

Fraud Proofs

Maximising Light Client Security and Scaling Blockchains with Dishonest Majorities


Mustafa Al-Bassam

8 November 2018

Motivation

@musalbas


- ▶ Currently there is a huge trade off between **decentralisation** and on-chain **scalability**.



Motivation

@musalbas

- ▶ Currently there is a huge trade off between **decentralisation** and on-chain **scalability**.
- ▶ Because light clients (non-fully validating nodes) will accept invalid blocks.
 - ▶ They assume that the majority of the consensus is honest.



Can we reduce this tradeoff?

@musalbas

- ▶ The big question: how can we make non-fully validating nodes (light clients) reject invalid blocks, so that they don't have to trust miners?

Can we reduce this tradeoff?

@musalbas

- ▶ The big question: how can we make non-fully validating nodes (light clients) reject invalid blocks, so that they don't have to trust miners?
- ▶ Let's use fraud proofs and data availability proofs!

Read the full paper (33 pages)

@musalbas

► Fraud Proofs: Maximising Light Client Security and Scaling Blockchains with Dishonest Majorities

Mustafa Al-Bassam (UCL)
Alberto Sonnino (UCL)
Vitalik Buterin (Ethereum)

► Paper:

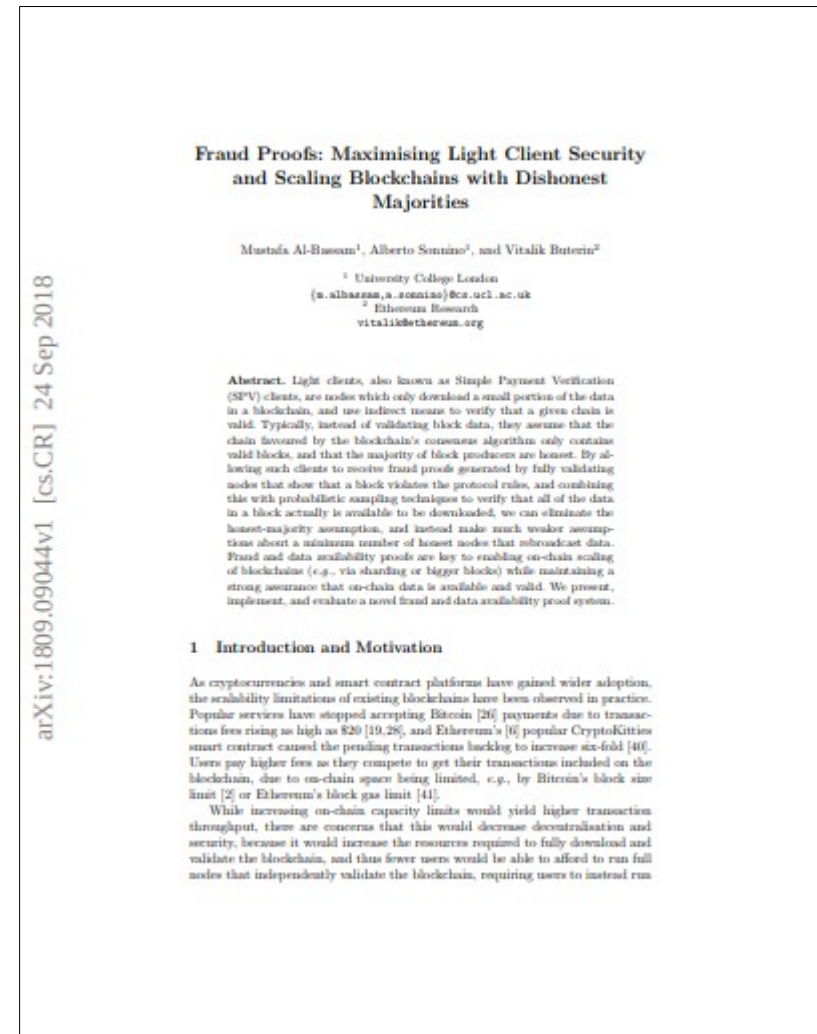
► arxiv.org/abs/1809.09044

► Code:

► github.com/musalbas/rsmt2d


► github.com/musalbas/smt

► github.com/asonnino/fraudproofs-prototype



Fraud proofs: the basic idea

@musalbas



Related discussions on fraud proofs

@musalbas

- ▶ The original Bitcoin whitepaper briefly mentions “**alerts**”; which are messages that full nodes can send to light clients to alert them the block is invalid. (Vulnerable to DoS.)

