

# Coins: A Billion Bitcoin Users

Robin Linus

January 12, 2020

## Abstract

Coins is a Bitcoin extension designed for payments at scale. We propose an efficient solution to the double-spending problem using a bitcoin-backed proof-of-stake. Validators vote on sidechain blocks with one-time signatures, forming a record that cannot be changed without destroying their collateral. Every user can become a validator by locking bitcoins. One-time signatures guarantee that validators lose their stake for publishing conflicting histories. Checkpoints in Bitcoin’s blockchain prevent long-range attacks. Assuming a rational majority of validators the sidechain provides safety and liveness. The sidechain’s footprint within bitcoin’s blockchain is minimal. The protocol is a generic consensus mechanism allowing for arbitrary sidechain assets. Spawning multiple, independent instances scales horizontally to a network of trustless sidechains.

## 1 Introduction

Bitcoin is a revolutionary alternative to the traditional, government-approved banking system [1]. However when Satoshi Nakamoto introduced it in 2008 the world’s first response was: *“We very, very much need such a system, but the way I understand your proposal, it does not seem to scale to the required size”* [2]. Ever since numerous scalability solutions have been proposed. Still, as of today, the Bitcoin network processes less than seven transactions per second because growing the chain faster reduces decentralization significantly. In contrast, centralized payment services such as PayPal or Visa serve billions of users at up to 50,000 TX/s.

Currently, off-chain payments via the Lightning Network are the most promising approach to scale Bitcoin [3]. They allow a much higher throughput, yet they hardly scale to billions of users. They still require too many on-chain transactions to open and close payment channels. Adoption is even further constrained by the inbound-capacity of payment channels and the need to lock funds for every new user. These constraints lead to many

layers of complexity and a tendency towards centralized and custodial solutions which contrast Bitcoin’s purpose of being permissionless, trustless and censorship-resistant.

Sidechains have been proposed as an alternative solution for scalability [4]. They introduce parallel blockchains enabling payments within a simplistic system similar to Bitcoin. Yet, their designs depend on federations or miners validating sidechain blocks which limits security, scalability and flexibility. We introduce a new sidechain consensus mechanism with a permissionless, bitcoin-backed proof-of-stake. Furthermore, we leverage the Lightning Network for efficient cross-chain communication. This results in a flexible and scalable network of sidechains.

## 2 Bitcoin Stake

It is impossible to produce distributed consensus except by consuming an external resource. This is because if block production has no ongoing costs, neither does attacking the chain [5] [6]. We anchor our proof-of-stake mechanism into Bitcoin’s proof-of-work consensus which is computationally, and therefore thermodynamically, very expensive. We use bitcoins as an external resource to produce sidechain consensus.

### 2.1 One-time Signatures

A characteristic of Bitcoin’s digital signature algorithm is that it needs to produce, for each signature generation, a fresh random value (hereafter designated as *nonce*). Reusing the nonce value on two signatures of different messages allows attackers to recover the private key algebraically.

The nonce reuse vulnerability can be used to discourage validators from participating in double-spending attacks [7] [8]. Each validator pre-commits to his sequence of nonces, such that the system can constrain a validator’s signature for the  $n$ -th block to be valid only if he signed it using his  $n$ -th nonce. This guarantees validators can not create valid signatures for conflicting blocks without leaking their private key and losing their collateral.




Figure 1: A validator votes for blocks with one-time signatures.

## 2.2 Collateral Contracts

For validators' one-time signatures to be scarce, each of them has to lock a collateral such that the bitcoin network can penalize malicious actors for signing conflicting histories. In the following we discuss *collateral contracts* which serve as an adaptor to produce a sidechain consensus from Bitcoin's consensus. Our contract expresses: *If Alice leaks her key she loses her bitcoins*. So for Alice to become a validator she locks bitcoins in an output such that:

- Option A: One year later she gets her money back.
- Option B: She can destroy her money right now.

## 2.3 Collateral Contracts in Bitcoin Script

Currently, burning funds is not supported in Bitcoin script. Yet, upcoming Bitcoin features such as `SIGHASH_NOINPUT` or `OP_CHECKTEMPLATEVERIFY` allow trustless covenants. For example, `SIGHASH_NOINPUT` allows to pre-commit to a particular signature and thus, a particular spending transaction. In combination with a 2-of-2 multi-signature this gives a simple, trustless solution for our collateral contract.

For now, we need a workaround to implement the collateral contract. In the following we discuss solutions. All of these trust-minimizing workarounds are highly undesirable additional complexity. Fortunately, in the long term, we will certainly have some clean solution for covenants.

### 2.3.1 Trust-minimized Workaround 1

A simple solution is to introduce a trusted party, say, Bob:

- Alice creates a funding transaction with the following output:
  - Alice can spend her money in one year.
  - Alice and Bob can always spend her money collaboratively.
- Bob pre-signs and publishes a punishment transaction:
  - burn her collateral now to address `0x000...000` ( if Alice agrees ).
- Bob immediately deletes his secret key ( in case his machine gets compromised ).

