Bitcoin Staking: Unlocking 21M Bitcoins to Secure the Proof-of-Stake Economy

The Babylon Team*

Version 1.0: 2023-07-13

Abstract

Proof-of-Stake (PoS) chains are secured by capital but capital can be very expensive. Bitcoin is a Proof-of-Work chain but it is also a $600 Billion asset and most of it is idle capital. We propose the concept of *Bitcoin staking* which allows bitcoin holders to stake their idle bitcoins to increase the security of PoS chains and in the process earn yield. We present a Bitcoin staking protocol which allows bitcoin holders to trustlessly stake their bitcoins *without* bridging them to the PoS chain but yet provides the chain with *full* slashable security guarantees. The protocol supports fast stake unbonding to maximize the liquidity for bitcoin holders. Moreover, the protocol is designed as a modular plug-in for use on top of many different PoS consensus algorithms and provides a primitive upon which restaking protocols can be built. A system architecture is proposed for scaling the protocol to many stakers and many PoS chains with a Bitcoin-staked Babylon chain acting as a control plane to synchronize between Bitcoin and the PoS chains. Bitcoin staking enables an important new use case for Bitcoin and takes a significant step towards integrating Bitcoin and the Proof-of-Stake economy.

1 Proof-of-stake security needs capital

The past few years have witnessed a shift in the blockchain industry from Proof-of-Work (PoW) to Proof-of-Stake (PoS) as the sybil resistance mechanism. A defining event in this trend is the Merge in September 2022, the migration of Ethereum from PoW to PoS consensus. While PoW blockchains are secured by miners solvinfg difficult mathematical puzzles, PoS blockchains are secured by validators holding stake. The stake of a validator serves as a *deposit* which can be slashed when the validator violates the protocol. Slashability is a feature that PoW chains do not have, and is a main motivation for the migration of Ethereum from PoW to PoS [18]. The larger the market capitalization of the securing stake, the more costly it is to attack the chain, and the stronger the economic security of the chain. Thus, while PoW chains are secured by work, PoS chains are secured by *capital*.

Such capital is oftentimes difficult to attract, especially for small chains or chains in their bootstrapping phase. High inflation rates are required to pay high yields to attract such capital. For example, in the Cosmos ecosystem, consisting of over 60 application specific chains, initial annualized inflation rates of 20% to 100% are quite common. Such high inflation hampers the long term

*Contact: contact@babylonchain.io. This litepaper is based on a forthcoming technical paper [9], a collaboration between the Babylon team, Robin Linus at ZeroSync and Orfeas Stefanos Thyfronitis Litos at Common Prefix and Imperial College.

growth of the chain. The high cost also puts security in tension of the *utility* of the chain: the inflation could have been used to incentivize the applications on the chain

Akash [1], a Cosmos SDK chain running a decentralized AI computing platform, provides a good case study. An extremely high initial inflation rate of 100% in the AKT token pays for both security and for incentivizing providers to rent out high-quality compute hardware [30]. As the inflation rate decreases over time, the tension between security and utility becomes even more acute.

2 Bitcoin: a $600 billion asset

Despite the shift to PoS, the largest crypto asset, Bitcoin, with more than half the total crypto asset as of this writing, is still secured by a PoW chain. Compared to PoS assets, the Bitcoin asset has several significant differences:

1. **Unencumbered**. Since Bitcoin is secured by work, the Bitcoin asset itself is not used to secure the Bitcoin chain. In contrast, each PoS asset is used to secure its own chain.

2. **More idle**. Most of the Bitcoin asset sits idle and is not deployed. Since most of yield-generating activities such as DeFi lending and security staking happen on PoS chains, bitcoins need to be *bridged* to other chains and/or sent to a third party centralized custodian in order to extract yield. Such bridges and centralized custodians are considered too risky for many bitcoin holders, For example, wBTC, one of the largest wrapped Bitcoin assets, only has a market capitalization of less than $5B, a fraction of a percent of the total Bitcoin market capitalization.

3. **More decentralized**. Bitcoin, being the oldest blockchain, has arguably the most decentralized set of token holders: miners, early adopters and developers, project founders, individual investors, institutional investors, exchanges, etc. In contrast, the assets of many PoS chains are concentrated in the hands of early investors, founders and team members and foundations, at least in the early stage of the projects. A concentrated asset exposes the network to centralization when the asset is staked to validate the network.

4. **Less volatile**. Bitcoin, being the largest crypto asset, has significantly lower volatility than most PoS assets. Volatilty of a PoS asset is a significant issue to the security of a PoS chain because security is a direct function of the market capitalization of the staked asset and a sharp drop in value of the asset provides an opportunity for an attacker to strike.

3 Bitcoin staking

Given these properties, why not put bitcoins to work by staking them to help secure PoS chains? This concept of *Bitcoin staking* is the focus of this work.

Bitcoin staking is a two-sided market place (Figure 1). On the one side are PoS chains which need security and are willing to pay yields for it. On the other side are bitcoin holders who have the capital and want to earn yield on it. A *Bitcoin staking protocol* is a *security-sharing protocol* that realizes this two-sided market place. A good Bitcoin staking protocol should provide strong security guarantees to both the consumer PoS chains (so that they are willing to pay yield for it) and to the provider bitcoin holders (so that they are willing to stake).