Related discussions on fraud proofs


@musalbas

- ▶ The original Bitcoin whitepaper briefly mentions “**alerts**”; which are messages that full nodes can send to light clients to alert them the block is invalid. (Vulnerable to DoS.)
- ▶ G. Maxwell and P. Todd have done some work on “compact fraud proofs”. Early proposals require a different fraud proof for each way to violate the rules. We improve on this.

Generalising the blockchain as a state transition system

@musalbas


- ▶ Each transaction reads and modifies the state of the blockchain.
 - ▶ $\text{transition}(\text{state}, \text{tx}) = \text{state}'$ or *error*



Generalising the blockchain as a state transition system

@musalbas

- ▶ Each transaction reads and modifies the state of the blockchain.
 - ▶ $\text{transition}(\text{state}, \text{tx}) = \text{state}'$ or *error*
- ▶ But what if one of the transitions is erroneous?
 - ▶ We need a way to prove this.



Representing the entire state as a Merkle root


@musalbas

- ▶ We can store the state as a key-value store.
 - ▶ Key: UTXO ID. Value: 1 if unspent, 0 otherwise.
- ▶ Ethereum presently uses a **Patricia tree** to do this.

Representing the entire state as a Merkle root

@musalbas


- ▶ We can store the state as a key-value store.
 - ▶ Key: UTXO ID. Value: 1 if unspent, 0 otherwise.
- ▶ We can also use a **Sparse Merkle tree** to do this.
 - ▶ A tree with 2^{256} leaves (every possible SHA-256 hash).



Representing the entire state as a Merkle root

@musalbas

- ▶ Merkle proofs are still $O(\log(n))$ as most branches will only contain leaves with default values (0).
- ▶ To access key K in the tree, access the $hash(K)$ th item.
- ▶ We can thus represent the entire state of the blockchain as a single Merkle root.



Generalising the blockchain as a state transition system

@musalbas

- ▶ Each transaction changes the state root of the blockchain.
 - ▶ $\text{transitionRoot}(\text{stateRoot}, \text{tx}, \text{witnesses}) = \text{stateRoot}'$ or *error*
 - ▶ The witnesses of a transaction are Merkle proofs of all the parts of the state the transactions accesses.
- ▶ **Execution trace:** we need to include the post-state root of transactions in the block. (e.g. every few transactions)



Generalising the blockchain as a state transition system


@musalbas

- ▶ The fraud proofs consists of the transaction and its pre-state root, post-state root, witnesses (+ Merkle proofs).

- ▶ If the Merkle proofs are valid, and

$\text{rootTransition}(\text{stateRoot}_{i+1}, \text{tx}_{i+1}, \text{witnesses}_{i+1}) \neq \text{stateRoot}_{i+2}$

then the fraud proof is valid.




Generalising the blockchain as a state transition system

@musalbas

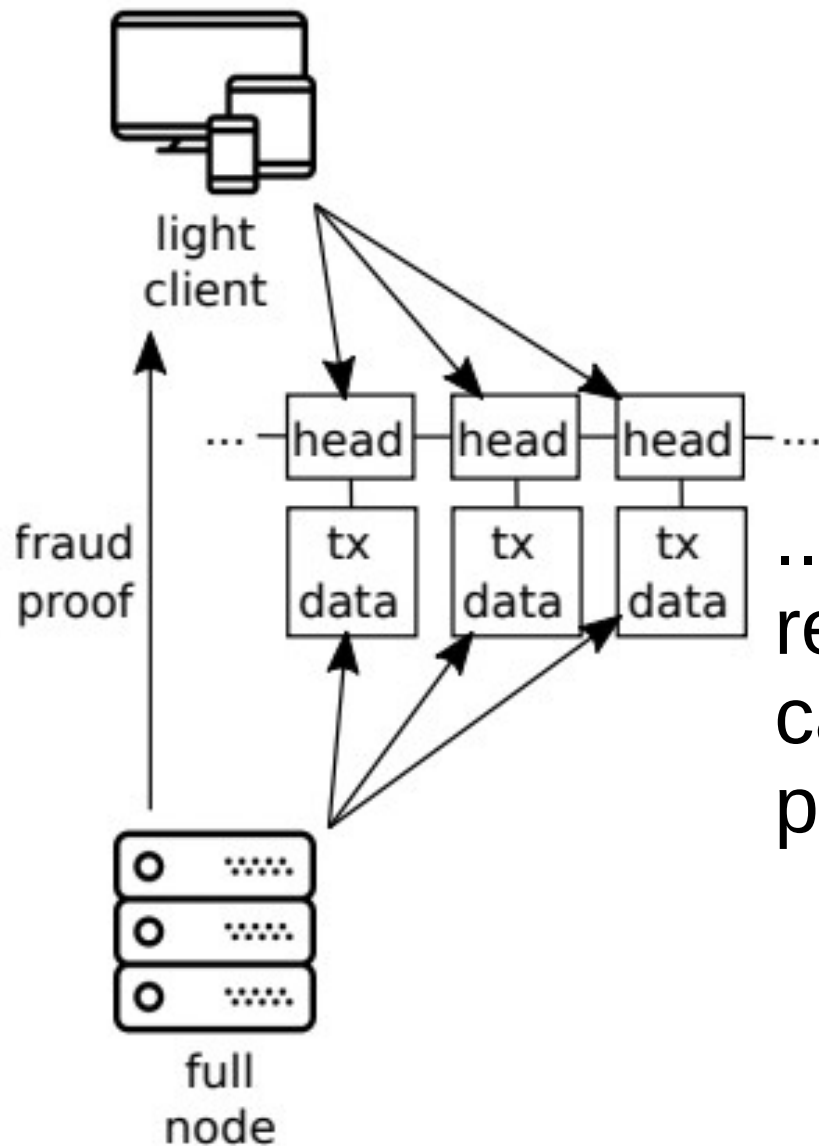
- ▶ You don't have to include the state root after every transaction – this saves block space but the fraud proof gets bigger.

