Hooks: A Simple and Modular Checkpointing Protocol for Blockchains

Pedro Antonino, Antoine Durand, Namrata Jain, Garry Lancaster, Jonathan Lawrence, A. W. Roscoe

The Blockhouse Technology Ltd., Oxford, UK

October 24, 2024

BLOCKHOUSE
The Blockhouse Technology Ltd.

Summary

BLOCKHOUSE
The Blockhouse Technology Ltd.

Table of contents

BLOCKHOUSE
The Blockhouse Technology Ltd.

Finality issues in blockchain

BLOCKHOUSE
The Blockhouse Technology Ltd.

Many issues recognized in the literature:

- Non-instant finality.
- Long-range attacks.
- Posterior corruptions.
- Trusted bootstrap.

Finality issues in blockchain

BLOCKHOUSE
The Blockhouse Technology Ltd.

- All related and solvable, but requires careful considerations !
- Our proposed solution: A checkpointing layer.
  - $\rightarrow$ Simple and modular.

Table of contents

The blockchain model

BLOCKHOUSE
The Blockhouse Technology Ltd.

The standard State Machine Replication-style definition:

Transaction Ledger (informal)

Nodes propose transactions and they output the final ones.

- Safety : the final transactions are the same for all honest nodes.
- Liveness : proposed transactions eventually becomes final.

The checkpoint layer properties

BLOCKHOUSE
The Blockhouse Technology Ltd.

- Takes an underlying blockchain $\mathcal{B}$.
- $\mathcal{B}$ may have weak Safety properties.
- The checkpoint layer:
  - same structure as a blockchain.
  - provides stronger Safety properties.

What is a "weak" blockchain?

BLOCKHOUSE
The Blockhouse Technology Ltd.

- What is a weak transaction ledger?
  - *E.g.*, for Bitcoin, the (probabilistic) security bound depends on the network delay.
  - For Proof-of-Stake protocols, past participants may cause Safety issues.
  - More generally, Safety could be broken in unexpected situations.

What is a "weak" blockchain?

BLOCKHOUSE
The Blockhouse Technology Ltd.

We sidestep the problem and define the weakest Safety property that works.

TLBS - Time-Limited Block Safety (informal)

$TLBS(h)$ holds iff, honest nodes agrees on the L blocks from height h to height $h + L$, from the time that the first block is known until the last block is known by all.

Our assumptions

BLOCKHOUSE
The Blockhouse Technology Ltd.

With input blockchain $\mathcal{B}$:

- TLBS *only for Liveness*.
- Sybil-resistance through *L*-Chain Quality.
  - Within *L* consecutive blocks, there is less than a third of malicious block authors.
- $\mathcal{B}$'s Liveness.
- The Secure Deletion assumption.
  - Honest nodes can irrevocably delete their state.
- The execution model taken from $\mathcal{B}$.

Table of contents

BLOCKHOUSE
The Blockhouse Technology Ltd.

Hooks as a checkpointing layer

BLOCKHOUSE
The Blockhouse Technology Ltd.

Checkpoint Safety

Only a single checkpoint will ever be created for every block height.

Checkpoint Liveness

Hooks is live as long as *TLBS* holds.

Per-block overhead is $O(1)$.

Safety Improvements

BLOCKHOUSE
The Blockhouse Technology Ltd.

Our Safety property is stronger than the Transaction Ledger Safety.

- Mitigates long-range attacks, *e.g.*, in case of Posterior Corruptions.
- Online nodes are immune because they have time-related information.
  - Only joining nodes are concerned.
  - The checkpointing proofs are sufficient for nodes to join.

  $\rightarrow$ Free property : Trustless Bootstrap.

Table of contents

The birds eye

BLOCKHOUSE
The Blockhouse Technology Ltd.

The algorithm, in short.

1. Block b author include its public key in b.
2. If b's L-th descendant becomes final, the author signs it.
   - This signature is called a *hook*.
3. Submit the hook and delete the key
4. A blocks is checkpointed if its L-depth subtree has $\frac{2}{3}$rd of blocks hooked.
   $\rightarrow$ The hooks set is a checkpoint proof !

The birds eye

BLOCKHOUSE
The Blockhouse Technology Ltd.

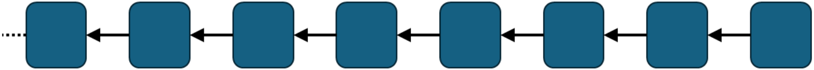

The birds eye

BLOCKHOUSE
The Blockhouse Technology Ltd.