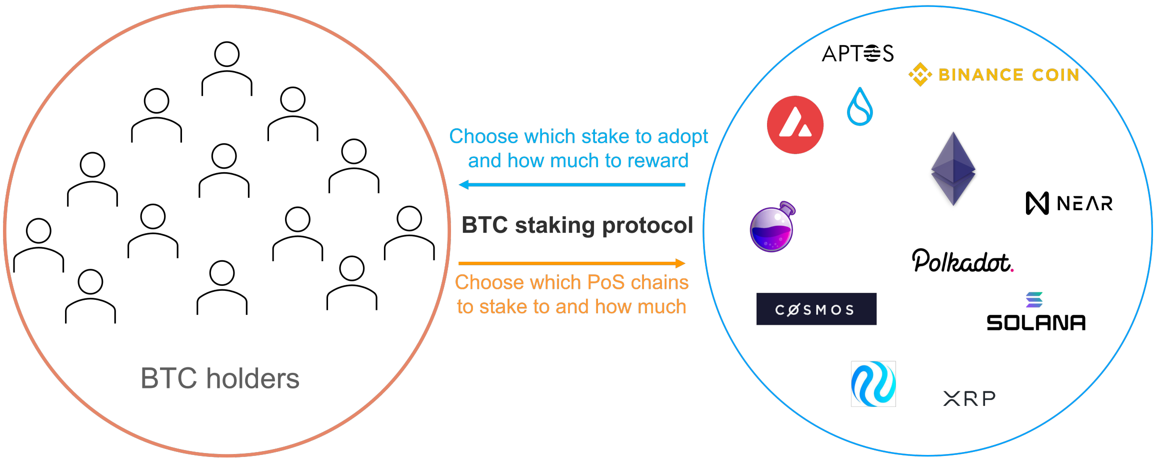
BTC holders
Choose which stake to adopt and how much to reward
BTC staking protocol
Choose which PoS chains to stake to and how much
APTOS
BINANCE COIN
NEAR
Polkadot.
COSMOS
SOLANA
XRP

Figure 1: Bitcoin staking is a two-sided market place.

4 Babylon Bitcoin staking protocol: security properties

We present here a Bitcoin staking protocol which, when used in conjunction with an off-the-shelf consumer PoS chain, has three important security properties:

1. **Fully slashabable PoS security**. Whenever there is a safety violation, 1/3 of the Bitcoin stake is guaranteed to be slashed. As long as 2/3 of the Bitcoin stake follows the PoS protocol honestly, the PoS chain is live.

2. **Staker security**. Each Bitcoin staker is guaranteed to be able to withdraw its funds, or unbond, as long as the staker follows the PoS protocol honestly.

3. **Staker liquidity**. Unbonding of the staked bitcoins is guaranteed to be secure and fast without the need of social consensus.

Property 1 says that protocol violators are slashed. Property 2 says that the only ones that lose their stakes are protocol violators. Taken together, Property 1 and 2 reflect the gold standard of PoS security: full slashability, as advocated by Buterin and Griffith [18]. Indeed, full slashability is the central design goal of PoS Ethereum [19] as well as that of Tendermint [13, 15], one of the most widely used consensus engines for building PoS blockchains including Cosmos SDK chains, Polygon, BNB Chain, etc. In fact, Property 2 is even stronger than in stand-alone PoS protocols: the staker can still unbond even if *all* other stakers on the PoS chain are dishonest. Withdrawal censorship is not possible in our Bitcoin staking protocol. Thus, our protocol provides *trustless staking*.

Stand-alone PoS chains like PoS Ethereum or Cosmos SDK chains suffer from a long unbonding time, of the order of several weeks, due to the use of social consensus to counter long range attacks [16, 21, 12, 22], a fundamental "nothing-at-stake" attack vector in PoS chains. In contrast, in our Bitcoin staking protocol, the bitcoin stake distribution is maintained on the Bitcoin chain, which is immune to such long range attacks. By a proper design of the staking protocol, we show that Property 3 is achievable.

5 Challenges

We consider two basic approaches to Bitcoin staking, each of which has its challenge.

1. **Bridging to PoS chain.** One approach to Bitcoin staking is to first bridge bitcoins from the Bitcoin chain to the consumer PoS chain, and enforce slashing rules there. While this approach can give slashable security to the PoS chain (Property 1), a fundamental limitation is the security of the bridging solution itself. The security of most existing Bitcoin bridges is based on trusting either a centralized custodian (eg. Bitgo for wBTC) or a multisig bridge committee. (See Section 9.8 for more discussions.) Even an ideal Bitcoin bridge relies on trusting the stakers of the destination chain. Hence, with a bridging solution, it is not possible to achieve Property 2, i.e. trustless staking,

2. **Remote staking from Bitcoin chain.** To avoid bridging bitcoins, an alternative approach is *remote staking*: lock the staked bitcoins in a contract on the Bitcoin chain and then slash the stake when there is a protocol violation on the consumer PoS chain. This is the approach used in security sharing solutions such as Eigenlayer's Ethereum restaking protocol [36] and Cosmos ecosystem's mesh security [11, 4]. In both settings, the provider chains, i.e. the sources of security, have Turing-complete smart contract layers. This makes the implementation of slashing on the provider chain technically straightforward, because protocol violation evidences can be sent back from the consumer chain to the provider chain and slashing done by smart contracts on the provider chain. However, the provider chain in our setting is Bitcoin, which does not support smart contracts but only has a scripting language with limited expressiveness. So while we obtain trustless staking (Property 2) by virtue of the fact that the bitcoins remain on the Bitcoin chain, a key challenge now is how to perform slashing to achieve Property 1, i.e. fully slashable PoS security.

Our Bitcoin staking protocol follows the remote staking approach but we overcame the lack of smart contracts by a combination of advanced cryptography, consensus protocol innovations and optimized use of the Bitcoin scripting language. Before we go into the details of these technologies, let us first illustrate the high-level functionalities of the Bitcoin staking protocol by going through a staker's journey.

6 A Bitcoin staker's journey

Alice has 1 bitcoin and she wants to stake it on a PoS chain. First she enters into a staking contract by sending a staking transaction to the Bitcoin chain to lock her bitcoin into a self-custodian vault. The bitcoin can be unlocked, only by using Alice's private key, in one of two possible ways:

1. Alice issues an *unbonding transaction* and the bitcoin will be unlocked and returned to Alice in 3 days.

2. Alice issues a *slashing transaction* which sends the bitcoin to a burn address.

Once this staking transaction appears in the Bitcoin chain, Alice can start validating for the PoS chain by signing blocks using her key. During her validation duty, there are two possible paths.