Fraud proof



The data availability problem

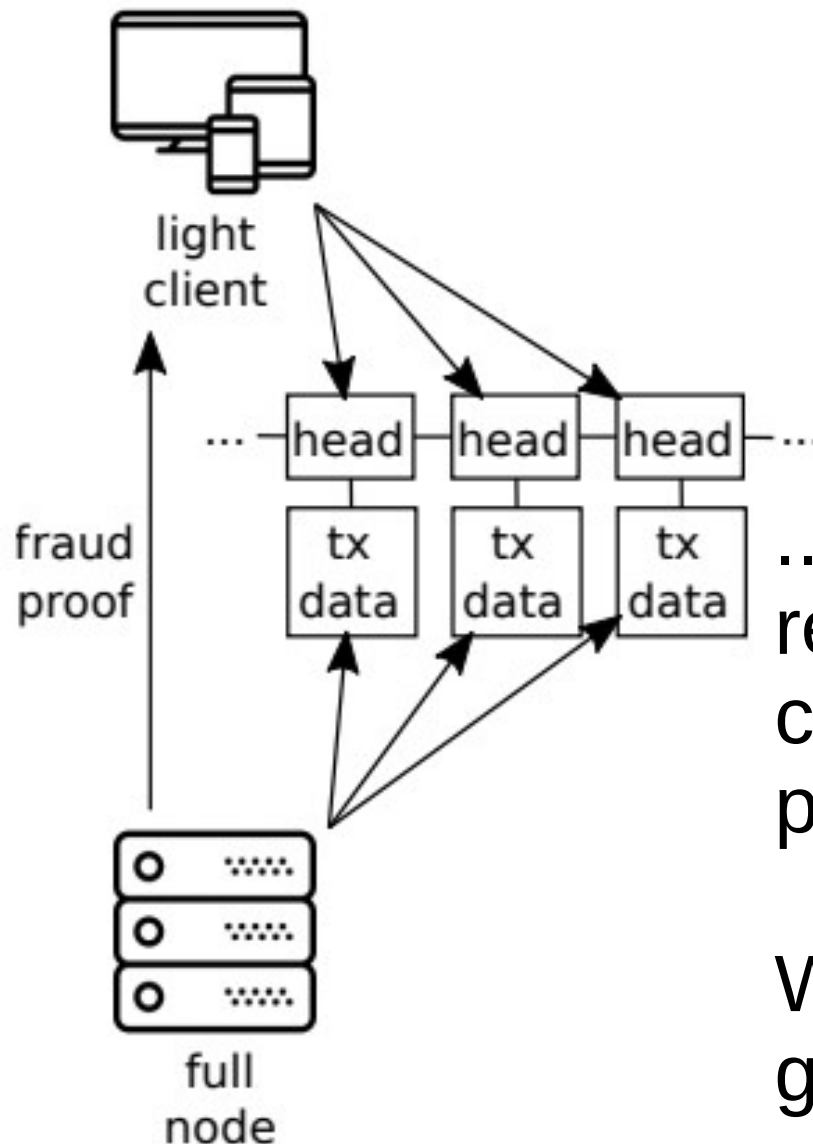
@musalbas



...but if the miner doesn't release the tx data, we can't generate a fraud proof!

