Dfinity: Internet computer: under partial synchrony Yu Tang, Shuhan Chen

Background: Dfinity; Internet Computer Blockchain

Internet Computer Blockchain:

"the world's first web-speed, web-serving public blockchain network that can scale"



Background: Dfinity; Internet Computer Blockchain

The Internet Computer hosts special smart contracts, called "canisters"



Background: Dfinity; Internet Computer Blockchain



Background: Related Work

• PBFT

- Optimistic responsive
- HotStuff
 - Eliminates all-to-all communication steps of PBFT
 - Does not rely on view-change; rely on a "pacemaker" subprotocol
 - Optimistic responsive
- Tendermint / Algorand
 - Rely on a peer-to-peer gossip sub-layer for communication
 - Not optimistically responsive
- MirBFT
 - Instances of PBFT run concurrently

Background

- The Internet Computer Consensus (ICC) protocols
 - Dynamic collection of intercommunicating replicated state machines
 - ICC0, ICC1, ICC2
 - Features:
 - Fully specified: no unspecified protocols
 - Extremely simple: no complicated subprotocols
 - Robust consensus: peak performance is partially sacrificed to ensure reasonable performance
 - ICC1: peer-to-peer gossip sub-layer; ICC2: subprotocol for reliable broadcast
 - Optimistically responsive

Assumptions

- n parties, at most t < n/3 corrupt parties
- Static corruption
- Partial synchronous
- Blockchain based
- Proceed in rounds
 - A random beacon to assign to each party a rank; lowest rank is the leader of the round
- Local computation is negligible relative to network latency
- The only type of communication: broadcast
- Each party has a pool which holds the set of all messages received from all parties

Protocol ICC0: Primitives

Threshold signature scheme:

(t, h, n)-threshold signature scheme:

- n parties are initialized with a public-key/secret-key pair
- public keys for all n parties
- a global public key

- Signing algorithm
- Signature share verification algorithm
- Signature share combining algorithm
- Signature verification algorithm

Blocks

For k>=1, a round-k block B: (block, k, α, phash, payload)

- Authentic
- Valid
- Notarized
- Finalized

Root serves as its own authenticator, notarization, and finalization.

Protocol Components

- A collision resistant hash function, H
- A signature scheme S_auth
- An instance S_notary of a (t, n t, n)-threshold signature scheme
- An instance S_final of a (t, n-t, n)-threshold signature scheme
- An instance S_beacon of a (t, t + 1, n)-threshold signature scheme, used to implement a random beacon

High level description of the protocol

• Under cryptographic assumptions, it is guaranteed that in each round k >=1:

P1 (deadlock freeness): at least one notarized block of depth k will be added to the block-tree.

P2 (safety): if a notarized block of depth k is finalized, then there is no other notarized block of depth k.

P3 (liveness): if the network is synchronous over a short interval of time beginning at the point in time where any honest party first enters round k, and the leader in round k is honest, then the block proposed by the leader in round k will be finalized.

Invariants in Sync settings

I An honest leader's block will be uniquely certified in an iteration.

II In each iteration, at least one block, but possibly many, will be certified.

III At the end of iteration k, if all iteration (k-1) blocks extend the same B_{k-2} , then B_{k-2} is uniquely extendable. i.e., no other iteration (k-2) block can be extended from then on.

Protocol details

- Authentic
- Valid /consensus/src/consensus/validator.rs
- Notarized /consensus/src/consensus/notary.rs
- Finalized /consensus/src/consensus/finalizer.rs

- Tree Building Subprotocol
- Finalization Subprotocol

Protocol details: Tree Building Subprotocol

```
broadcast a share of the round-1 random beacon
For each round k = 1, 2, 3 \dots:
     wait for t + 1 shares of the round-k random beacon
     compute the round-k random beacon (which defines the permutation \pi for round k)
     broadcast a share of the random beacon for round k + 1
     let r_{\rm me} be the rank of P_{\alpha} according to the permutation \pi
     \mathcal{N} \leftarrow \emptyset // the set of blocks for which notarization shares have been broadcast by P_{\alpha}
     \mathcal{D} \leftarrow \emptyset // the set of ranks disgualified by P_{\alpha}
     done \leftarrow false
     proposed \leftarrow false
     t_0 \leftarrow \text{clock}()
     repeat
           wait for either:
                (a) a notarized round-k block B.
                         or a full set of notarization shares for some valid
                         but non-notarized round-k block B:
                      // Finish the round
                      combine the notarization shares into a notarization for B, if necessary
                      broadcast the notarization for B
                      done \leftarrow true
                      if \mathcal{N} \subseteq \{B\} then broadcast a finalization share for B
                (b) not proposed and clock() \geq t_0 + \Delta_{\text{prop}}(r_{\text{me}}):
                      // Propose a block
                      choose a notarized round-(k-1) block B_{\rm p}
                      payload \leftarrow getPayload(B_p)
                      create a new round-k block B = (block, k, \alpha, H(B_p), payload)
                      broadcast B, B's authenticator, and the notarization for B's parent
                      proposed \leftarrow true
                (c) a valid round-k block B of rank r such that
                         B \notin \mathcal{N}, r \notin \mathcal{D}, \operatorname{clock}() \geq t_0 + \Delta_{\operatorname{ntrv}}(r), \text{ and }
                         there is no valid round-k block B^* of rank r^* \in [r] \setminus \mathcal{D}:
                      // Echo block B
                            and either broadcast a notarization share for it or disqualify its rank
                      if r \neq r_{me} then
                           broadcast B, B's authenticator, and the notarization for B's parent
                      if some block in \mathcal{N} has rank r
                           then \mathcal{D} \leftarrow \mathcal{D} \cup \{r\}
                           else \mathcal{N} \leftarrow \mathcal{N} \cup \{B\}, broadcast a notarization share for B
     until done
```

Protocol details: Tree Building Subprotocol

Delay functions: [n] $\rightarrow R{\geq}0$,based on the rank of the proposer-

- Δ prop: delay proposing a block
- Δ ntry: delay generating a notarization share on a block

Liveness: $2\delta + \Delta_prop(0) \le \Delta_ntry(1)$, δ : network delay

Protocol details: Finalization Subprotocol



Consensus Properties

Messages are placed in blocks. We reach agreement using a blockchain.