One, the *happy* path (Figure 2(a)), is that Alice follows the protocol honestly, and when she wants to unstake, she issues an unbonding request by sending a unbonding transaction to the Bitcoin chain (Figure 2(b)). Once the unbonding transaction enters the Bitcoin chain, Alice's validation duty on

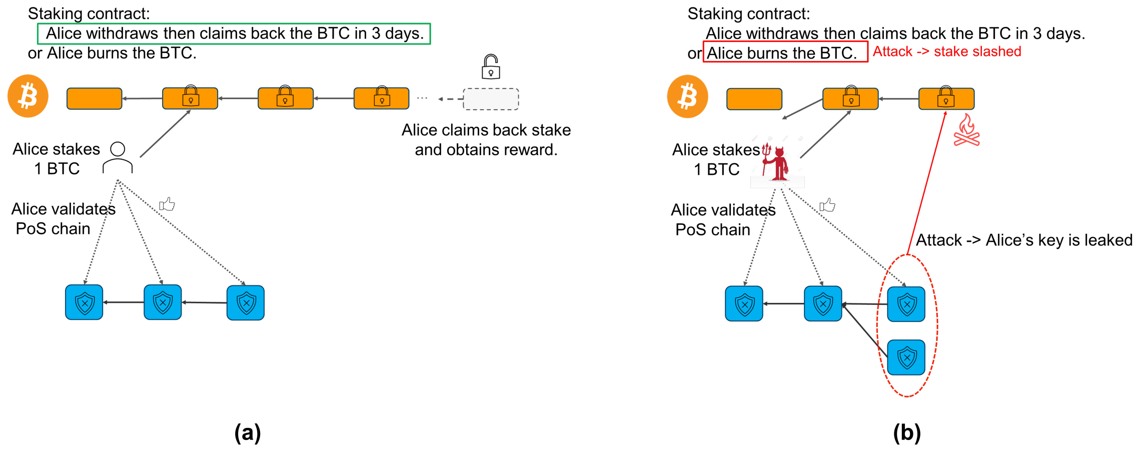

Figure 2: A Bitcoin staker's journey: (a) the happy path: Alice stakes, validates the PoS chain, requests unbonding, and unstakes in 3 days; (b) the unhappy path: Alice stakes, commits a safety offense to the PoS chain, and then her bitcoin got burned.

the PoS chain stops, and after 3 days, the withdrawal request is granted and the 1 bitcoin returns to Alice. The PoS chain will also grant a reward to Alice.

Two, the *unhappy* path (Figure 2(b)), is that Alice turns malicious and participates in a double-spend (safety) attack on the PoS chain (Figure 2(c)). In that case, the staking protocol ensures that Alice's private key will be leaked to the public. Anyone can now pretend to be Alice and send the slashing transaction to the Bitcoin chain and burns Alice's 1 bitcoin. The existence of this unhappy path makes sure safety violations are slashable and everyone will keep things happy.

7 Technology ingredients

The following are the key ingredients which enable the protocol functionalities described in Section 6 and overcome the challenges described in Section 5. Details will be found in the forthcoming full paper [9].

7.1 Staking contracts via Bitcoin convenant emulation

Since there is no smart contract layer on Bitcoin, staking contracts have to be expressed in terms of UTXO transactions written in the Bitcoin script [10]. Each UTXO transaction spends funds from the UTXO set and the Bitcoin Script provides a small number of opcodes to specify conditions for spending the funds. A stacking contract has 4 transactions:

- a staking transaction, where the input is the address of the staker's bitcoin and the output can be spent in one of two ways:
  - an unbonding transaction, which allows the output to be spent by the staker after a relative lock time (measuring the unbonding time) has expired (the relative lock time can be implemented by the opcode `OP_CHECKSEQUENCEVERIFY` [10]);

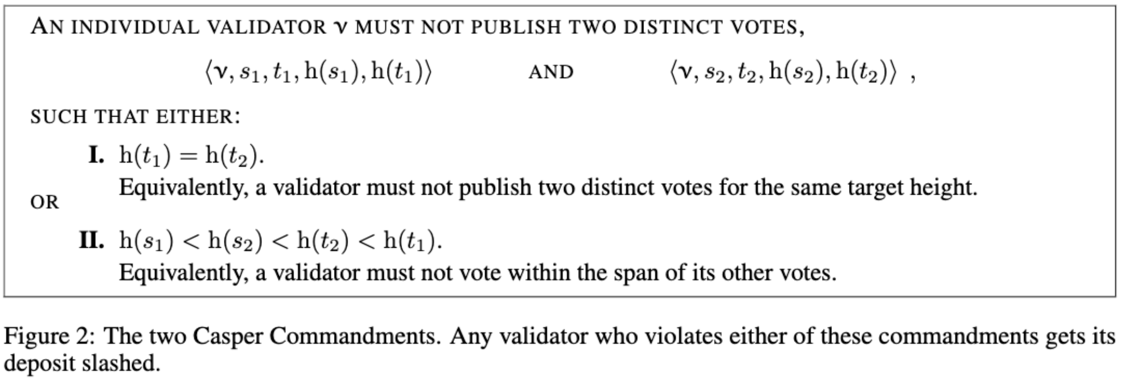

AN INDIVIDUAL VALIDATOR ν MUST NOT PUBLISH TWO DISTINCT VOTES,

$$\langle \nu, s_1, t_1, \mathrm{h}(s_1), \mathrm{h}(t_1) \rangle \qquad \text{AND} \qquad \langle \nu, s_2, t_2, \mathrm{h}(s_2), \mathrm{h}(t_2) \rangle \ ,$$

SUCH THAT EITHER:

I. $\mathrm{h}(t_1) = \mathrm{h}(t_2)$.
Equivalently, a validator must not publish two distinct votes for the same target height.

OR

II. $\mathrm{h}(s_1) < \mathrm{h}(s_2) < \mathrm{h}(t_2) < \mathrm{h}(t_1)$.
Equivalently, a validator must not vote within the span of its other votes.

Figure 2: The two Casper Commandments. Any validator who violates either of these commandments gets its deposit slashed.