The data availability problem

@musalbas




...but if the miner doesn't release the tx data, we can't generate a fraud proof!

We need a way to guarantee **data availability**.

Erasure coding

@musalbas


- ▶ Using erasure coding, you can extend data x pieces long to $2x$ pieces, such that you can recover the whole data from any x pieces.



Naive data availability scheme

@musalbas


- ▶ We can require miners to commit to the Merkle root of the erasure coded version of the block data. In order for a miner to hide any piece of the block, they must hide 50%.



Naive data availability scheme

@musalbas


- ▶ We can require miners to commit to the Merkle root of the erasure coded version of the block data. In order for a miner to hide any piece of the block, they must hide 50%.
- ▶ Thus clients can randomly sample parts of the block*, and if 50% is hidden, then there is a $1-2^{-s}$ chance of landing on an unavailable piece after s samples, in which case the block is rejected.



Naive data availability scheme

@musalbas


- ▶ Clients gossip pieces to full nodes, to enable full nodes to recover the data.
- ▶ There must be enough light clients making samples to reconstruct the whole data (at least x pieces).
 - ▶ We'll do some analysis on this in a moment.



Naive data availability scheme

@musalbas


- ▶ Problem: what if the miner incorrectly generates the erasure code?



Naive data availability scheme

@musalbas

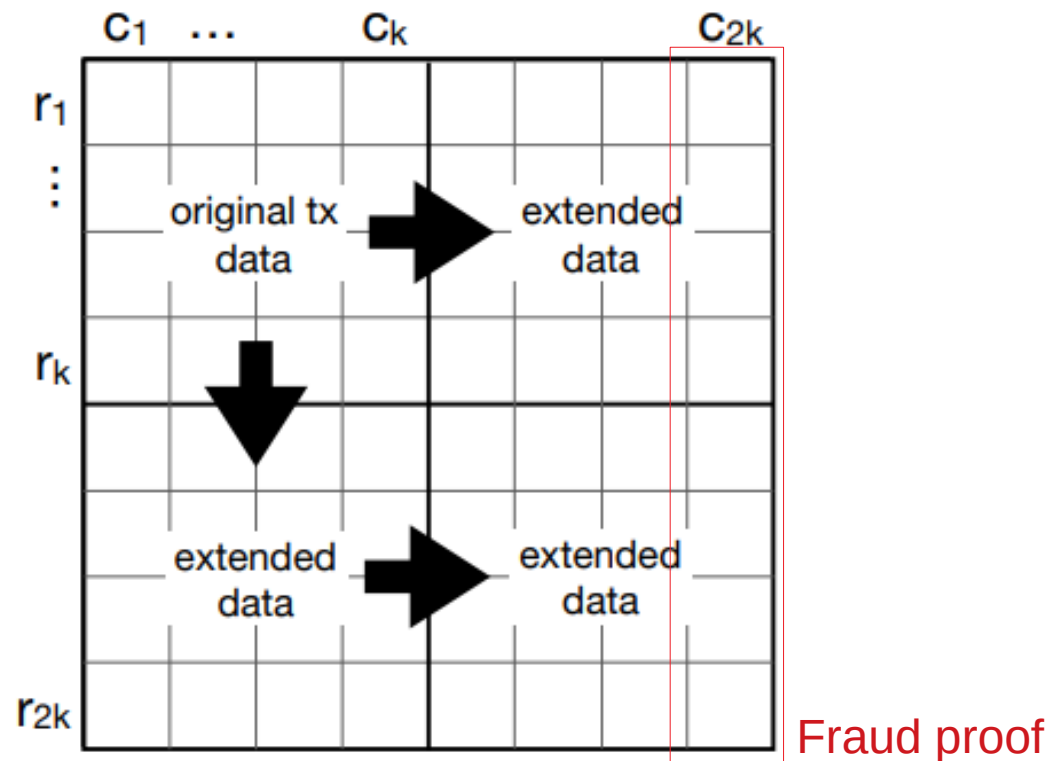
- ▶ Problem: what if the miner incorrectly generates the erasure code?
 - ▶ In order to prove this, the fraud proof consists of the entire block, as clients will need to download and regenerate the erasure code to check if it's correct.
 - ▶ That's $O(\text{blocksize})$. Back to square one!