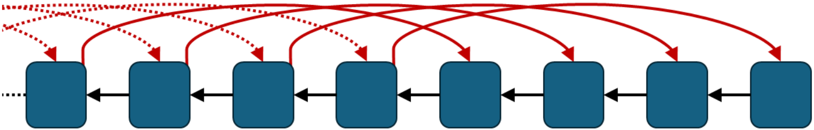

The birds eye

BLOCKHOUSE
The Blockhouse Technology Ltd.

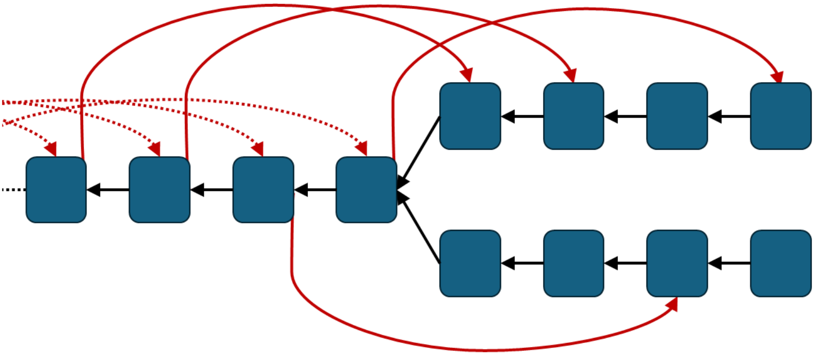

The formal arguments

BLOCKHOUSE
The Blockhouse Technology Ltd.

- When there is a fork at height h, there are L common block authors to vote on branches
- By quorum intersection, there is *at most* one branch checkpointed.
  - Any two sets of $\frac{2}{3}L$ hooks intersects at at least one *honest* node.
  - Honest nodes will never send two hooks.
- If $TLBS(h)$ does not holds, it might be none !

Some additional details

- *Ignore* non-checkpointed branches.
- Wait until your own block is checkpointed before sending a hook.

A weaker alternative version

BLOCKHOUSE
The Blockhouse Technology Ltd.

With $\frac{2}{3}$ honesty in Chain Quality, Hooks cannot be applied to honest majority blockchains.

- $\frac{1}{2}$ honesty also works, but we have weaker Safety.
- The algorithm must be modified to track *equivocating* hooks.
- Equivocating hooks may cause checkpoints to be (eventually) invalidated.

Weak Safety (informal)

If *TLBS* does not hold at some height, then there may be multiple checkpointed branches. In this case, eventually none of them will be checkpointed.

Some possible improvements

BLOCKHOUSE
The Blockhouse Technology Ltd.

- Hooks can be aggregated into a single signature for short checkpoint proofs.
- Avoid storing the node public key with key-evolving signatures.
- Make the analysis in the Universal Composability framework.

Conclusion

BLOCKHOUSE
The Blockhouse Technology Ltd.

In short, we take a weak blockchain, and,

- Prevent many safety issues when possible (*e.g.*, asynchrony, posterior corruption)
- Otherwise will only break Liveness.
- Offers trustless bootstrap/long-range attack resistance.
- Keep performance unaffected (experimentally confirmed).
- And possibly more (*e.g.* Quantum Resistance).

The end

BLOCKHOUSE
The Blockhouse Technology Ltd.

Thank you for your attention

$$
\begin{aligned}
TLBS(h) := \ & \text{Let } t_1 := \min\{t' \mid \exists i \in \mathcal{H},\ \mathcal{F}_i^{t'}.h = h\} \\
& \text{Let } t_2 := \min\{t' \mid \forall i \in \mathcal{H},\ \mathcal{F}_i^{t'}.h = h + l\} \\
& \forall h' \in [h, h + l], \\
& \# \bigcup_{\substack{t \in [t_1, t_2] \\ i \in \mathcal{H}}} \{\mathcal{F}_i^t[h]\} \leq 1
\end{aligned}
$$

$$
\begin{aligned}
TLBS(h) := \ & \text{Let } t_1 := \min\{t' \mid \exists i \in \mathcal{H},\ \mathcal{F}_i^{t'}.h = h\} \\
& \text{Let } t_2 := \min\{t' \mid \forall i \in \mathcal{H},\ \mathcal{F}_i^{t'}.h = h + l\} \\
& \forall h' \in [h, h+l], \\
& \# \bigcup_{\substack{t \in [t_1, t_2] \\ i \in \mathcal{H}}} \{\mathcal{F}_i^t[h]\} \leq 1
\end{aligned}
$$