Figure 3: Slashing conditions of Casper. Figure taken from [18].

  - a slashing transaction, which allows the staker to spend the output immediately to a burn address (unspendable output)

- an unstaking transaction, which can spend the output of the unbonding transaction after the relative timelock has expired,

The staking contract is an example of a *Bitcoin convenant* [25, 26], where the output of a transaction is constrained to be spent in certain ways. Convents can be implemented by `OP_CHECKTEMPLATEVERIFY` [8], a Bitcoin script opcode proposed for inclusion in a future upgraade of the Bitcoin Script. Multiple ways have been proposed to emulate convenants before the upgrade. One of the innovations in our work is a novel almost trustless emulation of convents. For more details, see [9].

7.2 Automated slashing via accountable assertions and finality gadgets

Because of the lack of smart contracts on Bitcoin, one cannot just send any safety violation evidences and rely on Bitcoin to process such evidence. Our protocol instead allows the sending of an evidence that can directly lead to slashing: the staker's private key. To ensure the private key of the staker is leaked whenever a safety violation occurs. we combine two ideas: (a) accountable assertions, from cryptography [32], and (b) finality gadgets, from blockchain consensus [18, 27, 28].

Extractable one-time signatures are ones in which when a signer signs two messages using the same private key, then the private key is leaked, i.e. can be extracted from the signatures of the two messages. EOTS has been proposed as a general method for penalizing equivocation [32], such as double-spending the same bitcoin. However, slashing conditions for consensus protocols are more complex than equivocation on two specific messages. For example, in Casper [18], the slashability module of the PoS Ethereum protocol, there are two sets of slashing conditions (Figure 3). The first set of slashing conditions, that of signing two blocks at the same height, is an equivocation, the second set is more complex and cannot be expressed as equivocations. Similarly, in Tendermint, there are again two sets of slashing conditions. One set of slashing conditions is for signing two blocks in the same round at the same height, but the other set of slashing conditions comes from so-called amnesia attacks [14], which again are not directly expressible as equivocations.

We bypass this issue by not changing the signature scheme for the basic consensus protocol itself, but instead add an extra signing round after the base consensus protocol has finalized a block, signed using extractable one-time signatures. A block is considered truly finalized if it is both finalized by the base protocol *and* receives EOTS signed by more than 2/3 of the stake. One can interpret this extra round of signing as a type of *finality gadget* [27], an EOTS finality gadget. It is shown that if there is a safety violation in this modified protocol, then more than 1/3 of the stake has signed two blocks at the same height using EOTS [9]. This leads to the extraction of the private keys of those stakers. Moreover, the EOTS signature scheme can be implemented by Schnorr signatures, the signature scheme used in Bitcoin. Hence these extracted private keys can be used to spend the slashing transactions.

One very important advantage of this finality-gadget-based solution is its *modular nature*: it can be used on top of all BFT consensus protocols with no change to the base consensus protocols themselves. This enables the technology to be PoS-chain-agnostic.

7.3 Fast unbonding via Bitcoin timestamping

In PoS chains with native staking, the unbonding time is very long (eg. three weeks in the Cosmos Hub), due to the need for social consensus to counteract long-range attacks. Long-range attacks, where the attacker builds an alternate fork after unbonding, is costless. On the other hand, long range attacks on a PoW chain like Bitcoin is very costly, as much energy is expended to build an alternate long fork to replace the canonical longest chain. Since the unbonding request in our Bitcoin staking protocol is submitted to the Bitcoin chain, the cost to remove that transaction from the Bitcoin chain will be immensely costly. This suggests that unbonding in our Bitcoin staking protocol can take place fast without the need of social consensus. However, the issue is that while the stake distribution is maintained on the Bitcoin chain, the voting for blocks is happening on the PoS chains. An attacker can slow down the PoS chain, resulting in an out-of-date staker set used to validate a PoS block. This would mean that an attacker can unstake on the Bitcoin chain but still has the voting power to fork on the PoS chain. Even though the private key of the staker is leaked, it is too late to slash it because the staker has already unstaked on the Bitcoin chain.

To avoid this attack and to enable fast unbonding, the PoS chain should be tightly synchronized with the Bitcoin chain. This can be achieved by a technique called *Bitcoin timestamping*, where the PoS block hashes and the staker set voting for the blocks are recorded on the Bitcoin chain. Interestingly, this Bitcoin timestamping technique is also very useful for safe and fast unbonding in PoS chains with native staking, allowing the use of Bitcoin as an external trust [35]. We are repositioning this technique for enabling fast unbonding for the Bitcoin staking protocol.

8 System architecture

Based on the above primitives, the core infrastructure of the Bitcoin staking protocol is a control plane between Bitcoin and the PoS chains (Figure 4). This control plane is responsible for various key functionalities, including

- provide Bitcoin timestamping service to the PoS chains to enable their synchronization with the Bitcoin network.
- act as a market place, match Bitcoin stakes and the PoS chains, and track the staking and validation information, such as EOTS key registration and refreshment;
- record the finality signatures for the PoS chains;