Multidimensional coding

@musalbas


- ▶ We can use multidimensional coding to fix this.
- ▶ If any row or column is incorrectly generated, a fraud proof that the code is incorrectly generated is limited to that specific row or column. That's $O(\sqrt{\text{blocksize}})$.



Multidimensional coding

@musalbas


- ▶ Miner has to hide roughly 25% of the square to hide any pieces.



Probability analysis

@musalbas


- ▶ What is the probability of a client landing on at least one available piece if the miner has hidden $\sim 25\%$ of the square (if sampling without replacement)?
- ▶ 60% after 3 samplings; 99% after 15 samplings.



Selective releasing of pieces

@musalbas

- ▶ What if a miner only releases pieces as clients ask for them?
 - ▶ Then the miner can always pass the sampling challenge of the first couple of hundred to thousand clients.
 - ▶ The exact number of how many clients can be fooled depends on how many samples they make each (s) and how wide is the square (k).



$p_e(Z \geq \gamma)$	$s = 2$	$s = 5$	$s = 10$	$s = 20$	$s = 50$
$k = 16$	692	277	138	69	28
$k = 32$	2805	1,122	561	280	112
$k = 64$	11,289	4,516	2,258	1,129	451
$k = 128$	>40,000	~18,000	~9,000	~4,500	1,811

Preventing selective releasing of pieces

@musalbas

- ▶ We can prevent this by assuming an enhanced network model.
 - ▶ Clients send requests anonymously.
 - ▶ The order in which requests are received by the network are uniformly random (i.e. client's sampled are interleaved). For example, using a mix net.
- ▶ This would mean that a miner would have same probability of fooling all clients, including the first ones to ask for samples.
 - ▶ Because requests are unlinkable to clients; sampling challenges cannot be satisfied on a per-client basis.

Block validity security assumptions comparison

@musalbas

- ▶ What additional security assumptions are necessary to only accept valid blocks?

Full nodes	SPV clients	SPV clients + fraud proofs
	+ 51% of consensus is honest	+ at least one honest full node in connected network graph + maximum network delay to receive proofs (e.g. 5 mins) + minimum number of light clients (few hundred)

Performance: space/bandwidth

@musalbas

Object	Space complexity	250KB block	1MB block
State fraud proof	$O(p + p \log(d) + w \log(s) + w)$	14,090b	14,410b
Availability fraud proof	$O(d^{0.5} + d^{0.5} \log(d^{0.5}))$	5,120b	12,288b
Single sample response	$O(\text{shareSize} + \log(d))$	320b	368b
Header	$O(1)$	128b	128b
Header + axis roots	$O(d^{0.5})$	2,176b	4,224b

Table 2: Worst case space complexity and illustrative sizes for various objects for 250KB and 1MB blocks. p represents the number of transactions in a period, w represents the number of witnesses for those transactions, d is short for `dataLength`, and s is the number of key-value pairs in the state tree. For the illustrative sizes, we assume that a period consists of 10 transactions, the average transaction size is 225 bytes, and that conservatively there are 2^{30} non-default nodes in the state tree.

Performance: computation

@musalbas

Action	Time complexity	250KB block	1MB block
[G] State fraud proof	$O(b + p \log(d) + w)$	289.78 ms	981.88 ms
[V] State fraud proof	$O(p + p \log(d) + w)$	1.50 ms	1.50 ms
[G] Availability fraud proof	$O(d^2 + d^{0.5} \log(d^{0.5}))$	7.96ms	50.88ms
[V] Availability fraud proof	$O(d + d^{0.5} \log(d^{0.5}))$	0.05ms	0.19ms
[G] Single sample response	$O(\log(d^{0.5}))$	< 0.00001ms	< 0.00001ms
[V] Single sample response	$O(\log(d^{0.5}))$	< 0.00001ms	< 0.00001ms

Table 3: Worst case time complexity and benchmarks for various actions for 250KB and 1MB blocks (mean over 10 repeats), where [G] means generate and [V] means verify. p represents the number of transactions in a period, b represents the number of transactions in the block, w represents the number of witnesses for those transactions, d is short for `dataLength`, and s is the number of key-value pairs in the state tree. For the benchmarks, we assume that a period consists of 10 transactions, the average transaction size is 225 bytes, and each transaction writes to one key in the state tree.

Thank you

@musalbas

▶ Questions?

▶ Twitter: @musalbas

▶ Email: mus@musalbas.com

▶ Paper:

▶ arxiv.org/abs/1809.09044

▶ Code:

▶ github.com/musalbas/rsmt2d

▶ github.com/musalbas/smt

▶ github.com/asonnino/fraudproofs-prototype