There is no counter-party risk here for Alice, because she executes her funding transaction only after she received Bob's signature for the punishment transaction and completed her setup.

For the system to minimize trust in Bob, multiple parties can participate and if only one is honest and deletes their key, then this scheme is secure. Bitcoin script currently supports more than 60 participants per multi-signature transaction [9].

Containing more than 60 signatures, the punishment transaction becomes large and expensive, but as long as the validator is honest, the transaction doesn't need to be included in Bitcoin's blockchain. To incentivize Bitcoin miners to execute the punishment transaction quickly in case of misbehavior, it pays a high miners fee.

It is possible to further minimize trust in single parties by using ECDSA signature aggregation to allow for thousands of Bobs participating in a single combined signature [10]. When Schnorr signatures become available in Bitcoin, such aggregated signatures will become more simple because of the linearity of Schnorr's scheme [11].

The Bobs could be a trusted federation or a percentage of the current validators. With aggregated signatures they could also be a significant percentage of current coin holders. The trusted parties sign blindly to reduce the risk of censorship. A simple scheme exploits that Bitcoin transactions use double SHA256. Thus, Alice can ask Bob to sign the single-round SHA256 hash of her transaction. The second round is Bob's hashing function for his signature algorithm. This way Bob doesn't learn what he signs until Alice publishes her contract.

### 2.3.2 Trust-minimized Workaround 2

Alternatively, we can implement the collateral contract with a simple

```
<one year> OP_CHECKSEQUENCEVERIFY OP_DROP <A> OP_CHECKSIG
```

OP\_CHECKSEQUENCEVERIFY ensures, as a side effect, that the spending transaction opts in to replace-by-fee. Thus, if the pubkey <A> is used in a single-sign signature scheme (which reveals the privkey if double-signed), then at the end of the period, anyone who saw the double-signing can claim that fund and thus act as "Bob". Indeed, many "Bob"s will act and claim this fund, increasing the fee each time to try to get their version onchain. Eventually, some "Bob" will just put the entire fund as fee and put a measly OP\_RETURN as single output. This "burns" the funds by donating it to miners.

From the point of view of Alice this is hardly distinguishable from losing the fund right now, since Alice will have a vanishingly low chance of spending it after the collateral period ends, and Alice still cannot touch the funds now anyway. Alice also cannot plausibly bribe a miner, since the miner could always get more funds by replacing the transaction internally with a spend-everything-on-fees OP\_RETURN output transaction, and can only persuade the miner not to engage in this behavior by offering more than the collateral is worth (which is always worse than just losing the collateral).

A OP\_CHECKTEMPLATEVERIFY would work better for this use-case, but even without it you do not need a trusted party to implement the collateral contract.

Drawback of this solution is that cheating validators are not immediately removed from Bitcoin’s UTXO set. Yet, this might be a decent tradeoff because we have a succinct proof to exclude a cheating validator: knowledge of his private key.

Another drawback is that it requires trust in miners. If Alice cooperates with a significant share of Bitcoin’s hash power, she has proportional chances of mining the transaction herself. Then she would not have any cost of attacking the chain.

### 3 Consensus Mechanism

The sidechain’s consensus is anchored into Bitcoin’s proof-of-work consensus. Validators are defined by collateral contracts included in Bitcoin’s UTXO set. Assuming all sidechain nodes are running a Bitcoin full node, they implicitly are in consensus about the exact validator set at every point in time without exchanging any messages. All randomness is determined by Bitcoin’s proof-of-work which is studied well [12].

The validator set is determined only by Bitcoin’s blockchain which prevents many classical problems of pure proof-of-stake protocols such as the nothing-at-stake problem, long-range attacks, stake grinding and costless simulation [5] [13]. Still, a malicious majority can halt the system.

#### 3.1 Voting for the next Block

The set of validators is known and so is the number of possible votes per sidechain block. An election is economically final as soon as there is a majority such that validators would have to burn deposits to change it. Minority votes “wasted” for conflicting blocks help finalize the result, too.

The leader is determined from randomness derived from bitcoin’s most recent block. The leader proposes a signed block and all other validators confirm the block by signing it too. Once 51% of validators signed a block it cannot be reverted without burning stake. If a leader does not propose a block within a given time then all other validators start broadcasting skip messages for that block. Once 51% of validators signed a skip message the block is skipped and the next round begins.

#### 3.2 Safety and Liveness

By definition, a decentralized cryptocurrency must be susceptible to 51% attacks whether by hashrate, stake, or other permissionlessly-acquirable resources. Our system always provides safety, but an attacker can halt the network as long as he controls a majority of

all bitcoins at stake. Halting the system costs at least the time value of locking sufficient amounts of bitcoins.