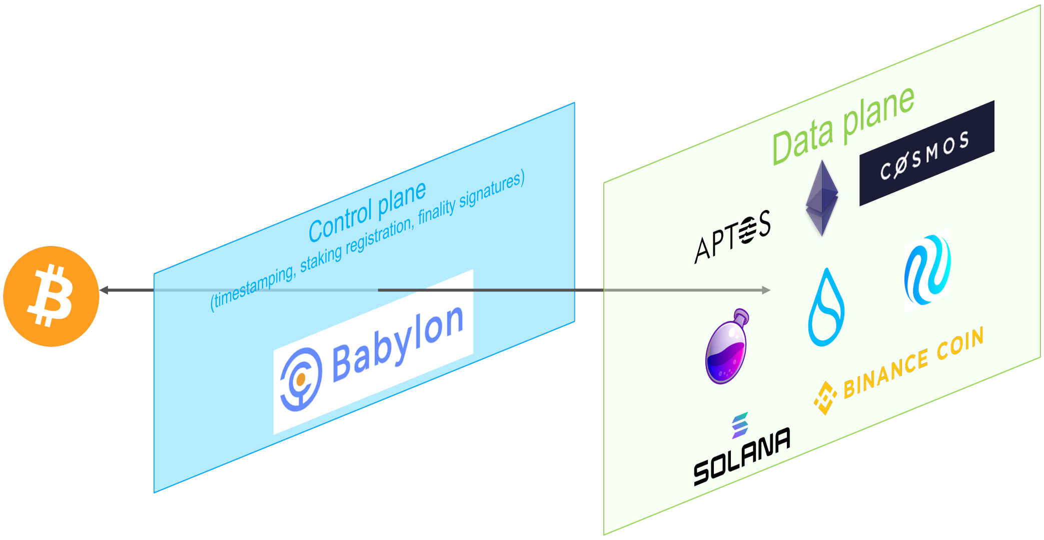

Figure 4: System architecture with control and data planes

On the other hand, validators of each PoS chain, in addition to generate and validate blocks as in the normal consensus protocol, also sign finality signatures on the finality gadget. Together, these validators run the data plane of the architecture.

The control plane is implemented as a chain to make sure that it is decentralized, secure, censorship resistant, and scalable. For example, Bitcoin network's limited and expensive block space makes it unsustainable and unscalable for every PoS chain to directly timestamp on to it, which hinders the adoption of Bitcoin staking. To solve this problem, the Babylon team has designed a secure Bitcoin timestamping protocol and implemented it as a Cosmos-SDK based Babylon chain. This chain enables efficient timestamp aggregation for any number of Cosmos SDK chains via the standard IBC (inter-blockchain communication) protocol. First launched in February 2023, its testnet has integrated with over 30 Cosmos SDK chains across different verticals (Figure 5).

This yields a 3-layer architecture where Babylon chain acts as the control plane, enabling the interaction between Bitcoin and the data plane, i.e., PoS chains. This architecture could also bring network effects and enable interoperability potentials. For example, one may settle cross-PoS-chain transactions on Babylon chain based on the finality status of the two PoS chains on Babylon chain.

9 Related Works

9.1 Cross-chain staking, restaking and mesh security

Every existing PoS chain is secured by the native asset that is maintained in the ledger of the chain. For example, PoS Ethereum is secured by ETH, Cosmos Hub is secured by ATOM, and BNB Chain is secured by BNB. Using only the native token, however, upper-bounds the economic security of the PoS chain by the market capitalization of the native token. Staking of *remote* crypto assets *instead of* or in addition to the native asset on a PoS chain provides an avenue to improve the chain's security by increasing the total market capitalization staked. An approach that is emerging in the blockchain industry can be termed *cross-chain staking*: the staked foreign asset remains in its own chain, but

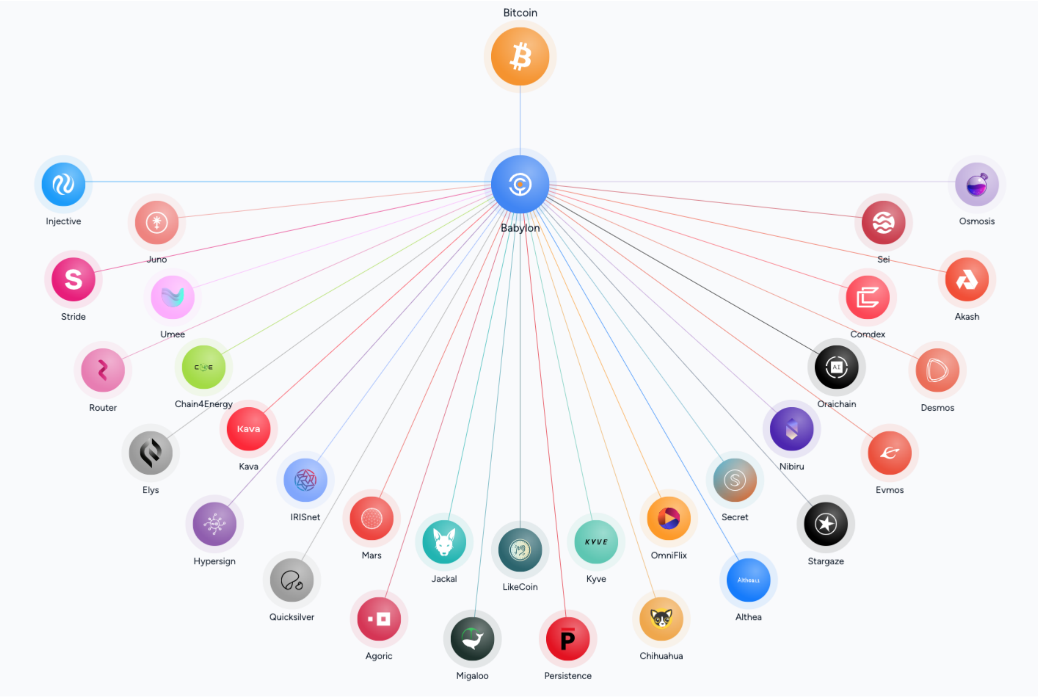

Figure 5: Bitcoin timestamping testnet with 31 IBC-enabled chains.

locked in a staking contract designated for a preferred validator of the chain it is securing. The staked asset is slashed only if the validator commits slashable offenses. This was the idea behind the concept of *mesh security* proposed for the Cosmos ecosystem [11, 4]. The asset of one Cosmos appchain (the provider chain) can be *cross-staked* to help secure another Cosmos appchain (the consumer chain). This cross-staking protocol is in turn inspired by Eigenlayer's Ethereum restaking concept [36]. The idea is to take staked ETH on PoS Ethereum and restake it to secure middleware (so called AVS, Actively Verified Systems) such as data availability layers, bridges, oracle services, etc. With these projects, a generalized form of PoS is emerging where a crypto asset can be used to secure chains and services other than its own chain.