Suppose we have *n* replicas running a subnet together.

Safety: If two (honest) replicas think that the *i*-th block is agreed upon, they must have the same block Liveness: at some point, every (honest) replica will think that the *i*-th block is agreed upon, for any *i* Validity: all agreed upon blocks are valid

> We want this even if up to f of the n nodes misbehave, with n = 3f+1

In examples on these slides, we use n = 4, f = 1

Block Maker



Note: We need sufficiently many replicas to take the role of block maker each round. If we were to choose only one, we would not be fault tolerant.

Random Beacon

At every height, we have a Random Beacon, which is an unpredictable random value shared by the replicas



Notarization

Block Proposals may be invalid. Every block in the chain must be *notarized*. The notarization process ensures that every round a *valid* block proposal is published.

Step 1

Replica 1 receives a block proposal for height 30, building on some notarized height 29 block Step 2

Replica 1 sees that the block is valid, signs it, and broadcasts its *notarization* share

Step 3

Replica 1 sees that replicas 3 and 4 also published their notarization shares on the block

Step 4

3 notarization shares are sufficient approval: the shares are aggregated into a single full notarization. Block 30 is now notarized, and notaries wait for height 31 blocks









A Notary May Notary-Sign Multiple Blocks

To ensure that at least one block becomes fully notarized, notaries may notary-sign multiple blocks



Notarization with Block Maker Ranking

We can reduce the number of notarized blocks per round by letting notaries use the block maker ranking in the notarization process.

- The time following a round start is divided into time slots
- Initially, in time slot 0, a notary only notary-signs a block proposal from the rank-0 block maker
- After some timeout, if they still haven't seen a valid rank 0 block proposal, slot 1 starts, and the notary is also willing to notary-sign a rank 1 block proposal. After more timeout, slot 2 starts, and notaries would accept rank 2 blocks if no better block proposal was observed, etc.
- If the network works well and the rank 0 block maker sends a valid block that reaches notaries before they are willing to notary-sign other blocks, this will be the only notarized block at that height



Notarization with Block Rank

By taking block rank into consideration, we can reduce the amount of notarized blocks



Finalization

Notaries create "finalization shares" on a block if they did not notary-sign any other block at that height.



n=4, f=1:

P1: at least one notarized block of depth k will be added to the block-tree.

P2: if a notarized block of depth k is finalized, then there is no other notarized block of depth k.

P3: if the network is synchronous over a short interval of time beginning at the point in time where any honest party first enters round k, and the leader in round k is honest, then the block proposed by the leader in round k will be finalized.

Analysis: Expected Message Complexity and Latency

- Partial synchrony (depend on value of h, the rank of the leading honest party in round k)
 - Message complexity in round k is: $O((h+1)n^2)=O(n^3)$
 - Latency:

(iv) the communication network is δ -synchronous at all times throughout the interval

 $[T, T + \Delta_0(h, \delta) + 2h\delta],$

where

$$\Delta_0(h,\delta) \coloneqq \max(2\delta + \Delta_{\text{prop}}(h), \delta + \Delta_{\text{ntry}}(h)).$$

Then the all honest parties finish round k by time T plus

 $\Delta_0(h,\delta) + (2h+1)\delta.$

Analysis: Expected Message Complexity and Latency

$$\Pr[h \ge i] = \frac{t}{n} \cdot \frac{t-1}{n-1} \cdot \dots \cdot \frac{t-(i-1)}{n-(i-1)} < \frac{1}{3^i}.$$

$$\mathbf{E}[h] = \sum_{i \ge 1} \Pr[h \ge i] \le \sum_{i \ge 1} \frac{1}{3^i} = \frac{1}{2}.$$

- With probabilistic analysis, the expected message complexity is $O(n^2)$.
- With probabilistic analysis, the expected latency is bounded by:

 $\Delta_{\text{bnd}} + 3\delta + \max(\epsilon, \delta).$

Analysis: Comparison

- Sync Dfinity (static adversary)
 - Message complexity:
 - basic design: unbounded
 - With equivocation check: O(n²)
 - Latency: $O(\Delta)$
 - optimistic case (c << Δ): 8 Δ ; pessimistic case (c = Δ): 14 Δ
- Partial sync Dfinity: ICC0
 - Message complexity:
 - O(n²) with overwhelming probability
 - Latency:

$$\Delta_{\text{bnd}} + 3\delta + \max(\epsilon, \delta).$$

Comparison to [HMW18, AMNR18]

- Proposing step
- only guaranteed safety in a synchronous setting
- not optimistically responsive
- potentially unbounded communication complexity.

Shortcomings

- Difficulty on adaptive adversaries scenario (round k rank determined in round k-1)
- 2. Need to disseminate secret keys to parties