### 3.3 Ongoing Cost to Attack the Chain

One might argue that time value of locked bitcoins is too cheap to protect the chain. That is a fundamental misunderstanding of the underlying market mechanism.

In Bitcoin, the proof of work security is not determined by the price of electricity, but by the Dollar value of the block rewards. Miners have an incentive to spend about  $1\text{reward} = 12.5 \text{ BTC} \approx \$100000$  to produce a block. The cost of block production equals the cost of attacking the chain. So more Bitcoin users lead to a higher bitcoin price and thus, a higher value of the block rewards and subsequently, to more security. The exact same market forces align the incentives for our proof of stake. The total time value of all locked bitcoins converges against 1 sidechain revenue per block. Whenever there is less time value locked, someone will stake their bitcoins to earn the cheap sidechain rewards. The ongoing cost for an attack is about one sidechain revenue per block time. This is independent of the consumed resource. Security depends only on the exchange rate of the sidechain asset. Sustainable price is mostly driven by usage and resulting network effects. Thus, a sidechain provides security proportional to the number of its users. The seemingly low investment of time value of bitcoins is priced in.

The above argumentation becomes even more obvious if we introduce an additional cost and let validators burn bitcoins for every on-chain vote. Proof-of-burn has recently been formally analysed [14]. The economic implications of burning significant amounts of Bitcoin are questionable. Additionally, a level of security comparable to Bitcoin requires the system's BTC burn rate to be equal to Bitcoin's inflation rate. Still, people would certainly burn 12.5 BTC per sidechain block as long as the sidechain block rewards make it profitable. Fortunately, burning bitcoins is not necessary though. It just makes the cost of block production more obvious.

For decentralized consensus it is sufficient to consume any limited resource for block production. Just like with any other resource, the time value of bitcoins can provide security proportionally to the value of the sidechain block rewards.

## 4 Scaling and Sharding

Coins can scale horizontally almost indefinitely by spawning multiple fully independent instances in parallel. All instances are naturally synchronized by Bitcoin's blockchain. Each instance manages a separate currency, but in a standardized format to facilitate

cross-chain transactions via atomic swaps. Cross-chain communication requires no cross-chain validation. A swap is validated only between the two swapping clients.

This scales well. Assuming each instance grows approximately like Bitcoin’s blockchain. Assuming further, optimizations such as Schnorr Signatures make the sidechain about three times more compact than Bitcoin, such that it can process about 20 transactions per second at the same network load. Then to achieve transaction rates as high as centralized services with billions of users we need about 2,500 independent instances with about a million users each. For 200 validators per instance we need about 500,000 on-chain transactions per month, which the current Bitcoin network could process in less than a day.

## 4.1 Overhead of Sharding

Transactions within an instance are simple, but cross-shard communication is more complex and expensive. Shards might be off-sync because our clock is Bitcoin’s consensus which has a high latency. This is a difficulty for cross-chain atomic swaps. Also every shard manages a separate currency which implies varying exchange rates, but also creates a market for balancing shards and flexibility to tolerate local failures. The overhead of naive cross-chain communication scales quadratically because each shard needs a channel into each other shard. HTLCs allow for multi-hop communication such that shards can share their channels. It is efficient to centralize the communication around Bitcoin as a settlement layer. This allows payments between arbitrary shards within 2 hops. Theoretically every shard requires only one payment channel – the one with bitcoin. With the same argument each currency requires only one exchange rate. This enhances usability because wallets can abstract away all underlying currencies and display to users only a single currency: BTC.

## 4.2 Limitation: The Altcoin Problem

The system requires one independent asset per sidechain. A sidechain’s security is limited by its asset’s value. Only if a sidechain’s revenue is sufficiently valuable, it can motivate many Bitcoin holders to protect it. Therefore, only large sidechains with lots of users and valuable assets can provide security. Small sidechains are insecure.

Furthermore, these sidechain assets are only a means of transfer. The global unit of account is BTC. Yet, it is not easy to abstract away all hops seamlessly. For usability it’s important to pay invoices denominated in the central settlement currency BTC. The exchange rates can be hidden in transaction fees. Nevertheless, sidechains’ assets must converge against correlating very strongly with Bitcoin’s price. Otherwise, people leave the sidechain because they lose trust in its asset. This can result in the collapse of a

sidechain’s security incentives.

Additionally, atomic swaps with varying exchange rates introduce the free option problem [15] which might become practical to exploit in situations when there is a high price volatility. All of these assumptions are practical because they also underlie the existing PoW altcoin economy.

A sidechain’s security depends on its usefulness.

It is possible to peg sidechain assets to Bitcoin. An eternal one-way peg [16] can limit the price fluctuation of a sidechain’s asset. It is possible to demand from validators to burn a certain amount of BTC to become a validator. Those burned BTC can be issued in the sidechain as block subsidy.

