue. Such a connection handover is hard to implement correctly with minimal application-layer service interruption. Also, handling connection handover often complicates application-layer protocols (such as base consensus protocols).

Intelchain's network layer, in order to solve this problem, introduces a clean separation between node identity (cryptographic key pair possessed by the node) and node locators (network/transport-layer locator where the node can be reached) using the industry-standard Host Identity Protocol Version 2 (HIPv2). HIPv2 lets locators of a node change over time while keeping the node identity, by providing mechanisms for locator discovery, node-to-node security association, and tunneling of upper-layer traffic associated with local/remote node identity as endpoints.

7. Incentive Model

7.1 Consensus Rewards

After the successful commitment of a block, a protocol-defined number of new tokens will be rewarded to all validators who signed the block in proportion to their voting shares. The transactions fees are rewarded to validators similarly.

7.2 Stake Slashing

For any misbehaviors detected by the network, a certain amount of staked tokens will be slashed. For example, if a leader failed to finish the consensus process and triggered the leader change process, P_{vote} staked tokens will be slashed. If validators are proven to sign a dishonest block, all of their stake under the same shard will be slashed. This severe punishment is meant to strongly discourage any dishonest behavior and make the network as secure as possible. A proof of misbehavior can be two signed blocks that conflict with each other. Any validator can submit a transaction to prove the misbehavior of other validator and if verified, the slashed token will be rewarded to the prover(s).

7.3 Stake withdrawal

Long-range Attacks

Proof-of-stake blockchains, unlike proof-of-work blockchains, tend to suffer from *long-range attacks* . These are attacks that leverage the fact that proofs are based on signatures rather than on resource-intensive tasks. In a long-range attack, the private keys of honest validators are stolen long after they have been used, and the attacker is able to create a forked blockchain by signing fake blocks with those keys. When this happens, new validators joining the network have no way to distinguish between the original, legitimate chain and the attacker's simulated chain.

Long-range attacks happen in the following two scenarios. Private key can be compromised either by a lack of security on validators, or more commonly, by the fact, after a validator withdraw their token, he could financially benefits if an attacker which would be looking to buy its private key. Also, by design each set of validators is trusted to approve the block of transactions that also determines the next set of validators. After enough private key (i.e. those that collectively hold

more than $\frac{2}{3}$ voting shares in a shard) has been compromised, an attacker has total control on who the subsequent validators is.

Long-range Defense: Resonant Quorums

Proof-of-work blockchain protects against the above attacks by giving honest validators an objective method of fork choice. In a proof-of-work blockchain, the fork choice to select the canonical chain is the accumulated amount of work done in terms of hashes computed.

In a proof-of-stake blockchain, the only objective measure that can be used to select between forks is the total weighting of signatures used to approve each block. If we use these weighted signatures to compare two different blocks, we come to the following equation to determine when a chain may be forked:

Safety = “Block approval key weight” - “Compromised key weight”

The “Block approval key weight” means the voting power of the keys that signed on the block. If, by stake weight, more private keys are compromised than were used to approve of a block, then the block can be forked. Until then, validators will always prefer the original, legitimate version of the block.

Intelchain maximizes the safety of each block in its proof-of-stake blockchain by maximizing this equation. It is infeasible to disincentivize leaking private keys in the long term. Intelchain instead incentivizes validators to maximize the approval weight of each block after a quorum has been achieved. This is done by requiring validators to sign each quorum-approved block before allowing those validators to withdraw their stake. These new additional signatures only need to exist within the blockchain, and they do not need to be generated at consensus time for each block. Because of this, the new signatures can be added to subsequent blocks when validators decide to withdraw their stake, and so they may freely improve the safety of the chain without impacting its liveness.

8. Future Research

8.1 Fraud Proofs

The capability of proving the misbehavior of validators is important for a light client to trust the block data they received. In the case of cross-shard communication, each shard is a light client of other shards. Ensuring that messages sent between shards are trustable is crucial for inter-shard data consistency. We are actively researching the topic of data availability and fraud proofs to securitize our protocol.

8.2 Stateless Validators

In a high throughput blockchain, the size of the blockchain data will grow faster than existing chains, which is a major problem for new validators to sync up quickly. This makes the resharding process problematic because if new validators can’t sync up in time, then the quorum of validators may not be reached for a new block to be approved, and even if the quorum is met, the security of

the protocol would be reduced. State block pruning is one mitigation to the problem, but it's not optimal since the state itself can grow large. We are actively looking into enabling stateless client where validators doesn't have to sync up the full state to validate transactions.

References

[1] J.R. Douceur, The Sybil attack, in: 1st International Workshop on Peer-to-Peer Systems (IPTPS 02), 2002.

[2] Al-Bassam, M., Sonnino, A., & Buterin, V. (2018). Fraud Proofs: Maximising Light Client Security and Scaling Blockchains with Dishonest Majorities. CoRR, abs/1809.09044.

[3] Vasin, P. (2014) Blackcoin's Proof-of-Stake Protocol v2, https://blackcoin.co/blackcoin-pos-protocolv2-whitepaper.pdf

[4] A. Kiayias, I. Konstantinou, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. Cryptology ePrint Archive, Report 2016/889, 2016. http://eprint.iacr.org/.

[5] P. Daian, R. Pass and E. Shi, Snow White: Robustly reconfigurable consensus and applications to provably secure proofs of stake, Cryptology ePrint Archive, Report 2016/919, 2017.

[6] Rafael Pass and Elaine Shi. Thunderella: Blockchains with optimistic instant confirmation. https://eprint.iacr.org/2017/913.pdf.

[7] The Zilliqa Team. The zilliqa technical whitepaper. https://docs.zilliqa.com/whitepaper.pdf, August 2017.

[8] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. Available at https://bitcoin.org/ bitcoin.pdf.

[9] Miguel Castro and Barbara Liskov. Practical Byzantine Fault Tolerance. In Proceedings of the 3rd Symposium on Operating Systems Design and Implementation (OSDI '99), New Orleans, Louisiana, February 1999.

[10] E. Kokoris-Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford. Enhancing Bitcoin Security and Performance with Strong Consistency via Collective Signing. In Proceedings of the 25th USENIX Conference on Security Symposium, 2016.