Our Bitcoin staking protocol can be viewed as an example of a cross-chain staking protocol, but there are two important differences with Cosmos mesh security and Ethereum restaking. First, in Ethereum restaking and Cosmos mesh security, the asset is already staked to secure the provider chain. In contrast, Bitcoin is secured by PoW rather than by the Bitcoin asset itself, so the Bitcoin asset is unencumbered. This reduces the over-leveraged risks that arise from restaking [17, 36]. Second, there are no smart contracts on Bitcoin to implement the stake slashing. Instead, we optimize the use of the Bitcoin scripting language and use an advanced cryptographic mechanism to achieve the same goal.

9.2 Accountability and slashability

An important property of many PoS chains is their ability to hold protocol violators accountable in a provable way [18, 20, 33]. This property does not exist in PoW chains because miners do not have on-chain identities. Indeed, the property of *accountable safety*, i.e. the ability to hold 1/3 of the validators accountable in the event of safety violations, is central to the design of PoS Ethereum [18, 20]. However, there is a gap between accountability and on-chain slashing, i.e., actually taking away the stake of the violators on-chain using the protocol violation evidence. In particular, in the case of safety violation, more than 1/3 of the validators are already adversarial and can censor the violation evidence from getting on the chain and effecting the slashing. In such a case, a complex social consensus process has to happen off-chain so that the violators can be slashed and kicked out of the validator set, and the remaining honest validators can restart the chain [35]. In contrast, our Bitcoin staking protocol does not suffer from this issue as the Bitcoin stake resides on the Bitcoin chain, not on the PoS chain, and it is automatically slashed once a safety violation occurs on the PoS chain.

9.3 Accountable assertions and Stakechain

The general concept of using extractable one-time signatures to penalize equivocations originated in [32]. In that work, a deposit has to be time-locked on the Bitcoin chain as a prerequisite for a party in a distributed protocol to make *accountable assertions*. Whenever two different assertions are made in the same context, the private key of the party is leaked, and anyone can use the private key to take the deposit. [24] builds upon this concept to create a Proof-of-Stake sidechain backed by bitcoins. However, the proof-of-stake protocol presented involves only one phase of voting at each height of the blockchain. While this allows safety violations to be straightforwardly modeled as conflicting accountable assertions (with the block height as the context, and a safety violation as an equivocation on two blocks at the same height), the protocol is not live even when the attacker has very small stake. In contrast, known BFT protocol designs all involve *multiple* phases of voting to ensure liveness. In contrast, our work does not attempt to design a PoS protocol from scratch, but instead the Bitcoin staking protocol is used *in conjunction* with any PoS consensus protocol as an add-on finality gadget. This guarantees liveness of the overall protocol as long as the underlying consensus protocol is live, but yet slashability is achieved because safety violations are all equivocations on the same block height when using EOTS to sign on the finality gadget. Moreover, the deposit contracts in our protocol allows on-demand withdrawals of funds (after a certain delay), while the contracts in [32] allows only fixed-term deposits.

9.4 Finality gadgets

Broadly speaking, a finality gadget can be thought of as overlay protocol to be used on top of an existing consensus protocol to provide additional security guarantees. The first finality gadget is the Casper FFG [18], to be used on top of a longest chain protocol to guarantee safety under network partition (a security property that longest chain protocol does not have). Another finality gadget is GRANDPA [34], used in Polkadot. PoS Ethereum's beacon chain consensus protocol, Gasper [19], uses Casper FFG as a finality gadget on top of a LMD (Latest Message Driven) GHOST protocol. However, [27] shows that Gasper is vulnerable to a liveness attack. The first finality gadget construction that is formally proven to be secured is the snap-and-chat protocol [27]. The accountability gadget proposed in [28] allows the additional of accountability properties to a longest chain protocol. The EOTS finality gadget construction in our Bitcoin staking protocol follows a similar philosophy. It adds the property of slashability of Bitcoin stake to an existing BFT consensus protocol.

9.5 Bitcoin merge mining

Merge mining is the first technique, invented in 2010 by Satoshi Nakamoto, to share the security of Bitcoin. It is used to secure the first Bitcoin sidechain, Namecoin. The current largest Bitcoin sidechain supported by merge mining is Rootstock [5]. Using the merge mining technique, Bitcoin miners can simultaneously mine Bitcoin and another PoW chain without using extra energy. However, as a security sharing protocol, merge mining is threatened by a "Nothing-at-Stake" problem: in principle, miners could attack the sidechain while mining the Bitcoin chain honestly. As Bitcoin is the major source of revenue for the miners, there may not be sufficient deterrent against malicious behaviors on the sidechain. In contrast, with the Bitcoin staking, everything is at stake: malicious behavior on the PoS chain is slashable. Thus, Bitcoin staking is a much stronger security sharing technique than merge mining.

9.6 Bitcoin timestamping

Another technique to share the security of Bitcoin is timestamping [35]: hashes and signatures of PoS blocks are submitted as transactions and recorded on the Bitcoin chain. This gives an extra layer of ordering of the PoS blocks which can be used to break ties in case there are forks on the PoS chain. This technique is the basis of the Babylon Bitcoin timestamping testnet. Since Bitcoin takes a long time to confirm transactions, the secure recording of these timestamps on the Bitcoin chain is a slow process. As such , Bitcoin timestamping is effective in providing long-range security, such as against long-range attacks. In contrast, Bitcoin staking adds to the economic security of the PoS chain, which protects it against short-range attacks. Moreover, as discussed earlier, Bitcoin timestamping is also an integral part of the Bitcoin staking protocol, serving the role of synchronizing between the PoS chain and Bitcoin.

9.7 Proof-of-transfer and Stacks

Stacks [7] develops a proof-of-transfer (PoX) consensus mechanism where miners compete against each other to become the next block proposer by sending bitcoins to specific addresses on the Bitcoin chain, the higher the amount sent, the greater the chance. This is a fundamentally different mechanism than proof-of-stake protocols, and thus slashability and stakers' security properties do not apply to Stacks' PoX.