## 5 Validator Signatures

In this chapter we discuss details of the validators signatures. We explain how to pre-commit to a particular nonce per block efficiently. Finally, we explain a scheme for better hot key security.

### 5.1 Compact Nonce Commitments

Constructed naively, validators would have to pre-commit to millions of nonces within their funding transaction. A more efficient construction is to let validators commit to their next nonce by signing it with the previous signature. This way only one nonce has to get stored at each point in time.

### 5.2 Hot Key Security

Hardware wallets are stateless and therefore they’re incompatible with nonce commitments. To be compatible they would need to keep track of every nonce they have signed. Otherwise an attacker can extract the private key by requesting two different signatures with the same nonce.

Therefore, to protect the validators’ nodes, they can use a *validation key* to sign sidechain blocks and a separate *redeem key* to spend the bitcoin deposit once it is unlocked. Until then, the redeem key remains in a cold wallet. Nodes having access only to the validation key are a much less attractive target for attackers.



## 6 Conclusion

We have proposed a trustless Bitcoin extension which scales to a global payment system. We started with the usual framework of a sidechain, which provides strong control of ownership, but is incomplete without a trust-minimized way to prevent double-spending. To solve this, we proposed a bitcoin-backed proof-of-stake consensus mechanism that quickly becomes economically impractical for an attacker to change if honest nodes control a majority of stakes. Validators and leaders are elected via Bitcoin’s consensus with little coordination. Nodes can leave and rejoin the network at will, accepting checkpoints in the bitcoin blockchain as proof of what happened while they were gone. Validators can not create conflicting histories without losing their Bitcoin collateral. They vote with their signatures, expressing their acceptance of valid blocks by signing them and rejecting invalid blocks by refusing to sign them. Any needed rules and incentives can be enforced with this consensus mechanism. Spawning multiple instance scales horizontally. Cross-chain communication is facilitated via atomic swaps.

## References

- [1] Satoshi Nakamoto et al. Bitcoin: A peer-to-peer electronic cash system. 2008.
- [2] James A. Donald. The cryptography mailing list - bitcoin p2p e-cash paper. Bitcoin Stack Exchange.
- [3] Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments, 2016.
- [4] Adam Back, Matt Corallo, Luke Dashjr, Mark Friedenbach, Gregory Maxwell, Andrew Miller, Andrew Poelstra, Jorge Timón, and Pieter Wuille. Enabling blockchain innovations with pegged sidechains. *URL: [http://www. opensciencereview. com/papers/123/enablingblockchain-innovations-with-pegged-sidechains](http://www.opensciencereview.com/papers/123/enablingblockchain-innovations-with-pegged-sidechains), page 72*, 2014.
- [5] Andrew Poelstra et al. Distributed consensus from proof of stake is impossible. 2014.
- [6] Cristina Pérez-Solà, Sergi Delgado-Segura, Guillermo Navarro-Arribas, and Jordi Herrera-Joancomartí. Double-spending prevention for bitcoin zero-confirmation transactions. *International Journal of Information Security*, 18(4):451–463, 2019.
- [7] Tim Ruffing, Aniket Kate, and Dominique Schröder. Liar, liar, coins on fire!: Penalizing equivocation by loss of bitcoins. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 219–230. ACM, 2015.

- [8] Malte Möser, Ittay Eyal, and Emin Gün Sirer. Bitcoin covenants. In *International Conference on Financial Cryptography and Data Security*, pages 126–141. Springer, 2016.
- [9] Pieter Wuille. Will segwit allow for m of n multisig with very large n and m? Bitcoin Stack Exchange.
- [10] Rosario Gennaro and Steven Goldfeder. Fast multiparty threshold ecdsa with fast trustless setup. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1179–1194. ACM, 2018.
- [11] Gregory Maxwell, Andrew Poelstra, Yannick Seurin, and Pieter Wuille. Simple schnorr multi-signatures with applications to bitcoin. *Designs, Codes and Cryptography*, 87(9):2139–2164, 2019.
- [12] Joseph Bonneau, Jeremy Clark, and Steven Goldfeder. On bitcoin as a public randomness source. *IACR Cryptology ePrint Archive*, 2015:1015, 2015.
- [13] Jonah Brown-Cohen, Arvind Narayanan, Alexandros Psomas, and S Matthew Weinberg. Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 459–473. ACM, 2019.
- [14] Kostis Karantias, Aggelos Kiayias, and Dionysis Zindros. Proof-of-burn.
- [15] ZmnSCPxj. [lightning-dev] an argument for single-asset lightning network. Lightning Developers Mailing List, 2018.
- [16] Ruben Somsen. 21 million bitcoins to rule all sidechains: The perpetual one-way peg. *Medium*, 2020.