Nevertheless, to bridge bitcoins to Stacks in order to allow Stacks' smart contracts to access the asset, Stacks proposes a way to mint and burn a synthetic Bitcoin token named sBTC secured by STX token stakers, called "Stackers" [6]. The stackers work as a 70% threshold signature group, and are tasked with two responsibilities: 1) sBTC minting and redemption, and 2) approving forking of the Stacks ledger that is already finalized. Therefore, the security of the sBTC bridge is safe if more than 30% of stackers are honest, and live if at least 70% of them honestly sign transactions. A key advantage of our work is that we make Bitcoin staking possible without requiring the bridging of bitcoins, whose security is often capped by the overall locked asset value of the token issued by the bridge project.

Compared to Stacks, our Bitcoin staking protocol does not require spending of bitcoins, but preserves the Bitcoin staked as long as no security violations occur. This enables more efficient and scalable utilization of the asset for security applications.

9.8 Bitcoin bridging

Broadly speaking, there are three categories of Bitcoin bridges today: centralized, collateral-based, and sidechain-based, as well as potential security hardening provided by hardware solutions. We leave out atomic swaps in our discussion as they have not been adopted by prevailing Bitcoin bridges, possibly due to challenges in usability, latency and liquidity sourcing.

Centralized bridges are run by a central party trusted by its users, and a typical example is a centralized exchange that allows users to deposit Bitcoin and its wrapped tokens from other chains, and withdraw to any of such supported chains. For instance, a Binance user could deposit her native Bitcoin into her Binance account and then withdraw a wrapped Bitcoin token on the BNB chain. Another example is wBTC where Bitgo serves as the custodian of the native Bitcoin [37]. These solutions work under the strong assumption that the centralized parties will not do harm, either intentionally, or when they are under attack, they will sufficiently compensate users for their losses.

Interlay is a solution to bring Bitcoin to the Polkadot ecosystem via overcollateralized vaults that provide the peg-in (creating a wrapped Bitcoin upon receiving a native Bitcoin) and peg-out (redeeming the native Bitcoin after destroying the wrapped Bitcoin) of Bitcoin [23]. A key trade-off here is security (i.e., greater collateralization ratio in case vaults go rogue and steal Bitcoin) vs capacity (the amount of Bitcoin bridged is throttled by the amount of collateral and the collateralization ratio). Similarly, sBTC [6] proposed by Stacks in its Nakamoto upgrade [7] tasks the collective group of “stacked" STX token holders to perform the peg-in and peg-out operations between Bitcoin and the sBTC token on the Stacks chain.

Alternatively, Nomic is a Tendermint-based chain that provides a way for Bitcoin to be bridged into nBTC, which can in turn, via the Inter-Blockchain Communication Protocol (IBC) [3], be used on Osmosis and other Cosmos Zones [29]. The limitation of such bridging solutions is that the overall security of bridged tokens depends on the security of the Nomic chain and is loosely capped by the total staked value of the Nomic token. Similarly, Rootstock runs Bitcoin light clients on their miners and rely on the latter for the peg-in and peg-out between native Bitcoin and the synthetic Bitcoin token on their chain [5, 31]. Without taking into account the additional security provided by hardware as elaborated below, the security of their wrapped Bitcoin token is on a par with the security of the proof-of-work chain of Rootstock.

In addition, Rootstock's Bitcoin pegging mechanism named PowerPeg [31] leverages secure hardware to strengthen the security of the pegging of Bitcoin. Similar hardware-based security enhancement is also adopted by Avalanche's Bitcoin bridge using Intel SGX [2]. The usage of hardware root of trust could in principle reduce the attack surface of such bridges, especially when code integrity can be verified at run time. However, practical software security considerations apply here: a) if the bridging logic running inside the secure hardware relies on critical information obtained from external sources, the security of such a bridge is reduced to that of the external component; and b) if a security vulnerability in the code running inside the secure hardware is exploited, the security enhancement provided by the hardware could be rendered ineffective.

As we mentioned previously, a key risk in the prevailing Bitcoin bridges is due to the redeemability of wrapped Bitcoin tokens that are protected by a chain much less secure than Bitcoin. Fortunately, to enable Bitcoin staking to secure external chains and systems, we do not need the full transferability of the locked Bitcoin. Our proposed Bitcoin staking scheme circumvents the security and capacity challenges of prevailing Bitcoin bridges by restricting spending operations on the locked Bitcoin to only slashing safety violations. Consequently, our scheme provides strong safety guarantees as stated in Section 4.

10 Conclusions

Bitcoin is the first and still the top blockchain in terms of market capitalization. However, beyond a store of value, its utility has been limited by its small blockspace, high latency, and limited programmability. In particular, previous efforts in scaling Bitcoin and broadening its use cases by building sidechains and other layer 2 projects are hampered by the inability to bridge a large quantity of bitcoins to these chains. The bridges are limited either by security or by capacity or by both.

Our work brings a new and important use case for the Bitcoin asset: staking to provide security to the PoS world. What we showed is that for at least this use case, bridging of the Bitcoin asset to other chains is not needed, yet full economic security can be provided to the PoS chains. The biggest challenge to achieve this is to be able to remotely slash all safety violations without having a smart contract on the Bitcoin chain. We achieve this by synthesizing four concepts into a single protocol: 1) *accountable assertions* to leak private keys upon equivocations, 2) *finality gadgets* to convert all safety violations into equivocations of accountable assertions, 3) *Bitcoin convenant emulation* to enforce burning of funds upon leaking of private key, and 4) *Bitcoin timestamping* to ensure the slashing transaction can be spent before unstaking. Our construction is modular, and can be used on top of all PoS consensus protocols. No soft or hard fork of Bitcoin is needed to implement our Bitcoin staking protocol.

With new use cases like Ordinals, Bitcoin has recently undergone a renaissance of sorts. We believe that the use case of Bitcoin staking will add further momentum to this renaissance, and will inspire efforts to find other trustless use cases of the enormous Bitcoin asset.

What happens on Bitcoin stays on Bitcoin.

11 Acknowledgements

This litepaper draws from a collaboration between the Babylon team, Robin Linus at ZeroSync and Orfeas Stefanos Thyfronitis Litos at Common Prefix and Imperial College. We met Robin Linus at Bitcoin Miami 2023 where he told us about his work Stakechain. At that time we had already been working intensely on Bitcoin staking, so we joined forces to complete this work.

References

[1] Akash networks. https://akash.network/. Accessed: 2023-07-10.

[2] How does the avalanche bridge™ work? https://support.avax.network/en/articles/6349640-how-does-the-avalanche-bridge-work.

[3] Inter-blockchain communication protocol. https://ibcprotocol.org/.

[4] Mesh security. https://github.com/osmosis-labs/mesh-security.

[5] Rsk. https://www.rsk.co/. Accessed: 2021-11-3.

[6] sbtc: Design of a trustless two-way peg for bitcoin. https://stx.is/sbtc-pdf.

[7] Stacks: A bitcoin layer for smart contracts. https://stx.is/nakamoto.

[8] Bip 119: `OP_CHECKTEMPLATEVERIFY`, 2023. https://github.com/bitcoin/bips/blob/master/bip-0119.mediawiki.

[9] Bitcoin staking. In progress., 2023.

[10] Script, 2023.

[11] Sunny Aggarwal. Mesh security talk at cosmoverse 2022. https://youtu.be/Z2ZBKo9-iRs?t=4937.

[12] Christian Badertscher, Peter Gaži, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros Genesis: Composable proof-of-stake blockchains with dynamic availability. In *Conference on Computer and Communications Security*, CCS '18, pages 913–930. ACM, 2018.

[13] Ethan Buchman. *Tendermint: Byzantine fault tolerance in the age of blockchains.* PhD thesis, University of Guelph, 2016.

[14] Ethan Buchman, Rachid Guerraoui, Jovan Komatovic, Zarko Milosevic, Dragos-Adrian Seredinschi, and Josef Widder. Revisiting tendermint: Design tradeoffs, accountability, and practical use. In *DSN (Supplements)*, pages 11–14. IEEE, 2022.

[15] Ethan Buchman, Jae Kwon, and Zarko Milosevic. The latest gossip on BFT consensus. *arXiv preprint arXiv:1807.04938*, 2018.

[16] Vitalik Buterin. Proof of stake: How i learned to love weak subjectivity, 2014.

[17] Vitalik Buterin. Don't overload Ethereum consensus. Available at: https://vitalik.ca/general/2023/05/21/dont_overload.html, May 2023.

[18] Vitalik Buterin and Virgil Griffith. Casper the friendly finality gadget. *arXiv:1710.09437*, 2019.

[19] Vitalik Buterin, Diego Hernandez, Thor Kamphefner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. Combining ghost and casper. *arXiv:2003.03052*, 2020.

[20] Vitalik Buterin, Diego Hernandez, Thor Kamphefner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. Combining GHOST and Casper. *arXiv preprint arXiv:2003.03052*, 2020.

[21] Phil Daian, Rafael Pass, and Elaine Shi. Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. In *Financial Cryptography and Data Security*, FC '19, pages 23–41. Springer, 2019.

[22] Evangelos Deirmentzoglou, Georgios Papakyriakopoulos, and Constantinos Patsakis. A survey on long-range attacks for proof of stake protocols. *IEEE Access*, 7:28712–28725, 2019.

[23] Interlay Labs. Interlay v2: Bitcoin finance, unbanked, February 2023. https://gateway.pinata.cloud/ipfs/QmWp62gdLssFpAoG2JqK8sy3m3rTRUa8LyzoSY8ZFisYNB.

[24] Robin Linus. Stakechain: A bitcoin-backed proof-of-stake, December 2021. https://coins.github.io/stakechains.pdf.

[25] Gregory Maxwell. Coincovenants using scip signatures, an amusingly bad idea, 2013. https://bitcointalk.org/index.php?topic=278122.0.

[26] Malte Moser, Ittay Eyal, and Emin Gun Sirer. Bitcoin convenants. In *Financial Cryptography*, 2015.

[27] Joachim Neu, Ertem Nusret Tas, and David Tse. Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In *Symposium on Security and Privacy*, S&P '21. IEEE, 2021.

[28] Joachim Neu, Ertem Nusret Tas, and David Tse. The availability-accountability dilemma and its resolution via accountability gadgets. In *Financial Cryptography*, 2022.

[29] Nomic. Nomic bitcoin bridge. https://www.nomic.io/.

[30] Greg Osouri and dam Bozanich. Akt: Akash network token mining economics, January 2020. https://ipfs.io/ipfs/QmdV52bF7j4utynJ6L11RgG93FuJiUmBH1i7pRD6NjUt6B.

[31] Rootstock. Powpeg: Building the most secure, permissionless and uncensorable bitcoin peg. https://dev.rootstock.io/rsk/architecture/powpeg/.

[32] 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*, CCS '15, page 219–230, New York, NY, USA, 2015. Association for Computing Machinery.

[33] Peiyao Sheng, Gerui Wang, Kartik Nayak, Sreeram Kannan, and Pramod Viswanath. BFT protocol forensics. In *CCS*, pages 1722–1743. ACM, 2021.

[34] Alistair Stewart and Eleftherios Kokoris-Kogia. GRANDPA: A Byzantine finality gadget. *arXiv:2007.01560*, 2020.

[35] Ertem Nusret Tas, David Tse, Fisher Yu, Sreeram Kannan, and Mohammad Ali Maddah-Ali. Bitcoin-enhanced proof-of-stake security: Possibilities and impossibilities. In *IEEE Symposium on Security and Privacy*. IEEE, IEEE, 2023.

[36] EigenLayer Team. Eigenlayer: The restaking collective. https://docs.eigenlayer.xyz/overview/whitepaper.

[37] WBTC. Wrapped bitcoin (wbtc). https://wbtc.network/.